

INTERNATIONAL STANDARD



**OPC unified architecture –
Part 4: Services**

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV



THIS PUBLICATION IS COPYRIGHT PROTECTED

Copyright © 2020 IEC, Geneva, Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either IEC or IEC's member National Committee in the country of the requester. If you have any questions about IEC copyright or have an enquiry about obtaining additional rights to this publication, please contact the address below or your local IEC member National Committee for further information.

IEC Central Office
3, rue de Varembe
CH-1211 Geneva 20
Switzerland

Tel.: +41 22 919 02 11
info@iec.ch
www.iec.ch

About the IEC

The International Electrotechnical Commission (IEC) is the leading global organization that prepares and publishes International Standards for all electrical, electronic and related technologies.

About IEC publications

The technical content of IEC publications is kept under constant review by the IEC. Please make sure that you have the latest edition, a corrigendum or an amendment might have been published.

IEC publications search - webstore.iec.ch/advsearchform

The advanced search enables to find IEC publications by a variety of criteria (reference number, text, technical committee,...). It also gives information on projects, replaced and withdrawn publications.

IEC Just Published - webstore.iec.ch/justpublished

Stay up to date on all new IEC publications. Just Published details all new publications released. Available online and once a month by email.

IEC Customer Service Centre - webstore.iec.ch/csc

If you wish to give us your feedback on this publication or need further assistance, please contact the Customer Service Centre: sales@iec.ch.

Electropedia - www.electropedia.org

The world's leading online dictionary on electrotechnology, containing more than 22 000 terminological entries in English and French, with equivalent terms in 16 additional languages. Also known as the International Electrotechnical Vocabulary (IEV) online.

IEC Glossary - std.iec.ch/glossary

67 000 electrotechnical terminology entries in English and French extracted from the Terms and Definitions clause of IEC publications issued since 2002. Some entries have been collected from earlier publications of IEC TC 37, 77, 86 and CISPR.

IECNORM.COM : Click to view the full text of IEC 60364-4:2020 RVV

INTERNATIONAL STANDARD



**OPC unified architecture –
Part 4: Services**

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 25.040.40; 35.100.05

ISBN 978-2-8322-8664-7

Warning! Make sure that you obtained this publication from an authorized distributor.

CONTENTS

FOREWORD	13
1 Scope	16
2 Normative references	16
3 Terms, definitions, abbreviated terms and conventions	17
3.1 Terms and definitions	17
3.2 Abbreviated terms	18
3.3 Conventions for Service definitions	18
4 Overview	20
4.1 Service Set model	20
4.2 Request/response Service procedures	23
5 Service Sets	23
5.1 General	23
5.2 Service request and response header	24
5.3 Service results	24
5.4 Discovery Service Set	25
5.4.1 Overview	25
5.4.2 FindServers	27
5.4.3 FindServersOnNetwork	28
5.4.4 GetEndpoints	30
5.4.5 RegisterServer	32
5.4.6 RegisterServer2	36
5.5 SecureChannel Service Set	38
5.5.1 Overview	38
5.5.2 OpenSecureChannel	39
5.5.3 CloseSecureChannel	43
5.6 Session Service Set	44
5.6.1 Overview	44
5.6.2 CreateSession	44
5.6.3 ActivateSession	49
5.6.4 CloseSession	52
5.6.5 Cancel	53
5.7 NodeManagement Service Set	54
5.7.1 Overview	54
5.7.2 AddNodes	54
5.7.3 AddReferences	56
5.7.4 DeleteNodes	57
5.7.5 DeleteReferences	59
5.8 View Service Set	60
5.8.1 Overview	60
5.8.2 Browse	60
5.8.3 BrowseNext	63
5.8.4 TranslateBrowsePathsToNodeIds	65
5.8.5 RegisterNodes	67
5.8.6 UnregisterNodes	68
5.9 Query Service Set	69
5.9.1 Overview	69

5.9.2	Querying Views	69
5.9.3	QueryFirst	70
5.9.4	QueryNext	73
5.10	Attribute Service Set	74
5.10.1	Overview	74
5.10.2	Read	75
5.10.3	HistoryRead	76
5.10.4	Write	79
5.10.5	HistoryUpdate	82
5.11	Method Service Set	84
5.11.1	Overview	84
5.11.2	Call	84
5.12	MonitoredItem Service Set	87
5.12.1	MonitoredItem model	87
5.12.2	CreateMonitoredItems	93
5.12.3	ModifyMonitoredItems	95
5.12.4	SetMonitoringMode	97
5.12.5	SetTriggering	98
5.12.6	DeleteMonitoredItems	100
5.13	Subscription Service Set	101
5.13.1	Subscription model	101
5.13.2	CreateSubscription	110
5.13.3	ModifySubscription	111
5.13.4	SetPublishingMode	113
5.13.5	Publish	114
5.13.6	Republish	116
5.13.7	TransferSubscriptions	117
5.13.8	DeleteSubscriptions	119
6	Service behaviours	120
6.1	Security	120
6.1.1	Overview	120
6.1.2	Obtaining and installing an Application Instance Certificate	120
6.1.3	Determining if a Certificate is trusted	121
6.1.4	Creating a SecureChannel	124
6.1.5	Creating a Session	126
6.1.6	Impersonating a User	127
6.2	Authorization Services	127
6.2.1	Overview	127
6.2.2	Indirect handshake with an Identity Provider	128
6.2.3	Direct handshake with an Identity Provider	128
6.3	Session-less Service invocation	129
6.3.1	Description	129
6.3.2	Parameters	130
6.3.3	Service results	131
6.4	Software Certificates	131
6.5	Auditing	133
6.5.1	Overview	133
6.5.2	General audit logs	133
6.5.3	General audit Events	133

6.5.4	Auditing for Discovery Service Set	134
6.5.5	Auditing for SecureChannel Service Set	134
6.5.6	Auditing for Session Service Set	134
6.5.7	Auditing for NodeManagement Service Set	135
6.5.8	Auditing for Attribute Service Set	135
6.5.9	Auditing for Method Service Set	135
6.5.10	Auditing for View, Query, MonitoredItem and Subscription Service Set	136
6.6	Redundancy	136
6.6.1	Redundancy overview	141
6.6.2	Server Redundancy	141
6.6.3	Client Redundancy	153
6.6.4	Network Redundancy	153
6.6.5	Manually forcing Failover	155
6.7	Re-establishing connections	155
6.8	Durable Subscriptions	157
7	Common parameter type definitions	158
7.1	ApplicationDescription	158
7.2	ApplicationInstanceCertificate	159
7.3	BrowseResult	160
7.4	ContentFilter	161
7.4.1	ContentFilter structure	161
7.4.2	ContentFilterResult	161
7.4.3	FilterOperator	162
7.4.4	FilterOperand parameters	169
7.5	Counter	171
7.6	ContinuationPoint	171
7.7	DataValue	172
7.7.1	General	172
7.7.2	PicoSeconds	172
7.7.3	SourceTimestamp	173
7.7.4	ServerTimestamp	173
7.7.5	StatusCodes assigned to a value	174
7.8	DiagnosticInfo	174
7.9	DiscoveryConfiguration parameters	175
7.9.1	Overview	175
7.9.2	MdnsDiscoveryConfiguration	176
7.10	EndpointDescription	176
7.11	ExpandedNodeId	177
7.12	ExtensibleParameter	177
7.13	Index	177
7.14	IntegerId	178
7.15	MessageSecurityMode	178
7.16	MonitoringParameters	178
7.17	MonitoringFilter parameters	179
7.17.1	Overview	179
7.17.2	DataChangeFilter	180
7.17.3	EventFilter	181
7.17.4	AggregateFilter	183
7.18	MonitoringMode	184

7.19	NodeAttributes parameters	185
7.19.1	Overview	185
7.19.2	ObjectAttributes parameter	186
7.19.3	VariableAttributes parameter	186
7.19.4	MethodAttributes parameter	187
7.19.5	ObjectTypeAttributes parameter	187
7.19.6	VariableTypeAttributes parameter	188
7.19.7	ReferenceTypeAttributes parameter	188
7.19.8	DataTypeAttributes parameter	189
7.19.9	ViewAttributes parameter	189
7.19.10	GenericAttributes parameter	190
7.20	NotificationData parameters	190
7.20.1	Overview	190
7.20.2	DataChangeNotification parameter	191
7.20.3	EventNotificationList parameter	191
7.20.4	StatusChangeNotification parameter	192
7.21	NotificationMessage	192
7.22	NumericRange	192
7.23	QueryDataSet	193
7.24	ReadValueId	194
7.25	ReferenceDescription	195
7.26	RelativePath	196
7.27	RegisteredServer	197
7.28	RequestHeader	197
7.29	ResponseHeader	199
7.30	ServiceFault	199
7.31	SessionAuthenticationToken	200
7.32	SignatureData	201
7.33	SignedSoftwareCertificate	202
7.34	StatusCode	203
7.34.1	General	203
7.34.2	Common StatusCodes	205
7.35	TimestampsToReturn	209
7.36	UserIdentityToken parameters	209
7.36.1	Overview	209
7.36.2	Token Encryption and Proof of Possession	210
7.36.3	AnonymousIdentityToken	214
7.36.4	UserNameIdentityToken	214
7.36.5	X509IdentityTokens	216
7.36.6	IssuedIdentityToken	216
7.37	UserTokenPolicy	217
7.38	VersionTime	218
7.39	ViewDescription	218
Annex A	(informative) BNF definitions	220
A.1	Overview over BNF	220
A.2	BNF of RelativePath	220
A.3	BNF of NumericRange	221
Annex B	(informative) ContentFilter and Query examples	222
B.1	Simple ContentFilter examples	222

B.1.1	Overview	222
B.1.2	Example 1	222
B.1.3	Example 2	223
B.2	Complex examples of Query filters	224
B.2.1	Overview	224
B.2.2	Used type model	224
B.2.3	Example Notes	227
B.2.4	Example 1	228
B.2.5	Example 2	229
B.2.6	Example 3	230
B.2.7	Example 4	233
B.2.8	Example 5	234
B.2.9	Example 6	235
B.2.10	Example 7	237
B.2.11	Example 8	239
B.2.12	Example 9	241
Figure 1	– Discovery Service Set	20
Figure 2	– SecureChannel Service Set	20
Figure 3	– Session Service Set	21
Figure 4	– NodeManagement Service Set	21
Figure 5	– View Service Set	21
Figure 6	– Attribute Service Set	22
Figure 7	– Method Service Set	22
Figure 8	– MonitoredItem and Subscription Service Sets	23
Figure 9	– Discovery process	26
Figure 10	– Using a Gateway Server	31
Figure 11	– The registration process – Manually launched servers	33
Figure 12	– The registration process – Automatically launched servers	33
Figure 13	– SecureChannel and Session Services	39
Figure 14	– Multiplexing users on a Session	46
Figure 15	– MonitoredItem model	88
Figure 16	– Typical delay in change detection	90
Figure 17	– Queue overflow handling	91
Figure 18	– Triggering model	92
Figure 19	– Obtaining and installing an Application Instance Certificate	121
Figure 20	– Determining if an Application Instance Certificate is trusted	124
Figure 21	– Establishing a SecureChannel	125
Figure 22	– Establishing a Session	126
Figure 23	– Impersonating a User	127
Figure 24	– Indirect handshake with an Identity Provider	128
Figure 25	– Direct handshake with an Identity Provider	129
Figure 26	– Transparent Redundancy setup example	143
Figure 27	– Non-Transparent Redundancy setup	144
Figure 28	– Client Start-up steps	148

Figure 29 – Cold Failover.....	149
Figure 30 – Warm Failover.....	150
Figure 31 – Hot Failover	151
Figure 32 – HotAndMirrored Failover	152
Figure 33 – Server proxy for Redundancy	153
Figure 34 – Transparent network Redundancy	154
Figure 35 – Non-transparent network Redundancy.....	155
Figure 36 – Reconnect sequence.....	156
Figure 37 – Logical layers of a Server.....	200
Figure 38 – Obtaining a SessionAuthenticationToken	201
Figure 39 – EncryptedSecret layout	211
Figure B.1 – Filter logic tree example.....	222
Figure B.2 – Filter logic tree example.....	223
Figure B.3 – Example Type Nodes	226
Figure B.4 – Example Instance Nodes	227
Figure B.5 – Example 1 Filter.....	228
Figure B.6 – Example 2 Filter logic tree	230
Figure B.7 – Example 3 Filter logic tree	231
Figure B.8 – Example 4 Filter logic tree	233
Figure B.9 – Example 5 Filter logic tree	235
Figure B.10 – Example 6 Filter logic tree	236
Figure B.11 – Example 7 Filter logic tree	238
Figure B.12 – Example 8 Filter logic tree	240
Figure B.13 – Example 9 Filter logic tree	241
Table 1 – Service definition table.....	19
Table 2 – Parameter Types defined in IEC 62541-3	19
Table 3 – FindServers Service parameters	28
Table 4 – FindServersOnNetwork Service parameters	29
Table 5 – GetEndpoints Service parameters	32
Table 6 – RegisterServer Service parameters	35
Table 7 – RegisterServer Service result codes.....	36
Table 8 – RegisterServer2	37
Table 9 – RegisterServer2 Service result codes	37
Table 10 – RegisterServer2 Operation Level result codes	37
Table 11 – OpenSecureChannel Service parameters	41
Table 12 – OpenSecureChannel Service result codes	43
Table 13 – CloseSecureChannel Service parameters.....	44
Table 14 – CloseSecureChannel Service result codes	44
Table 15 – CreateSession Service parameters.....	47
Table 16 – CreateSession Service result codes	49
Table 17 – ActivateSession Service parameters.....	51
Table 18 – ActivateSession Service result codes	52

Table 19 – CloseSession Service parameters	53
Table 20 – CloseSession Service result codes	53
Table 21 – Cancel Service parameters	53
Table 22 – AddNodes Service parameters	54
Table 23 – AddNodes Service result codes	55
Table 24 – AddNodes Operation Level result codes	55
Table 25 – AddReferences Service parameters	56
Table 26 – AddReferences Service result codes	57
Table 27 – AddReferences Operation Level result codes	57
Table 28 – DeleteNodes Service parameters	58
Table 29 – DeleteNodes Service result codes	58
Table 30 – DeleteNodes Operation Level result codes	59
Table 31 – DeleteReferences Service parameters	59
Table 32 – DeleteReferences Service result codes	60
Table 33 – DeleteReferences Operation Level result codes	60
Table 34 – Browse Service parameters	61
Table 35 – Browse Service result codes	62
Table 36 – Browse Operation Level result codes	63
Table 37 – BrowseNext Service parameters	64
Table 38 – BrowseNext Service result codes	64
Table 39 – BrowseNext Operation Level result codes	65
Table 40 – TranslateBrowsePathsToNodeIds Service parameters	66
Table 41 – TranslateBrowsePathsToNodeIds Service result codes	66
Table 42 – TranslateBrowsePathsToNodeIds Operation Level result codes	67
Table 43 – RegisterNodes Service parameters	68
Table 44 – RegisterNodes Service result codes	68
Table 45 – UnregisterNodes Service parameters	69
Table 46 – UnregisterNodes Service result codes	69
Table 47 – QueryFirst Request parameters	71
Table 48 – QueryFirst Response parameters	72
Table 49 – QueryFirst Service result codes	73
Table 50 – QueryFirst Operation Level result codes	73
Table 51 – QueryNext Service parameters	74
Table 52 – QueryNext Service result codes	74
Table 53 – Read Service parameters	75
Table 54 – Read Service result codes	76
Table 55 – Read Operation Level result codes	76
Table 56 – HistoryRead Service parameters	77
Table 57 – HistoryRead Service result codes	79
Table 58 – HistoryRead Operation Level result codes	79
Table 59 – Write Service parameters	81
Table 60 – Write Service result codes	82
Table 61 – Write Operation Level result codes	82

Table 62 – HistoryUpdate Service parameters	83
Table 63 – HistoryUpdate Service result codes	83
Table 64 – HistoryUpdate Operation Level result codes	84
Table 65 – Call Service parameters	85
Table 66 – Call Service result codes	86
Table 67 – Call Operation Level result codes	87
Table 68 – Call Input Argument Result Codes	87
Table 69 – CreateMonitoredItems Service parameters	94
Table 70 – CreateMonitoredItems Service result codes	95
Table 71 – CreateMonitoredItems Operation Level result codes	95
Table 72 – ModifyMonitoredItems Service parameters	96
Table 73 – ModifyMonitoredItems Service result codes	97
Table 74 – ModifyMonitoredItems Operation Level result codes	97
Table 75 – SetMonitoringMode service parameters	98
Table 76 – SetMonitoringMode Service result codes	98
Table 77 – SetMonitoringMode Operation Level result codes	98
Table 78 – SetTriggering Service parameters	99
Table 79 – SetTriggering Service result codes	99
Table 80 – SetTriggering Operation Level result codes	100
Table 81 – DeleteMonitoredItems Service parameters	100
Table 82 – DeleteMonitoredItems Service result codes	101
Table 83 – DeleteMonitoredItems Operation Level result codes	101
Table 84 – Subscription states	104
Table 85 – Subscription state table	105
Table 86 – State variables and parameters	108
Table 87 – Functions	109
Table 88 – CreateSubscription Service parameters	110
Table 89 – CreateSubscription Service result codes	111
Table 90 – ModifySubscription Service parameters	112
Table 91 – ModifySubscription Service result codes	113
Table 92 – SetPublishingMode Service parameters	113
Table 93 – SetPublishingMode Service result codes	113
Table 94 – SetPublishingMode Operation Level result codes	114
Table 95 – Publish Service parameters	115
Table 96 – Publish Service result codes	115
Table 97 – Publish Operation Level Result Codes	116
Table 98 – Republish Service parameters	116
Table 99 – Republish Service result codes	116
Table 100 – TransferSubscriptions Service parameters	118
Table 101 – TransferSubscriptions Service result codes	118
Table 102 – TransferSubscriptions Operation Level result codes	119
Table 103 – DeleteSubscriptions Service parameters	119
Table 104 – DeleteSubscriptions Service result codes	120

Table 105 – DeleteSubscriptions Operation Level result codes	120
Table 106 – Certificate validation steps	122
Table 107 – SessionlessInvoke Service parameters.....	130
Table 108 – SessionlessInvoke Service result codes	131
Table 109 – ServiceLevel ranges	146
Table 110 – Server Failover modes.....	147
Table 111 – Redundancy Failover actions.....	148
Table 112 – ApplicationDescription	159
Table 113 – ApplicationInstanceCertificate	160
Table 114 – BrowseResult	160
Table 115 – ContentFilter structure.....	161
Table 116 – ContentFilterResult structure	161
Table 117 – ContentFilterResult result codes.....	162
Table 118 – ContentFilterResult Operand result codes	162
Table 119 – Basic FilterOperator definition	162
Table 120 – Complex FilterOperator definition	165
Table 121 – Wildcard characters.....	166
Table 122 – Conversion rules	167
Table 123 – Data Precedence rules	168
Table 124 – Logical AND Truth table	169
Table 125 – Logical OR Truth table.....	169
Table 126 – FilterOperand parameter Typelds	169
Table 127 – ElementOperand	170
Table 128 – LiteralOperand	170
Table 129 – AttributeOperand	170
Table 130 – SimpleAttributeOperand	171
Table 131 – DataValue	172
Table 132 – DiagnosticInfo	175
Table 133 – DiscoveryConfiguration parameterTypelds	175
Table 134 – MdnsDiscoveryConfiguration	176
Table 135 – EndpointDescription	176
Table 136 – ExpandedNodeId	177
Table 137 – ExtensibleParameter base type	177
Table 138 – MessageSecurityMode values	178
Table 139 – MonitoringParameters	178
Table 140 – MonitoringFilter parameterTypelds	179
Table 141 – DataChangeFilter	180
Table 142 – EventFilter structure	182
Table 143 – EventFilterResult structure	182
Table 144 – EventFilterResult result codes	183
Table 145 – AggregateFilter structure	184
Table 146 – AggregateFilterResult structure	184
Table 147 – MonitoringMode values.....	185

Table 148 – NodeAttributes parameterTypelds	185
Table 149 – Bit mask for specified Attributes	186
Table 150 – ObjectAttributes	186
Table 151 – VariableAttributes	187
Table 152 – MethodAttributes	187
Table 153 – ObjectTypeAttributes	188
Table 154 – VariableTypeAttributes	188
Table 155 – ReferenceTypeAttributes	189
Table 156 – DataTypeAttributes	189
Table 157 – ViewAttributes	190
Table 158 – GenericAttributes	190
Table 159 – NotificationData parameterTypelds	191
Table 160 – DataChangeNotification	191
Table 161 – EventNotificationList	192
Table 162 – StatusChangeNotification	192
Table 163 – NotificationMessage	192
Table 164 – NumericRange	193
Table 165 – QueryDataSet	194
Table 166 – ReadValueId	195
Table 167 – ReferenceDescription	196
Table 168 – RelativePath	196
Table 169 – RegisteredServer	197
Table 170 – RequestHeader	198
Table 171 – ResponseHeader	199
Table 172 – ServiceFault	200
Table 173 – SignatureData	201
Table 174 – SignedSoftwareCertificate	203
Table 175 – StatusCode bit assignments	204
Table 176 – DataValue InfoBits	205
Table 177 – Common Service result codes	206
Table 178 – Common Operation Level result codes	208
Table 179 – TimestampsToReturn values	209
Table 180 – UserIdentityToken parameterTypelds	210
Table 181 – Legacy UserIdentityToken Encrypted Token Secret Format	211
Table 182 – EncryptedSecret layout	213
Table 183 – EncryptedSecret DataTypes	213
Table 184 – RsaEncryptedSecret structure	214
Table 185 – AnonymousIdentityToken	214
Table 186 – UserNameIdentityToken	215
Table 187 – EncryptionAlgorithm selection	216
Table 188 – X.509 v3 Identity Token	216
Table 189 – IssuedIdentityToken	217
Table 190 – UserTokenPolicy	218

Table 191 – ViewDescription.....	219
Table A.1 – RelativePath	220
Table A.2 – <i>RelativePath</i> Examples	221
Table B.1 – ContentFilter example	223
Table B.2 – ContentFilter example	223
Table B.3 – Example 1 NodeTypeDescription	228
Table B.4 – Example 1 ContentFilter.....	228
Table B.5 – Example 1 QueryDataSets	229
Table B.6 – Example 2 NodeTypeDescription	229
Table B.7 – Example 2 ContentFilter.....	230
Table B.8 – Example 2 QueryDataSets	230
Table B.9 – Example 3 NodeTypeDescription	231
Table B.10 – Example 3 ContentFilter.....	232
Table B.11 – Example 3 QueryDataSets	233
Table B.12 – Example 4 NodeTypeDescription.....	233
Table B.13 – Example 4 ContentFilter	234
Table B.14 – Example 4 QueryDataSets	234
Table B.15 – Example 5 NodeTypeDescription.....	234
Table B.16 – Example 5 ContentFilter.....	235
Table B.17 – Example 5 QueryDataSets	235
Table B.18 – Example 6 NodeTypeDescription.....	236
Table B.19 – Example 6 ContentFilter.....	236
Table B.20 – Example 6 QueryDataSets.....	237
Table B.21 – Example 6 QueryDataSets without additional information.....	237
Table B.22 – Example 7 NodeTypeDescription.....	238
Table B.23 – Example 7 ContentFilter.....	239
Table B.24 – Example 7 QueryDataSets	239
Table B.25 – Example 8 NodeTypeDescription.....	240
Table B.26 – Example 8 ContentFilter.....	240
Table B.27 – Example 8 QueryDataSets	240
Table B.28 – Example 9 NodeTypeDescription.....	241
Table B.29 – Example 9 ContentFilter.....	242
Table B.30 – Example 9 QueryDataSets	242

INTERNATIONAL ELECTROTECHNICAL COMMISSION

OPC UNIFIED ARCHITECTURE –

Part 4: Services

FOREWORD

- 1) The International Electrotechnical Commission (IEC) is a worldwide organization for standardization comprising all national electrotechnical committees (IEC National Committees). The object of IEC is to promote international co-operation on all questions concerning standardization in the electrical and electronic fields. To this end and in addition to other activities, IEC publishes International Standards, Technical Specifications, Technical Reports, Publicly Available Specifications (PAS) and Guides (hereafter referred to as "IEC Publication(s)"). Their preparation is entrusted to technical committees; any IEC National Committee interested in the subject dealt with may participate in this preparatory work. International, governmental and non-governmental organizations liaising with the IEC also participate in this preparation. IEC collaborates closely with the International Organization for Standardization (ISO) in accordance with conditions determined by agreement between the two organizations.
- 2) The formal decisions or agreements of IEC on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC National Committees.
- 3) IEC Publications have the form of recommendations for international use and are accepted by IEC National Committees in that sense. While all reasonable efforts are made to ensure that the technical content of IEC Publications is accurate, IEC cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 4) In order to promote international uniformity, IEC National Committees undertake to apply IEC Publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any IEC Publication and the corresponding national or regional publication shall be clearly indicated in the latter.
- 5) IEC itself does not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. IEC is not responsible for any services carried out by independent certification bodies.
- 6) All users should ensure that they have the latest edition of this publication.
- 7) No liability shall attach to IEC or its directors, employees, servants or agents including individual experts and members of its technical committees and IEC National Committees for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication, use of, or reliance upon, this IEC Publication or any other IEC Publications.
- 8) Attention is drawn to the Normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.
- 9) Attention is drawn to the possibility that some of the elements of this IEC Publication may be the subject of patent rights. IEC shall not be held responsible for identifying any or all such patent rights.

This redline version of the official IEC Standard allows the user to identify the changes made to the previous edition. A vertical bar appears in the margin wherever a change has been made. Additions are in green text, deletions are in strikethrough red text.

International Standard IEC 62541-4 has been prepared by subcommittee 65E: Devices and integration in enterprise systems, of IEC technical committee 65: Industrial-process measurement, control and automation.

This third edition cancels and replaces the second edition published in 2015. This edition constitutes a technical revision.

This edition includes the following significant technical changes with respect to the previous edition:

- a) Added ability to resend all data of monitored items in a Subscription using the ResendData Method.
- b) Added support for durable Subscriptions (lifetime of hours or days).
- c) Added Register2 and FindServersOnNetwork Services to support network-wide discovery using capability filters.
- d) Removed definition of software certificates. Will be defined in a future edition.
- e) Extended and partially revised the redundancy definition. Added sub-range definitions for ServiceLevel and added more terms for redundancy.
- f) Added a section on how to use Authorization Services to request user access tokens.
- g) Added JSON Web Tokens (JWTs) as a new user token.
- h) Added the concept of session-less service invocation.
- i) Added a generic structure that allows passing any number of attributes to the AddNodes Service.
- j) Added requirement to protect against user identity token attacks.
- k) Added new EncryptedSecret format for user identity tokens.

The text of this standard is based on the following documents:

FDIS	Report on voting
65E/716/FDIS	65E/732/RVD

Full information on the voting for the approval of this International Standard can be found in the report on voting indicated in the above table.

This document has been drafted in accordance with the ISO/IEC Directives, Part 2.

Throughout this document and the other parts of the IEC 62541 series, certain document conventions are used:

Italics are used to denote a defined term or definition that appears in Clause 3 in one of the parts of the series.

Italics are also used to denote the name of a service input or output parameter or the name of a structure or element of a structure that are usually defined in tables.

The *italicized terms and names* are also, with a few exceptions, written in camel-case (the practice of writing compound words or phrases in which the elements are joined without spaces, with each element's initial letter capitalized within the compound). For example the defined term is *AddressSpace* instead of Address Space. This makes it easier to understand that there is a single definition for *AddressSpace*, not separate definitions for Address and Space.

A list of all parts of the IEC 62541 series, published under the general title *OPC Unified Architecture*, can be found on the IEC website.

The committee has decided that the contents of this document will remain unchanged until the stability date indicated on the IEC website under "<http://webstore.iec.ch>" in the data related to the specific document. At this date, the document will be

- reconfirmed,
- withdrawn,
- replaced by a revised edition, or
- amended.

IMPORTANT – The 'colour inside' logo on the cover page of this publication indicates that it contains colours which are considered to be useful for the correct understanding of its contents. Users should therefore print this document using a colour printer.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

OPC UNIFIED ARCHITECTURE –

Part 4: Services

1 Scope

This part of IEC 62541 defines the OPC Unified Architecture (OPC UA) *Services*. The *Services* ~~described~~ *defined* are the collection of abstract Remote Procedure Calls (RPC) that are implemented by OPC UA *Servers* and called by OPC UA *Clients*. All interactions between OPC UA *Clients* and *Servers* occur via these *Services*. The defined *Services* are considered abstract because no particular RPC mechanism for implementation is defined in this document. IEC 62541-6 specifies one or more concrete mappings supported for implementation. For example, one mapping in IEC 62541-6 is to XML Web Services. In that case the *Services* described in this document appear as the Web service methods in the WSDL contract.

Not all OPC UA *Servers* will need to implement all of the defined *Services*. IEC 62541-7 defines the *Profiles* that dictate which *Services* need to be implemented in order to be compliant with a particular *Profile*.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

IEC TR 62541-1, *OPC Unified Architecture – Part 1: Overview and Concepts*

IEC TR 62541-2, *OPC Unified Architecture – Part 2: Security Model*

IEC 62541-3, *OPC Unified Architecture – Part 3: Address Space Model*

IEC 62541-5, *OPC Unified Architecture – Part 5: Information Model*

IEC 62541-6, *OPC Unified Architecture – Part 6: Mappings*

IEC 62541-7, *OPC Unified Architecture – Part 7: Profiles*

IEC 62541-8, *OPC Unified Architecture – Part 8: Data Access*

IEC 62541-11, *OPC Unified Architecture – Part 11: Historical Access*

IEC 62541-12¹, *OPC Unified Architecture – Part 12: Discovery and Global Services*

IEC 62541-13, *OPC Unified Architecture – Part 13: Aggregates*

¹ Under preparation. Stage at the time of publication: IEC CDV 62541-12:2018.

3 Terms, definitions, abbreviated terms and conventions

3.1 Terms and definitions

For the purposes of this document, the terms and definitions given in IEC TR 62541-1, IEC TR 62541-2, IEC 62541-3 and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1.1

Active Server

Server which is currently sourcing information

Note 1 to entry: In OPC UA redundant systems, an *Active Server* is the *Server* that a *Client* is using as the source of data.

3.1.2

Deadband

permitted range for value changes that will not trigger a data change *Notification*

Note 1 to entry: *Deadband* can be applied as a filter when subscribing to *Variables* and is used to keep noisy signals from updating the *Client* unnecessarily. This document defines *AbsoluteDeadband* as a common filter. IEC 62541-8 defines an additional *Deadband* filter.

3.1.3

DiscoveryEndpoint

Endpoint that allows *Clients* access to *Discovery Services* without security

Note 1 to entry: A *DiscoveryEndpoint* allows access to *Discovery Services* without a *Session* and without message security.

3.1.4

Endpoint

physical address available on a network that allows *Clients* to access one or more *Services* provided by a *Server*

Note 1 to entry: Each *Server* may have multiple *Endpoints*. ~~The address of an *Endpoint* shall include~~ Each *Endpoint* includes a *HostName*.

3.1.5

Failed Server

Server that is not operational.

Note 1 to entry: In OPC UA redundant systems, a *Failed Server* is a *Server* that is unavailable or is not able to serve data.

3.1.6

Failover

act of switching the source or target of information

Note 1 to entry: In OPC UA redundant systems, a *Failover* is the act of a *Client* switching away from a failed or degraded *Server* to another *Server* in the redundant set (*Server Failover*). In some cases a *Client* may have no knowledge of a *Failover* action occurring (transparent redundancy). A *Client Failover* is the act of an alternate *Client* replacing an existing failed or degraded *Client* connection to a *Server*.

3.1.7

Gateway Server

Server that acts as an intermediary for one or more *Servers*

Note 1 to entry: *Gateway Servers* may be deployed to limit external access, provide protocol conversion or to provide features that the underlying *Servers* do not support.

3.1.8

HostName

unique identifier for a machine on a network

Note 1 to entry: This identifier ~~shall be~~ is unique within a local network; however, it may also be globally unique. The identifier can be an IP address.

3.1.9

Redundancy

presence of duplicate components enabling the continued operation after a failure of an OPC UA component

Note 1 to entry: This may apply to *Servers*, *Clients* or networks.

3.1.10

Redundant server set

two or more *Servers* that are redundant with each other

Note 1 to entry: A *Redundant Server Set* is a group of *Servers* that are configured to provide *Redundancy*. These *Servers* have requirements related to the address space and provide *Failovers*.

3.1.11

SecurityToken

identifier for a cryptographic key set

Note 1 to entry: ~~All Security Tokens belong to a security context which is in case of OPC UA the Secure Channel.~~ All *SecurityTokens* belong to a security context. For OPC UA the security context is the *SecureChannel*.

3.1.12

SoftwareCertificate

digital certificate for a software product that can be installed on several hosts to describe the capabilities of the software product

Note 1 to entry: Different installations of one software product could have the same software certificate. Software certificates are not relevant for security. They are used to identify a software product and its supported features. *SoftwareCertificates* are described in 6.4.

3.2 Abbreviated terms

API	Application Programming Interface
BNF	Backus-Naur Form
CA	Certificate Authority
CRL	Certificate Revocation List
CTL	Certificate Trust List
DA	Data Access
NAT	Network Address Translation
UA	Unified Architecture
URI	Uniform Resource Identifier
URL	Uniform Resource Locator

3.3 Conventions for Service definitions

OPC UA *Services* contain parameters that are conveyed between the *Client* and the *Server*. The OPC UA *Service* specifications use tables to describe *Service* parameters, as shown in Table 1. Parameters are organized in this table into request parameters and response parameters.

Table 1 – Service definition table

Name	Type	Description
Request		Defines the request parameters of the <i>Service</i>
Simple Parameter Name		Description of this parameter
Constructed Parameter Name		Description of the constructed parameter
Component Parameter Name		Description of the component parameter
Response		Defines the response parameters of the <i>Service</i>

The Name, Type and Description columns contain the name, data type and description of each parameter. All parameters are mandatory, although some may be unused under certain circumstances. The Description column specifies the value to be supplied when a parameter is unused.

Two types of parameters are defined in these tables, simple and constructed. Simple parameters have a simple data type, such as *Boolean* or *String*.

Constructed parameters are composed of two or more component parameters, which can be simple or constructed. Component parameter names are indented below the constructed parameter name.

The data types used in these tables may be base types, common types to multiple *Services* or *Service*-specific types. Base data types are defined in IEC 62541-3. The base types used in *Services* are listed in Table 2. Data types that are common to multiple *Services* are defined in Clause 7. Data types that are *Service*-specific are defined in the parameter table of the *Service*.

Table 2 – Parameter Types defined in IEC 62541-3

Parameter Type
BaseDataType
Boolean
ByteString
Double
Duration
Guid
Int32
LocaleId
NodeId
QualifiedName
String
UInt16
UInt32
UInteger
UtcTime
XmlElement

The parameters of the Request and Indication service primitives are represented in Table 1 as Request parameters. Likewise, the parameters of the Response and Confirmation service primitives are represented in Table 1 as Response parameters. All request and response parameters are conveyed between the sender and receiver without change. Therefore, separate columns for request, indication, response and confirmation parameter values are not needed and have been intentionally omitted to improve readability.

4 Overview

4.1 Service Set model

This clause specifies the OPC UA *Services*. The OPC UA *Service* definitions are abstract descriptions and do not represent a specification for implementation. The mapping between the abstract descriptions and the *Communication Stack* derived from these *Services* are defined in IEC 62541-6. In the case of an implementation as web services, the OPC UA *Services* correspond to the web service and an OPC UA *Service* corresponds to an operation of the web service.

These *Services* are organized into *Service Sets*. Each *Service Set* defines a set of related *Services*. The organization in *Service Sets* is a logical grouping used in this document and is not used in the implementation.

The *Discovery Service Set*, illustrated in Figure 1, defines *Services* that allow a *Client* to discover the *Endpoints* implemented by a *Server* and to read the security configuration for each of those *Endpoints*.




Figure 1 – Discovery Service Set

The *SecureChannel Service Set*, illustrated in Figure 2, defines *Services* that allow a *Client* to establish a communication channel to ensure the *Confidentiality* and *Integrity* of *Messages* exchanged with the *Server*.




Figure 2 – SecureChannel Service Set

The *Session Service Set*, illustrated in Figure 3, defines *Services* that allow the *Client* to authenticate the user on whose behalf it is acting and to manage *Sessions*.




Figure 3 – Session Service Set

The *NodeManagement Service Set*, illustrated in Figure 4, defines *Services* that allow the *Client* to add, modify and delete *Nodes* in the *AddressSpace*.




Figure 4 – NodeManagement Service Set

The *View Service Set*, illustrated in Figure 5, defines *Services* that allow *Clients* to browse through the *AddressSpace* or subsets of the *AddressSpace* called *Views*. The *Query Service Set* allows *Clients* to get a subset of data from the *AddressSpace* or the *View*.




Figure 5 – View Service Set

The *Attribute Service Set* is illustrated in Figure 6. It defines *Services* that allow *Clients* to read and write *Attributes* of *Nodes*, including their historical values. Since the value of a *Variable* is modelled as an *Attribute*, these *Services* allow *Clients* to read and write the values of *Variables*.




Figure 6 – Attribute Service Set

The *Method Service Set* is illustrated in Figure 7. It defines *Services* that allow *Clients* to call methods. Methods run to completion when called. They may be called with method-specific input parameters and may return method-specific output parameters.




Figure 7 – Method Service Set

The *MonitoredItem Service Set* and the *Subscription Service Set*, illustrated in Figure 8, are used together to subscribe to *Nodes* in the OPC UA AddressSpace.

The *MonitoredItem Service Set* defines *Services* that allow *Clients* to create, modify, and delete *MonitoredItems* used to monitor *Attributes* for value changes and *Objects* for *Events*.

These *Notifications* are queued for transfer to the *Client* by *Subscriptions*.

The *Subscription Service Set* defines *Services* that allow *Clients* to create, modify and delete *Subscriptions*. *Subscriptions* send *Notifications* generated by *MonitoredItems* to the *Client*. *Subscription Services* also provide for *Client* recovery from missed *Messages* and communication failures.




Figure 8 – MonitoredItem and Subscription Service Sets

4.2 Request/response Service procedures

Request/response *Service* procedures describe the processing of requests received by the *Server*, and the subsequent return of responses. The procedures begin with the requesting *Client* submitting a *Service* request *Message* to the *Server*.

Upon receipt of the request, the *Server* processes the *Message* in two steps. In the first step, it attempts to decode and locate the *Service* to execute. The error handling for this step is specific to the communication technology used and is described in IEC 62541-6.

If it succeeds, then it attempts to access each operation identified in the request and perform the requested operation. For each operation in the request, it generates a separate success/failure code that it includes in a positive response *Message* along with any data that is to be returned.

To perform these operations, both the *Client* and the *Server* may make use of the API of a *Communication Stack* to construct and interpret *Messages* and to access the requested operation.

The implementation of each service request or response handling shall check that each service parameter lies within the specified range for that parameter.

5 Service Sets

5.1 General

This clause defines the OPC UA *Service Sets* and their *Services*. Clause 7 contains the definitions of common parameters used by these *Services*. Subclause 6.5 describes auditing requirements for all services.

Whether or not a *Server* supports a *Service Set*, or a *Service* within a *Service Set*, is defined by its *Profile*. *Profiles* are described in IEC 62541-7.

5.2 Service request and response header

Each *Service* request has a *RequestHeader* and each *Service* response has a *ResponseHeader*.

The *RequestHeader* structure is defined in 7.28 and contains common request parameters such as *authenticationToken*, *timestamp* and *requestHandle*.

The *ResponseHeader* structure is defined in 7.29 and contains common response parameters such as *serviceResult* and *diagnosticInfo*.

5.3 Service results

Service results are returned at two levels in OPC UA responses, one that indicates the status of the *Service* call, and the other that indicates the status of each operation requested by the *Service*.

Service results are defined via the *StatusCode* (see 7.34).

The status of the *Service* call is represented by the *serviceResult* contained in the *ResponseHeader* (see 7.29). The mechanism for returning this parameter is specific to the communication technology used to convey the *Service* response and is defined in IEC 62541-6.

The status of individual operations in a request is represented by individual *StatusCodes*.

The following cases define the use of these parameters.

- a) A bad code is returned in *serviceResult* if the *Service* itself failed. In this case, a *ServiceFault* is returned. The *ServiceFault* is defined in 7.30.
- b) The good code is returned in *serviceResult* if the *Service* fully or partially succeeded. In this case, other response parameters are returned. The *Client* shall always check the response parameters, especially all *StatusCodes* associated with each operation. These *StatusCodes* may indicate bad or uncertain results for one or more operations requested in the *Service* call.

All *Services* with arrays of operations in the request shall return a bad code in the *serviceResult* if the array is empty.

The *Services* define various specific *StatusCodes* and a *Server* shall use these specific *StatusCodes* as described in the *Service*. A *Client* should be able to handle these *Service* specific *StatusCodes*. In addition, a *Client* shall expect other common *StatusCodes* defined in Table 177 and Table 178. Additional details for *Client* handling of specific *StatusCodes* may be defined in IEC 62541-7.

If the *Server* discovers, through some out-of-band mechanism that the application or user credentials used to create a *Session* or *SecureChannel* have been compromised, then the *Server* should immediately terminate all sessions and channels that use those credentials. In this case, the *Service* result code should be either *Bad_IdentityTokenRejected* or *Bad_CertificateUntrusted*.

Message parsing can fail due to syntax errors or if data contained within the message exceeds ranges supported by the receiver. When this happens messages shall be rejected by the receiver. If the receiver is a *Server* then it shall return a *ServiceFault* with result code of

Bad_DecodingError or *Bad_EncodingLimitsExceeded*. If the receiver is the *Client* then the *Communication Stack* should report these errors to the *Client* application.

Many applications will place limits on the size of messages and/or data elements contained within these messages. For example, a *Server* may reject requests containing string values longer than a certain length. These limits are typically set by administrators and apply to all connections between a *Client* and a *Server*.

Clients that receive *Bad_EncodingLimitsExceeded* faults from the *Server* will likely ~~have~~ need to reformulate their requests. The administrator may need to increase the limits for the *Client* if it receives a response from the *Server* with this fault.

In some cases, parsing errors are fatal and it is not possible to return a fault. For example, the incoming message could exceed the buffer capacity of the receiver. In these cases, these errors may be treated as a communication fault which requires the *SecureChannel* to be re-established (see 5.5).

The *Client* and *Server* reduce the chances of a fatal error by exchanging their message size limits in the *CreateSession* service. This will allow either party to avoid sending a message that causes a communication fault. The *Server* should return a *Bad_ResponseTooLarge* fault if a serialized response message exceeds the message size specified by the *Client*. Similarly, the *Client Communication Stack* should report a *Bad_RequestTooLarge* error to the application before sending a message that exceeds the *Server's* limit.

Note that the message size limits only apply to the raw message body and do not include headers or the effect of applying any security. This means that a message body that is smaller than the specified maximum could still cause a fatal error.

5.4 Discovery Service Set

5.4.1 Overview

This *Service Set* defines *Services* used to discover the *Endpoints* implemented by a *Server* and to read the security configuration for those *Endpoints*. The *Discovery Services* are implemented by individual *Servers* and by dedicated *Discovery Servers*. IEC 62541-12² describes how to use the *Discovery Services* with dedicated *Discovery Servers*.

Every *Server* shall have a *DiscoveryEndpoint* that *Clients* can access without establishing a *Session*. This *Endpoint* may or may not be the same *Session Endpoint* that *Clients* use to establish a *SecureChannel*. *Clients* read the security information necessary to establish a *SecureChannel* by calling the *GetEndpoints* *Service* on the *DiscoveryEndpoint*.

In addition, *Servers* may register themselves with a well-known *Discovery Server* using the *RegisterServer* *Service*. *Clients* can later discover any registered *Servers* by calling the *FindServers* *Service* on the *Discovery Server*.

The ~~complete~~ discovery process using *FindServers* is illustrated in Figure 9. The establishment of a *SecureChannel* (with *MessageSecurityMode* NONE) for *FindServers* and *GetEndpoints* is omitted from the figure for clarity.

² ~~Under consideration.~~




Figure 9 – Discovery process

The URL for a *DiscoveryEndpoint* shall provide all of the information that the *Client* needs to connect to the *DiscoveryEndpoint*.

Once a *Client* retrieves the *Endpoints*, the *Client* can save this information and use it to connect directly to the *Server* again without going through the discovery process. If the *Client* finds that it cannot connect then the *Server* configuration may have changed and the *Client* needs to go through the discovery process again.

DiscoveryEndpoints shall not require any message security, but it may require transport layer security. In production systems, Administrators may disable discovery for security reasons and *Clients* shall rely on cached *EndpointDescriptions*. To provide support for systems with disabled *Discovery Services* *Clients* shall allow *Administrators* to manually update the *EndpointDescriptions* used to connect to a *Server*. *Servers* shall allow *Administrators* to disable the *DiscoveryEndpoint*.

A *Client* shall be careful when using the information returned from a *DiscoveryEndpoint* since it has no security. A *Client* does this by comparing the information returned from the *DiscoveryEndpoint* to the information returned in the *CreateSession* response. A *Client* shall verify that:

- The *ApplicationUri* specified in the *Server Certificate* is the same as the *ApplicationUri* provided in the *EndpointDescription*.
- The *Server Certificate* returned in *CreateSession* response is the same as the *Certificate* used to create the *SecureChannel*.
- The *EndpointDescriptions* returned from the *DiscoveryEndpoint* are the same as the *EndpointDescriptions* returned in the *CreateSession* response.

If the *Client* detects that one of the above requirements is not fulfilled, then the *Client* shall close the *SecureChannel* and report an error.

A *Client* shall verify the *HostName* specified in the *Server Certificate* is the same as the *HostName* contained in the *endpointUrl* provided in the *EndpointDescription* returned by *CreateSession*. If there is a difference then the *Client* shall report the difference and may close the *SecureChannel*. *Servers* shall add all possible *HostNames* like *MyHost* and *MyHost.local* into the *Server Certificate*. This includes IP addresses of the host or the *HostName* exposed by a NAT router used to connect to the *Server*.

5.4.2 FindServers

5.4.2.1 Description

This *Service* returns the *Servers* known to a *Server* or *Discovery Server*. The behaviour of *Discovery Servers* is described in detail in IEC 62541-12³.

The *Client* may reduce the number of results returned by specifying filter criteria. A *Discovery Server* returns an empty list if no *Servers* match the criteria specified by the client. The filter criteria supported by this *Service* are described in 5.4.2.2.

Every *Server* shall provide a *DiscoveryEndpoint* that supports this *Service*. The *Server* shall always return a record that describes itself, however in some cases more than one record may be returned. *Gateway Servers* shall return a record for each *Server* that they provide access to plus (optionally) a record that allows the *Gateway Server* to be accessed as an ordinary OPC UA *Server*. Non-transparent redundant *Servers* shall provide a record for each *Server* in the *Redundant Server Set*.

Every *Server* shall have a globally unique identifier called the *ServerUri*. This identifier should be a fully qualified domain name; however, it may be a GUID or similar construct that ensures global uniqueness. The *ServerUri* returned by this *Service* shall be the same value that appears in index 0 of the *ServerArray* property (see IEC 62541-5). The *ServerUri* is returned as the *applicationUri* field in the *ApplicationDescription* (see 7.1)

Every *Server* shall also have a human readable identifier called the *ServerName* which is not necessarily globally unique. This identifier may be available in multiple locales.

A *Server* may have multiple *HostNames*. For this reason, the *Client* shall pass the URL it used to connect to the *Endpoint* to this *Service*. The implementation of this *Service* shall use this information to return responses that are accessible to the *Client* via the provided URL.

This *Service* shall not require ~~any~~ message security but it may require transport layer security.

Some *Servers* may be accessed via a *Gateway Server* and shall have a value specified for *gatewayServerUri* in their *ApplicationDescription* (see 7.1). The *discoveryUrls* provided in *ApplicationDescription* shall belong to the *Gateway Server*. Some *Discovery Servers* may return multiple records for the same *Server* if that *Server* can be accessed via multiple paths.

This *Service* can be used without security and it is therefore vulnerable to Denial of Service (DOS) attacks. A *Server* should minimize the amount of processing required to send the response for this *Service*. This can be achieved by preparing the result in advance. The *Server* should also add a short delay before starting processing of a request during high traffic conditions.

~~The *DiscoveryUri* returned by this *Service* is ambiguous if there are multiple *TransportProfiles* (e.g. UA XML or UA Binary encoding) associated with the URL scheme. *Clients* that support multiple *TransportProfiles* should attempt to use alternate *TransportProfiles* if the first choice does not succeed.~~

5.4.2.2 Parameters

Table 3 defines the parameters for the *Service*.

³ ~~Under consideration.~~

Table 3 – FindServers Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always ~~omitted~~ null. The <i>authenticationToken</i> shall be ignored if it is provided. The type <i>RequestHeader</i> is defined in 7.28.
endpointUrl	String	The network address that the <i>Client</i> used to access the <i>DiscoveryEndpoint</i> . The <i>Server</i> uses this information for diagnostics and to determine what URLs to return in the response. The <i>Server</i> should return a suitable default URL if it does not recognize the <i>HostName</i> in the URL.
localeIds []	LocaleId	List of locales to use. The <i>Server</i> should return the ~~ServerName~~ <i>applicationName</i> in the <i>ApplicationDescription</i> defined in 7.1 using one of locales specified. If the <i>Server</i> supports more than one of the requested locales then the <i>Server</i> shall use the locale that appears first in this list. If the <i>Server</i> does not support any of the requested locales it chooses an appropriate default locale. The <i>Server</i> chooses an appropriate default locale if this list is empty.
serverUris []	String	List of <i>Servers</i> to return. All known <i>Servers</i> are returned if the list is empty. A <i>serverUri</i> matches the <i>applicationUri</i> from the <i>ApplicationDescription</i> defined in 7.1.
Response		
responseHeader	ResponseHeader	Common response parameters. The <i>ResponseHeader</i> type is defined in 7.29.
servers []	ApplicationDescription	List of <i>Servers</i> that meet criteria specified in the request. This list is empty if no <i>Servers</i> meet the criteria. The <i>ApplicationDescription</i> type is defined in 7.1.

5.4.2.3 Service results

Common *StatusCodes* are defined in Table 177.

5.4.3 FindServersOnNetwork

5.4.3.1 Description

This *Service* returns the *Servers* known to a *Discovery Server*. Unlike *FindServers*, this *Service* is only implemented by *Discovery Servers*.

The *Client* may reduce the number of results returned by specifying filter criteria. An empty list is returned if no *Server* matches the criteria specified by the *Client*.

This *Service* shall not require message security but it may require transport layer security.

Each time the *Discovery Server* creates or updates a record in its cache, it shall assign a monotonically increasing identifier to the record. This allows *Clients* to request records in batches by specifying the identifier for the last record received in the last call to *FindServersOnNetwork*. To support this the *Discovery Server* shall return records in numerical order starting from the lowest record identifier. The *Discovery Server* shall also return the last time the counter was reset, for example due to a restart of the *Discovery Server*. If a *Client*

detects that this time is more recent than the last time the *Client* called the *Service*, it shall call the *Service* again with a startingRecordId of 0.

This *Service* can be used without security and it is therefore vulnerable to denial of service (DOS) attacks. A *Server* should minimize the amount of processing required to send the response for this *Service*. This can be achieved by preparing the result in advance.

5.4.3.2 Parameters

Table 4 defines the parameters for the *Service*.

Table 4 – FindServersOnNetwork Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The <i>authenticationToken</i> shall be ignored if it is provided. The type <i>RequestHeader</i> is defined in 7.28.
startingRecordId	Counter	Only records with an identifier greater than this number will be returned. Specify 0 to start with the first record in the cache.
maxRecordsToReturn	UInt32	The maximum number of records to return in the response. 0 indicates that there is no limit.
serverCapabilityFilter[]	String	List of <i>Server</i> capability filters. The set of allowed <i>Server</i> capabilities are defined in IEC 62541-12. Only records with all of the specified <i>Server</i> capabilities are returned. The comparison is case insensitive. If this list is empty then no filtering is performed.
Response		
responseHeader	ResponseHeader	Common response parameters. The <i>ResponseHeader</i> type is defined in 7.29.
lastCounterResetTime	UtcTime	The last time the counters were reset.
servers[]	ServerOnNetwork	List of DNS service records that meet criteria specified in the request. This list is empty if no <i>Servers</i> meet the criteria.
recordId	UInt32	A unique identifier for the record. This can be used to fetch the next batch of <i>Servers</i> in a subsequent call to <i>FindServersOnNetwork</i> .
serverName	String	The name of the <i>Server</i> specified in the mDNS announcement (see IEC 62541-12). This may be the same as the <i>ApplicationName</i> for the <i>Server</i> .
discoveryUrl	String	The URL of the <i>DiscoveryEndpoint</i> .
serverCapabilities	String[]	The set of <i>Server</i> capabilities supported by the <i>Server</i> . The set of allowed <i>Server</i> capabilities are defined in IEC 62541-12.

5.4.3.3 Service results

Common *StatusCodes* are defined in Table 177.

5.4.4 GetEndpoints

5.4.4.1 Description

This *Service* returns the *Endpoints* supported by a *Server* and all of the configuration information required to establish a *SecureChannel* and a *Session*.

This *Service* shall not require ~~any~~ message security but it may require transport layer security.

A *Client* may reduce the number of results returned by specifying filter criteria based on *LocaleIds* and *Transport Profile* URIs. The *Server* returns an empty list if no *Endpoints* match the criteria specified by the client. The filter criteria supported by this *Service* are described in 5.4.4.2.

A *Server* may support multiple security configurations for the same *Endpoint*. In this situation, the *Server* shall return separate *EndpointDescription* records for each available configuration. *Clients* should treat each of these configurations as distinct *Endpoints* even if the physical URL happens to be the same.

The security configuration for an *Endpoint* has four components:

- Server Application Instance Certificate
- Message Security Mode
- Security Policy
- Supported User Identity Tokens

The *ApplicationInstanceCertificate* is used to secure the *OpenSecureChannel* request (see 5.5.2). The *MessageSecurityMode* and the *SecurityPolicy* tell the *Client* how to secure messages sent via the *SecureChannel*. The *UserIdentityTokens* tell the *Client* which type of user credentials shall be passed to the *Server* in the *ActivateSession* request (see 5.6.3).

If the *securityPolicyUri* is NONE and none of the *UserTokenPolicies* requires encryption, the *Client* shall ignore the *ApplicationInstanceCertificate*.

Each *EndpointDescription* also specifies a URI for the *Transport Profile* that the *Endpoint* supports. The *Transport Profiles* specify information such as message encoding format and protocol version and are defined in IEC 62541-7. ~~Clients shall fetch the Server's SoftwareCertificates if they want to discover the complete list of Profiles supported by the Server (see 7.30).~~

Messages are secured by applying standard cryptography algorithms to the messages before they are sent over the network. The exact set of algorithms used depends on the *SecurityPolicy* for the *Endpoint*. IEC 62541-7 defines *Profiles* for common *SecurityPolicies* and assigns a unique URI to them. It is expected that applications have built in knowledge of the *SecurityPolicies* that they support, as a result, only the Profile URI for the *SecurityPolicy* is specified in the *EndpointDescription*. A *Client* cannot connect to an *Endpoint* that does not support a *SecurityPolicy* that it recognizes.

An *EndpointDescription* may specify that the message security mode is NONE. This configuration is not recommended unless the applications are communicating on a physically isolated network where the risk of intrusion is extremely small. If the message security is NONE then it is possible for *Clients* to deliberately or accidentally hijack *Sessions* created by other *Clients*.

A *Server* may have multiple *HostNames*. For this reason, the *Client* shall pass the URL it used to connect to the *Endpoint* to this *Service*. The implementation of this *Service* shall use this information to return responses that are accessible to the *Client* via the provided URL.

This *Service* can be used without security and it is therefore vulnerable to Denial of Service (DOS) attacks. A *Server* should minimize the amount of processing required to send the response for this *Service*. This can be achieved by preparing the result in advance. The *Server* should also add a short delay before starting processing of a request during high traffic conditions.

Some of the *EndpointDescriptions* returned in a response shall specify the *Endpoint* information for a *Gateway Server* that can be used to access another *Server*. In these situations, the *gatewayServerUri* is specified in the *EndpointDescription* and all security checks used to verify *Certificates* shall use the *gatewayServerUri* (see 6.1.3) instead of the *serverUri*.

To connect to a *Server* via the gateway the *Client* shall first establish a *SecureChannel* with the *Gateway Server*. Then the *Client* shall call the *CreateSession* service and pass the *serverUri* specified in the *EndpointDescription* to the *Gateway Server*. The *Gateway Server* shall then connect to the underlying *Server* on behalf of the *Client*. The process of connecting to a *Server* via a *Gateway Server* is illustrated in Figure 10.




Figure 10 – Using a Gateway Server

5.4.4.2 Parameters

Table 5 defines the parameters for the *Service*.

Table 5 – GetEndpoints Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always ~~omitted~~ null. The <i>authenticationToken</i> shall be ignored if it is provided. The type <i>RequestHeader</i> is defined in 7.28.
endpointUrl	String	The network address that the <i>Client</i> used to access the <i>DiscoveryEndpoint</i> . The <i>Server</i> uses this information for diagnostics and to determine what URLs to return in the response. The <i>Server</i> should return a suitable default URL if it does not recognize the <i>HostName</i> in the URL.
localeIds []	LocaleId	List of locales to use. Specifies the locale to use when returning human readable strings. This parameter is described in 5.4.2.2.
profileUris []	String	List of <i>Transport Profile</i> that the returned <i>Endpoints</i> shall support. IEC 62541-7 defines URIs for the <i>Transport Profiles</i> . All <i>Endpoints</i> are returned if the list is empty. If the URI is a URL, this URL may have a query string appended. The <i>Transport Profiles</i> that support query strings are defined in IEC 62541-7.
Response		
responseHeader	ResponseHeader	Common response parameters. The <i>ResponseHeader</i> type is defined in 7.29.
Endpoints []	EndpointDescription	List of <i>Endpoints</i> that meet criteria specified in the request. This list is empty if no <i>Endpoints</i> meet the criteria. The <i>EndpointDescription</i> type is defined in 7.10.

5.4.4.3 Service results

Common *StatusCodes* are defined in Table 177.

5.4.5 RegisterServer**5.4.5.1 Description**

This Service is implemented by *Discovery Servers*.

This Service registers a *Server* with a *Discovery Server*. This Service will be called by a *Server* or a separate configuration utility. *Clients* will not use this Service.

A *Server* shall establish a *SecureChannel* with the *Discovery Server* before calling this Service. The *SecureChannel* is described in 5.5. The *Administrator* of the *Server* shall provide the *Server* with an *EndpointDescription* for the *Discovery Server* as part of the configuration process. *Discovery Servers* shall reject registrations if the *serverUri* provided does not match the *applicationUri* in *Server Certificate* used to create the *SecureChannel*.

This Service can only be invoked via *SecureChannels* that support *Client* authentication (i.e. HTTPS cannot be used to call this Service).

A *Server* only provides its *serverUri* and the URLs of the *DiscoveryEndpoints* to the *Discovery Server*. *Clients* shall use the *GetEndpoints Service* to fetch the most up to date configuration information directly from the *Server*.

The *Server* shall provide a localized name for itself in all locales that it supports.

Servers shall be able to register themselves with a *Discovery Server* running on the same machine. The exact mechanisms depend on the *Discovery Server* implementation and are described in IEC 62541-6.

There are two types of *Server* applications: those which are manually launched including a start by the operating system at boot and those that are automatically launched when a *Client* attempts to connect. The registration process that a *Server* shall use depends on which category it falls into.

The registration process for manually launched *Servers* is illustrated in Figure 11.




Figure 11 – The registration process – Manually launched servers

The registration process for automatically launched *Servers* is illustrated in Figure 12.




Figure 12 – The registration process – Automatically launched servers

The registration process is designed to be platform independent, robust and able to minimize problems created by configuration errors. For that reason, *Servers* shall register themselves more than once.

Under normal conditions, manually launched *Servers* shall periodically register with the *Discovery Server* as long as they are able to receive connections from *Clients*. If a *Server* goes offline then it shall register itself once more and indicate that it is going offline. The period of the recurring registration-frequency should be configurable; however, the maximum is 10 minutes. If an error occurs during registration (e.g. the *Discovery Server* is not running) then the *Server* shall periodically re-attempt registration. The frequency of these attempts should start at 1 second but gradually increase until the registration frequency is the same as what it would be if no errors occurred. The recommended approach would be to double the period of each attempt until reaching the maximum.

When an automatically launched *Server* (or its install program) registers with the *Discovery Server* it shall provide a path to a semaphore file which the *Discovery Server* can use to determine if the *Server* has been uninstalled from the machine. The *Discovery Server* shall have read access to the file system that contains the file.

5.4.5.2 Parameters

Table 6 defines the parameters for the *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 6 – RegisterServer Service parameters

Name	Type	Description
Request		
—requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always omitted. The type <i>RequestHeader</i> is defined in 7.28.
—Server	RegisteredServer	The server to register. This structure is defined in-line with the following indented items.
—serverUri	String	The globally unique identifier for the <i>Server</i> instance. The <i>serverUri</i> matches the <i>applicationUri</i> from the <i>ApplicationDescription</i> defined in 7.1.
—productUri	String	The globally unique identifier for the <i>Server</i> product.
—serverNames []	LocalizedText	A list of localized descriptive names for the <i>Server</i> . The list shall have at least one valid entry.
—serverType	Enum ApplicationType	The type of application. The enumeration values are defined in Table 103. The value "CLIENT_1" (The application is a <i>Client</i>) is not allowed. The <i>Service</i> result shall be <i>Bad_InvalidArgument</i> in this case.
—gatewayServerUri	String	The URI of the <i>Gateway Server</i> associated with the <i>discoveryUrls</i> . This value is only specified by <i>Gateway Servers</i> that wish to register the <i>Servers</i> that they provide access to. For <i>Servers</i> that do not act as a <i>Gateway Server</i> this parameter shall be null.
—discoveryUrls []	String	A list of <i>Discovery Endpoints</i> for the <i>Server</i> . The list shall have at least one valid entry.
—semaphoreFilePath	String	The path to the semaphore file used to identify an automatically-launched server instance; Manually-launched servers will not use this parameter. If a Semaphore file is provided, the <i>isOnline</i> flag is ignored. If a Semaphore file is provided and exists, the <i>LocalDiscoveryServer</i> shall save the registration information in a persistent data store that it reads whenever the <i>LocalDiscoveryServer</i> starts. If a Semaphore file is specified but does not exist the <i>Discovery Server</i> shall remove the registration from any persistent data store. If the <i>Server</i> has registered with a <i>semaphoreFilePath</i> , the <i>DiscoveryServer</i> shall check that this file exists before returning the <i>ApplicationDescription</i> to the client. If the <i>Server</i> did not register with a <i>semaphoreFilePath</i> (it is null or empty) then the <i>DiscoveryServer</i> does not attempt to verify the existence of the file before returning the <i>ApplicationDescription</i> to the client.
—isOnline	Boolean	True if The <i>Server</i> is currently able to accept connections from <i>Clients</i> . The <i>DiscoveryServer</i> shall return <i>ApplicationDescriptions</i> to the <i>Client</i> . The <i>Server</i> is expected to periodically re-register with the <i>DiscoveryServer</i> . False if the <i>Server</i> is currently unable to accept connections from <i>Clients</i> . The <i>DiscoveryServer</i> shall NOT return <i>ApplicationDescriptions</i> to the <i>Client</i> . This parameter is ignored if a <i>semaphoreFilePath</i> is provided.
Response		
—ResponseHeader	ResponseHeader	Common response parameters. The type <i>ResponseHeader</i> is defined in 7.27.

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
Server	RegisteredServer	The <i>Server</i> to register. The type <i>RegisteredServer</i> is defined in 7.27.
Response		
ResponseHeader	ResponseHeader	Common response parameters. The type <i>ResponseHeader</i> is defined in 7.29.

5.4.5.3 Service results

Table 7 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 7 – RegisterServer Service result codes

Symbolic Id	Description
Bad_InvalidArgument	See Table 177 for the description of this result code.
Bad_ServerUriInvalid	See Table 177 for the description of this result code.
Bad_ServerNameMissing	No <i>ServerName</i> was specified.
Bad_DiscoveryUriMissing	No discovery URL was specified.
Bad_SemaphoreFileMissing	The semaphore file specified is not valid.

5.4.6 RegisterServer2

5.4.6.1 Description

This *Service* is implemented by *Discovery Servers*.

This *Service* allows a *Server* to register its *DiscoveryUrls* and capabilities with a *Discovery Server*. It extends the registration information from *RegisterServer* with information necessary for *FindServersOnNetwork*. This *Service* will be called by a *Server* or a separate configuration utility. *Clients* will not use this *Service*.

Servers that support *RegisterServer2* shall try to register with the *Discovery Server* using this *Service* and shall fall back to *RegisterServer* if *RegisterServer2* fails with the status *Bad_ServiceUnsupported*.

A *Discovery Server* that implements this *Service* needs to assign unique record ids each time this *Service* is called. See 5.4.3 for more details.

This *Service* can only be invoked via *SecureChannels* that support *Client* authentication (i.e. HTTPS cannot be used to call this *Service*).

5.4.6.2 Parameters

Table 8 defines the parameters for the *Service*.

Table 8 – RegisterServer2

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
Server	RegisteredServer	The <i>Server</i> to register. The type <i>RegisteredServer</i> is defined in 7.27.
discoveryConfiguration []	ExtensibleParameter DiscoveryConfiguration	Additional configuration settings for the <i>Server</i> to register. The <i>discoveryConfiguration</i> is an extensible parameter type defined in 7.9. Discovery <i>Servers</i> that do not understand a configuration shall return <i>Bad_NotSupported</i> for this configuration.
Response		
ResponseHeader	ResponseHeader	Common response parameters. The type <i>ResponseHeader</i> is defined in 7.29.
configurationResults []	StatusCode	List of results for the <i>discoveryConfiguration</i> parameters.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>discoveryConfiguration</i> parameters. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.4.6.3 Service results

Table 9 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 9 – RegisterServer2 Service result codes

Symbolic Id	Description
Bad_InvalidArgument	See Table 177 for the description of this result code.
Bad_ServerUriInvalid	See Table 177 for the description of this result code.
Bad_ServerNameMissing	No <i>ServerName</i> was specified.
Bad_DiscoveryUriMissing	No discovery URL was specified.
Bad_SemaphoreFileMissing	The semaphore file specified is not valid.
Bad_ServiceUnsupported	See Table 177 for the description of this result code.

5.4.6.4 StatusCodes

Table 10 defines values for the operation level *configurationResults* parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 10 – RegisterServer2 Operation Level result codes

Symbolic Id	Description
Bad_NotSupported	See Table 178 for the description of this result code.

5.5 SecureChannel Service Set

5.5.1 Overview

This *Service Set* defines *Services* used to open a communication channel that ensures the confidentiality and *Integrity* of all *Messages* exchanged with the *Server*. The base concepts for OPC UA security are defined in IEC TR 62541-2.

The *SecureChannel Services* are unlike other *Services* because they are not implemented directly by the *OPC UA Application*. Instead, they are provided by the *Communication Stack* on which the *OPC UA Application* is built. For example, an OPC UA Server may be built on a SOAP stack that allows applications to establish a *SecureChannel* using ~~the WS-SecureConversation specification~~ HTTPS. In these cases, the *OPC UA Application* shall verify that the *Message* it received was in the context of ~~a WS-SecureConversation~~ an HTTPS connection. IEC 62541-6 describes how the *SecureChannel Services* are implemented.

A *SecureChannel* is a long-running logical connection between a single *Client* and a single *Server*. This channel maintains a set of keys known only to the *Client* and *Server*, which are used to authenticate and encrypt *Messages* sent across the network. The *SecureChannel Services* allow the *Client* and *Server* to securely negotiate the keys to use.

Logical connections may be initiated by the *Client* or by the *Server* as described in IEC 62541-6. After the connection is initiated, the *SecureChannel* is opened and closed by the *Client* using the *SecureChannel Services*.

An *EndpointDescription* tells a *Client* how to establish a *SecureChannel* with a given *Endpoint*. A *Client* may obtain the *EndpointDescription* from a *Discovery Server*, via some non-UA defined directory server or from its own configuration.

The exact algorithms used to authenticate and encrypt *Messages* are described in the *SecurityPolicy* field of the *EndpointDescription*. A *Client* shall use these algorithms when it creates a *SecureChannel*.

It should be noted that some *SecurityPolicies* defined in IEC 62541-7 will turn off authentication and encryption resulting in a *SecureChannel* that provides no security.

When a *Client* and *Server* are communicating via a *SecureChannel*, they shall verify that all incoming *Messages* have been signed and encrypted according to the requirements specified in the *EndpointDescription*. An *OPC UA Application* shall not process any *Message* that does not conform to these requirements.

The relationship between the *SecureChannel* and the *OPC UA Application* depends on the implementation technology. IEC 62541-6 defines any requirements that depend on the technology used.

The correlation between the *OPC UA Application Session* and the *SecureChannel* is illustrated in Figure 13. The *Communication Stack* is used by the *OPC UA Applications* to exchange *Messages*. In the first step, the *SecureChannel Services* are used to establish a *SecureChannel* between the two *Communication Stacks* which allows the secure exchange of *Messages*. In the second step, the *OPC UA Applications* use the *Session Service Set* to establish an *OPC UA Application Session*.




Figure 13 – SecureChannel and Session Services

Once a *Client* has established a *Session* it may wish to access the *Session* from a different *SecureChannel*. The *Client* can do this by validating the new *SecureChannel* with the *ActivateSession Service* described in 5.6.3.

If a *Server* acts as a *Client* to other *Servers*, which is commonly referred to as *Server chaining*, then the *Server* shall be able to maintain user level security. By this we mean that the user identity should be passed to the underlying *Server* or it should be mapped to an appropriate user identity in the underlying server. It is unacceptable to ignore user level security. This is required to ensure that security is maintained and that a user does not obtain information that they should not have access to. Whenever possible a *Server* should impersonate the original *Client* by passing the original *Client's* user identity to the underlying *Server* when it calls the *ActivateSession Service*. If impersonation is not an option then the *Server* shall map the original *Client's* user identity onto a new user identity which the underlying *Server* does recognize.

5.5.2 OpenSecureChannel

5.5.2.1 Description

This *Service* is used to open or renew a *SecureChannel* that can be used to ensure *Confidentiality* and *Integrity* for *Message* exchange during a *Session*. This *Service* requires the *Communication Stack* to apply the various security algorithms to the *Messages* as they are sent and received. Specific implementations of this *Service* for different *Communication Stacks* are described in IEC 62541-6.

Each *SecureChannel* has a globally-unique identifier and is valid for a specific combination of *Client* and *Server* application instances. Each channel contains one or more *SecurityTokens* that identify a set of cryptography keys that are used to encrypt and authenticate *Messages*. *SecurityTokens* also have globally-unique identifiers which are attached to each *Message* secured with the token. This allows an authorized receiver to know how to decrypt and verify the *Message*.

SecurityTokens have a finite lifetime negotiated with this *Service*. However, differences between the system clocks on different machines and network latencies mean that valid *Messages* could arrive after the token has expired. To prevent valid *Messages* from being discarded, the applications should do the following:

- Clients* should request a new *SecurityToken* after 75 % of its lifetime has elapsed. This should ensure that *Clients* will receive the new *SecurityToken* before the old one actually expires.
- Servers* shall use the existing *SecurityToken* to secure outgoing *Messages* until the *SecurityToken* expires or the *Server* receives a *Message* secured with a new

SecurityToken. This should ensure that *Clients* do not reject *Messages* secured with the new *SecurityToken* that arrive before the *Client* receives the new *SecurityToken*.

- c) *Clients* should accept *Messages* secured by an expired *SecurityToken* for up to 25 % of the token lifetime. This should ensure that *Messages* sent by the *Server* before the token expired are not rejected because of network delays.

Each *SecureChannel* exists until it is explicitly closed or until the last token has expired and the overlap period has elapsed. A *Server* application should limit the number of *SecureChannels*. To protect against misbehaving *Clients* and denial of service attacks, the *Server* shall close the oldest *SecureChannel* that has no *Session* assigned before reaching the maximum number of supported *SecureChannels*.

The *OpenSecureChannel* request and response *Messages* shall be signed with the sender's *Certificate* private key. These *Messages* shall always be encrypted. If the transport layer does not provide encryption, then these *Messages* shall be encrypted with the receiver's *Certificate* public key. These requirements for *OpenSecureChannel* only apply if the *securityPolicyUri* is not None.

If the protocol defined in IEC 62541-6 requires that *Application Instance Certificates* are used in the *OpenSecureChannel* Service ~~shall be the Application Instance Certificates~~, then *Clients* and *Servers* shall verify that the same *Certificates* ~~were~~ are used in the *CreateSession* and *ActivateSession* Services. *Certificates* are not provided and shall not be verified if the *securityPolicyUri* is None.

If the *securityPolicyUri* is not None, a *Client* shall verify the *HostName* specified in the *Server Certificate* is the same as the *HostName* contained in the *endpointUrl*. If there is a difference then the *Client* shall report the difference and may choose to not open the *SecureChannel*. *Servers* shall add all possible *HostNames* like *MyHost* and *MyHost.local* into the *Server Certificate*. This includes IP addresses of the host or the *HostName* exposed by a NAT router used to connect to the *Server*.

Clients should be prepared to replace the *HostName* returned in the *EndpointDescription* with the *HostName* or the IP addresses they used to call *GetEndpoints*.

5.5.2.2 Parameters

Table 11 defines the parameters for the *Service*.

Unlike other *Services*, the parameters for this *Service* provide only an abstract definition. The concrete representation on the network depends on the mappings defined in IEC 62541-6.

Table 11 – OpenSecureChannel Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always omitted null. The type <i>RequestHeader</i> is defined in 7.28.
clientCertificate	ApplicationInstanceCertificate	A <i>Certificate</i> that identifies the <i>Client</i> . The <i>OpenSecureChannel</i> request shall be signed with the private key for this <i>Certificate</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2. If the <i>securityPolicyUri</i> is None, the <i>Server</i> shall ignore the <i>ApplicationInstanceCertificate</i> .
requestType	Enum SecurityToken RequestType	The type of <i>SecurityToken</i> request: An enumeration that shall be one of the following: ISSUE_0 creates a new <i>SecurityToken</i> for a new <i>SecureChannel</i> . RENEW_1 creates a new <i>SecurityToken</i> for an existing <i>SecureChannel</i> .
secureChannelId	ByteString BaseDataType	The identifier for the <i>SecureChannel</i> that the new token should belong to. This parameter shall be null when creating a new <i>SecureChannel</i> . The concrete security protocol definition in IEC 62541-6 chooses the concrete <i>DataType</i> .
securityMode	Enum MessageSecurityMode	The type of security to apply to the messages. The type <i>MessageSecurityMode</i> type is defined in 7.15. A <i>SecureChannel</i> may have need to be created even if the <i>securityMode</i> is NONE. The exact behaviour depends on the mapping used and is described in the IEC 62541-6.
securityPolicyUri	String	The URI for <i>SecurityPolicy</i> to use when securing messages sent over the <i>SecureChannel</i> . The set of known URIs and the <i>SecurityPolicies</i> associated with them are defined in IEC 62541-7.
clientNonce	ByteString	A random number that shall not be used in any other request. A new <i>clientNonce</i> shall be generated for each time a <i>SecureChannel</i> is renewed. This parameter shall have a length equal to key size used for the symmetric encryption algorithm that is identified by the securityPolicyUri. This parameter shall have a length equal to the <i>SecureChannelNonceLength</i> defined for the <i>SecurityPolicy</i> in IEC 62541-7. The <i>SecurityPolicy</i> is identified by the <i>securityPolicyUri</i> .
requestedLifetime	Duration	The requested lifetime, in milliseconds, for the new <i>SecurityToken</i> . It specifies when the <i>Client</i> expects to renew the <i>SecureChannel</i> by calling the <i>OpenSecureChannel Service</i> again. If a <i>SecureChannel</i> is not renewed, then all <i>Messages</i> sent using the current <i>SecurityTokens</i> shall be rejected by the receiver. Several cryptanalytic attacks become easier as more material encrypted with a specific key is available. By limiting the amount of data processed using a particular key, those attacks are made more difficult. Therefore the volume of data exchanged between <i>Client</i> and <i>Server</i> shall be limited by establishing a new <i>SecurityToken</i> after the lifetime. The setting of the requested lifetime depends on the expected number of exchanged messages and their size in the lifetime. A higher volume of data requires shorter lifetime.

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> type definition).
securityToken	ChannelSecurityToken	Describes the new <i>SecurityToken</i> issued by the <i>Server</i> . This structure is defined in-line with the following indented items.
channelId	~~ByteString~~ BaseDataType	A unique identifier for the <i>SecureChannel</i> . This is the identifier that shall be supplied whenever the <i>SecureChannel</i> is renewed. The concrete security protocol definition in IEC 62541-6 chooses the concrete <i>Data Type</i> .
tokenId	ByteString	A unique identifier for a single <i>SecurityToken</i> within the channel. This is the identifier that shall be passed with each <i>Message</i> secured with the <i>SecurityToken</i> .
createdAt	UtcTime	The time when the <i>SecurityToken</i> was created.
revisedLifetime	Duration	The lifetime of the <i>SecurityToken</i> in milliseconds. The UTC expiration time for the token may be calculated by adding the lifetime to the <i>createdAt</i> time.
serverNonce	ByteString	A random number that shall not be used in any other request. A new <i>serverNonce</i> shall be generated for each time a <i>SecureChannel</i> is renewed. ~~This parameter shall have a length equal to key size used for the symmetric encryption algorithm that is identified by the <i>securityPolicyUri</i>.~~ This parameter shall have a length equal to the <i>SecureChannelNonceLength</i> defined for the <i>SecurityPolicy</i> in IEC 62541-7. The <i>SecurityPolicy</i> is identified by the <i>securityPolicyUri</i> .

5.5.2.3 Service results

Table 12 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 12 – OpenSecureChannel Service result codes

Symbolic Id	Description
Bad_SecurityChecksFailed	See Table 177 for the description of this result code.
Bad_CertificateTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateIssuerTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateHostNameInvalid	See Table 177 for the description of this result code.
Bad_CertificateUriInvalid	See Table 177 for the description of this result code.
Bad_CertificateUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateIssuerUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateUntrusted	See Table 177 for the description of this result code.
Bad_CertificateRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateRevoked	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevoked	See Table 177 for the description of this result code.
Bad_RequestTypeInvalid	The security token request type is not valid.
Bad_SecurityModeRejected	The security mode does not meet the requirements set by the <i>Server</i> .
Bad_SecurityPolicyRejected	The security policy does not meet the requirements set by the <i>Server</i> .
Bad_SecureChannelIdInvalid	See Table 177 for the description of this result code.
Bad_NonceInvalid	See Table 177 for the description of this result code. A <i>Server</i> shall check the minimum length of the <i>Client</i> nonce and return this status if the length is below 32 bytes. A check for duplicated nonce can only be done in <i>OpenSecureChannel</i> calls with the request type RENEW_1.

5.5.3 CloseSecureChannel

5.5.3.1 Description

This *Service* is used to terminate a *SecureChannel*.

The request *Messages* shall be signed with the appropriate key associated with the current token for the *SecureChannel*.

5.5.3.2 Parameters

Table 13 defines the parameters for the *Service*.

Specific protocol mappings defined in IEC 62541-6 may choose to omit the response.

Table 13 – CloseSecureChannel Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The ~~sessionId~~ <i>authenticationToken</i> is always ~~emitted~~ null. The type <i>RequestHeader</i> is defined in 7.28.
secureChannelId	~~ByteString~~ BaseDataType	The identifier for the <i>SecureChannel</i> to close. The concrete security protocol definition in IEC 62541-6 chooses the concrete <i>Data Type</i> .
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

5.5.3.3 Service results

Table 14 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 14 – CloseSecureChannel Service result codes

Symbolic Id	Description
Bad_SecureChannelIdInvalid	See Table 177 for the description of this result code.

5.6 Session Service Set

5.6.1 Overview

This *Service Set* defines *Services* for an application layer connection establishment in the context of a *Session*.

5.6.2 CreateSession

5.6.2.1 Description

This *Service* is used by an OPC UA *Client* to create a *Session* and the *Server* returns two values which uniquely identify the *Session*. The first value is the *sessionId* which is used to identify the *Session* in the audit logs and in the *Server's AddressSpace*. The second is the *authenticationToken* which is used to associate an incoming request with a *Session*.

Before calling this *Service*, the *Client* shall create a *SecureChannel* with the *OpenSecureChannel Service* to ensure the *Integrity* of all *Messages* exchanged during a *Session*. This *SecureChannel* has a unique identifier which the *Server* shall associate with the *authenticationToken*. The *Server* may accept requests with the *authenticationToken* only if they are associated with the same *SecureChannel* that was used to create the *Session*. The *Client* may associate a new *SecureChannel* with the *Session* by calling the *ActivateSession* method.

The *SecureChannel* is always managed by the *Communication Stack* which means it shall provide APIs which the *Server* can use to find out information about the *SecureChannel* used for any given request. The *Communication Stack* shall, at a minimum, provide the *SecurityPolicy* and *SecurityMode* used by the *SecureChannel*. It shall also provide a *SecureChannelId* which uniquely identifies the *SecureChannel* or the *Client Certificate* used to establish the *SecureChannel*. The *Server* uses one of these to identify the *SecureChannel*.

used to send a request. Clause 7.31 describes how to create the *authenticationToken* for different types of *Communication Stack*.

Depending upon on the *SecurityPolicy* and the *SecurityMode* of the *SecureChannel*, the exchange of *ApplicationInstanceCertificates* and *Nonces* may be optional and the signatures may be empty. See IEC 62541-7 for the definition of *SecurityPolicies* and the handling of these parameters.

The *Server* returns its *EndpointDescriptions* in the response. *Clients* use this information to determine whether the list of *EndpointDescriptions* returned from the *DiscoveryEndpoint* matches the *Endpoints* that the *Server* has. If there is a difference then the *Client* shall close the *Session* and report an error. The *Server* returns all *EndpointDescriptions* for the *serverUri* specified by the *Client* in the request. The *Client* only verifies *EndpointDescriptions* with a *transportProfileUri* that matches the *profileUri* specified in the original *GetEndpoints* request. A *Client* may skip this check if the *EndpointDescriptions* were provided by a trusted source such as the *Administrator*.

The *Session* created with this *Service* shall not be used until the *Client* calls the *ActivateSession Service* and provides its *SoftwareCertificates* and proves possession of its *Application Instance Certificate* and any user identity token that it provided.

A *Server* application should limit the number of *Sessions*. To protect against misbehaving *Clients* and denial of service attacks, the *Server* shall close the oldest *Session* that is not activated before reaching the maximum number of supported *Sessions*.

The *SoftwareCertificates* parameter in the *Server* response is deprecated to reduce the message size for OPC UA Applications with limited resources. The *SoftwareCertificates* are provided in the *Server's AddressSpace* as defined in IEC 62541-5. A *SoftwareCertificate* identifies the capabilities of the *Server* and also contains the list of OPC UA *Profiles* supported by the *Server*. OPC UA *Profiles* are defined in IEC 62541-7.

Additional *Certificates* issued by other organizations may be included to identify additional *Server* capabilities. Examples of these *Profiles* include support for specific information models and support for access to specific types of devices.

When a *Session* is created, the *Server* adds an entry for the *Client* in its *SessionDiagnosticsArray Variable*. See IEC 62541-5 for a description of this *Variable*.

Sessions are created to be independent of the underlying communications connection. Therefore, if a communications connection fails, the *Session* is not immediately affected. The exact mechanism to recover from an underlying communication connection error depends on the *SecureChannel* mapping as described in IEC 62541-6.

Sessions are terminated by the *Server* automatically if the *Client* fails to issue a *Service* request on the *Session* within the timeout period negotiated by the *Server* in the *CreateSession Service* response. This protects the *Server* against *Client* failures and against situations where a failed underlying connection cannot be re-established. *Clients* shall be prepared to submit requests in a timely manner to prevent the *Session* from closing automatically. *Clients* may explicitly terminate *Sessions* using the *CloseSession Service*.

When a *Session* is terminated, all outstanding requests on the *Session* are aborted and *Bad_SessionClosed StatusCodes* are returned to the *Client*. In addition, the *Server* deletes the entry for the *Client* from its *SessionDiagnosticsArray Variable* and notifies any other *Clients* who were subscribed to this entry.

If a *Client* invokes the *CloseSession Service* then all *Subscriptions* associated with the *Session* are also deleted if the *deleteSubscriptions* flag is set to TRUE. If a *Server* terminates a *Session* for any other reason, *Subscriptions* associated with the *Session*, are not deleted.

Each *Subscription* has its own lifetime to protect against data loss in the case of a *Session* termination. In these cases, the *Subscription* can be reassigned to another *Client* before its lifetime expires.

Some *Servers*, such as aggregating *Servers*, also act as *Clients* to other *Servers*. These *Servers* typically support more than one system user, acting as their agent to the *Servers* that they represent. Security for these *Servers* is supported at two levels.

First, each OPC UA *Service* request contains a string parameter that is used to carry an audit record id. A *Client*, or any *Server* operating as a *Client*, such as an aggregating *Server*, can create a local audit log entry for a request that it submits. This parameter allows the *Client* to pass the identifier for this entry with the request. If the *Server* also maintains an audit log, then it can include this id in the audit log entry that it writes. When the log is examined and the entry is found, the examiner will be able to relate it directly to the audit log entry created by the *Client*. This capability allows for traceability across audit logs within a system. See IEC TR 62541-2 for additional information on auditing. A *Server* that maintains an audit log shall provide the information in the audit log entries via event *Messages* defined in this document. The *Server* may choose to only provide the *Audit* information via event *Messages*. The *Audit EventType* is defined in IEC 62541-3.

Second, these aggregating *Servers* may open independent *Sessions* to the underlying *Servers* for each *Client* that accesses data from them. Figure 14 illustrates this concept.




Figure 14 – Multiplexing users on a Session

5.6.2.2 Parameters

Table 15 defines the parameters for the *Service*.

Table 15 – CreateSession Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always ~~omitted~~ null. The type <i>RequestHeader</i> is defined in 7.28.
clientDescription	Application Description	Information that describes the <i>Client</i> application. The type <i>ApplicationDescription</i> is defined in 7.1.
serverUri	String	This value is only specified if the <i>EndpointDescription</i> has a <i>gatewayServerUri</i> . This value is the <i>applicationUri</i> from the <i>EndpointDescription</i> which is the <i>applicationUri</i> for the underlying <i>Server</i> . The type <i>EndpointDescription</i> is defined in 7.10.
endpointUrl	String	The network address that the <i>Client</i> used to access the <i>Session Endpoint</i> . The <i>HostName</i> portion of the URL should be one of the <i>HostNames</i> for the application that are specified in the <i>Server's ApplicationInstanceCertificate</i> (see 7.2). The <i>Server</i> shall raise an <i>AuditUrlMismatchEventType</i> event if the URL does not match the <i>Server's HostNames</i> . <i>AuditUrlMismatchEventType</i> event type is defined in IEC 62541-5. The <i>Server</i> uses this information for diagnostics and to determine the set of <i>EndpointDescriptions</i> to return in the response.
sessionName	String	Human readable string that identifies the <i>Session</i> . The <i>Server</i> makes this name and the <i>sessionId</i> visible in its <i>AddressSpace</i> for diagnostic purposes. The <i>Client</i> should provide a name that is unique for the instance of the <i>Client</i> . If this parameter is not specified the <i>Server</i> shall assign a value.
clientNonce	ByteString	A random number that should never be used in any other request. This number shall have a minimum length of 32 bytes. Profiles may increase the required length. The <i>Server</i> shall use this value to prove possession of its <i>Application Instance Certificate</i> in the response.
clientCertificate	ApplicationInstance Certificate	The <i>Application Instance Certificate</i> issued to the <i>Client</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2. If the <i>securityPolicyUri</i> is None, the <i>Server</i> shall ignore the <i>ApplicationInstanceCertificate</i> .
Requested SessionTimeout	Duration	Requested maximum number of milliseconds that a <i>Session</i> should remain open without activity. If the <i>Client</i> fails to issue a <i>Service</i> request within this interval, then the <i>Server</i> shall automatically terminate the <i>Client Session</i> .
maxResponse MessageSize	UInt32	The maximum size, in bytes, for the body of any response message. The <i>Server</i> should return a <i>Bad_ResponseTooLarge</i> service fault if a response message exceeds this limit. The value zero indicates that this parameter is not used. The transport protocols defined in IEC 62541-6 may imply minimum message sizes. More information on the use of this parameter is provided in 5.3.

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> type).
sessionId	NodeId	A unique <i>NodeId</i> assigned by the <i>Server</i> to the <i>Session</i> . This identifier is used to access the diagnostics information for the <i>Session</i> in the <i>Server AddressSpace</i> . It is also used in the audit logs and any events that report information related to the <i>Session</i> . The <i>Session</i> diagnostic information is described in IEC 62541-5. Audit logs and their related events are described in 6.5.
authentication Token	Session AuthenticationToken	A unique identifier assigned by the <i>Server</i> to the <i>Session</i> . This identifier shall be passed in the <i>RequestHeader</i> of each request and is used with the <i>SecureChannelId</i> to determine whether a <i>Client</i> has access to the <i>Session</i> . This identifier shall not be reused in a way that the <i>Client</i> or the <i>Server</i> has a chance of confusing them with a previous or existing <i>Session</i> . The <i>SessionAuthenticationToken</i> type is described in 7.31.
revisedSession Timeout	Duration	Actual maximum number of milliseconds that a <i>Session</i> shall remain open without activity. The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.
serverNonce	ByteString	A random number that should never be used in any other request. This number shall have a minimum length of 32 bytes. The <i>Client</i> shall use this value to prove possession of its <i>Application Instance Certificate</i> in the <i>ActivateSession</i> request. This value may also be used to prove possession of the <i>userIdentityToken</i> it specified in the <i>ActivateSession</i> request.
serverCertificate	ApplicationInstance Certificate	The <i>Application Instance Certificate</i> issued to the <i>Server</i> . A <i>Server</i> shall prove possession by using the private key to sign the <i>Nonce</i> provided by the <i>Client</i> in the request. The <i>Client</i> shall verify that this <i>Certificate</i> is the same as the one it used to create the <i>SecureChannel</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2. If the <i>securityPolicyUri</i> is NONE and none of the <i>UserTokenPolicies</i> requires encryption, the <i>Client</i> shall ignore the <i>ApplicationInstanceCertificate</i> .
serverEndpoints []	EndpointDescription	List of <i>Endpoints</i> that the <i>Server</i> supports. The <i>Server</i> shall return a set of <i>EndpointDescriptions</i> available for the <i>serverUri</i> specified in the request. The <i>EndpointDescription</i> type is defined in 7.10. The <i>Client</i> shall verify this list with the list from a <i>DiscoveryEndpoint</i> if it used a <i>DiscoveryEndpoint</i> to fetch the <i>EndpointDescriptions</i> . It is recommended that <i>Servers</i> only include the <i>server.applicationUri</i> , <i>endpointUrl</i> , <i>securityMode</i> , <i>securityPolicyUri</i> , <i>userIdentityTokens</i> , <i>transportProfileUri</i> and <i>securityLevel</i> with all other parameters set to null. Only the recommended parameters shall be verified by the client.
serverSoftware Certificates []	SignedSoftware Certificate	This parameter is deprecated and the array shall be empty. The <i>SoftwareCertificates</i> are provided in the <i>Server AddressSpace</i> as defined in IEC 62541-5.
serverSignature	SignatureData	This is a signature generated with the private key associated with the <i>serverCertificate</i> . This parameter is calculated by appending the <i>clientNonce</i> to the <i>clientCertificate</i> and signing the resulting sequence of bytes. If the <i>clientCertificate</i> contains a chain, the signature calculation shall be done only with the leaf <i>Certificate</i> . For backward compatibility a <i>Client</i> shall check the signature with the full chain if the check with the leaf <i>Certificate</i> fails. The <i>SignatureAlgorithm</i> shall be the <i>AsymmetricSignatureAlgorithm</i> specified in the <i>SecurityPolicy</i> for the <i>Endpoint</i> . The <i>SignatureData</i> type is defined in 7.32.
maxRequest MessageSize	UInt32	The maximum size, in bytes, for the body of any request message. The <i>Client Communication Stack</i> should return a <i>Bad_RequestTooLarge</i> error to the application if a request message exceeds this limit. The value zero indicates that this parameter is not used. See IEC 62541-6 for protocol specific minimum or default values. 5.3 provides more information on the use of this parameter.

5.6.2.3 Service results

Table 16 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 16 – CreateSession Service result codes

Symbolic Id	Description
Bad_SecureChannelIdInvalid	See Table 177 for the description of this result code.
Bad_NonceInvalid	See Table 177 for the description of this result code. A <i>Server</i> shall check the minimum length of the <i>Client</i> nonce and return this status if the length is below 32 bytes. A check for a duplicated nonce is optional and requires access to the nonce used to create the secure channel.
Bad_SecurityChecksFailed	See Table 177 for the description of this result code.
Bad_CertificateTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateIssuerTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateHostNameInvalid	See Table 177 for the description of this result code.
Bad_CertificateUriInvalid	See Table 177 for the description of this result code.
Bad_CertificateUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateIssuerUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateUntrusted	See Table 177 for the description of this result code.
Bad_CertificateRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateRevoked	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevoked	See Table 177 for the description of this result code.
Bad_TooManySessions	The <i>Server</i> has reached its maximum number of <i>Sessions</i> .
Bad_ServerUriInvalid	See Table 177 for the description of this result code.

5.6.3 ActivateSession

5.6.3.1 Description

This *Service* is used by the *Client* ~~to submit its SoftwareCertificates to the Server for validation and~~ to specify the identity of the user associated with the *Session*. This *Service* request shall be issued by the *Client* before it issues any ~~other~~ *Service* request ~~other than~~ *CloseSession* after *CreateSession*. Failure to do so shall cause the *Server* to close the *Session*.

Whenever the *Client* calls this *Service* the *Client* shall prove that it is the same application that called the *CreateSession Service*. The *Client* does this by creating a signature with the private key associated with the *clientCertificate* specified in the *CreateSession* request. This signature is created by appending the last *serverNonce* provided by the *Server* to the *serverCertificate* and calculating the signature of the resulting sequence of bytes.

Once used, a *serverNonce* cannot be used again. For that reason, the *Server* returns a new *serverNonce* each time the *ActivateSession Service* is called.

When the *ActivateSession Service* is called for the first time then the *Server* shall reject the request if the *SecureChannel* is not same as the one associated with the *CreateSession* request. Subsequent calls to *ActivateSession* may be associated with different *SecureChannels*. If this is the case then the *Server* shall verify that the *Certificate* the *Client* used to create the new *SecureChannel* is the same as the *Certificate* used to create the original *SecureChannel*. In addition, the *Server* shall verify that the *Client* supplied a *UserIdentityToken* that is identical to the token currently associated with the *Session*. Once

the *Server* accepts the new *SecureChannel* it shall reject requests sent via the old *SecureChannel*.

The *ActivateSession Service* is used to associate a user identity with a *Session*. When a *Client* provides a user identity then it shall provide proof that it is authorized to use that user identity. The exact mechanism used to provide this proof depends on the type of the *UserIdentityToken*. If the token is a *UserNameIdentityToken* then the proof is the *password* that is included in the token. If the token is an *X509IdentityToken* then the proof is a signature generated with private key associated with the *Certificate*. The data to sign is created by appending the last *serverNonce* to the *serverCertificate* specified in the *CreateSession* response. If a token includes a secret then it should be encrypted using the public key from the *serverCertificate*.

Servers shall take proper measures to protect against attacks on user identity tokens. Such an attack is assumed if repeated connection attempts with invalid user identity tokens happen. One option is to lock out an OPC UA *Client* for a period of time if the user identity token validation fails several times. The OPC UA *Client* is either detected by IP address for unsecured connections or by the *ApplicationInstanceUri* for secured connections. Another option is delaying the *Service* response when the validation of a user identity fails. This delay time could be increased with repeated failures. Sporadic failures shall not delay connections with valid tokens.

Clients can change the identity of a user associated with a *Session* by calling the *ActivateSession Service*. The *Server* validates the signatures provided with the request and then validates the new user identity. If no errors occur the *Server* replaces the user identity for the *Session*. Changing the user identity for a *Session* may cause discontinuities in active *Subscriptions* because the *Server* may ~~have~~ need to tear down connections to an underlying system and re-establish them using the new credentials.

When a *Client* supplies a list of locale ids in the request, each locale id is required to contain the language component. It may optionally contain the <country/region> component. When the *Server* returns a *LocalizedText* in the context of the *Session*, it also may return both the language and the country/region or just the language as its default locale id.

When a *Server* returns a string to the *Client*, it first determines if there are available translations for it. If there are, then the *Server* returns the string whose locale id exactly matches the locale id with the highest priority in the *Client*-supplied list.

If there are no exact matches, then the *Server* ignores the <country/region> component of the locale id, and returns the string whose <language> component matches the <language> component of the locale id with the highest priority in the *Client* supplied list.

If there still are no matches, then the *Server* returns the string that it has along with the locale id.

A *Gateway Server* is expected to impersonate the user provided by the *Client* when it connects to the underlying *Server*. This means it shall re-calculate the signatures on the *UserIdentityToken* using the nonce provided by the underlying *Server*. The *Gateway Server* will ~~have~~ need to use its own user credentials if the *UserIdentityToken* provided by the *Client* does not support impersonation.

5.6.3.2 Parameters

Table 17 defines the parameters for the *Service*.

Table 17 – ActivateSession Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The type <i>RequestHeader</i> is defined in 7.28.
clientSignature	SignatureData	This is a signature generated with the private key associated with the <i>clientCertificate</i>. This parameter is calculated by appending the <i>serverNonce</i> to the <i>serverCertificate</i> and signing the resulting sequence of bytes. If the <i>serverCertificate</i> contains a chain, the signature calculation shall be done only with the leaf <i>Certificate</i>. For backward compatibility a <i>Server</i> shall check the signature with the full chain if the check with the leaf <i>Certificate</i> fails. The <i>SignatureAlgorithm</i> shall be the <i>AsymmetricSignatureAlgorithm</i> specified in the <i>SecurityPolicy</i> for the <i>Endpoint</i>. The <i>SignatureData</i> type is defined in 7.32.
clientSoftwareCertificates []	SignedSoftwareCertificate	~~These are the <i>SoftwareCertificates</i> which have been issued to the <i>Client</i> application.~~ ~~The <i>productUri</i> contained in the <i>SoftwareCertificates</i> shall match the <i>productUri</i> in the <i>ApplicationDescription</i> passed by the <i>Client</i> in the <i>CreateSession</i> requests. <i>Certificates</i> without matching <i>productUri</i> should be ignored.~~ ~~<i>Servers</i> may reject connections from <i>Clients</i> if they are not satisfied with the <i>SoftwareCertificates</i> provided by the <i>Client</i>.~~ ~~This parameter only needs to be specified in the first <i>ActivateSession</i> request after <i>CreateSession</i>. It shall always be omitted if the <i>maxRequestMessageSize</i> returned from the <i>Server</i> in the <i>CreateSession</i> response is less than one megabyte.~~ Reserved for future use. The <i>SignedSoftwareCertificate</i> type is defined in 7.33.
localeIds []	LocaleId	List of locale ids in priority order for localized strings. The first <i>LocaleId</i> in the list has the highest priority. If the <i>Server</i> returns a localized string to the <i>Client</i>, the <i>Server</i> shall return the translation with the highest priority that it can. If it does not have a translation for any of the locales identified in this list, then it shall return the string value that it has and include the locale id with the string. See IEC 62541-3 for more detail on locale ids. If the <i>Client</i> fails to specify at least one locale id, the <i>Server</i> shall use any that it has. This parameter only needs to be specified during the first call to <i>ActivateSession</i> during a single application <i>Session</i>. If it is not specified the <i>Server</i> shall keep using the current <i>localeIds</i> for the <i>Session</i>.
userIdentityToken	ExtensibleParameter UserIdentityToken	The credentials of the user associated with the <i>Client</i> application. The <i>Server</i> uses these credentials to determine whether the <i>Client</i> should be allowed to activate a <i>Session</i> and what resources the <i>Client</i> has access to during this <i>Session</i>. The <i>UserIdentityToken</i> is an extensible parameter type defined in 7.36. The <i>EndpointDescription</i> specifies what <i>UserIdentityTokens</i> the <i>Server</i> shall accept. Null or empty user token shall always be interpreted as anonymous.
userTokenSignature	SignatureData	If the <i>Client</i> specified a user identity token that supports digital signatures, then it shall create a signature and pass it as this parameter. Otherwise the parameter is ~~omitted~~ null. The <i>SignatureAlgorithm</i> depends on the identity token type. The <i>SignatureData</i> type is defined in 7.32.

Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
serverNonce	ByteString	A random number that should never be used in any other request. This number shall have a minimum length of 32 bytes. The <i>Client</i> shall use this value to prove possession of its <i>Application Instance Certificate</i> in the next call to <i>ActivateSession</i> request.
results []	StatusCode	List of validation results for the <i>SoftwareCertificates</i> (see 7.34 for <i>StatusCode</i> definition).
diagnosticInfos []	DiagnosticInfo	List of diagnostic information associated with <i>SoftwareCertificate</i> validation errors (see 7.8 for <i>DiagnosticInfo</i> definition). This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.6.3.3 Service results

Table 18 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 18 – ActivateSession Service result codes

Symbolic Id	Description
Bad_IdentityTokenInvalid	See Table 177 for the description of this result code.
Bad_IdentityTokenRejected	See Table 177 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_ApplicationSignatureInvalid	The signature provided by the <i>Client</i> application is missing or invalid.
Bad_UserSignatureInvalid	The user token signature is missing or invalid.
Bad_NoValidCertificates	The <i>Client</i> did not provide at least one <i>Software Certificate</i> that is valid and meets the profile requirements for the <i>Server</i> .
Bad_IdentityChangeNotSupported	The <i>Server</i> does not support changing the user identity assigned to the session.

5.6.4 CloseSession

5.6.4.1 Description

This *Service* is used to terminate a *Session*. The *Server* takes the following actions when it receives a *CloseSession* request:

- It stops accepting requests for the *Session*. All subsequent requests received for the *Session* are discarded.
- It returns negative responses with the *StatusCode* *Bad_SessionClosed* to all requests that are currently outstanding to provide for the timely return of the *CloseSession* response. *Clients* are urged to wait for all outstanding requests to complete before submitting the *CloseSession* request.
- It removes the entry for the *Client* in its *SessionDiagnosticsArray Variable*.

When the *CloseSession Service* is called before the *Session* is successfully activated, the *Server* shall reject the request if the *SecureChannel* is not the same as the one associated with the *CreateSession* request.

5.6.4.2 Parameters

Table 19 defines the parameters for the *Service*.

Table 19 – CloseSession Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
deleteSubscriptions	Boolean	If the value is TRUE, the Server deletes all Subscriptions associated with the Session. If the value is FALSE, the Server keeps the Subscriptions associated with the Session until they timeout based on their own lifetime.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

5.6.4.3 Service results

Table 20 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 20 – CloseSession Service result codes

Symbolic Id	Description
Bad_SessionIdInvalid	See Table 177 for the description of this result code.

5.6.5 Cancel

5.6.5.1 Description

This *Service* is used to cancel outstanding *Service* requests. Successfully cancelled service requests shall respond with *Bad_RequestCancelledByClient*.

5.6.5.2 Parameters

Table 21 defines the parameters for the *Service*.

Table 21 – Cancel Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
requestHandle	IntegerId	The <i>requestHandle</i> assigned to one or more requests that should be cancelled. All outstanding requests with the matching <i>requestHandle</i> shall be cancelled.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
cancelCount	UInt32	Number of cancelled requests.

5.6.5.3 Service results

Common *StatusCodes* are defined in Table 177.

5.7 NodeManagement Service Set

5.7.1 Overview

This *Service Set* defines *Services* to add and delete *AddressSpace Nodes* and *References* between them. All added *Nodes* continue to exist in the *AddressSpace* even if the *Client* that created them disconnects from the *Server*.

5.7.2 AddNodes

5.7.2.1 Description

This *Service* is used to add one or more *Nodes* into the *AddressSpace* hierarchy. Using this *Service*, each *Node* is added as the *TargetNode* of a *HierarchicalReference* to ensure that the *AddressSpace* is fully connected and that the *Node* is added as a child within the *AddressSpace* hierarchy (see IEC 62541-3).

When a *Server* creates an instance of a *TypeDefinitionNode* it shall create the same hierarchy of *Nodes* beneath the new *Object* or *Variable* depending on the *ModellingRule* of each *InstanceDeclaration*. All *Nodes* with a *ModellingRule* of *Mandatory* shall be created or an existing *Node* shall be referenced that conforms to the *InstanceDeclaration*. The creation of *Nodes* with other *ModellingRules* is *Server* specific.

5.7.2.2 Parameters

Table 22 defines the parameters for the *Service*.

Table 22 – AddNodes Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToAdd []	AddNodesItem	List of <i>Nodes</i> to add. All <i>Nodes</i> are added as a <i>Reference</i> to an existing <i>Node</i> using a hierarchical <i>ReferenceType</i> . This structure is defined in-line with the following indented items.
parentNodeId	Expanded NodeId	<i>ExpandedNodeId</i> of the parent <i>Node</i> for the <i>Reference</i> . The <i>ExpandedNodeId</i> type is defined in 7.11.
referenceTypeId	NodeId	<i>NodeId</i> of the hierarchical <i>ReferenceType</i> to use for the <i>Reference</i> from the parent <i>Node</i> to the new <i>Node</i> .
requestedNewNodeId	Expanded NodeId	<i>Client</i> requested expanded <i>NodeId</i> of the <i>Node</i> to add. The <i>serverIndex</i> in the expanded <i>NodeId</i> shall be 0. If the <i>Server</i> cannot use this <i>NodeId</i>, it rejects this <i>Node</i> and returns the appropriate error code. If the <i>Client</i> does not want to request a <i>NodeId</i>, then it sets the value of this parameter to the null expanded <i>NodeId</i>. If the <i>Node</i> to add is a <i>ReferenceType Node</i>, its <i>NodeId</i> should be a numeric id. See IEC 62541-3 for a description of <i>ReferenceType NodeIds</i>.
browseName	QualifiedName	The browse name of the <i>Node</i> to add.
nodeClass	NodeClass	<i>NodeClass</i> of the <i>Node</i> to add.
nodeAttributes	Extensible Parameter NodeAttributes	The <i>Attributes</i> that are specific to the <i>NodeClass</i>. The <i>NodeAttributes</i> parameter type is an extensible parameter type specified in 7.19. A <i>Client</i> is allowed to omit values for some or all <i>Attributes</i>. If an <i>Attribute</i> value is omitted null, the <i>Server</i> shall use the default values from the <i>TypeDefinitionNode</i>. If a <i>TypeDefinitionNode</i> was not provided the <i>Server</i> shall choose a suitable default value. The <i>Server</i> may still add an optional <i>Attribute</i> to the <i>Node</i> with an appropriate default value even if the <i>Client</i> does not specify a value.
typeDefinition	Expanded NodeId	<i>NodeId</i> of the <i>TypeDefinitionNode</i> for the <i>Node</i> to add. This parameter shall

		be null for all <i>NodeClasses</i> other than <i>Object</i> and <i>Variable</i> in which case it shall be provided.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	AddNodesResult	List of results for the <i>Nodes</i> to add. The size and order of the list matches the size and order of the <i>nodesToAdd</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the <i>Node</i> to add (see 7.34 for <i>StatusCode</i> definition).
addedNodeId	NodeId	<i>Server</i> assigned <i>NodeId</i> of the added <i>Node</i> . Null <i>NodeId</i> if the operation failed.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Nodes</i> to add (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>nodesToAdd</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.2.3 Service results

Table 23 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 23 – AddNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.2.4 StatusCodes

Table 24 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 24 – AddNodes Operation Level result codes

Symbolic Id	Description
Bad_ParentNodeIdInvalid	The parent node id does not to refer to a valid node.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_ReferenceNotAllowed	The reference could not be created because it violates constraints imposed by the data model.
Bad_NodeIdRejected	The requested node id was rejected either because it was invalid or because the <i>Server</i> does not allow node ids to be specified by the client.
Bad_NodeIdExists	The requested node id is already used by another node.
Bad_NodeClassInvalid	See Table 178 for the description of this result code.
Bad_BrowseNameInvalid	See Table 178 for the description of this result code.
Bad_BrowseNameDuplicated	The browse name is not unique among nodes that share the same relationship with the parent.
Bad_NodeAttributesInvalid	The node <i>Attributes</i> are not valid for the node class.
Bad_TypeDefinitionInvalid	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.

5.7.3 AddReferences

5.7.3.1 Description

This *Service* is used to add one or more *References* to one or more *Nodes*. The *NodeClass* is an input parameter that is used to validate that the *Reference* to be added matches the *NodeClass* of the *TargetNode*. This parameter is not validated if the *Reference* refers to a *TargetNode* in a remote *Server*.

In certain cases, adding new *References* to the *AddressSpace* shall require that the *Server* add new *Server* ids to the *Server's ServerArray Variable*. For this reason, remote *Servers* are identified by their URI and not by their *ServerArray* index. This allows the *Server* to add the remote *Server* URIs to its *ServerArray*.

5.7.3.2 Parameters

Table 25 defines the parameters for the *Service*.

Table 25 – AddReferences Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
referencesToAdd []	AddReferences Item	List of <i>Reference</i> instances to add to the <i>SourceNode</i> . The <i>targetNodeClass</i> of each <i>Reference</i> in the list shall match the <i>NodeClass</i> of the <i>TargetNode</i> . This structure is defined in-line with the following indented items.
sourceNodeid	Nodeid	<i>Nodeid</i> of the <i>Node</i> to which the <i>Reference</i> is to be added. The source <i>Node</i> shall always exist in the <i>Server</i> to add the <i>Reference</i> . The <i>isForward</i> parameter can be set to FALSE if the target <i>Node</i> is on the local <i>Server</i> and the source <i>Node</i> on the remote <i>Server</i> .
referenceTypeid	Nodeid	<i>Nodeid</i> of the <i>ReferenceType</i> that defines the <i>Reference</i> .
isForward	Boolean	If the value is TRUE, the <i>Server</i> creates a forward <i>Reference</i> . If the value is FALSE, the <i>Server</i> creates an inverse <i>Reference</i> .
targetServerUri	String	URI of the remote <i>Server</i> . If this parameter is not null, it overrides the <i>serverIndex</i> in the <i>targetNodeid</i> .
targetNodeid	Expanded Nodeid	Expanded <i>Nodeid</i> of the <i>TargetNode</i> . The <i>ExpandedNodeid</i> type is defined in 7.11.
targetNodeClass	NodeClass	<i>NodeClass</i> of the <i>TargetNode</i> . The <i>Client</i> shall specify this since the <i>TargetNode</i> might not be accessible directly by the <i>Server</i> .
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	Status Code	List of <i>StatusCodes</i> for the <i>References</i> to add (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>referencesToAdd</i> request parameter.
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>References</i> to add (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>referencesToAdd</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.3.3 Service results

Table 26 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 26 – AddReferences Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.3.4 StatusCodes

Table 27 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 27 – AddReferences Operation Level result codes

Symbolic Id	Description
Bad_SourceNodeIdInvalid	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_ServerUriInvalid	See Table 177 for the description of this result code.
Bad_TargetNodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeClassInvalid	See Table 178 for the description of this result code.
Bad_ReferenceNotAllowed	The reference could not be created because it violates constraints imposed by the data model on this <i>Server</i> .
Bad_ReferenceLocalOnly	The reference type is not valid for a reference to a remote <i>Server</i> .
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_DuplicateReferenceNotAllowed	The reference type between the nodes is already defined.
Bad_InvalidSelfReference	The <i>Server</i> does not allow this type of self reference on this node.

5.7.4 DeleteNodes

5.7.4.1 Description

This *Service* is used to delete one or more *Nodes* from the *AddressSpace*.

When any of the *Nodes* deleted by an invocation of this *Service* is the *TargetNode* of a *Reference*, then those *References* are left unresolved based on the *deleteTargetReferences* parameter.

When any of the *Nodes* deleted by an invocation of this *Service* is being monitored, then a *Notification* containing the status code *Bad_NodeIdUnknown* is sent to the monitoring *Client* indicating that the *Node* has been deleted.

5.7.4.2 Parameters

Table 28 defines the parameters for the *Service*.

Table 28 – DeleteNodes Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToDelete []	DeleteNodes Item	List of <i>Nodes</i> to delete. This structure is defined in-line with the following indented items.
nodeId	NodeId	<i>NodeId</i> of the <i>Node</i> to delete.
deleteTargetReferences	Boolean	A <i>Boolean</i> parameter with the following values: TRUE delete <i>References</i> in <i>TargetNodes</i> that <i>Reference</i> the <i>Node</i> to delete. FALSE delete only the <i>References</i> for which the <i>Node</i> to delete is the source. The <i>Server</i> cannot guarantee that it is able to delete all <i>References</i> from <i>TargetNodes</i> if this parameter is TRUE.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCodes	List of <i>StatusCodes</i> for the <i>Nodes</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the list of the <i>nodesToDelete</i> request parameter.
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>Nodes</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>nodesToDelete</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.4.3 Service results

Table 29 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 29 – DeleteNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.4.4 StatusCodes

Table 30 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 30 – DeleteNodes Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_NoDeleteRights	See Table 178 for the description of this result code.
Uncertain_ReferenceNotDeleted	The <i>Server</i> was not able to delete all target references.

5.7.5 DeleteReferences

5.7.5.1 Description

This *Service* is used to delete one or more *References* of a *Node*.

When any of the *References* deleted by an invocation of this *Service* are contained in a *View*, then the *ViewVersion Property* is updated if this *Property* is supported.

Table 31 defines the parameters for the *Service*.

Table 31 – DeleteReferences Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
referencesToDelete []	DeleteReferences Item	List of <i>References</i> to delete. This structure is defined in-line with the following indented items.
sourceNodeId	NodeId	<i>NodeId</i> of the <i>Node</i> that contains the <i>Reference</i> to delete.
referenceTypeId	NodeId	<i>NodeId</i> of the <i>ReferenceType</i> that defines the <i>Reference</i> to delete.
isForward	Boolean	If the value is TRUE, the <i>Server</i> deletes a forward <i>Reference</i> . If the value is FALSE, the <i>Server</i> deletes an inverse <i>Reference</i> .
targetNodeId	ExpandedNodeId	<i>NodeId</i> of the <i>TargetNode</i> of the <i>Reference</i> . If the <i>Server</i> index indicates that the <i>TargetNode</i> is a remote <i>Node</i> , then the <i>nodeId</i> shall contain the absolute namespace URI. If the <i>TargetNode</i> is a local <i>Node</i> the <i>nodeId</i> shall contain the namespace index.
deleteBidirectional	Boolean	A <i>Boolean</i> parameter with the following values: TRUE delete the specified <i>Reference</i> and the opposite <i>Reference</i> from the <i>TargetNode</i> . If the <i>TargetNode</i> is located in a remote <i>Server</i> , the <i>Server</i> is permitted to delete the specified <i>Reference</i> only. FALSE delete only the specified <i>Reference</i> .
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of <i>StatusCodes</i> for the <i>References</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>referencesToDelete</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>References</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>referencesToDelete</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.5.2 Service results

Table 32 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 32 – DeleteReferences Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.5.3 StatusCodes

Table 33 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 33 – DeleteReferences Operation Level result codes

Symbolic Id	Description
Bad_SourceNodeIdInvalid	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_ServerIndexInvalid	The <i>Server</i> index is not valid.
Bad_TargetNodeIdInvalid	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_NoDeleteRights	See Table 178 for the description of this result code.

5.8 View Service Set

5.8.1 Overview

Clients use the browse *Services* of the *View Service Set* to navigate through the *AddressSpace* or through a *View* which is a subset of the *AddressSpace*.

A *View* is a subset of the *AddressSpace* created by the *Server*. Future versions of IEC 62541-4 may also define services to create *Client*-defined *Views*. See IEC 62541-5 for a description of the organization of views in the *AddressSpace*.

5.8.2 Browse

5.8.2.1 Description

This *Service* is used to discover the *References* of a specified *Node*. The browse can be further limited by the use of a *View*. This Browse *Service* also supports a primitive filtering capability.

In some cases it may take longer than the *Client* timeout hint to process all nodes to browse. In this case the *Server* may return zero results with a continuation point for the affected nodes before the timeout expires.

5.8.2.2 Parameters

Table 34 defines the parameters for the *Service*.

Table 34 – Browse Service parameters

Name	Type	Description																		
Request																				
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).																		
View	ViewDescription	Description of the <i>View</i> to browse (see 7.39 for <i>ViewDescription</i> definition). An empty <i>ViewDescription</i> value indicates the entire <i>AddressSpace</i> . Use of the empty <i>ViewDescription</i> value causes all <i>References</i> of the <i>nodesToBrowse</i> to be returned. Use of any other <i>View</i> causes only the <i>References</i> of the <i>nodesToBrowse</i> that are defined for that <i>View</i> to be returned.																		
requestedMax ReferencesPerNode	Counter	Indicates the maximum number of references to return for each starting <i>Node</i> specified in the request. The value 0 indicates that the <i>Client</i> is imposing no limitation (see 7.5 for <i>Counter</i> definition).																		
nodesToBrowse []	BrowseDescription	A list of nodes to Browse. This structure is defined in-line with the following indented items.																		
nodeId	NodeId	<i>NodeId</i> of the <i>Node</i> to be browsed. If a <i>view</i> is provided, it shall include this <i>Node</i> .																		
browseDirection	Enum BrowseDirection	An enumeration that specifies the direction of <i>References</i> to follow. It has the following values: FORWARD_0 select only forward <i>References</i> . INVERSE_1 select only inverse <i>References</i> . BOTH_2 select forward and inverse <i>References</i> . The returned <i>References</i> do indicate the direction the <i>Server</i> followed in the <i>isForward</i> parameter of the <i>ReferenceDescription</i> . Symmetric <i>References</i> are always considered to be in forward direction therefore the <i>isForward</i> flag is always set to TRUE and symmetric <i>References</i> are not returned if <i>browseDirection</i> is set to <i>INVERSE_1</i> .																		
referenceTypeId	NodeId	Specifies the <i>NodeId</i> of the <i>ReferenceType</i> to follow. Only instances of this <i>ReferenceType</i> or its subtypes are returned. If not specified then all <i>References</i> are returned and <i>includeSubtypes</i> is ignored.																		
includeSubtypes	Boolean	Indicates whether subtypes of the <i>ReferenceType</i> should be included in the browse. If TRUE, then instances of <i>referenceTypeId</i> and all of its subtypes are returned.																		
nodeClassMask	UInt32	Specifies the <i>NodeClasses</i> of the <i>TargetNodes</i> . Only <i>TargetNodes</i> with the selected <i>NodeClasses</i> are returned. The <i>NodeClasses</i> are assigned the following bits: <table><tr><th>Bit</th><th>NodeClass</th></tr><tr><td>0</td><td>Object</td></tr><tr><td>1</td><td>Variable</td></tr><tr><td>2</td><td>Method</td></tr><tr><td>3</td><td>ObjectType</td></tr><tr><td>4</td><td>VariableType</td></tr><tr><td>5</td><td>ReferenceType</td></tr><tr><td>6</td><td>DataType</td></tr><tr><td>7</td><td>View</td></tr></table> If set to zero, then all <i>NodeClasses</i> are returned. If the <i>NodeClass</i> is unknown for a remote <i>Node</i> , the <i>nodeClassMask</i> is ignored.	Bit	NodeClass	0	Object	1	Variable	2	Method	3	ObjectType	4	VariableType	5	ReferenceType	6	DataType	7	View
Bit	NodeClass																			
0	Object																			
1	Variable																			
2	Method																			
3	ObjectType																			
4	VariableType																			
5	ReferenceType																			
6	DataType																			
7	View																			

Name	Type	Description														
resultMask	UInt32	Specifies the fields in the <i>ReferenceDescription</i> structure that should be returned. The fields are assigned the following bits: <table><tr><th>Bit</th><th>Result</th></tr><tr><td>0</td><td>ReferenceType</td></tr><tr><td>1</td><td>IsForward</td></tr><tr><td>2</td><td>NodeClass</td></tr><tr><td>3</td><td>BrowseName</td></tr><tr><td>4</td><td>DisplayName</td></tr><tr><td>5</td><td>TypeDefinition</td></tr></table> The <i>ReferenceDescription</i> type is defined in 7.25.	Bit	Result	0	ReferenceType	1	IsForward	2	NodeClass	3	BrowseName	4	DisplayName	5	TypeDefinition
Bit	Result															
0	ReferenceType															
1	IsForward															
2	NodeClass															
3	BrowseName															
4	DisplayName															
5	TypeDefinition															
Response																
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).														
results []	BrowseResult	A list of <i>BrowseResults</i>. The size and order of the list matches the size and order of the <i>nodesToBrowse</i> specified in the request. The <i>BrowseResult</i> type is defined in 7.3.														
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>results</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>results</i> response parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.														

5.8.2.3 Service results

Table 35 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 35 – Browse Service result codes

Symbolic Id	Description
Bad_ViewIdUnknown	See Table 177 for the description of this result code.
Bad_ViewTimestampInvalid	See Table 177 for the description of this result code.
Bad_ViewParameterMismatchInvalid	See Table 177 for the description of this result code.
Bad_ViewVersionInvalid	See Table 177 for the description of this result code.
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.8.2.4 StatusCodes

Table 36 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 36 – Browse Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_BrowseDirectionInvalid	See Table 178 for the description of this result code.
Bad_NodeNotInView	See Table 178 for the description of this result code.
Bad_NoContinuationPoints	See Table 178 for the description of this result code.
Uncertain_NotAllNodesAvailable	Browse results may be incomplete because of the unavailability of a subsystem.

5.8.3 BrowseNext

5.8.3.1 Description

This *Service* is used to request the next set of *Browse* or *BrowseNext* response information that is too large to be sent in a single response. "Too large" in this context means that the *Server* is not able to return a larger response or that the number of results to return exceeds the maximum number of results to return that was specified by the *Client* in the original *Browse* request. The *BrowseNext* shall be submitted on the same *Session* that was used to submit the *Browse* or *BrowseNext* that is being continued.

5.8.3.2 Parameters

Table 37 defines the parameters for the *Service*.

Table 37 – BrowseNext Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
releaseContinuationPoints	Boolean	A <i>Boolean</i> parameter with the following values: TRUE passed <i>continuationPoints</i> shall be reset to free resources in the <i>Server</i>. The continuation points are released and the results and diagnosticInfos arrays are empty. FALSE passed <i>continuationPoints</i> shall be used to get the next set of browse information. A <i>Client</i> shall always use the continuation point returned by a <i>Browse</i> or <i>BrowseNext</i> response to free the resources for the continuation point in the <i>Server</i>. If the <i>Client</i> does not want to get the next set of browse information, <i>BrowseNext</i> shall be called with this parameter set to TRUE.
continuationPoints []	Continuation Point	A list of <i>Server</i>-defined opaque values that represent continuation points. The value for a continuation point was returned to the <i>Client</i> in a previous <i>Browse</i> or <i>BrowseNext</i> response. These values are used to identify the previously processed <i>Browse</i> or <i>BrowseNext</i> request that is being continued and the point in the result set from which the browse response is to continue. Clients may mix continuation points from different <i>Browse</i> or <i>BrowseNext</i> responses. The <i>ContinuationPoint</i> type is described in 7.6.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	BrowseResult	A list of references that met the criteria specified in the original <i>Browse</i> request. The size and order of this list matches the size and order of the <i>continuationPoints</i> request parameter. The <i>BrowseResult</i> type is defined in 7.3.
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>results</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>results</i> response parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.8.3.3 Service results

Table 38 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 38 – BrowseNext Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.8.3.4 StatusCodes

Table 39 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 39 – BrowseNext Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_BrowseDirectionInvalid	See Table 178 for the description of this result code.
Bad_NodeNotInView	See Table 178 for the description of this result code.
Bad_ContinuationPointInvalid	See Table 178 for the description of this result code.

5.8.4 TranslateBrowsePathsToNodeIds

5.8.4.1 Description

This *Service* is used to request that the *Server* translates one or more browse paths to *NodeIds*. Each browse path is constructed of a starting *Node* and a *RelativePath*. The specified starting *Node* identifies the *Node* from which the *RelativePath* is based. The *RelativePath* contains a sequence of *ReferenceTypes* and *BrowseNames*.

One purpose of this *Service* is to allow programming against type definitions. Since *BrowseNames* shall be unique in the context of type definitions, a *Client* may create a browse path that is valid for a type definition and use this path on instances of the type. For example, an *ObjectType* "Boiler" may have a "HeatSensor" *Variable* as *InstanceDeclaration*. A graphical element programmed against the "Boiler" may need to display the *Value* of the "HeatSensor". If the graphical element would be called on "Boiler1", an instance of "Boiler", it would need to call this *Service* specifying the *NodeId* of "Boiler1" as starting *Node* and the *BrowseName* of the "HeatSensor" as browse path. The *Service* would return the *NodeId* of the "HeatSensor" of "Boiler1" and the graphical element could subscribe to its *Value Attribute*.

If a *Node* has multiple targets with the same *BrowseName*, the *Server* shall return a list of *NodeIds*. However, since one of the main purposes of this *Service* is to support programming against type definitions, the *NodeId* of the *Node* based on the type definition of the starting *Node* is returned as the first *NodeId* in the list.

5.8.4.2 Parameters

Table 40 defines the parameters for the *Service*.

Table 40 – TranslateBrowsePathsToNodeIds Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
browsePaths []	BrowsePath	List of browse paths for which <i>NodeIds</i> are being requested. This structure is defined in-line with the following indented items.
startingNode	NodeId	<i>NodeId</i> of the starting <i>Node</i> for the browse path.
relativePath	RelativePath	The path to follow from the <i>startingNode</i> . The last element in the <i>relativePath</i> shall always have a <i>targetName</i> specified. This further restricts the definition of the <i>RelativePath</i> type. The <i>Server</i> shall return <i>Bad_BrowseNameInvalid</i> if the <i>targetName</i> is missing. The <i>RelativePath</i> structure is defined in 7.26.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	BrowsePathResult	List of results for the list of browse paths. The size and order of the list matches the size and order of the <i>browsePaths</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the browse path (see 7.34 for <i>StatusCode</i> definition).
targets []	BrowsePathTarget	List of targets for the <i>relativePath</i> from the <i>startingNode</i> . This structure is defined in-line with the following indented items. A <i>Server</i> may encounter a <i>Reference</i> to a <i>Node</i> in another <i>Server</i> which it cannot follow while it is processing the <i>RelativePath</i> . If this happens the <i>Server</i> returns the <i>NodeId</i> of the external <i>Node</i> and sets the <i>remainingPathIndex</i> parameter to indicate which <i>RelativePath</i> elements still need to be processed. To complete the operation the <i>Client</i> shall connect to the other <i>Server</i> and call this service again using the target as the <i>startingNode</i> and the unprocessed elements as the <i>relativePath</i> .
targetId	ExpandedNodeId	The identifier for a target of the <i>RelativePath</i> .
remainingPathIndex	Index	The index of the first unprocessed element in the <i>RelativePath</i> . This value shall be equal to the maximum value of <i>Index</i> data type if all elements were processed (see 7.13 for <i>Index</i> definition).
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the list of browse paths (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>browsePaths</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.8.4.3 Service results

Table 41 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in 7.34.

Table 41 – TranslateBrowsePathsToNodeIds Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.8.4.4 StatusCodes

Table 42 defines values for the operation level *statusCode* parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 42 – TranslateBrowsePathsToNodeIds Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_NothingToDo	See Table 177 for the description of this result code. This code indicates that the relativePath contained an empty list.
Bad_BrowseNameInvalid	See Table 178 for the description of this result code. This code indicates that a TargetName was missing in a RelativePath.
Uncertain_ReferenceOutOfServer	The path element has targets which are in another server.
Bad_TooManyMatches	The requested operation has too many matches to return. Users should use queries for large result sets. <i>Servers</i> should allow at least 10 matches before returning this error code.
Bad_QueryTooComplex	The requested operation requires too many resources in the server.
Bad_NoMatch	The requested relativePath cannot be resolved to a target to return.

5.8.5 RegisterNodes

5.8.5.1 Description

A *Server* often has no direct access to the information that it manages. Variables or services might be in underlying systems where additional effort is required to establish a connection to these systems. The *RegisterNodes Service* can be used by *Clients* to register the *Nodes* that they know they will access repeatedly (e.g. Write, Call). It allows *Servers* to set up anything needed so that the access operations will be more efficient. Clients can expect performance improvements when using registered *NodeIds*, but the optimization measures are vendor-specific. For *Variable Nodes Servers* shall concentrate their optimization efforts on the *Value Attribute*.

Registered *NodeIds* are only guaranteed to be valid within the current *Session*. *Clients* shall unregister unneeded *Ids* immediately to free up resources.

RegisterNodes does not validate the *NodeIds* from the request. *Servers* will simply copy unknown *NodeIds* in the response. Structural *NodeId* errors (size violations, invalid id types) will cause the complete *Service* to fail.

For the purpose of *Auditing*, *Servers* shall not use the registered *NodeIds* but only the canonical *NodeIds* which is the value of the *NodeId Attribute*.

5.8.5.2 Parameters

Table 43 defines the parameters for the *Service*.

Table 43 – RegisterNodes Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToRegister []	NodeId	List of <i>NodeIds</i> to register that the <i>Client</i> has retrieved through browsing, querying or in some other manner.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
registeredNodeIds []	NodeId	A list of <i>NodeIds</i> which the <i>Client</i> shall use for subsequent access operations. The size and order of this list matches the size and order of the <i>nodesToRegister</i> request parameter. The <i>Server</i> may return the <i>NodeId</i> from the request or a new (an alias) <i>NodeId</i>. It is recommended that the <i>Server</i> return a numeric <i>NodeIds</i> for aliasing. In case no optimization is supported for a <i>Node</i>, the <i>Server</i> shall return the <i>NodeId</i> from the request.

5.8.5.3 Service results

Table 44 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 44 – RegisterNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_NodeIdInvalid	See Table 178 for the description of this result code. <i>Servers</i> shall completely reject the <i>RegisterNodes</i> request if any of the <i>NodeIds</i> in the <i>nodesToRegister</i> parameter are structurally invalid.

5.8.6 UnregisterNodes

5.8.6.1 Description

This *Service* is used to unregister *NodeIds* that have been obtained via the *RegisterNodes* service.

UnregisterNodes does not validate the *NodeIds* from the request. *Servers* shall simply unregister *NodeIds* that are known as registered *NodeIds*. Any *NodeIds* that are in the list, but are not registered *NodeIds* are simply ignored.

5.8.6.2 Parameters

Table 50 defines the parameters for the *Service*.

Table 45 – UnregisterNodes Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToUnregister []	NodeId	A list of <i>NodeIds</i> that have been obtained via the <i>RegisterNodes</i> service.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

5.8.6.3 Service results

Table 51 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 46 – UnregisterNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.9 Query Service Set

5.9.1 Overview

This *Service Set* is used to issue a *Query* to a *Server*. OPC UA *Query* is generic in that it provides an underlying storage mechanism independent *Query* capability that can be used to access a wide variety of OPC UA data stores and information management systems. OPC UA *Query* permits a *Client* to access data maintained by a *Server* without any knowledge of the logical schema used for internal storage of the data. Knowledge of the *AddressSpace* is sufficient.

An *OPC UA Application* is expected to use the OPC UA *Query Services* as part of an initialization process or an occasional information synchronization step. For example, OPC UA *Query* would be used for bulk data access of a persistent store to initialise an analysis application with the current state of a system configuration. A *Query* may also be used to initialise or populate data for a report.

A *Query* defines what instances of one or more *TypeDefinitionNodes* in the *AddressSpace* should supply a set of *Attributes*. Results returned by a *Server* are in the form of an array of *QueryDataSets*. The selected *Attribute* values in each *QueryDataSet* come from the definition of the selected *TypeDefinitionNodes* or related *TypeDefinitionNodes* and appear in results in the same order as the *Attributes* that were passed into the *Query*. *Query* also supports *Node* filtering on the basis of *Attribute* values, as well as relationships between *TypeDefinitionNodes*.

See Annex B for example queries.

5.9.2 Querying Views

A *View* is a subset of the *AddressSpace* available in the *Server*. See IEC 62541-5 for a description of the organization of *Views* in the *AddressSpace*.

For any existing *View*, a *Query* may be used to return a subset of data from the *View*. When an application issues a *Query* against a *View*, only data defined by the *View* is returned. Data not included in the *View* but included in the original *AddressSpace* is not returned.

The *Query Services* supports access to current and historical data. The *Service* supports a *Client* querying a past version of the *AddressSpace*. Clients may specify a *ViewVersion* or a *Timestamp* in a *Query* to access past versions of the *AddressSpace*. OPC UA *Query* is complementary to Historical Access in that the former is used to *Query* an *AddressSpace* that existed at a time and the latter is used to *Query* for the value of *Attributes* over time. In this way, a *Query* can be used to retrieve a portion of a past *AddressSpace* so that *Attribute* value history may be accessed using Historical Access even if the *Node* is no longer in the current *AddressSpace*.

Servers that support *Query* are expected to be able to access the *AddressSpace* that is associated with the local *Server* and any *Views* that are available on the local *Server*. If a *View* or the *AddressSpace* also references a remote *Server*, query may be able to access the *AddressSpace* of the remote *Server*, but it is not required. If a *Server* does access a remote *Server* the access shall be accomplished using the user identity of the *Client* as described in 5.5.1.

5.9.3 QueryFirst

5.9.3.1 Description

This *Service* is used to issue a *Query* request to the *Server*. The complexity of the *Query* can range from very simple to highly sophisticated. The *Query* can simply request data from instances of a *TypeDefinitionNode* or *TypeDefinitionNode* subject to restrictions specified by the filter. On the other hand, the *Query* can request data from instances of related *Node* types by specifying a *RelativePath* from an originating *TypeDefinitionNode*. In the filter, a separate set of paths can be constructed for limiting the instances that supply data. A filtering path can include multiple *RelatedTo* operators to define a multi-hop path between source instances and target instances. For example, one could filter on students that attend a particular school, but return information about students and their families. In this case, the student school relationship is traversed for filtering, but the student family relationship is traversed to select data. For a complete description of *ContentFilter* see 7.4, also see Clause B.1 for simple examples and Clause B.2 for more complex examples of content filter and queries.

The *Client* provides an array of *NodeTypeDescription* which specify the *NodeId* of a *TypeDefinitionNode* and selects what *Attributes* are to be returned in the response. A *Client* can also provide a set of *RelativePaths* through the type system starting from an originating *TypeDefinitionNode*. Using these paths, the *Client* selects a set of *Attributes* from *Nodes* that are related to instances of the originating *TypeDefinitionNode*. Additionally, the *Client* can request the *Server* return instances of subtypes of *TypeDefinitionNodes*. If a selected *Attribute* does not exist in a *TypeDefinitionNode* but does exist in a subtype, it is assumed to have a null value in the *TypeDefinitionNode* in question. Therefore, this does not constitute an error condition and a null value is returned for the *Attribute*.

The *Client* can use the filter parameter to limit the result set by restricting *Attributes* and *Properties* to certain values. Another way the *Client* can use a filter to limit the result set is by specifying how instances should be related, using *RelatedTo* operators. In this case, if an instance at the top of the *RelatedTo* path cannot be followed to the bottom of the path via specified hops, no *QueryDataSets* are returned for the starting instance or any of the intermediate instances.

When querying for related instances in the *RelativePath*, the *Client* can optionally ask for *References*. A *Reference* is requested via a *RelativePath* that only includes a *ReferenceType*. If all *References* are desired then the root *ReferenceType* is listed. These *References* are returned as part of the *QueryDataSets*.

Query Services allow a special handling of the `targetName` field in the `RelativePath`. In several Query use cases a type `NodeId` is necessary in the path instead of a `QualifiedName`. Therefore the Client is allowed to specify a `NodeId` in the `QualifiedName`. This is done by setting the `namespaceIndex` of the `targetName` to zero and the name part of the `targetName` to the XML representation of the `NodeId`. The XML representation is defined in IEC 62541-6. When matching instances are returned as the target node, the target node shall be an instance of the specified type or subtype of the specified type.

Table 47 defines the request parameters and Table 48 the response parameters for the *QueryFirst Service*.

Table 47 – QueryFirst Request parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
View	ViewDescription	Specifies a <i>View</i> and temporal context to a <i>Server</i> (see 7.39 for <i>ViewDescription</i> definition).
nodeTypes []	NodeTypeDescription	This is the <i>Node</i> type description. This structure is defined in-line with the following indented items.
typeDefinitionNode	ExpandedNodeId	<i>NodeId</i> of the originating <i>TypeDefinitionNode</i> of the instances for which data is to be returned.
includeSubtypes	Boolean	A flag that indicates whether the <i>Server</i> should include instances of subtypes of the <i>TypeDefinitionNode</i> in the list of instances of the <i>Node</i> type.
dataToReturn []	QueryDataDescription	Specifies an <i>Attribute</i> or <i>Reference</i> from the originating <i>typeDefinitionNode</i> along a given <i>relativePath</i> for which to return data. This structure is defined in-line with the following indented items.
relativePath	RelativePath	Browse path relative to the originating <i>Node</i> that identifies the <i>Node</i> which contains the data that is being requested, where the originating <i>Node</i> is an instance <i>Node</i> of the type defined by the type definition <i>Node</i>. The instance <i>Nodes</i> are further limited by the filter provided as part of this call. For a definition of <i>relativePath</i> see 7.26. This relative path could end on a <i>Reference</i>, in which case the <i>ReferenceDescription</i> of the <i>Reference</i> would be returned as its value. The <code>targetName</code> field of the <i>relativePath</i> may contain a type <i>NodeId</i>. This is done by setting the <code>namespaceIndex</code> of the <code>targetName</code> to zero and the <code>name</code> part of the <code>targetName</code> to the XML representation of the <i>NodeId</i>. The XML representation is defined in IEC 62541-6. When matching instances are returned as the target node, the target node shall be an instance of the specified type or subtype of the specified type.
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>Attribute</i> Id. The <i>IntegerId</i> is defined in 7.14. The <i>IntegerId</i> for <i>Attributes</i> are defined in IEC 62541-6. If the <i>RelativePath</i> ended in a <i>Reference</i> then this parameter is 0 and ignored by the server.
indexRange	NumericRange	This parameter is used to identify a single element of a structure or an array, or a single range of indexes for arrays. If a range of elements are specified, the values are returned as a composite. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is null if the specified <i>Attribute</i> is not an array or a structure. However, if the specified <i>Attribute</i> is an array or a structure, and this parameter is null, then all elements are to be included in the range.
Filter	ContentFilter	Resulting <i>Nodes</i> shall be limited to the <i>Nodes</i> matching the criteria defined by the filter. <i>ContentFilter</i> is discussed in 7.4. If an empty filter is provided then the entire <i>AddressSpace</i> shall be examined and all <i>Nodes</i> that contain a matching requested <i>Attribute</i> or <i>Reference</i> are returned.

Name	Type	Description
maxDataSetsToReturn	Counter	The number of <i>QueryDataSets</i> that the <i>Client</i> wants the <i>Server</i> to return in the response and on each subsequent continuation call response. The <i>Server</i> is allowed to further limit the response, but shall not exceed this limit. A value of 0 indicates that the <i>Client</i> is imposing no limitation.
maxReferencesToReturn	Counter	The number of <i>References</i> that the <i>Client</i> wants the <i>Server</i> to return in the response for each <i>QueryDataSet</i> and on each subsequent continuation call response. The <i>Server</i> is allowed to further limit the response, but shall not exceed this limit. A value of 0 indicates that the <i>Client</i> is imposing no limitation. For example a result where 4 <i>Nodes</i> are being returned, but each has 100 <i>References</i> , if this limit were set to 50 then only the first 50 <i>References</i> for each <i>Node</i> would be returned on the initial call and a continuation point would be set indicating additional data.

Table 48 – QueryFirst Response parameters

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
queryDataSets []	QueryDataSet	The array of <i>QueryDataSets</i> . This array is empty if no <i>Nodes</i> or <i>References</i> met the <i>nodeTypes</i> criteria. In this case the continuationPoint parameter shall be empty. The <i>QueryDataSet</i> type is defined in 7.23.
continuationPoint	ContinuationPoint	Server-defined opaque value that identifies the continuation point. The continuation point is used only when the <i>Query</i> results are too large to be returned in a single response. "Too large" in this context means that the <i>Server</i> is not able to return a larger response or that the number of <i>QueryDataSets</i> to return exceeds the maximum number of <i>QueryDataSets</i> to return that was specified by the <i>Client</i> in the request. The continuation point is used in the <i>QueryNext Service</i> . When not used, the value of this parameter is null. If a continuation point is returned, the <i>Client</i> shall call <i>QueryNext</i> to get the next set of <i>QueryDataSets</i> or to free the resources for the continuation point in the <i>Server</i> . A continuation point shall remain active until the <i>Client</i> passes the continuation point to <i>QueryNext</i> or the session is closed. If the maximum continuation points have been reached the oldest continuation point shall be reset. The <i>ContinuationPoint</i> type is described in 7.6.
parsingResults[]	ParsingResult	List of parsing results for <i>QueryFirst</i> . The size and order of the list matches the size and order of the <i>NodeTypes</i> request parameter. This structure is defined in-line with the following indented items. This list is populated with any status codes that are related to the processing of the node types that are part of the query. The array can be empty if no errors were encountered. If any node type encountered an error all node types shall have an associated status code.
statusCode	StatusCode	Parsing result for the requested <i>NodeTypeDescription</i> .
dataStatusCodes []	StatusCode	List of results for <i>dataToReturn</i> . The size and order of the list matches the size and order of the <i>dataToReturn</i> request parameter. The array can be empty if no errors were encountered.
dataDiagnosticInfos []	DiagnosticInfo	List of diagnostic information <i>dataToReturn</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>dataToReturn</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the query request.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the requested <i>NodeTypeDescription</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the query request.
filterResult	ContentFilter Result	A structure that contains any errors associated with the filter. This structure shall be empty if no errors occurred. The <i>ContentFilterResult</i> type is defined in 7.4.2.

5.9.3.2 Service results

If the *Query* is invalid or cannot be processed, then *QueryDataSets* are not returned and only a *Service* result, *filterResult*, *parsingResults* and optional *DiagnosticInfo* is returned. Table 49 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 49 – QueryFirst Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_ContentFilterInvalid	See Table 178 for the description of this result code.
Bad_ViewIdUnknown	See Table 177 for the description of this result code.
Bad_ViewTimestampInvalid	See Table 177 for the description of this result code.
Bad_ViewParameterMismatchInvalid	See Table 177 for the description of this result code.
Bad_ViewVersionInvalid	See Table 177 for the description of this result code.
Bad_InvalidFilter	The provided filter is invalid, see the <i>filterResult</i> for specific errors
Bad_NodelistError	The <i>NodeTypes</i> provided contain an error, see the <i>parsingResults</i> for specific errors
Bad_InvalidView	The provided <i>ViewDescription</i> is not a valid <i>ViewDescription</i> .
Good_ResultsMayBeIncomplete	The <i>Server</i> should have followed a reference to a node in a remote <i>Server</i> but did not. The result set may be incomplete.

5.9.3.3 StatusCodes

Table 50 defines values for the *parsingResults statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 50 – QueryFirst Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_NotTypeDefinition	The provided <i>NodeId</i> was not a type definition <i>NodeId</i> .
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.

5.9.4 QueryNext

5.9.4.1 Descriptions

This *Service* is used to request the next set of *QueryFirst* or *QueryNext* response information that is too large to be sent in a single response. "Too large" in this context means that the *Server* is not able to return a larger response or that the number of *QueryDataSets* to return exceeds the maximum number of *QueryDataSets* to return that was specified by the *Client* in the original request. The *QueryNext* shall be submitted on the same session that was used to submit the *QueryFirst* or *QueryNext* that is being continued.

5.9.4.2 Parameters

Table 51 defines the parameters for the *Service*.

Table 51 – QueryNext Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
releaseContinuationPoint	Boolean	A <i>Boolean</i> parameter with the following values: TRUE passed <i>continuationPoint</i> shall be reset to free resources for the continuation point in the <i>Server</i>. FALSE passed <i>continuationPoint</i> shall be used to get the next set of <i>QueryDataSets</i>. A <i>Client</i> shall always use the continuation point returned by a <i>QueryFirst</i> or <i>QueryNext</i> response to free the resources for the continuation point in the <i>Server</i>. If the <i>Client</i> does not want to get the next set of <i>Query</i> information, <i>QueryNext</i> shall be called with this parameter set to TRUE. If the parameter is set to TRUE all array parameters in the response shall contain empty arrays.
continuationPoint	ContinuationPoint	Server defined opaque value that represents the continuation point. The value of the continuation point was returned to the <i>Client</i> in a previous <i>QueryFirst</i> or <i>QueryNext</i> response. This value is used to identify the previously processed <i>QueryFirst</i> or <i>QueryNext</i> request that is being continued, and the point in the result set from which the browse response is to continue. The <i>ContinuationPoint</i> type is described in 7.6.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
queryDataSets []	QueryDataSet	The array of <i>QueryDataSets</i>. The <i>QueryDataSet</i> type is defined in 7.23.
revisedContinuationPoint	ContinuationPoint	Server-defined opaque value that represents the continuation point. It is used only if the information to be returned is too large to be contained in a single response. When not used or when <i>releaseContinuationPoint</i> is set, the value of this parameter is null. The <i>ContinuationPoint</i> type is described in 7.6.

5.9.4.3 Service results

Table 52 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 52 – QueryNext Service result codes

Symbolic Id	Description
Bad_ContinuationPointInvalid	See Table 178 for the description of this result code.

5.10 Attribute Service Set

5.10.1 Overview

This *Service Set* provides *Services* to access *Attributes* that are part of *Nodes*.

5.10.2 Read

5.10.2.1 Description

This *Service* is used to read one or more *Attributes* of one or more *Nodes*. For constructed *Attribute* values whose elements are indexed, such as an array, this *Service* allows *Clients* to read the entire set of indexed values as a composite, to read individual elements or to read ranges of elements of the composite.

The *maxAge* parameter is used to direct the *Server* to access the value from the underlying data source, such as a device, if its copy of the data is older than that which the *maxAge* specifies. If the *Server* cannot meet the requested maximum age, it returns its "best effort" value rather than rejecting the request.

5.10.2.2 Parameters

Table 53 defines the parameters for the *Service*.

Table 53 – Read Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
maxAge	Duration	Maximum age of the value to be read in milliseconds. The age of the value is based on the difference between the <i>ServerTimestamp</i> and the time when the <i>Server</i> starts processing the request. For example if the <i>Client</i> specifies a <i>maxAge</i> of 500 milliseconds and it takes 100 milliseconds until the <i>Server</i> starts processing the request, the age of the returned value could be 600 milliseconds prior to the time it was requested. If the <i>Server</i> has one or more values of an <i>Attribute</i> that are within the maximum age, it can return any one of the values or it can read a new value from the data source. The number of values of an <i>Attribute</i> that a <i>Server</i> has depends on the number of <i>MonitoredItems</i> that are defined for the <i>Attribute</i>. In any case, the <i>Client</i> can make no assumption about which copy of the data will be returned. If the <i>Server</i> does not have a value that is within the maximum age, it shall attempt to read a new value from the data source. If the <i>Server</i> cannot meet the requested <i>maxAge</i>, it returns its "best effort" value rather than rejecting the request. This may occur when the time it takes the <i>Server</i> to process and return the new data value after it has been accessed is greater than the specified maximum age. If <i>maxAge</i> is set to 0, the <i>Server</i> shall attempt to read a new value from the data source. If <i>maxAge</i> is set to the max Int32 value or greater, the <i>Server</i> shall attempt to get a cached value. Negative values are invalid for <i>maxAge</i>.
timestampsToReturn	Enum TimestampsToReturn	An enumeration that specifies the <i>Timestamps</i> to be returned for each requested <i>Variable Value Attribute</i> . The <i>TimestampsToReturn</i> enumeration is defined in 7.35.
nodesToRead []	ReadValueId	List of <i>Nodes</i> and their <i>Attributes</i> to read. For each entry in this list, a <i>StatusCode</i> is returned, and if it indicates success, the <i>Attribute Value</i> is also returned. The <i>ReadValueId</i> parameter type is defined in 7.24.

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	DataValue	List of <i>Attribute</i> values (see 7.7 for <i>DataValue</i> definition). The size and order of this list matches the size and order of the <i>nodesToRead</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToRead</i> parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of this list matches the size and order of the <i>nodesToRead</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToRead</i> parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.2.3 Service results

Table 54 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 54 – Read Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_MaxAgeInvalid	The max age parameter is invalid.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.

5.10.2.4 StatusCodes

Table 55 defines values for the operation level *statusCode* contained in the *DataValue* structure of each *values results* element. Common *StatusCodes* are defined in Table 178.

Table 55 – Read Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeNoData	See Table 178 for the description of this result code.
Bad_DataEncodingInvalid	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_NotReadable	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_SecurityModelInsufficient	See Table 178 for the description of this result code.

5.10.3 HistoryRead

5.10.3.1 Description

This *Service* is used to read historical values or *Events* of one or more *Nodes*. For constructed *Attribute* values whose elements are indexed, such as an array, this *Service*

allows *Clients* to read the entire set of indexed values as a composite, to read individual elements or to read ranges of elements of the composite. *Servers* may make historical values available to *Clients* using this *Service*, although the historical values themselves are not visible in the *AddressSpace*.

The *AccessLevel Attribute* defined in IEC 62541-3 indicates a *Node's* support for historical values. Several request parameters indicate how the *Server* is to access values from the underlying history data source. The *EventNotifier Attribute* defined in IEC 62541-3 indicates a *Node's* support for historical *Events*.

The *continuationPoint* parameter in the *HistoryRead* is used to mark a point from which to continue the read if not all values could be returned in one response. The value is opaque for the *Client* and is only used to maintain the state information for the *Server* to continue from. A *Server* may use the timestamp of the last returned data item if the timestamp is unique. This can reduce the need in the *Server* to store state information for the continuation point.

In some cases it may take longer than the *Client* timeout hint to read the data for all nodes to read. Then the *Server* may return zero results with a continuation point for the affected nodes before the timeout expires. That allows the *Server* to resume the data acquisition on the next *Client* read call.

For additional details on reading historical data and historical *Events* see IEC 62541-11.

5.10.3.2 Parameters

Table 56 defines the parameters for the *Service*.

Table 56 – HistoryRead Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
historyReadDetails	Extensible Parameter HistoryReadDetails	The details define the types of history reads that can be performed. The <i>HistoryReadDetails</i> parameter type is an extensible parameter type formally defined in IEC 62541-11. The <i>ExtensibleParameter</i> type is defined in 7.12.
timestampsToReturn	Enum TimestampsToReturn	An enumeration that specifies the timestamps to be returned for each requested <i>Variable Value Attribute</i> . The <i>TimestampsToReturn</i> enumeration is defined in 7.35. Specifying a <i>TimestampsToReturn</i> of NEITHER is not valid. A <i>Server</i> shall return a <i>Bad_InvalidTimestampArgument StatusCode</i> in this case. IEC 62541-11 defines exceptions where this parameter shall be ignored.
releaseContinuation Points	Boolean	A <i>Boolean</i> parameter with the following values: TRUE passed <i>continuationPoints</i> shall be reset to free resources in the <i>Server</i> . FALSE passed <i>continuationPoints</i> shall be used to get the next set of historical information. A <i>Client</i> shall always use the continuation point returned by a <i>HistoryRead</i> response to free the resources for the continuation point in the <i>Server</i> . If the <i>Client</i> does not want to get the next set of historical information, <i>HistoryRead</i> shall be called with this parameter set to TRUE.
nodesToRead []	HistoryReadValueId	This parameter contains the list of items upon which the historical retrieval is to be performed. This structure is defined in-line with the following indented items.

Name	Type	Description
nodeId	NodeId	If the <i>HistoryReadDetails</i> is RAW, PROCESSED, MODIFIED or ATTIME: The <i>nodeId</i> of the <i>Nodes</i> whose historical values are to be read. The value returned shall always include a timestamp. If the <i>HistoryReadDetails</i> is EVENTS: The <i>NodeId</i> of the <i>Node</i> whose <i>Event</i> history is to be read. If the <i>Node</i> does not support the requested access for historical values or historical <i>Events</i> the appropriate error response for the given <i>Node</i> shall be generated.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for arrays. If a range of elements is specified, the values are returned as a composite. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is null if the value is not an array. However, if the value is an array, and this parameter is null, then all elements are to be included in the range.
dataEncoding	QualifiedName	A <i>QualifiedName</i> that specifies the data encoding to be returned for the <i>Value</i> to be read (see 7.24 for definition how to specify the data encoding). This parameter only applies if the <i>DataType</i> of the <i>Variable</i> is a subtype of <i>Structure</i>. It is an error to specify this parameter if the <i>DataType</i> of the <i>Variable</i> is not a subtype of <i>Structure</i>. The parameter is ignored when reading history of <i>Events</i>.
continuationPoint	ByteString ContinuationPoint	For each <i>NodesToRead</i> item this parameter specifies a continuation point returned from a previous <i>HistoryRead</i> call, allowing the <i>Client</i> to continue that read from the last value received. The <i>HistoryRead</i> is used to select an ordered sequence of historical values or events. A continuation point marks a point in that ordered sequence, such that the <i>Server</i> returns the subset of the sequence that follows that point. A null value indicates that this parameter is not used. See 7.6 for a general description of continuation points. This continuation point is described in more detail in IEC 62541-11.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> type).
results []	HistoryReadResult	List of read results. The size and order of the list matches the size and order of the <i>nodesToRead</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the <i>NodesToRead</i> item (see 7.34 for <i>StatusCode</i> definition).
continuationPoint	ByteString ContinuationPoint	This parameter is used only if the number of values to be returned is too large to be returned in a single response or if the timeout provided as hint by the <i>Client</i> is close to expiring and not all nodes have been processed. When this parameter is not used, its value is null. <i>Servers</i> shall support at least one continuation point per <i>Session</i>. <i>Servers</i> specify a max history continuation points per <i>Session</i> in the <i>Server capabilities Object</i> defined in IEC 62541-5. A continuation point shall remain active until the <i>Client</i> passes the continuation point to <i>HistoryRead</i> or the <i>Session</i> is closed. If the max continuation points have been reached the oldest continuation point shall be reset.
historyData	Extensible Parameter HistoryData	The history data returned for the <i>Node</i>. The <i>HistoryData</i> parameter type is an extensible parameter type formally defined in IEC 62541-11. It specifies the types of history data that can be returned. The <i>ExtensibleParameter</i> base type is defined in 7.12.
diagnosticInfos []	Diagnostic Info	List of diagnostic information. The size and order of the list matches the size and order of the <i>nodesToRead</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToRead</i> parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.3.3 Service results

Table 57 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 57 – HistoryRead Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.
Bad_HistoryOperationInvalid	See Table 178 for the description of this result code.
Bad_HistoryOperationUnsupported	See Table 178 for the description of this result code. The requested history operation is not supported by the server.

5.10.3.4 StatusCodes

Table 58 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178. History access specific *StatusCodes* are defined in IEC 62541-11.

Table 58 – HistoryRead Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_DataEncodingInvalid	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_ContinuationPointInvalid	See Table 177 for the description of this result code.
Bad_InvalidTimestampArgument	The defined timestamp to return was invalid.
Bad_HistoryOperationUnsupported	See Table 178 for the description of this result code. The requested history operation is not supported for the requested node.
Bad_NoContinuationPoints	See Table 178 for the description of this result code. See 7.6 for the rules to apply this status code.

5.10.4 Write

5.10.4.1 Description

This *Service* is used to write values to one or more *Attributes* of one or more *Nodes*. For constructed *Attribute* values whose elements are indexed, such as an array, this *Service* allows *Clients* to write the entire set of indexed values as a composite, to write individual elements or to write ranges of elements of the composite.

The values are written to the data source, such as a device, and the *Service* does not return until it writes the values or determines that the value cannot be written. In certain cases, the *Server* will successfully write to an intermediate system or *Server*, and will not know if the data source was updated properly. In these cases, the *Server* should report a success code that indicates that the write was not verified. In the cases where the *Server* is able to verify that it has successfully written to the data source, it reports an unconditional success.

The order the operations are processed in the *Server* is not defined and depends on the different data sources and the internal *Server* logic. If an *Attribute* and *Node* combination is contained in more than one operation, the order of the processing is undefined. If a *Client* requires sequential processing the *Client* needs separate *Service* calls.

It is possible that the *Server* may successfully write some *Attributes*, but not others. Rollback is the responsibility of the *Client*.

If a *Server* allows writing of *Attributes* with the *DataType LocalizedText*, the *Client* can add or overwrite the text for a locale by writing the text with the associated *LocaleId*. Writing a null *String* for the text for a locale shall delete the *String* for that locale. Writing a null *String* for the *locale* and a non-null *String* for the *text* is setting the *text* for an invariant locale. Writing a null *String* for the *text* and a null *String* for the *locale* shall delete the entries for all locales. If a *Client* attempts to write a *locale* that is either syntactically invalid or not supported, the *Server* returns *Bad_LocaleNotSupported*. The *Write* behaviour for *Value Attributes* with a *LocalizedText DataType* is *Server* specific but it is recommended to follow the same rules.

5.10.4.2 Parameters

Table 59 defines the parameters for the *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 59 – Write Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToWrite []	WriteValue	List of <i>Nodes</i> and their <i>Attributes</i> to write. This structure is defined in-line with the following indented items.
nodeId	NodeId	<i>NodeId</i> of the <i>Node</i> that contains the <i>Attributes</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>Attribute</i> id. The <i>IntegerId</i> is defined in 7.14. The <i>IntegerIds</i> for the <i>Attributes</i> are defined in IEC 62541-6.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for arrays. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is not used if the specified <i>Attribute</i> is not an array. However, if the specified <i>Attribute</i> is an array and this parameter is not used, then all elements are to be included in the range. The parameter is null if not used. A <i>Server</i> shall return a <i>Bad_WriteNotSupported</i> error if an <i>indexRange</i> is provided and writing of <i>indexRange</i> is not possible for the <i>Node</i>.
value	DataValue	The <i>Node's Attribute</i> value (see 7.7 for <i>DataValue</i> definition). If the <i>indexRange</i> parameter is specified then the <i>Value</i> shall be an array even if only one element is being written. If the <i>SourceTimestamp</i> or the <i>ServerTimestamp</i> is specified, the <i>Server</i> shall use these values. The <i>Server</i> returns a <i>Bad_WriteNotSupported</i> error if it does not support writing of timestamps. A <i>Server</i> shall return a <i>Bad_TypeMismatch</i> error if the data type of the written value is not the same type or subtype of the <i>Attribute's DataType</i>. Based on the <i>DataType</i> hierarchy, subtypes of the <i>Attribute DataType</i> shall be accepted by the <i>Server</i>. <i>Servers</i> may reject subtypes defined in newer specification versions than supported by the <i>Server</i> with <i>Bad_TypeMismatch</i>. For the <i>Value Attribute</i> the <i>DataType</i> is defined through the <i>DataType Attribute</i>. A <i>ByteString</i> is structurally the same as a one dimensional array of <i>Byte</i>. A <i>Server</i> shall accept a <i>ByteString</i> if an array of <i>Byte</i> is expected. The <i>Server</i> returns a <i>Bad_DataEncodingUnsupported</i> error if it does not support the provided data encoding. <i>Simple DataTypes</i> (see IEC 62541-3) use the same representation on the wire as their super types and therefore writing a value of a simple <i>DataType</i> cannot be distinguished from writing a value of its super type. The <i>Server</i> shall assume that by receiving the correct wire representation for a simple <i>DataType</i> the correct type was chosen. <i>Servers</i> are allowed to impose additional data validations on the value independent of the encoding (e.g. having an image in GIF format in a <i>ByteString</i>). In this case the <i>Server</i> shall return a <i>Bad_TypeMismatch</i> error if the validation fails.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of results for the <i>Nodes</i> to write (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>nodesToWrite</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToWrite</i> parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Nodes</i> to write (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>nodesToWrite</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.4.3 Service results

Table 60 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 60 – Write Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.10.4.4 StatusCodes

Table 61 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 61 – Write Operation Level result codes

Symbolic Id	Description
Good_CompletesAsynchronously	See Table 177 for the description of this result code. The value was successfully written to an intermediate system but the <i>Server</i> does not know if the data source was updated properly.
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeNoData	See Table 178 for the description of this result code.
Bad_WriteNotSupported	The requested write operation is not supported. If a <i>Client</i> attempts to write any value, <i>quality status code</i> , timestamp combination and the <i>Server</i> does not support the requested combination (which could be a single quantity such as just timestamp); then the <i>Server</i> shall not perform any write on this <i>Node</i> and shall return this <i>StatusCode</i> for this <i>Node</i> . It is also used if writing an <i>IndexRange</i> is not supported for a <i>Node</i> .
Bad_NotWritable	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code. The current user does not have permission to write the attribute.
Bad_OutOfRange	See Table 178 for the description of this result code. If a <i>Client</i> attempts to write a value outside the valid range like a value not contained in the enumeration data type of the <i>Node</i> , the <i>Server</i> shall return this <i>StatusCode</i> for this <i>Node</i> .
Bad_TypeMismatch	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_NoCommunication	See Table 178 for the description of this result code.
Bad_LocaleNotSupported	The locale in the requested write operation is not supported.

5.10.5 HistoryUpdate

5.10.5.1 Description

This *Service* is used to update historical values or *Events* of one or more *Nodes*. Several request parameters indicate how the *Server* is to update the historical value or *Event*. Valid actions are Insert, Replace or Delete.

5.10.5.2 Parameters

Table 62 defines the parameters for the *Service*.

Table 62 – HistoryUpdate Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
historyUpdateDetails []	Extensible Parameter HistoryUpdate Details	The details defined for this update. The <i>HistoryUpdateDetails</i> parameter type is an extensible parameter type formally defined in IEC 62541-11. It specifies the types of history updates that can be performed. The <i>ExtensibleParameter</i> type is defined in 7.12.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	HistoryUpdate Result	List of update results for the history update details. The size and order of the list matches the size and order of the details element of the <i>historyUpdateDetails</i> parameter specified in the request. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the update of the <i>Node</i> (see 7.34 for <i>StatusCode</i> definition).
operationResults []	StatusCode	List of <i>StatusCodes</i> for the operations to be performed on a <i>Node</i> . The size and order of the list matches the size and order of any list defined by the details element being reported by this result entry.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the operations to be performed on a <i>Node</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of any list defined by the details element being reported by this <i>updateResults results</i> entry. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the history update details. The size and order of the list matches the size and order of the details element of the <i>historyUpdateDetails</i> parameter specified in the request. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.5.3 Service results

Table 63 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 63 – HistoryUpdate Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.10.5.4 StatusCodes

Table 64 defines values for the *statusCode* and *operationResults* parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178. History access specific *StatusCodes* are defined in IEC 62541-11.

Table 64 – HistoryUpdate Operation Level result codes

Symbolic Id	Description
Bad_NotWritable	See Table 178 for the description of this result code.
Bad_HistoryOperationInvalid	See Table 178 for the description of this result code.
Bad_HistoryOperationUnsupported	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code. The current user does not have permission to update the history.

5.11 Method Service Set

5.11.1 Overview

Methods represent the function calls of *Objects*. They are defined in IEC 62541-3. *Methods* are invoked and return only after completion (successful or unsuccessful). Execution times for *Methods* may vary, depending on the function that they perform.

The *Method Service Set* defines the means to invoke *Methods*. A *Method* shall be a component of an *Object*. Discovery is provided through the *Browse* and *Query Services*. *Clients* discover the *Methods* supported by a *Server* by browsing for the owning *Objects References* that identify their supported *Methods*.

Because *Methods* may control some aspect of plant operations, *Method* invocation may depend on environmental or other conditions. This may be especially true when attempting to re-invoke a *Method* immediately after it has completed execution. Conditions that are required to invoke the *Method* might not yet have returned to the state that permits the *Method* to start again.

5.11.2 Call

5.11.2.1 Description

This *Service* is used to call (invoke) a list of *Methods*. ~~Each method call is invoked within the context of an existing Session. If the Session is terminated, the results of the method's execution cannot be returned to the Client and are discarded. This is independent of the task actually performed at the Server.~~

This *Service* provides for passing input and output arguments to/from a *Method*. These arguments are defined by *Properties* of the *Method*.

If the *Method* is invoked in the context of a *Session* and the *Session* is terminated, the results of the *Method's* execution cannot be returned to the *Client* and are discarded. This is independent of the task actually performed at the *Server*.

The order the operations are processed in the *Server* is not defined and depends on the different tasks and the internal *Server* logic. If a *Method* is contained in more than one operation, the order of the processing is undefined. If a *Client* requires sequential processing the *Client* needs separate *Service* calls.

~~This Service provides for passing input and output arguments to/from a method. These arguments are defined by Properties of the method.~~

5.11.2.2 Parameters

Table 65 defines the parameters for the *Service*.

Table 65 – Call Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
methodsToCall []	CallMethodRequest	List of <i>Methods</i> to call. This structure is defined in-line with the following indented items.
objectId	NodeId	~~The <i>NodeId</i> shall be that of the <i>Object</i> or <i>ObjectType</i> that is the source of a <i>HasComponent Reference</i> (or subtype of <i>HasComponent Reference</i>) to the <i>Method</i> specified in <i>methodId</i>.~~ The <i>NodeId</i> shall be that of the <i>Object</i> or <i>ObjectType</i> on which the <i>Method</i> is invoked. In case of an <i>ObjectType</i>, the <i>ObjectType</i> or a super type of the <i>ObjectType</i> shall be the source of a <i>HasComponent Reference</i> (or subtype of <i>HasComponent Reference</i>) to the <i>Method</i> specified in <i>methodId</i>. In case of an <i>Object</i>, the <i>Object</i> or the <i>ObjectType</i> of the <i>Object</i> or a super type of that <i>ObjectType</i> shall be the source of a <i>HasComponent Reference</i> (or subtype of <i>HasComponent Reference</i>) to the <i>Method</i> specified in <i>methodId</i>. See IEC 62541-3 for a description of <i>Objects</i> and their <i>Methods</i>.
methodId	NodeId	<i>NodeId</i> of the <i>Method</i> to invoke. If the <i>objectId</i> is the <i>NodeId</i> of an <i>Object</i>, it is allowed to use the <i>NodeId</i> of a <i>Method</i> that is the target of a <i>HasComponent Reference</i> from the <i>ObjectType</i> of the <i>Object</i>.
inputArguments []	BaseDataType	List of input argument values. An empty list indicates that there are no input arguments. The size and order of this list matches the size and order of the input arguments defined by the input <i>InputArguments Property</i> of the <i>Method</i>. The name, a description and the data type of each argument are defined by the <i>Argument</i> structure in each element of the method's <i>InputArguments Property</i>.

Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	CallMethodResult	Result for the <i>Method</i> calls. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> of the <i>Method</i> executed in the server. This <i>StatusCode</i> is set to the <i>Bad_InvalidArgument StatusCode</i> if at least one input argument broke a constraint (e.g. wrong data type, value out of range). This <i>StatusCode</i> is set to a bad <i>StatusCode</i> if the <i>Method</i> execution failed in the server, e.g. based on an exception.
inputArgumentResults []	StatusCode	List of <i>StatusCodes</i> corresponding to the <i>inputArguments</i> . This list is empty unless the operation level result is <i>Bad_InvalidArgument</i> . If this list is populated, it has the same length as the <i>inputArguments</i> list.
inputArgumentDiagnosticInfos []	DiagnosticInfo	List of diagnostic information corresponding to the <i>inputArguments</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.
outputArguments []	BaseDataType	List of output argument values. An empty list indicates that there are no output arguments. The size and order of this list matches the size and order of the output arguments defined by the <i>OutputArguments Property</i> of the <i>Method</i> . The name, a description and the data type of each argument are defined by the <i>Argument</i> structure in each element of the methods <i>OutputArguments Property</i> .
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>statusCode</i> of the <i>results</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.11.2.3 Service results

Table 66 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 66 – Call Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.11.2.4 StatusCodes

Table 67 defines values for the *statusCode* parameter and Table 68 defines values for the *inputArgumentResults* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Server vendors or OPC UA companion specifications may reuse existing *StatusCodes* for application specific error information. This is valid as long as the canonical description of the *StatusCode* does not have a different meaning than the application specific description. To eliminate any vagueness, the *Server* should include the application specific description in the *DiagnosticInfo*.

Good *StatusCodes* with sub-status shall not be used as *statusCode* since many programming language bindings would cause such codes to throw an exception.

Table 67 – Call Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code. Used to indicate that the specified <i>Object</i> is not valid.
Bad_NodeIdUnknown	See Table 178 for the description of this result code. Used to indicate that the specified <i>Object</i> is not valid.
Bad_NotExecutable	The executable <i>Attribute</i> does not allow the execution of the <i>Method</i> .
Bad_ArgumentsMissing	The <i>Client</i> did not specify all of the input arguments for the <i>Method</i> .
Bad_TooManyArguments	The <i>Client</i> specified more input arguments than defined for the <i>Method</i> .
Bad_InvalidArgument	See Table 177 for the description of this result code. Used to indicate in the operation level results that one or more of the input arguments are invalid. The <i>inputArgumentResults</i> contain the specific status code for each invalid argument.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_SecurityModelInsufficient	See Table 178 for the description of this result code.
Bad_MethodInvalid	The method id does not refer to a <i>Method</i> for the specified <i>Object</i> .
Bad_NoCommunication	See Table 178 for the description of this result code.

Table 68 – Call Input Argument Result Codes

Symbolic Id	Description
Bad_OutOfRange	See Table 178 for the description of this result code. Used to indicate that an input argument is outside the acceptable range.
Bad_TypeMismatch	See Table 178 for the description of this result code. Used to indicate that an input argument does not have the correct data type. ~~A <i>ByteString</i> is structurally the same as a one dimensional array of <i>Byte</i>. A server shall accept a <i>ByteString</i> if an array of <i>Byte</i> is expected.~~ A <i>ByteString</i> is structurally the same as a one dimensional array of <i>Byte</i> . A <i>Server</i> shall accept a <i>ByteString</i> if an array of <i>Byte</i> is expected.

5.12 MonitoredItem Service Set

5.12.1 MonitoredItem model

5.12.1.1 Overview

Clients define *MonitoredItems* to subscribe to data and *Events*. Each *MonitoredItem* identifies the item to be monitored and the *Subscription* to use to send *Notifications*. The item to be monitored may be any *Node Attribute*.

Notifications are data structures that describe the occurrence of data changes and *Events*. They are packaged into *NotificationMessages* for transfer to the *Client*. The *Subscription* periodically sends *NotificationMessages* at a user-specified publishing interval, and the cycle during which these messages are sent is called a publishing cycle.

Four primary parameters are defined for *MonitoredItems* that tell the *Server* how the item is to be sampled, evaluated and reported. These parameters are the sampling interval, the monitoring mode, the filter and the queue parameter. Figure 15 illustrates these concepts.




Figure 15 – MonitoredItem model

Attributes, other than the *Value Attribute*, are only monitored for a change in value. The filter is not used for these *Attributes*. Any change in value for these *Attributes* causes a *Notification* to be generated.

The *Value Attribute* is used when monitoring *Variables*. *Variable* values are monitored for a change in value or a change in their status. The filters defined in this document (see 7.17.2) and in IEC 62541-8 are used to determine if the value change is large enough to cause a *Notification* to be generated for the *Variable*.

Objects and *views* can be used to monitor *Events*. *Events* are only available from *Nodes* where the *SubscribeToEvents* bit of the *EventNotifier Attribute* is set. The filter defined in this document (see 7.17.3) is used to determine if an *Event* received from the *Node* is sent to the *Client*. The filter also allows selecting fields of the *EventType* that will be contained in the *Event* such as *EventId*, *EventType*, *SourceNode*, *Time* and *Description*.

IEC 62541-3 describes the *Event* model and the base *EventTypes*.

The *Properties* of the base *EventTypes* and the representation of the base *EventTypes* in the *AddressSpace* are specified in IEC 62541-5.

5.12.1.2 Sampling interval

Each *MonitoredItem* created by the *Client* is assigned a sampling interval that is either inherited from the publishing interval of the *Subscription* or that is defined specifically to override that rate. A negative number indicates that the default sampling interval defined by the publishing interval of the *Subscription* is requested. The sampling interval indicates the fastest rate at which the *Server* should sample its underlying source for data changes.

A *Client* shall define a sampling interval of 0 if it subscribes for *Events*.

The assigned sampling interval defines a "best effort" cyclic rate that the *Server* uses to sample the item from its source. "Best effort" in this context means that the *Server* does its best to sample at this rate. Sampling at rates faster than this rate is acceptable, but not necessary to meet the needs of the *Client*. How the *Server* deals with the sampling rate and how often it actually polls its data source internally is a *Server* implementation detail. However, the time between values returned to the *Client* shall be greater or equal to the sampling interval.

The *Client* may also specify 0 for the sampling interval, which indicates that the *Server* should use the fastest practical rate. It is expected that *Servers* will support only a limited set of sampling intervals to optimize their operation. If the exact interval requested by the *Client* is

not supported by the *Server*, then the *Server* assigns to the *MonitoredItem* the most appropriate interval as determined by the *Server*. It returns this assigned interval to the *Client*. The *Server Capabilities Object* defined in IEC 62541-5 identifies the sampling intervals supported by the *Server*.

The *Server* may support data that is collected based on a sampling model or generated based on an exception-based model. The fastest supported sampling interval may be equal to 0, which indicates that the data item is exception-based rather than being sampled at some period. An exception-based model means that the underlying system does not require sampling and reports data changes.

The *Client* may use the revised sampling interval values as a hint for setting the publishing interval as well as the keep-alive count of a *Subscription*. If, for example, the smallest revised sampling interval of the *MonitoredItems* is 5 seconds, then the time before a keep-alive is sent should be longer than 5 seconds.

Note that, in many cases, the OPC UA *Server* provides access to a decoupled system and therefore has no knowledge of the data update logic. In this case, even though the OPC UA *Server* samples at the negotiated rate, the data might be updated by the underlying system at a much slower rate. In this case, changes can only be detected at this slower rate.

If the behaviour by which the underlying system updates the item is known, it will be available via the *MinimumSamplingInterval Attribute* defined in IEC 62541-3. If the *Server* specifies a value for the *MinimumSamplingInterval Attribute*, it shall always return a *revisedSamplingInterval* that is equal or higher than the *MinimumSamplingInterval* if the *Client* subscribes to the *Value Attribute*.

Clients should also be aware that the sampling by the OPC UA *Server* and the update cycle of the underlying system are usually not synchronized. This can cause additional delays in change detection, as illustrated in Figure 16.




Figure 16 – Typical delay in change detection

5.12.1.3 Monitoring mode

The monitoring mode parameter is used to enable and disable the sampling of a *MonitoredItem*, and also to provide for independently enabling and disabling the reporting of *Notifications*. This capability allows a *MonitoredItem* to be configured to sample, sample and report, or neither. Disabling sampling does not change the values of any of the other *MonitoredItem* parameter, such as its sampling interval.

When a *MonitoredItem* is enabled (i.e. when the *MonitoringMode* is changed from *DISABLED* to *SAMPLING* or *REPORTING*) or it is created in the enabled state, the *Server* shall report the first sample as soon as possible and the time of this sample becomes the starting point for the next sampling interval.

5.12.1.4 Filter

Each time a *MonitoredItem* is sampled, the *Server* evaluates the sample using the filter defined for the *MonitoredItem*. The filter parameter defines the criteria that the *Server* uses to determine if a *Notification* should be generated for the sample. The type of filter is dependent on the type of the item that is being monitored. For example, the *DataChangeFilter* and the *AggregateFilter* are used when monitoring *Variable Values* and the *EventFilter* is used when monitoring *Events*. Sampling and evaluation, including the use of filters, are described in this document. Additional filters may be defined in other parts of IEC 62541.

5.12.1.5 Queue parameters

If the sample passes the filter criteria, a *Notification* is generated and queued for transfer by the *Subscription*. The size of the queue is defined when the *MonitoredItem* is created. When the queue is full and a new *Notification* is received, the *Server* either discards the oldest *Notification* and queues the new one, or it replaces the last value added to the queue with the new one. The *MonitoredItem* is configured for one of these discard policies when the *MonitoredItem* is created. If a *Notification* is discarded for a *DataValue* and the size of the queue is larger than one, then the *Overflow* bit (flag) in the *InfoBits* portion of the *DataValue statusCode* is set. If *discardOldest* is TRUE, the oldest value gets deleted from the queue and the next value in the queue gets the flag set. If *discardOldest* is FALSE, the last value added to the queue gets replaced with the new value. The new value gets the flag set to indicate the lost values in the next *NotificationMessage*. Figure 17 illustrates the queue overflow handling.




Figure 17 – Queue overflow handling

If the queue size is one, the queue becomes a buffer that always contains the newest *Notification*. In this case, if the sampling interval of the *MonitoredItem* is faster than the publishing interval of the *Subscription*, the *MonitoredItem* will be over sampling and the *Client* will always receive the most up-to-date value. The discard policy is ignored if the queue size is one.

On the other hand, the *Client* may want to subscribe to a continuous stream of *Notifications* without any gaps, but does not want them reported at the sampling interval. In this case, the *MonitoredItem* would be created with a queue size large enough to hold all *Notifications* generated between two consecutive publishing cycles. Then, at each publishing cycle, the *Subscription* would send all *Notifications* queued for the *MonitoredItem* to the *Client*. The *Server* shall return *Notifications* for any particular item in the same order they are in the queue.

The *Server* may be sampling at a faster rate than the sampling interval to support other *Clients*; the *Client* should only expect values at the negotiated sampling interval. The *Server* may deliver fewer values than dictated by the sampling interval, based on the filter and implementation constraints. If a *DataChangeFilter* is configured for a *MonitoredItem*, it is always applied to the newest value in the queue compared to the current sample.

If, for example, the *AbsoluteDeadband* in the *DataChangeFilter* is "10", the queue could consist of values in the following order:


- 100
- 111
- 100
- 89
- 100

Queuing of data may result in unexpected behaviour when using a *Deadband* filter and the number of encountered changes is larger than the number of values that can be maintained. The new first value in the queue may not exceed the *Deadband* limit of the previous value sent to the *Client*.

The queue size is the maximum value supported by the *Server* when monitoring *Events*. In this case, the *Server* is responsible for the *Event* buffer. If *Events* are lost, an *Event* of the type *EventQueueOverflowEventType* is placed in the queue. This *Event* is generated when the first *Event* ~~has to be~~ is discarded on a *MonitoredItem* subscribing for *Events*. It is put into the Queue of the *MonitoredItem* in addition to the size of the Queue defined for this *MonitoredItem* without discarding any other *Event*. If *discardOldest* is set to TRUE it is put at the beginning of the queue and is never discarded, otherwise at the end. An aggregating *Server* shall not pass on such an *Event*. It shall be handled like other connection error scenarios.

5.12.1.6 Triggering model

The *MonitoredItems Service* allows the addition of items that are reported only when some other item (the triggering item) triggers. This is done by creating links between the triggered items and the items to report. The monitoring mode of the items to report is set to sampling-only so that it will sample and queue *Notifications* without reporting them. Figure 18 illustrates this concept.



IEC

Figure 18 – Triggering model

The triggering mechanism is a useful feature that allows *Clients* to reduce the data volume on the wire by configuring some items to sample frequently but only report when some other *Event* happens.

The following triggering behaviours are specified.

- If the monitoring mode of the triggering item is *SAMPLING*, then it is not reported when the triggering item triggers the items to report.
- If the monitoring mode of the triggering item is *REPORTING*, then it is reported when the triggering item triggers the items to report.
- If the monitoring mode of the triggering item is *DISABLED*, then the triggering item does not trigger the items to report.
- If the monitoring mode of the item to report is *SAMPLING*, then it is reported when the triggering item triggers the items to report.
- If the monitoring mode of the item to report is *REPORTING*, this effectively causes the triggering item to be ignored. All notifications of the items to report are sent after the publishing interval expires.
- If the monitoring mode of the item to report is *DISABLED*, then there will be no sampling of the item to report and therefore no notifications to report.
- The first trigger shall occur when the first notification is queued for the triggering item after the creation of the link.

Clients create and delete triggering links between a triggering item and a set of items to report. If the *MonitoredItem* that represents an item to report is deleted before its associated triggering link is deleted, the triggering link is also deleted, but the triggering item is otherwise unaffected.

Deletion of a *MonitoredItem* should not be confused with the removal of the *Attribute* that it monitors. If the *Node* that contains the *Attribute* being monitored is deleted, the *MonitoredItem* generates a *Notification* with a *StatusCode* *Bad_NodeIdUnknown* that indicates the deletion, but the *MonitoredItem* is not deleted.

5.12.2 CreateMonitoredItems

5.12.2.1 Description

This *Service* is used to create and add one or more *MonitoredItems* to a *Subscription*. A *MonitoredItem* is deleted automatically by the *Server* when the *Subscription* is deleted. Deleting a *MonitoredItem* causes its entire set of triggered item links to be deleted, but has no effect on the *MonitoredItems* referenced by the triggered items.

Calling the *CreateMonitoredItems Service* repetitively to add a small number of *MonitoredItems* each time may adversely affect the performance of the *Server*. Instead, *Clients* should add a complete set of *MonitoredItems* to a *Subscription* whenever possible.

When a *MonitoredItem* is added, the *Server* performs initialization processing for it. The initialization processing is defined by the *Notification* type of the item being monitored. *Notification* types are specified in this document and in the Access Type Specification parts of IEC 62541, such as IEC 62541-8. See IEC TR 62541-1 for a description of the Access Type Parts.

When a user adds a monitored item that the user is denied read access to, the add operation for the item shall succeed and the bad status *Bad_NotReadable* or *Bad_UserAccessDenied* shall be returned in the *Publish* response. This is the same behaviour for the case where the access rights are changed after the call to *CreateMonitoredItems*. If the access rights change to read rights, the *Server* shall start sending data for the *MonitoredItem*. The same procedure shall be applied for an *IndexRange* that does not deliver data for the current value but could deliver data in the future.

Monitored *Nodes* can be removed from the *AddressSpace* after the creation of a *MonitoredItem*. This does not affect the validity of the *MonitoredItem* but a *Bad_NodeIdUnknown* shall be returned in the *Publish* response. It is possible that the *MonitoredItem* becomes valid again if the *Node* is added again to the *AddressSpace* and the *MonitoredItem* still exists.

If a *NodeId* is known to be valid by a *Server* but the corresponding *Node Attributes* are currently not available, the *Server* may allow the creation of a *MonitoredItem* and return an appropriate *Bad_StatusCode* in the *Publish* response.

The return diagnostic info setting in the request header of the *CreateMonitoredItems* or the last *ModifyMonitoredItems Service* is applied to the *Monitored Items* and is used as the diagnostic information settings when sending *Notifications* in the *Publish* response.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

It is strongly recommended by OPC UA that a *Client* reuses a *Subscription* after a short network interruption by activating the existing *Session* on a new *SecureChannel* as described in 6.7. If a *Client* called *CreateMonitoredItems* during the network interruption and the call succeeded in the *Server* but did not return to the *Client*, then the *Client* does not know if the call succeeded. The *Client* may receive data changes for these monitored items but is not able to remove them since it does not know the *Server* handle for each monitored item. There is also no way for the *Client* to detect if the create succeeded. To delete and recreate the *Subscription* is also not an option since there may be several monitored items operating normally that should not be interrupted. To resolve this situation, the *Server Object* provides a *Method GetMonitoredItems* that returns the list of server and client handles for the monitored

items in a *Subscription*. This *Method* is defined in IEC 62541-5. The *Server* shall verify that the *Method* is called within the *Session* context of the *Session* that owns the *Subscription*.

5.12.2.2 Parameters

Table 69 defines the parameters for the *Service*.

Table 69 – CreateMonitoredItems Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> that will report <i>Notifications</i> for this <i>MonitoredItem</i> (see 7.14 for <i>IntegerId</i> definition).
timestampsToReturn	Enum TimestampsToReturn	An enumeration that specifies the timestamp <i>Attributes</i> to be transmitted for each <i>MonitoredItem</i> . The <i>TimestampsToReturn</i> enumeration is defined in 7.35. When monitoring <i>Events</i> , this applies only to <i>Event</i> fields that are of type <i>DataValue</i> .
itemsToCreate []	MonitoredItem CreateRequest	A list of <i>MonitoredItems</i> to be created and assigned to the specified <i>Subscription</i> . This structure is defined in-line with the following indented items.
itemToMonitor	ReadValueId	Identifies an item in the <i>AddressSpace</i> to monitor. To monitor for <i>Events</i> , the <i>attributeId</i> element of the <i>ReadValueId</i> structure is the id of the <i>EventNotifierAttribute</i> . The <i>ReadValueId</i> type is defined in 7.24.
monitoringMode	Enum MonitoringMode	The monitoring mode to be set for the <i>MonitoredItem</i> . The <i>MonitoringMode</i> enumeration is defined in 7.18.
requestedParameters	Monitoring Parameters	The requested monitoring parameters. <i>Servers</i> negotiate the values of these parameters based on the <i>Subscription</i> and the capabilities of the <i>Server</i> . The <i>MonitoringParameters</i> type is defined in 7.16.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	MonitoredItem CreateResult	List of results for the <i>MonitoredItems</i> to create. The size and order of the list matches the size and order of the <i>itemsToCreate</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the <i>MonitoredItem</i> to create (see 7.34 for <i>StatusCode</i> definition).
monitoredItemId	IntegerId	<i>Server</i> -assigned id for the <i>MonitoredItem</i> (see 7.14 for <i>IntegerId</i> definition). This id is unique within the <i>Subscription</i> , but might not be unique within the <i>Server</i> or <i>Session</i> . This parameter is present only if the <i>statusCode</i> indicates that the <i>MonitoredItem</i> was successfully created.
revisedSamplingInterval	Duration	The actual sampling interval that the <i>Server</i> will use. This value is based on a number of factors, including capabilities of the underlying system. The <i>Server</i> shall always return a <i>revisedSamplingInterval</i> that is equal or higher than the requested <i>samplingInterval</i> . If the requested <i>samplingInterval</i> is higher than the maximum sampling interval supported by the <i>Server</i> , the maximum sampling interval is returned.
revisedQueueSize	Counter	The actual queue size that the <i>Server</i> will use.
filterResult	Extensible Parameter MonitoringFilterResult	Contains any revised parameter values or error results associated with the <i>MonitoringFilter</i> specified in <i>requestedParameters</i> . This parameter may be emitted null if no errors occurred. The <i>MonitoringFilterResult</i> parameter type is an extensible parameter type specified in 7.17.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to create (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>itemsToCreate</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.2.3 Service results

Table 70 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 70 – CreateMonitoredItems Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.12.2.4 StatusCodes

Table 71 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 71 – CreateMonitoredItems Operation Level result codes

Symbolic Id	Description
Bad_MonitoringModelInvalid	See Table 178 for the description of this result code.
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeNoData	See Table 178 for the description of this result code. If the <i>ArrayDimensions</i> have a fixed length that cannot change and no data exists within the range of indexes specified, <i>Bad_IndexRangeNoData</i> is returned in <i>CreateMonitoredItems</i> . Otherwise if the length of the array is dynamic, the <i>Server</i> shall return this status in a <i>Publish</i> response for the <i>MonitoredItem</i> if no data exists within the range.
Bad_DataEncodingInvalid	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterInvalid	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterUnsupported	See Table 178 for the description of this result code.
Bad_FilterNotAllowed	See Table 177 for the description of this result code.
Bad_TooManyMonitoredItems	The <i>Server</i> has reached its maximum number of monitored items.

5.12.3 ModifyMonitoredItems

5.12.3.1 Description

This *Service* is used to modify *MonitoredItems* of a *Subscription*. Changes to the *MonitoredItem* settings shall be applied immediately by the *Server*. They take effect as soon as practical but not later than twice the new *revisedSamplingInterval*.

The return diagnostic info setting in the request header of the *CreateMonitoredItems* or the last *ModifyMonitoredItems* *Service* is applied to the *Monitored Items* and is used as the diagnostic information settings when sending Notifications in the *Publish* response.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

5.12.3.2 Parameters

Table 72 defines the parameters for the *Service*.

Table 72 – ModifyMonitoredItems Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> used to qualify the <i>monitoredItemId</i> (see 7.14 for <i>IntegerId</i> definition).
timestampsToReturn	Enum Timestamps ToReturn	An enumeration that specifies the timestamp <i>Attributes</i> to be transmitted for each <i>MonitoredItem</i> to be modified. The <i>TimestampsToReturn</i> enumeration is defined in 7.35. When monitoring <i>Events</i> , this applies only to <i>Event</i> fields that are of type <i>DataValue</i> .
itemsToModify []	MonitoredItemModifyRequest	The list of <i>MonitoredItems</i> to modify. This structure is defined in-line with the following indented items.
monitoredItemId	IntegerId	<i>Server</i> -assigned id for the <i>MonitoredItem</i> .
requestedParameters	Monitoring Parameters	The requested values for the monitoring parameters. The <i>MonitoringParameters</i> type is defined in 7.16. If the number of notifications in the queue exceeds the new queue size, the notifications exceeding the size shall be discarded following the configured discard policy.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	MonitoredItemModifyResult	List of results for the <i>MonitoredItems</i> to modify. The size and order of the list matches the size and order of the <i>itemsToModify</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>Status Code</i> for the <i>MonitoredItem</i> to be modified (see 7.34 for <i>Status Code</i> definition).
revisedSampling Interval	Duration	The actual sampling interval that the <i>Server</i> will use. The <i>Server</i> returns the value it will actually use for the sampling interval. This value is based on a number of factors, including capabilities of the underlying system. The <i>Server</i> shall always return a <i>revisedSamplingInterval</i> that is equal or higher than the requested <i>samplingInterval</i> . If the requested <i>samplingInterval</i> is higher than the maximum sampling interval supported by the <i>Server</i> , the maximum sampling interval is returned.
revisedQueueSize	Counter	The actual queue size that the <i>Server</i> will use.
filterResult	Extensible Parameter MonitoringFilter Result	Contains any revised parameter values or error results associated with the <i>MonitoringFilter</i> specified in the request. This parameter may be omitted null if no errors occurred. The <i>MonitoringFilterResult</i> parameter type is an extensible parameter type specified in 7.17.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to modify (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>itemsToModify</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.3.3 Service results

Table 73 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 73 – ModifyMonitoredItems Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.12.3.4 StatusCodes

Table 74 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 74 – ModifyMonitoredItems Operation Level result codes

Symbolic Id	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterInvalid	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterUnsupported	See Table 178 for the description of this result code.
Bad_FilterNotAllowed	See Table 177 for the description of this result code.

5.12.4 SetMonitoringMode

5.12.4.1 Description

This *Service* is used to set the monitoring mode for one or more *MonitoredItems* of a *Subscription*. Setting the mode to DISABLED causes all queued *Notifications* to be deleted.

5.12.4.2 Parameters

Table 75 defines the parameters for the *Service*.

Table 75 – SetMonitoringMode service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> used to qualify the <i>monitoredItemIds</i> (see 7.14 for <i>IntegerId</i> definition).
monitoringMode	Enum MonitoringMode	The monitoring mode to be set for the <i>MonitoredItems</i> . The <i>MonitoringMode</i> enumeration is defined in 7.18.
monitoredItemIds []	IntegerId	List of <i>Server</i> -assigned ids for the <i>MonitoredItems</i> .
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of <i>StatusCodes</i> for the <i>MonitoredItems</i> to enable/disable (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to enable/disable (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.4.3 Service results

Table 76 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 76 – SetMonitoringMode Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_MonitoringModeInvalid	See Table 178 for the description of this result code.

5.12.4.4 StatusCodes

Table 77 defines values for the operation level *statusCode results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 77 – SetMonitoringMode Operation Level result codes

Value	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.12.5 SetTriggering

5.12.5.1 Description

This *Service* is used to create and delete triggering links for a triggering item. The triggering item and the items to report shall belong to the same *Subscription*.

Each triggering link links a triggering item to an item to report. Each link is represented by the *MonitoredItem* id for the item to report. An error code is returned if this id is invalid.

See 5.12.1.6 for a description of the triggering model.

5.12.5.2 Parameters

Table 78 defines the parameters for the *Service*.

Table 78 – SetTriggering Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> that contains the triggering item and the items to report (see 7.14 for <i>IntegerId</i> definition).
triggeringItemId	IntegerId	<i>Server</i> -assigned id for the <i>MonitoredItem</i> used as the triggering item.
linksToAdd []	IntegerId	The list of <i>Server</i> -assigned ids of the items to report that are to be added as triggering links. The list of <i>linksToRemove</i> is processed before the <i>linksToAdd</i> .
linksToRemove []	IntegerId	The list of <i>Server</i> -assigned ids of the items to report for the triggering links to be deleted. The list of <i>linksToRemove</i> is processed before the <i>linksToAdd</i> .
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
addResults []	Status Code	List of <i>StatusCodes</i> for the items to add (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>linksToAdd</i> parameter specified in the request.
addDiagnosticInfos []	Diagnostic Info	List of diagnostic information for the links to add (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>linksToAdd</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.
removeResults []	Status Code	List of <i>StatusCodes</i> for the items to delete. The size and order of the list matches the size and order of the <i>linksToRemove</i> parameter specified in the request.
removeDiagnosticInfos []	Diagnostic Info	List of diagnostic information for the links to delete. The size and order of the list matches the size and order of the <i>linksToRemove</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.5.3 Service results

Table 79 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in 7.34.

Table 79 – SetTriggering Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.12.5.4 StatusCodes

Table 80 defines values for the results parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 80 – SetTriggering Operation Level result codes

Symbolic Id	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.12.6 DeleteMonitoredItems

5.12.6.1 Description

This *Service* is used to remove one or more *MonitoredItems* of a *Subscription*. When a *MonitoredItem* is deleted, its triggered item links are also deleted.

Successful removal of a *MonitoredItem*, however, might not remove *Notifications* for the *MonitoredItem* that are in the process of being sent by the *Subscription*. Therefore, *Clients* may receive *Notifications* for the *MonitoredItem* after they have received a positive response that the *MonitoredItem* has been deleted.

5.12.6.2 Parameters

Table 81 defines the parameters for the *Service*.

Table 81 – DeleteMonitoredItems Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> that contains the <i>MonitoredItems</i> to be deleted (see 7.14 for <i>IntegerId</i> definition).
monitoredItemIds []	IntegerId	List of <i>Server</i> -assigned ids for the <i>MonitoredItems</i> to be deleted.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of <i>StatusCodes</i> for the <i>MonitoredItems</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.6.3 Service results

Table 82 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 82 – DeleteMonitoredItems Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.12.6.4 StatusCodes

Table 83 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 83 – DeleteMonitoredItems Operation Level result codes

Symbolic Id	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.13 Subscription Service Set

5.13.1 Subscription model

5.13.1.1 Description

Subscriptions are used to report *Notifications* to the *Client*. Their general behaviour is summarized below. Their precise behaviour is described in 5.13.1.2.

- Subscriptions* have a set of *MonitoredItems* assigned to them by the *Client*. *MonitoredItems* generate *Notifications* that are to be reported to the *Client* by the *Subscription* (see 5.12.1 for a description of *MonitoredItems*).
- Subscriptions* have a publishing interval. The publishing interval of a *Subscription* defines the cyclic rate at which the *Subscription* executes. Each time it executes, it attempts to send a *NotificationMessage* to the *Client*. *NotificationMessages* contain *Notifications* that have not yet been reported to *Client*.
- NotificationMessages* are sent to the *Client* in response to *Publish* requests. *Publish* requests are normally queued to the *Session* as they are received, and one is de-queued and processed by a subscription related to this *Session* for each publishing cycle, if there are *Notifications* to report. When there are not, the *Publish* request is not de-queued from the *Session*, and the *Server* waits until the next cycle and checks again for *Notifications*.
- At the beginning of a cycle, if there are *Notifications* to send but there are no *Publish* requests queued, the *Server* enters a wait state for a *Publish* request to be received. When one is received, it is processed immediately without waiting for the next publishing cycle.
- NotificationMessages* are uniquely identified by sequence numbers that enable *Clients* to detect missed *Messages*. The publishing interval also defines the default sampling interval for its *MonitoredItems*.
- Subscriptions* have a keep-alive counter that counts the number of consecutive publishing cycles in which there have been no *Notifications* to report to the *Client*. When the maximum keep-alive count is reached, a *Publish* request is de-queued and used to return a keep-alive *Message*. This keep-alive *Message* informs the *Client* that the *Subscription* is still active. Each keep-alive *Message* is a response to a *Publish* request in which the *notificationMessage* parameter does not contain any *Notifications* and that contains the sequence number of the next *NotificationMessage* that is to be sent. In the clauses that follow, the term *NotificationMessage* refers to a response to a *Publish* request in which the *notificationMessage* parameter actually contains one or more *Notifications*, as opposed to a keep-alive *Message* in which this parameter contains no *Notifications*. The maximum

keep-alive count is set by the *Client* during *Subscription* creation and may be subsequently modified using the *ModifySubscription Service*. Similar to *Notification* processing described in (c) above, if there are no *Publish* requests queued, the *Server* waits for the next one to be received and sends the keep-alive immediately without waiting for the next publishing cycle.

- g) Publishing by a *Subscription* may be enabled or disabled by the *Client* when created, or subsequently using the *SetPublishingMode Service*. Disabling causes the *Subscription* to cease sending *NotificationMessages* to the *Client*. However, the *Subscription* continues to execute cyclically and continues to send keep-alive *Messages* to the *Client*.
- h) *Subscriptions* have a lifetime counter that counts the number of consecutive publishing cycles in which there have been no *Publish* requests available to send a *Publish* response for the *Subscription*. Any *Service* call that uses the *SubscriptionId* or the processing of a *Publish* response resets the lifetime counter of this *Subscription*. When this counter reaches the value calculated for the lifetime of a *Subscription* based on the *MaxKeepAliveCount* parameter in the *CreateSubscription Service* (5.13.2), the *Subscription* is closed. Closing the *Subscription* causes its *MonitoredItems* to be deleted. In addition the *Server* shall issue a *StatusChangeNotification notificationMessage* with the status code *Bad_Timeout*. The *StatusChangeNotification notificationMessage* type is defined in 7.20.4.
- i) *Sessions* maintain a retransmission queue of sent *NotificationMessages*. *NotificationMessages* are retained in this queue until they are acknowledged. The *Session* shall maintain a retransmission queue size of at least two times the number of *Publish* requests per *Session* the *Server* supports. The minimum size of the retransmission queue may be changed by a *Profile* in IEC 62541-7. The minimum number of *Publish* requests per *Session* the *Server* shall support is defined in IEC 62541-7. *Clients* are required to acknowledge *NotificationMessages* as they are received. In the case of a retransmission queue overflow, the oldest sent *NotificationMessage* gets deleted. If a *Subscription* is transferred to another *Session*, the queued *NotificationMessages* for this *Subscription* are moved from the old to the new *Session*.

The sequence number is an unsigned 32-bit integer that is incremented by one for each *NotificationMessage* sent. The value 0 is never used for the sequence number. The first *NotificationMessage* sent on a *Subscription* has a sequence number of 1. If the sequence number rolls over, it rolls over to 1.

When a *Subscription* is created, the first *Message* is sent at the end of the first publishing cycle to inform the *Client* that the *Subscription* is operational. A *NotificationMessage* is sent if there are *Notifications* ready to be reported. If there are none, a keep-alive *Message* is sent instead that contains a sequence number of 1, indicating that the first *NotificationMessage* has not yet been sent. This is the only time a keep-alive *Message* is sent without waiting for the maximum keep-alive count to be reached, as specified in (f) above.

A *Client* shall be prepared for receiving *Publish* responses for a *Subscription* more frequently than the corresponding publishing interval. One example is the situation where the number of available notifications exceeds the *Subscription* setting *maxNotificationsPerPublish*. A *Client* is always able to control the timing of the *Publish* responses by not queueing *Publish* requests. If a *Client* does not queue *Publish* requests in the *Server*, the *Server* can only send a *Publish* response if it receives a new *Publish* request. This would increase latency for delivery of notifications but allows a *Client* to throttle the number of received *Publish* responses in high load situations.

The value of the sequence number is never reset during the lifetime of a *Subscription*. Therefore, the same sequence number shall not be reused on a *Subscription* until over four billion *NotificationMessages* have been sent. At a continuous rate of one thousand *NotificationMessages* per second on a given *Subscription*, it would take roughly fifty days for the same sequence number to be reused. This allows *Clients* to safely treat sequence numbers as unique.

Sequence numbers are also used by *Clients* to acknowledge the receipt of *NotificationMessages*. *Publish* requests allow the *Client* to acknowledge all *Notifications* up to a specific sequence number and to acknowledge the sequence number of the last *NotificationMessage* received. One or more gaps may exist in between. Acknowledgements allow the *Server* to delete *NotificationMessages* from its retransmission queue.

Clients may ask for retransmission of selected *NotificationMessages* using the *Republish Service*. This *Service* returns the requested *Message*.

Subscriptions are designed to work independent of the actual communication connection between OPC UA *Client* and *Server* and independent of a *Session*. Short communication interruptions can be handled without losing data or events. To make sure that longer communication interruptions or planned disconnects can be handled without losing data or events, an OPC UA *Server* may support durable *Subscriptions*. If this feature is supported, the *Server* accepts a high *Subscription RequestedLifetimeCount* and large *MonitoredItem QueueSize* parameter settings. Subclause 6.8 describes how durable *Subscriptions* can be created and used.

5.13.1.2 State table

The state table formally describes the operation of the *Subscription*. The following model of operations is described by this state table. This description applies when publishing is enabled or disabled for the *Subscription*.

After creation of the *Subscription*, the *Server* starts the publishing timer and restarts it whenever it expires. If the timer expires the number of times defined for the *Subscription* lifetime without having received a *Subscription Service* request from the *Client*, the *Subscription* assumes that the *Client* is no longer present, and terminates.

Clients send *Publish* requests to *Servers* to receive *Notifications*. *Publish* requests are not directed to any one *Subscription* and, therefore, may be used by any *Subscription*. Each contains acknowledgements for one or more *Subscriptions*. These acknowledgements are processed when the *Publish* request is received. The *Server* then queues the request in a queue shared by all *Subscriptions*, except in the following cases.

- a) The previous *Publish* response indicated that there were still more *Notifications* ready to be transferred and there were no more *Publish* requests queued to transfer them.
- b) The publishing timer of a *Subscription* expired and there were either *Notifications* to be sent or a keep-alive *Message* to be sent.

In these cases, the newly received *Publish* request is processed immediately by the first *Subscription* to encounter either case (a) or case (b).

Each time the publishing timer expires, it is immediately reset. If there are *Notifications* or a keep-alive *Message* to be sent, it de-queues and processes a *Publish* request. When a *Subscription* processes a *Publish* request, it accesses the queues of its *MonitoredItems* and de-queues its *Notifications*, if any. It returns these *Notifications* in the response, setting the *moreNotifications* flag if it was not able to return all available *Notifications* in the response.

If there were *Notifications* or a keep-alive *Message* to be sent but there were no *Publish* requests queued, the *Subscription* assumes that the *Publish* request is late and waits for the next *Publish* request to be received, as described in case (b).

If the *Subscription* is disabled when the publishing timer expires or if there are no *Notifications* available, it enters the keep-alive state and sets the keep-alive counter to its maximum value as defined for the *Subscription*.

While in the keep-alive state, it checks for *Notifications* each time the publishing timer expires. If one or more *Notifications* have been generated, a *Publish* request is de-queued

and a *NotificationMessage* is returned in the response. However, if the publishing timer expires without a *Notification* becoming available, a *Publish* request is de-queued and a keep-alive *Message* is returned in the response. The *Subscription* then returns to the normal state of waiting for the publishing timer to expire again. If, in either of these cases, there are no *Publish* requests queued, the *Subscription* waits for the next *Publish* request to be received, as described in case (b).

The *Subscription* states are defined in Table 84.

Table 84 – Subscription states

State	Description
CLOSED	The <i>Subscription</i> has not yet been created or has terminated.
CREATING	The <i>Subscription</i> is being created.
NORMAL	The <i>Subscription</i> is cyclically checking for <i>Notifications</i> from its <i>MonitoredItems</i> . The keep-alive counter is not used in this state.
LATE	The publishing timer has expired and there are <i>Notifications</i> available or a keep-alive <i>Message</i> is ready to be sent, but there are no <i>Publish</i> requests queued. When in this state, the next <i>Publish</i> request is processed when it is received. The keep-alive counter is not used in this state.
KEEPAIVE	The <i>Subscription</i> is cyclically checking for <i>Notifications</i> from its <i>MonitoredItems</i> or for the keep-alive counter to count down to 0 from its maximum.

The state table is described in Table 85. The following rules and conventions apply.

- Events* represent the receipt of *Service* requests and the occurrence internal *Events*, such as timer expirations.
- Service* requests *Events* may be accompanied by conditions that test *Service* parameter values. Parameter names begin with a lower case letter.
- Internal *Events* may be accompanied by conditions that test state *Variable* values. State *Variables* are defined in 5.13.1.3. They begin with an upper case letter.
- Service* request and internal *Events* may be accompanied by conditions represented by functions whose return value is tested. Functions are identified by "()" after their name. They are described in 5.13.1.4.
- When an *Event* is received, the first transition for the current state is located and the transitions are searched sequentially for the first transition that meets the *Event* or conditions criteria. If none are found, the *Event* is ignored.
- Actions are described by functions and state *Variable* manipulations.
- The LifetimeTimerExpires *Event* is triggered when its corresponding counter reaches zero.

Table 85 – Subscription state table

#	Current State	Event/Conditions	Action	Next State
1	CLOSED	Receive CreateSubscription Request	CreateSubscription()	CREATING
2	CREATING	CreateSubscription fails	ReturnNegativeResponse()	CLOSED
3	CREATING	CreateSubscription succeeds	InitializeSubscription() MessageSent = FALSE ReturnResponse()	NORMAL
4	NORMAL	Receive <i>Publish</i> Request && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && MoreNotifications == FALSE))	DeleteAkedNotificationMsgs() EnqueuePublishingReq()	NORMAL
5	NORMAL	Receive <i>Publish</i> Request && PublishingEnabled == TRUE && MoreNotifications == TRUE	ResetLifetimeCounter() DeleteAkedNotificationMsgs() ReturnNotifications() MessageSent = TRUE	NORMAL
6	NORMAL	PublishingTimer Expires && PublishingReqQueued == TRUE && PublishingEnabled == TRUE && NotificationsAvailable == TRUE	ResetLifetimeCounter() StartPublishingTimer() DequeuePublishReq() ReturnNotifications() MessageSent == TRUE	NORMAL
7	NORMAL	PublishingTimer Expires && PublishingReqQueued == TRUE && MessageSent == FALSE && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	ResetLifetimeCounter() StartPublishingTimer() DequeuePublishReq() ReturnKeepAlive() MessageSent == TRUE	NORMAL
8	NORMAL	PublishingTimer Expires && PublishingReqQueued == FALSE && (MessageSent == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == TRUE))	StartPublishingTimer()	LATE

#	Current State	Event/Conditions	Action	Next State
9	NORMAL	PublishingTimer Expires && MessageSent == TRUE && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	StartPublishingTimer() ResetKeepAliveCounter()	KEEPALIVE
10	LATE	Receive <i>Publish</i> Request && PublishingEnabled == TRUE && (NotificationsAvailable == TRUE MoreNotifications == TRUE)	ResetLifetimeCounter() DeleteAkedNotificationMsgs() ReturnNotifications() MessageSent = TRUE	NORMAL
11	LATE	Receive <i>Publish</i> Request && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE && MoreNotifications == FALSE))	ResetLifetimeCounter() DeleteAkedNotificationMsgs() ReturnKeepAlive() MessageSent = TRUE	KEEPALIVE
12	LATE	PublishingTimer Expires	StartPublishingTimer()	LATE
13	KEEPALIVE	Receive <i>Publish</i> Request	DeleteAkedNotificationMsgs() EnqueuePublishingReq()	KEEPALIVE
14	KEEPALIVE	PublishingTimer Expires && PublishingEnabled == TRUE && NotificationsAvailable == TRUE && PublishingReqQueued == TRUE	ResetLifetimeCounter() StartPublishingTimer() DequeuePublishReq() ReturnNotifications() MessageSent == TRUE	NORMAL
15	KEEPALIVE	PublishingTimer Expires && PublishingReqQueued == TRUE && KeepAliveCounter == 1 && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	StartPublishingTimer() DequeuePublishReq() ReturnKeepAlive() ResetKeepAliveCounter()	KEEPALIVE

#	Current State	Event/Conditions	Action	Next State
16	KEEPALIVE	PublishingTimer Expires && KeepAliveCounter > 1 && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	StartPublishingTimer() KeepAliveCounter--	KEEPALIVE
17	KEEPALIVE	PublishingTimer Expires && PublishingReqQueued == FALSE && (KeepAliveCounter == 1 (KeepAliveCounter > 1 && PublishingEnabled == TRUE && NotificationsAvailable == TRUE))	StartPublishingTimer()	LATE
18	NORMAL LATE KEEPALIVE	Receive ModifySubscription Request	ResetLifetimeCounter() UpdateSubscriptionParams() ReturnResponse()	SAME
19	NORMAL LATE KEEPALIVE	Receive SetPublishingMode Request	ResetLifetimeCounter() SetPublishingEnabled() MoreNotifications = FALSE ReturnResponse()	SAME
20	NORMAL LATE KEEPALIVE	Receive Republish Request && RequestedMessageFound == TRUE	ResetLifetimeCounter() ReturnResponse()	SAME
21	NORMAL LATE KEEPALIVE	Receive Republish Request && RequestedMessageFound == FALSE	ResetLifetimeCounter() ReturnNegativeResponse()	SAME
22	NORMAL LATE KEEPALIVE	Receive TransferSubscriptions Request && SessionChanged() == FALSE	ResetLifetimeCounter() ReturnNegativeResponse ()	SAME
23	NORMAL LATE KEEPALIVE	Receive TransferSubscriptions Request && SessionChanged() == TRUE && ClientValidated() ==TRUE	SetSession() ResetLifetimeCounter() ReturnResponse() IssueStatusChangeNotification()	SAME
24	NORMAL LATE KEEPALIVE	Receive TransferSubscriptions Request && SessionChanged() == TRUE && ClientValidated() == FALSE	ReturnNegativeResponse()	SAME
25	NORMAL LATE KEEPALIVE	Receive DeleteSubscriptions Request && SubscriptionAssignedToClient ==TRUE	DeleteMonitoredItems() DeleteClientPublReqQueue()	CLOSED

#	Current State	Event/Conditions	Action	Next State
26	NORMAL LATE KEEPALIVE	Receive DeleteSubscriptions Request && SubscriptionAssignedToClient ==FALSE	ResetLifetimeCounter() ReturnNegativeResponse()	SAME
27	NORMAL LATE KEEPALIVE	LifetimeCounter == 1 The LifetimeCounter is decremented if PublishingTimer expires and PublishingReqQueued == FALSE The LifetimeCounter is reset if PublishingReqQueued == TRUE.	DeleteMonitoredItems() IssueStatusChangeNotification()	CLOSED

5.13.1.3 State variables and parameters

The state variables are defined alphabetically in Table 86.

Table 86 – State variables and parameters

State Variable	Description
MoreNotifications	A boolean value that is set to TRUE only by the CreateNotificationMsg() when there were too many <i>Notifications</i> for a single <i>NotificationMessage</i> .
LatePublishRequest	A boolean value that is set to TRUE to reflect that, the last time the publishing timer expired, there were no <i>Publish</i> requests queued.
LifetimeCounter	A value that contains the number of consecutive publishing timer expirations without <i>Client</i> activity before the <i>Subscription</i> is terminated.
MessageSent	A boolean value that is set to TRUE to mean that either a <i>NotificationMessage</i> or a keep-alive <i>Message</i> has been sent on the <i>Subscription</i> . It is a flag that is used to ensure that either a <i>NotificationMessage</i> or a keep-alive <i>Message</i> is sent out the first time the publishing timer expires.
NotificationsAvailable	A boolean value that is set to TRUE only when there is at least one <i>MonitoredItem</i> that is in the reporting mode and that has a <i>Notification</i> queued or there is at least one item to report whose triggering item has triggered and that has a <i>Notification</i> queued. The transition of this state <i>Variable</i> from FALSE to TRUE creates the "New <i>Notification</i> Queued" <i>Event</i> in the state table.
PublishingEnabled	The parameter that requests publishing to be enabled or disabled.
PublishingReqQueued	A boolean value that is set to TRUE only when there is a <i>Publish</i> request <i>Message</i> queued to the <i>Subscription</i> .
RequestedMessageFound	A boolean value that is set to TRUE only when the <i>Message</i> requested to be retransmitted was found in the retransmission queue.
SeqNum	The value that records the value of the sequence number used in <i>NotificationMessages</i> .
SubscriptionAssignedToClient	A boolean value that is set to TRUE only when the <i>Subscription</i> requested to be deleted is assigned to the <i>Client</i> that issued the request. A <i>Subscription</i> is assigned to the <i>Client</i> that created it. That assignment can only be changed through successful completion of the <i>TransferSubscriptions Service</i> .

5.13.1.4 Functions

The action functions are defined alphabetically in Table 87.

Table 87 – Functions

Function	Description
ClientValidated()	A boolean function that returns TRUE only when the <i>Client</i> that is submitting a <i>TransferSubscriptions</i> request is operating on behalf of the same user and supports the same <i>Profiles</i> as the <i>Client</i> of the previous <i>Session</i> .
CreateNotificationMsg()	Increment the SeqNum and create a <i>NotificationMessage</i> from the <i>MonitoredItems</i> assigned to the <i>Subscription</i> . Save the newly-created <i>NotificationMessage</i> in the retransmission queue. If all available <i>Notifications</i> can be sent in the <i>Publish</i> response, the <i>MoreNotifications</i> state <i>Variable</i> is set to FALSE. Otherwise, it is set to TRUE.
CreateSubscription()	Attempt to create the <i>Subscription</i> .
DeleteAkedNotificationMsgs()	Delete the <i>NotificationMessages</i> from the retransmission queue that were acknowledged by the request.
DeleteClientPublReqQueue()	Clear the <i>Publish</i> request queue for the <i>Client</i> that is sending the <i>DeleteSubscriptions</i> request, if there are no more <i>Subscriptions</i> assigned to that <i>Client</i> .
DeleteMonitoredItems()	Delete all <i>MonitoredItems</i> assigned to the <i>Subscription</i> .
DequeuePublishReq()	De-queue a publishing request in first-in first-out order. Validate if the publish request is still valid by checking the <i>timeoutHint</i> in the <i>RequestHeader</i> . If the request timed out, send a <i>Bad_Timeout</i> service result for the request and de-queue another publish request. ResetLifetimeCounter()
EnqueuePublishingReq()	Enqueue the publishing request.
InitializeSubscription()	ResetLifetimeCounter() MoreNotifications = FALSE PublishRateChange = FALSE PublishingEnabled = value of publishingEnabled parameter in the <i>CreateSubscription</i> request PublishingReqQueued = FALSE SeqNum = 0 SetSession() StartPublishingTimer()
IssueStatusChangeNotification()	Issue a <i>StatusChangeNotification notificationMessage</i> with a status code for the status change of the <i>Subscription</i> . The <i>StatusChangeNotification notificationMessage</i> type is defined in 7.20.4. <i>Bad_Timeout</i> status code is used if the lifetime expires and <i>Good_SubscriptionTransferred</i> is used if the <i>Subscriptions</i> was transferred to another <i>Session</i> .
ResetKeepAliveCounter()	Reset the keep-alive counter to the maximum keep-alive count of the <i>Subscription</i> . The maximum keep-alive count is set by the <i>Client</i> when the <i>Subscription</i> is created and may be modified using the <i>ModifySubscription Service</i> .
ResetLifetimeCounter()	Reset the <i>LifetimeCounter Variable</i> to the value specified for the lifetime of a <i>Subscription</i> in the <i>CreateSubscription Service</i> (5.13.2).
ReturnKeepAlive()	CreateKeepAliveMsg() ReturnResponse()
ReturnNegativeResponse ()	Return a <i>Service</i> response indicating the appropriate <i>Service</i> level error. No parameters are returned other than the responseHeader that contains the <i>Service</i> level <i>StatusCode</i> .
ReturnNotifications()	CreateNotificationMsg() ReturnResponse() If (MoreNotifications == TRUE) && (PublishingReqQueued == TRUE) { DequeuePublishReq() Loop through this function again }

Function	Description
ReturnResponse()	Return the appropriate response, setting the appropriate parameter values and <i>StatusCodes</i> defined for the <i>Service</i> .
SessionChanged()	A boolean function that returns TRUE only when the <i>Session</i> used to send a <i>TransferSubscriptions</i> request is different from the <i>Client Session</i> currently associated with the <i>Subscription</i> .
SetPublishingEnabled ()	Set the <i>PublishingEnabled</i> state <i>Variable</i> to the value of the <i>publishingEnabled</i> parameter received in the request.
SetSession	Set the <i>Session</i> information for the <i>Subscription</i> to match the <i>Session</i> on which the <i>TransferSubscriptions</i> request was issued.
StartPublishingTimer()	Start or restart the publishing timer and decrement the <i>LifetimeCounter Variable</i> .
UpdateSubscriptionParams()	Negotiate and update the <i>Subscription</i> parameters. If the new keep-alive interval is less than the current value of the keep-alive counter, perform <i>ResetKeepAliveCounter()</i> and <i>ResetLifetimeCounter()</i> .

5.13.2 CreateSubscription

5.13.2.1 Description

This *Service* is used to create a *Subscription*. *Subscriptions* monitor a set of *MonitoredItems* for *Notifications* and return them to the *Client* in response to *Publish* requests.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

5.13.2.2 Parameters

Table 88 defines the parameters for the *Service*.

Table 88 – CreateSubscription Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
requestedPublishingInterval	Duration	This interval defines the cyclic rate that the <i>Subscription</i> is being requested to return <i>Notifications</i> to the <i>Client</i> . This interval is expressed in milliseconds. This interval is represented by the publishing timer in the <i>Subscription</i> state table (see 5.13.1.2). The negotiated value for this parameter returned in the response is used as the default sampling interval for <i>MonitoredItems</i> assigned to this <i>Subscription</i> . If the requested value is 0 or negative, the <i>Server</i> shall revise with the fastest supported publishing interval.
requestedLifetimeCount	Counter	Requested lifetime count (see 7.5 for <i>Counter</i> definition). The lifetime count shall be a minimum of three times the keep keep-alive count. When the publishing timer has expired this number of times without a <i>Publish</i> request being available to send a <i>NotificationMessage</i> , then the <i>Subscription</i> shall be deleted by the <i>Server</i> .
requestedMaxKeepAliveCount	Counter	Requested maximum keep-alive count (see 7.5 for <i>Counter</i> definition). When the publishing timer has expired this number of times without requiring any <i>NotificationMessage</i> to be sent, the <i>Subscription</i> sends a keep-alive <i>Message</i> to the <i>Client</i> . The negotiated value for this parameter is returned in the response. If the requested value is 0, the <i>Server</i> shall revise with the smallest supported keep-alive count.
maxNotificationsPerPublish	Counter	The maximum number of notifications that the <i>Client</i> wishes to receive in a single <i>Publish</i> response. A value of zero indicates that there is no limit. The number of notifications per <i>Publish</i> is the sum of monitoredItems in the <i>DataChangeNotification</i> and events in the <i>EventNotificationList</i> .

Name	Type	Description
publishingEnabled	Boolean	A <i>Boolean</i> parameter with the following values: TRUE publishing is enabled for the <i>Subscription</i> . FALSE publishing is disabled for the <i>Subscription</i> . The value of this parameter does not affect the value of the monitoring mode <i>Attribute of MonitoredItems</i> .
priority	Byte	Indicates the relative priority of the <i>Subscription</i> . When more than one <i>Subscription</i> needs to send <i>Notifications</i> , the <i>Server</i> should de-queue a Publish request to the <i>Subscription</i> with the highest <i>priority</i> number. For <i>Subscriptions</i> with equal <i>priority</i> the <i>Server</i> should de-queue Publish requests in a round-robin fashion. A Client that does not require special priority settings should set this value to zero.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition). This identifier shall be unique for the entire <i>Server</i> , not just for the <i>Session</i> , in order to allow the <i>Subscription</i> to be transferred to another <i>Session</i> using the TransferSubscriptions service. After <i>Server</i> start-up the generation of <i>subscriptionIds</i> should start from a random <i>IntegerId</i> or continue from the point before the restart.
revisedPublishingInterval	Duration	The actual publishing interval that the <i>Server</i> will use, expressed in milliseconds. The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.
revisedLifetimeCount	Counter	The lifetime of the <i>Subscription</i> shall be a minimum of three times the keep-alive interval negotiated by the <i>Server</i> .
revisedMaxKeepAliveCount	Counter	The actual maximum keep-alive count (see 7.5 for <i>Counter</i> definition). The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.

5.13.2.3 Service results

Table 89 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 89 – CreateSubscription Service result codes

Symbolic Id	Description
Bad_TooManySubscriptions	The <i>Server</i> has reached its maximum number of subscriptions.

5.13.3 ModifySubscription

5.13.3.1 Description

This *Service* is used to modify a *Subscription*.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

Changes to the *Subscription* settings shall be applied immediately by the *Server*. They take effect as soon as practical but not later than twice the new *revisedPublishingInterval*.

5.13.3.2 Parameters

Table 90 defines the parameters for the *Service*.

Table 90 – ModifySubscription Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition).
requestedPublishingInterval	Duration	This interval defines the cyclic rate at which the <i>Subscription</i> is being requested to return <i>Notifications</i> to the <i>Client</i>. This interval is expressed in milliseconds. This interval is represented by the publishing timer in the <i>Subscription</i> state table (see 5.13.1.2). The negotiated value for this parameter returned in the response is used as the default sampling interval for <i>MonitoredItems</i> assigned to this <i>Subscription</i>. If the requested value is 0 or negative, the <i>Server</i> shall revise with the fastest supported publishing interval.
requestedLifetimeCount	Counter	Requested lifetime count (see 7.5 for <i>Counter</i> definition). The lifetime count shall be a minimum of three times the keep-alive count. When the publishing timer has expired this number of times without a <i>Publish</i> request being available to send a <i>NotificationMessage</i>, then the <i>Subscription</i> shall be deleted by the <i>Server</i>.
requestedMaxKeepAliveCount	Counter	Requested maximum keep-alive count (see 7.5 for <i>Counter</i> definition). When the publishing timer has expired this number of times without requiring any <i>NotificationMessage</i> to be sent, the <i>Subscription</i> sends a keep-alive <i>Message</i> to the <i>Client</i>. The negotiated value for this parameter is returned in the response. If the requested value is 0, the <i>Server</i> shall revise with the smallest supported keep-alive count.
maxNotificationsPerPublish	Counter	The maximum number of notifications that the <i>Client</i> wishes to receive in a single <i>Publish</i> response. A value of zero indicates that there is no limit.
priority	Byte	Indicates the relative priority of the <i>Subscription</i>. When more than one <i>Subscription</i> needs to send <i>Notifications</i>, the <i>Server</i> should de-queue a <i>Publish</i> request to the <i>Subscription</i> with the highest <i>priority</i> number. For <i>Subscriptions</i> with equal <i>priority</i> the <i>Server</i> should de-queue <i>Publish</i> requests in a round-robin fashion. ~~Any <i>Subscription</i> that needs to send a keep-alive <i>Message</i> shall take precedence regardless of its <i>priority</i>, in order to prevent the <i>Subscription</i> from expiring.~~ A <i>Client</i> that does not require special priority settings should set this value to zero.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
revisedPublishingInterval	Duration	The actual publishing interval that the <i>Server</i> will use, expressed in milliseconds. The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.
revisedLifetimeCount	Counter	The lifetime of the <i>Subscription</i> shall be a minimum of three times the keep-alive interval negotiated by the <i>Server</i> .
revisedMaxKeepAliveCount	Counter	The actual maximum keep-alive count (see 7.5 for <i>Counter</i> definition). The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.

5.13.3.3 Service results

Table 91 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 91 – ModifySubscription Service result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.13.4 SetPublishingMode

5.13.4.1 Description

This *Service* is used to enable sending of *Notifications* on one or more *Subscriptions*.

5.13.4.2 Parameters

Table 92 defines the parameters for the *Service*.

Table 92 – SetPublishingMode Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
publishingEnabled	Boolean	A <i>Boolean</i> parameter with the following values: TRUE publishing of <i>NotificationMessages</i> is enabled for the <i>Subscription</i> . FALSE publishing of <i>NotificationMessages</i> is disabled for the <i>Subscription</i> . The value of this parameter does not affect the value of the monitoring mode <i>Attribute</i> of <i>MonitoredItems</i> . Setting this value to FALSE does not discontinue the sending of keep-alive <i>Messages</i> .
subscriptionIds []	IntegerId	List of <i>Server</i> -assigned identifiers for the <i>Subscriptions</i> to enable or disable (see 7.14 for <i>IntegerId</i> definition).
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCodes	List of <i>StatusCodes</i> for the <i>Subscriptions</i> to enable/disable (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Subscriptions</i> to enable/disable (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.4.3 Service results

Table 93 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 93 – SetPublishingMode Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.13.4.4 StatusCodes

Table 94 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 94 – SetPublishingMode Operation Level result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.13.5 Publish

5.13.5.1 Description

This *Service* is used for two purposes. First, it is used to acknowledge the receipt of *NotificationMessages* for one or more *Subscriptions*. Second, it is used to request the *Server* to return a *NotificationMessage* or a keep-alive *Message*. Since *Publish* requests are not directed to a specific *Subscription*, they may be used by any *Subscription*. 5.13.1.2 describes the use of the *Publish Service*.

Client strategies for issuing *Publish* requests may vary depending on the networking delays between the *Client* and the *Server*. In many cases, the *Client* may wish to issue a *Publish* request immediately after creating a *Subscription*, and thereafter, immediately after receiving a *Publish* response.

In other cases, especially in high latency networks, the *Client* may wish to pipeline *Publish* requests to ensure cyclic reporting from the *Server*. Pipelining involves sending more than one *Publish* request for each *Subscription* before receiving a response. For example, if the network introduces a delay between the *Client* and the *Server* of 5 seconds and the publishing interval for a *Subscription* is one second, then the *Client* ~~will have to~~ shall issue *Publish* requests every second instead of waiting for a response to be received before sending the next request.

A *Server* should limit the number of active *Publish* requests to avoid an infinite number since it is expected that the *Publish* requests are queued in the *Server*. But a *Server* shall accept more queued *Publish* requests than created *Subscriptions*. It is expected that a *Server* supports several *Publish* requests per *Subscription*. When a *Server* receives a new *Publish* request that exceeds its limit it shall de-queue the oldest *Publish* request and return a response with the result set to *Bad_TooManyPublishRequests*. If a *Client* receives this *Service* result for a *Publish* request it shall not issue another *Publish* request before one of its outstanding *Publish* requests is returned from the *Server*.

Clients can limit the size of *Publish* responses with the *maxNotificationsPerPublish* parameter passed to the *CreateSubscription Service*. However, this could still result in a message that is too large for the *Client* or *Server* to process. In this situation, the *Client* will find that either the *SecureChannel* goes into a fault state and needs to be re-established or the *Publish* response returns an error and calling the *Republish Service* also returns an error. If either situation occurs then the *Client* ~~will have to~~ shall adjust its message processing limits or the parameters for the *Subscription* and/or *MonitoredItems*.

The return diagnostic info setting in the request header of the *CreateMonitoredItems* or the last *ModifyMonitoredItems Service* is applied to the *Monitored Items* and is used as the diagnostic information settings when sending *Notifications* in the *Publish* response.

5.13.5.2 Parameters

Table 95 defines the parameters for the *Service*.

Table 95 – Publish Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscription Acknowledgements []	Subscription Acknowledgement	The list of acknowledgements for one or more <i>Subscriptions</i> . This list may contain multiple acknowledgements for the same <i>Subscription</i> (multiple entries with the same <i>subscriptionId</i>). This structure is defined in-line with the following indented items.
subscriptionId	IntegerId	The <i>Server</i> assigned identifier for a <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition).
sequenceNumber	Counter	The sequence number being acknowledged (see 7.5 for <i>Counter</i> definition). The <i>Server</i> may delete the <i>Message</i> with this sequence number from its retransmission queue.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> for which <i>Notifications</i> are being returned (see 7.14 for <i>IntegerId</i> definition). The value 0 is used to indicate that there were no <i>Subscriptions</i> defined for which a response could be sent.
availableSequence Numbers []	Counter	A list of sequence number ranges that identify unacknowledged <i>NotificationMessages</i> that are available for retransmission from the <i>Subscription</i> 's retransmission queue. This list is prepared after processing the acknowledgements in the request (see 7.5 for <i>Counter</i> definition).
moreNotifications	Boolean	A <i>Boolean</i> parameter with the following values: TRUE the number of <i>Notifications</i> that were ready to be sent could not be sent in a single response. FALSE all <i>Notifications</i> that were ready are included in the response.
notificationMessage	Notification Message	The <i>NotificationMessage</i> that contains the list of <i>Notifications</i> . The <i>NotificationMessage</i> parameter type is specified in 7.21.
results []	StatusCode	List of results for the acknowledgements (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>subscriptionAcknowledgements</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the acknowledgements (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionAcknowledgements</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.5.3 Service results

Table 96 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 96 – Publish Service result codes

Symbolic Id	Description
Bad_TooManyPublishRequests	The <i>Server</i> has reached the maximum number of queued <i>Publish</i> requests.
Bad_NoSubscription	There is no <i>Subscription</i> available for this session.

5.13.5.4 StatusCodes

Table 97 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 97 – Publish Operation Level Result Codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_SequenceNumberUnknown	The sequence number is unknown to the <i>Server</i> .

5.13.6 Republish

5.13.6.1 Description

This *Service* requests the *Subscription* to republish a *NotificationMessage* from its retransmission queue. If the *Server* does not have the requested *Message* in its retransmission queue, it returns an error response.

See 5.13.1.2 for the detail description of the behaviour of this *Service*.

See 6.7 for a description of the issues and strategies regarding reconnect handling and *Republish*.

5.13.6.2 Parameters

Table 98 defines the parameters for the *Service*.

Table 98 – Republish Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> assigned identifier for the <i>Subscription</i> to be republished (see 7.14 for <i>IntegerId</i> definition).
retransmitSequence Number	Counter	The sequence number of a specific <i>NotificationMessage</i> to be republished (see 7.5 for <i>Counter</i> definition).
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
notificationMessage	Notification Message	The requested <i>NotificationMessage</i> . The <i>NotificationMessage</i> parameter type is specified in 7.21.

5.13.6.3 Service results

Table 99 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 99 – Republish Service result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_MessageNotAvailable	The requested message is no longer available.

5.13.7 TransferSubscriptions

5.13.7.1 Description

This *Service* is used to transfer a *Subscription* and its *MonitoredItems* from one *Session* to another. For example, a *Client* may need to reopen a *Session* and then transfer its *Subscriptions* to that *Session*. It may also be used by one *Client* to take over a *Subscription* from another *Client* by transferring the *Subscription* to its *Session*.

The *authenticationToken* contained in the request header identifies the *Session* to which the *Subscription* and *MonitoredItems* shall be transferred. The *Server* shall validate that the *Client* of that *Session* is operating on behalf of the same user and that the potentially new *Client* supports the *Profiles* that are necessary for the *Subscription*. If the *Server* transfers the *Subscription*, it returns the sequence numbers of the *NotificationMessages* that are available for retransmission. The *Client* should acknowledge all *Messages* in this list for which it will not request retransmission.

If the *Server* transfers the *Subscription* to the new *Session*, the *Server* shall issue a *StatusChangeNotification notificationMessage* with the status code *Good_SubscriptionTransferred* to the old *Session*. The *StatusChangeNotification notificationMessage* type is defined in 7.20.4.

5.13.7.2 Parameters

Table 100 defines the parameters for the *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 100 – TransferSubscriptions Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionIds []	IntegerId	List of identifiers for the <i>Subscriptions</i> to be transferred to the new <i>Client</i> (see 7.14 for <i>IntegerId</i> definition). These identifiers are transferred from the primary <i>Client</i> to a backup <i>Client</i> via external mechanisms.
sendInitialValues	Boolean	A <i>Boolean</i> parameter with the following values: TRUE the first Publish response(s) after the <i>TransferSubscriptions</i> call shall contain the current values of all Monitored Items in the Subscription where the Monitoring Mode is set to Reporting. If a value is queued for a data <i>MonitoredItem</i>, the next value in the queue is sent in the <i>Publish</i> response. If no value is queued for a data <i>MonitoredItem</i>, the last value sent is repeated in the <i>Publish</i> response. FALSE the first Publish response after the <i>TransferSubscriptions</i> call shall contain only the value changes since the last Publish response was sent. This parameter only applies to MonitoredItems used for monitoring Attribute changes.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	TransferResult	List of results for the <i>Subscriptions</i> to transfer. The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for each <i>Subscription</i> to be transferred (see 7.34 for <i>StatusCode</i> definition).
availableSequence Numbers []	Counter	A list of sequence number ranges that identify <i>NotificationMessages</i> that are in the <i>Subscription</i> 's retransmission queue. This parameter is null if the transfer of the <i>Subscription</i> failed. The <i>Counter</i> type is defined in 7.5.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Subscriptions</i> to transfer (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.7.3 Service results

Table 101 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 101 – TransferSubscriptions Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_InsufficientClientProfile	The <i>Client</i> of the current <i>Session</i> does not support one or more <i>Profiles</i> that are necessary for the <i>Subscription</i> .

5.13.7.4 StatusCodes

Table 102 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 102 – TransferSubscriptions Operation Level result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code. The <i>Client</i> of the current <i>Session</i> is not operating on behalf of the same user as the <i>Session</i> that owns the <i>Subscription</i> .

5.13.8 DeleteSubscriptions

5.13.8.1 Description

This *Service* is invoked to delete one or more *Subscriptions* that belong to the *Client's Session*.

Successful completion of this *Service* causes all *MonitoredItems* that use the *Subscription* to be deleted. If this is the last *Subscription* for the *Session*, then all *Publish* requests still queued for that *Session* are de-queued and shall be returned with *Bad_NoSubscription*.

Subscriptions that were transferred to another *Session* shall be deleted by the *Client* that owns the *Session*.

5.13.8.2 Parameters

Table 103 defines the parameters for the *Service*.

Table 103 – DeleteSubscriptions Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionIds []	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition).
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCodes	List of <i>StatusCodes</i> for the <i>Subscriptions</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Subscriptions</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.8.3 Service results

Table 104 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 104 – DeleteSubscriptions Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.13.8.4 StatusCodes

Table 105 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 105 – DeleteSubscriptions Operation Level result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

6 Service behaviours

6.1 Security

6.1.1 Overview

The OPC UA *Services* define a number of mechanisms to meet the security requirements outlined in IEC TR 62541-2. This clause describes a number of important security-related procedures that *OPC UA Applications* shall follow.

6.1.2 Obtaining and installing an Application Instance Certificate

All *OPC UA Applications* require an *Application Instance Certificate* which shall contain the following information:

- The network name or address of the computer where the application runs;
- The name of the organization that administers or owns the application;
- The name of the application;
- The URI of the application instance;
- The name of the *Certificate Authority* that issued the *Certificate*;
- The issue and expiry date for the *Certificate*;
- The public key issued to the application by the *Certificate Authority* (CA);
- A digital signature created by the *Certificate Authority* (CA).

In addition, each *Application Instance Certificate* has a private key which should be stored in a location that can only be accessed by the application. If this private key is compromised, the administrator shall assign a new *Application Instance Certificate* and private key to the application.

This *Certificate* may be generated automatically when the application is installed. In this situation the private key assigned to the *Certificate* shall be used to create the *Certificate* signature. *Certificates* created in this way are called self-signed *Certificates*.

If the administrator responsible for the application decides that a self-signed *Certificate* does not meet the security requirements of the organization, then the administrator should install a *Certificate* issued by a *Certificate Authority*. The steps involved in requesting an *Application Instance Certificate* from a *Certificate Authority* are shown in Figure 19.




Figure 19 – Obtaining and installing an Application Instance Certificate

Figure 19 illustrates the interactions between the application, the *Administrator* and the *Certificate Authority*. The *Application* is as *OPC UA Application* installed on a single machine. The *Administrator* is the person responsible for managing the machine and the *OPC UA Application*. The *Certificate Authority* is an entity that can issue digital *Certificates* that meet the requirements of the organization deploying the *OPC UA Application*.

If the *Administrator* decides that a self-signed *Certificate* meets the security requirements for the organization, then the *Administrator* may skip Steps 3 through 5. Application vendors shall ~~always create a default self-signed Certificate during~~ ensure that a *Certificate* is available after the installation process. Every *OPC UA Application* shall allow the *Administrators* to replace *Application Instance Certificates* with *Certificates* that meet their requirements.

When the *Administrator* requests a new *Certificate* from a *Certificate Authority*, the *Certificate Authority* may require that the *Administrator* provide proof of authorization to request *Certificates* for the organization that will own the *Certificate*. The exact mechanism used to provide this proof depends on the *Certificate Authority*.

Vendors may choose to automate the process of acquiring *Certificates* from an authority. If this is the case, the *Administrator* would still go through the steps illustrated in Figure 19, however, the installation program for the application would do them automatically and only prompt the *Administrator* to provide information about the application instance being installed.

6.1.3 Determining if a Certificate is trusted

Applications shall never communicate with another application that they do not trust. An *Application* decides if another application is trusted by checking whether the *Application Instance Certificate* for the other application is trusted. Applications shall rely on lists of *Certificates* provided by the *Administrator* to determine trust. There are two separate lists: a list of trusted *Applications* and a list of trusted *Certificate Authorities* (CAs). If an application is not directly trusted (i.e. its *Certificate* is not in the list of trusted applications) then the application shall build a chain of *Certificates* back to a trusted CA.

When building a chain each *Certificate* in the chain shall be validated. If any validation error occurs then the trust check fails. Some validation errors are non-critical which means they can be suppressed by a user of an *Application* with the appropriate privileges. Suppressed validation errors are always reported via auditing (i.e. an appropriate Audit event is raised).

Building a trust chain requires access to all *Certificates* in the chain. These *Certificates* may be stored locally or they may be provided with the application *Certificate*. Processing fails with *Bad_SecurityChecksFailed* if a CA *Certificate* cannot be found.

Table 106 specifies the steps used to validate a *Certificate* in the order that they shall be followed. These steps are repeated for each *Certificate* in the chain. Each validation step has a unique error status and audit event type that shall be reported if the check fails. The audit event is in addition to any audit event that was generated for the particular *Service* that was invoked. The *Service* audit event in its message text shall include the audit *EventId* of the *AuditCertificateEventType* (for more details, see 6.5). Processing halts if an error occurs, unless it is non-critical and it has been suppressed.

ApplicationInstanceCertificates shall not be used in a *Client* or *Server* until they have been evaluated and marked as trusted. This can happen automatically by a PKI trust chain or in an offline manner where the *Certificate* is marked as trusted by an administrator after evaluation.

Table 106 – Certificate validation steps

Step	Error/AuditEvent	Description
Certificate Structure	<i>Bad_CertificateInvalid</i> <i>Bad_SecurityChecksFailed</i> <i>AuditCertificateInvalidEventType</i>	The <i>Certificate</i> structure is verified. This error may not be suppressed. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .
Build Certificate Chain	<i>Bad_CertificateChainIncomplete</i> <i>Bad_SecurityChecksFailed</i> <i>AuditCertificateInvalidEventType</i>	The trust chain for the <i>Certificate</i> is created. An error during the chain creation may not be suppressed. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .
Signature	<i>Bad_CertificateInvalid</i> <i>Bad_SecurityChecksFailed</i> <i>AuditCertificateInvalidEventType</i>	A <i>Certificate</i> with an invalid signature shall always be rejected. A <i>Certificate</i> signature is invalid if the <i>Issuer Certificate</i> is unknown. A self-signed <i>Certificate</i> is its own issuer. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .
Security Policy Check	<i>Bad_CertificatePolicyCheckFailed</i> <i>Bad_SecurityChecksFailed</i> <i>AuditCertificateInvalidEventType</i>	A <i>Certificate</i> signature shall comply with the <i>CertificateSignatureAlgorithm</i> , <i>MinAsymmetricKeyLength</i> and <i>MaxAsymmetricKeyLength</i> requirements for the used <i>SecurityPolicy</i> defined in IEC 62541-7. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> . This error may be suppressed.
Trust List Check	<i>Bad_CertificateUntrusted</i> <i>Bad_SecurityChecksFailed</i> <i>AuditCertificateUntrustedEventType</i>	If the <i>Application Instance Certificate</i> is not trusted and none of the CA <i>Certificates</i> in the chain is trusted, the result of the <i>Certificate</i> validation shall be <i>Bad_CertificateUntrusted</i> . If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .

Step	Error/AuditEvent	Description
Validity Period	Bad_CertificateTimeInvalid Bad_CertificateIssuerTimeInvalid AuditCertificateExpiredEventType	The current time shall be after the start of the validity period and before the end. This error may be suppressed.
Host Name	Bad_CertificateHostNameInvalid AuditCertificateDataMismatchEventType	The HostName in the URL used to connect to the Server shall be the same as one of the HostNames specified in the <i>Certificate</i> . This check is skipped for CA <i>Certificates</i> . <i>This check is skipped for Server side validation.</i> This error may be suppressed.
URI	Bad_CertificateUriInvalid AuditCertificateDataMismatchEventType	<i>Application and Software Certificates</i> contain an application or product URI that shall match the URI specified in the <i>ApplicationDescription</i> provided with the <i>Certificate</i> . This check is skipped for CA <i>Certificates</i> . This error may not be suppressed. The <i>gatewayServerUri</i> is used to validate an <i>Application Certificate</i> when connecting to a <i>Gateway Server</i> (see 7.1).
Certificate Usage	Bad_CertificateUseNotAllowed Bad_CertificateIssuerUseNotAllowed AuditCertificateMismatchEventType	Each <i>Certificate</i> has a set of uses for the <i>Certificate</i> (see IEC 62541-6). These uses shall match use requested for the <i>Certificate</i> (i.e. Application, Software or CA). This error may be suppressed unless the <i>Certificate</i> indicates that the usage is mandatory.
Find Revocation List	Bad_CertificateRevocationUnknown Bad_CertificateIssuerRevocationUnknown AuditCertificateRevokedEventType	Each CA <i>Certificate</i> may have a revocation list. This check fails if this list is not available (i.e. a network interruption prevents the application from accessing the list). No error is reported if the <i>Administrator</i> disables revocation checks for a CA <i>Certificate</i> . This error may be suppressed.
Revocation Check	Bad_CertificateRevoked Bad_CertificateIssuerRevoked AuditCertificateRevokedEventType	The <i>Certificate</i> has been revoked and may not be used. This error may not be suppressed. <i>If this check fails on the Server side, the error Bad_SecurityChecksFailed shall be reported back to the Client.</i>

Certificates are usually placed in a central location called a *CertificateStore*. Figure 20 illustrates the interactions between the *Application*, the *Administrator* and the *CertificateStore*. The *CertificateStore* could be on the local machine or in some central server. The exact mechanisms used to access the *CertificateStore* depend on the application and PKI environment set up by the *Administrator*.




Figure 20 – Determining if an Application Instance Certificate is trusted

6.1.4 Creating a SecureChannel

All *OPC UA Applications* shall establish a *SecureChannel* before creating a *Session*. This *SecureChannel* requires that both applications have access to *Certificates* that can be used to encrypt and sign *Messages* exchange. The *Application Instance Certificates* installed by following the process described in 6.1.2 may be used for this purpose.

The steps involved in establishing a *SecureChannel* are shown in Figure 21.




Figure 21 – Establishing a SecureChannel

Figure 21 assumes *Client* and *Server* have online access to a *CertificateAuthority* (CA). If online access is not available and if the administrator has installed the CA public key on the local machine, then the *Client* and *Server* shall still validate the application *Certificates* using that key. The figure shows only one CA, however, there is no requirement that the *Client* and *Server* *Certificates* be issued by the same authority. A self-signed *Application Instance Certificate* does not need to be verified with a CA. Any *Certificate* shall be rejected if it is not in a trust list provided by the administrator.

Both the *Client* and *Server* shall have a list of *Certificates* that they have been configured to trust (sometimes called the *Certificate Trust List* or CTL). These trusted *Certificates* may be *Certificates* for *Certificate Authorities* or they may be *OPC UA Application Instance Certificates*. *OPC UA Applications* shall be configured to reject connections with applications that do not have a trusted *Certificate*.

Certificates can be compromised, which means they should no longer be trusted. Administrators can revoke a *Certificate* by removing it from the trust list for all applications or the CA can add the *Certificate* to the *Certificate Revocation List* (CRL) for the *Issuer Certificate*. Administrators may save a local copy of the CRL for each *Issuer Certificate* when online access is not available.

A *Client* does not need to call *GetEndpoints* each time it connects to the *Server*. This information should change rarely and the *Client* can cache it locally. If the *Server* rejects the *OpenSecureChannel* request the *Client* should call *GetEndpoints* and make sure the *Server* configuration has not changed.

There are two security risks which a *Client* shall be aware of when using the *GetEndpoints Service*. The first could come from a rogue *Discovery Server* that tries to direct the *Client* to a rogue *Server*. For this reason the *Client* shall verify that the *ServerCertificate* in the *EndpointDescription* is a trusted *Certificate* before it calls *CreateSession*.

The second security risk comes from a third party that alters the contents of the *EndpointDescriptions* as they are transferred over the network back to the *Client*. The *Client* protects itself against this by comparing the list of *EndpointDescriptions* returned from the *GetEndpoints Service* with list returned in the *CreateSession* response.

The exact mechanisms for using the *SecurityToken* to sign and encrypt *Messages* exchanged over the *SecureChannel* are described in IEC 62541-6. The process for renewing tokens is also described in detail in IEC 62541-6.

In many cases, the *Certificates* used to establish the *SecureChannel* will be the *Application Instance Certificates*. However, some *Communication Stacks* might not support *Certificates* that are specific to a single application. Instead, they expect all communication to be secured with a *Certificate* specific to a user or the entire machine. For this reason, *OPC UA Applications* will need to exchange their *Application Instance Certificates* when creating a *Session*.

6.1.5 Creating a Session

Once an *OPC UA Client* has established a *SecureChannel* with a *Server* it can create an *OPC UA Session*.

The steps involved in establishing a *Session* are shown in Figure 22.




Figure 22 – Establishing a Session

Figure 22 illustrates the interactions between a *Client*, a *Server*, a *Certificate Authority* (CA) and an ~~authentication service~~ identity provider. The CA is responsible for issuing the *Application Instance Certificates*. If the *Client* or *Server* does not have online access to the CA, then they shall validate the *Application Instance Certificates* using the CA public key that the administrator shall install on the local machine.

The ~~authentication service is~~ identity provider may be a central database that can verify that user token provided by the *Client*. This ~~authentication service~~ identity provider may also tell the *Server* which access rights the user has. The ~~authentication service~~ identity provider

depends on the user identity token. It could be a *Certificate Authority*, a Kerberos ticket granting service, a WS-Trust *Server* or a proprietary database of some sort.

The *Client* and *Server* shall prove possession of their *Application Instance Certificates* by signing the *Certificates* with a nonce appended. The exact mechanism used to create the proof of possession signatures is described in 5.6.2. Similarly, the *Client* shall prove possession ~~of some types of user identity tokens by creating signatures with the secret associated with the token~~ by either providing a secret like a password in the user identity token or by creating a signature with the secret associated with a user identity token like x.509 v3.

6.1.6 Impersonating a User

Once an OPC UA *Client* has established a *Session* with a *Server* it can change the user identity associated with the *Session* by calling the *ActivateSession* service.

The steps involved in impersonating a user are shown in Figure 23.




Figure 23 – Impersonating a User

6.2 Authorization Services

6.2.1 Overview

Authorization Services provide *Access Tokens* to *Clients* on behalf of *Users* that they pass to a *Server* to be granted access to resources.

In a basic model (as shown in Figure 22) the *Server* is responsible for authorization (i.e. deciding what a user can do) while a separate identity provider (e.g. the operating system) is responsible for authentication (deciding who the user is).

In more complex models, the *Server* relies on external *Authorization Services* to provide some of its authorization requirements. These *Authorization Services* act in concert with an external identity provider which validates the user credentials before the external *Authorization Service* creates an *Access Token* that tells the *Server* what the user is allowed to do. The *Client* interactions with these services may be indirect as shown in 6.2.2 or direct as shown in 6.2.3.

Even when the *Server* requires the *Client* to use an external *Authorization Service*, the *Server* is still responsible for managing and enforcing the *Permissions* assigned to *Nodes* in its *Address Space*. The subclauses below discuss the use of an external *Authorization Service* in more detail.


6.2.2 Indirect handshake with an Identity Provider

Authorization Services (AS) provide access to identity providers which can validate the credentials provided by *Clients*. They then provide tokens which can be passed to a *Server* instead of the credentials. These tokens are passed as an *IssuedIdentityToken* defined in 7.36.6.

The protocol to request tokens depends on the *Authorization Service* (AS). Common protocols include Kerberos and OAuth2. OAuth2 supports claims based authorization as described in IEC TR 62541-2.

Servers publish the *Authorization Services* (AS) they support in the *UserTokenPolicies* list return with *GetEndpoints*. The *IssuedTokenType* field specifies the protocol used to communicate with the AS. The *IssuerEndpointUrl* field contains the information needed by the *Client* to connect to the AS using the protocol required by the AS.

The basic handshake is shown in Figure 24.




IEC

Figure 24 – Indirect handshake with an Identity Provider

6.2.3 Direct handshake with an Identity Provider

Authorization Services require that *Servers* be registered with them because the *Access Tokens* can only be used with a single *Server*. This can introduce a lot of complexity for administrators. One way to reduce this complexity is to leverage the *Server* information that is already managed by a Global Discovery Service (GDS) described in IEC 62541-12. In this model the user identities are still managed by a central *Authorization Service*. The interactions are shown in Figure 25.



IEC

Figure 25 – Direct handshake with an Identity Provider

The *UserTokenPolicy* returned from the *Server* provides the URL of the *Authorization Service* and the identity provider. If the *Application Authorization Service* is linked with the GDS, it knows of all *Servers* which have been issued *Certificates*. The *ApplicationUri* is used as the identifier for the *Server* passed to the AS. The identity provider is responsible for managing users known to the system. It validates the credentials provided by the *Client* and returns an *Identity Access Token* which identifies the user. The *Identity Access Token* is passed to the *Application Authorization Service* which validates the *Client* and *Server* applications and creates a new *Access Token* that can be used to access the *Server*.

6.3 Session-less Service invocation

6.3.1 Description

The *Session-less Service* invocation is introduced for *Services*, such as *Read*, *Write* or *Call*, that do not require any caller specific state information. It is accessible through the *SessionlessInvoke Service* which provides the context information required to call *Services* without a *Session*.

Session-less invocation is limited to *Services* of the *View Service Set* (with exception of *RegisterNodes* and *UnregisterNodes*), *Attribute Service Set*, *Method Service Set*, *NodeManagement Service Set* and *Query Service Set*. All *Services* belonging to these *Service Sets* that are supported by a *Server* via a *Session* shall also be supported via the *SessionlessInvoke Service*.

Session-less Services can be invoked via a *SecureChannel* by using the *Access Token* returned from the *Authorization Service* as the *authenticationToken* in the *requestHeader*. The *SecureChannel* shall have encryption enabled to prevent eavesdroppers from seeing the *Access Token*. The *Access Token* provides the user authentication. If application authentication through the *SecureChannel* is sufficient, *Servers* may not require the *Access Token* and assume an anonymous user. In this case the *authenticationToken* shall be null.

The *SessionlessInvoke Messages* are just an envelope for the *Service* to invoke and do not have a *RequestHeader* and *ResponseHeader* like other *Services*. Those parameters are already part of the *body* which contains the *Message* for the *Service* to invoke.

Any *Endpoint* used for normal communication could be used for *Session-less* invocation provided the *Endpoint* supports encryption. The *Server* returns *Bad_ServiceUnsupported* if it does not support *Session-less* invocation for the request specified in the *body*. If it supports

invocation but not with the combination of *Endpoint* and security settings used, it returns *Bad_SecurityModelInsufficient*.

Servers may expose *Endpoints* which are only for use with *Session-less* invocation. These *Endpoints* shall support *GetEndpoints* and *FindServers* in addition to the *SessionlessInvokeService*. The *Server* returns *Bad_ServiceUnsupported* for the other *Services*.

A *Session* ensures that a namespace index or a server index does not change during the lifetime of a *Session*. This cannot be ensured between *Session-less Services* invocations. There are two options to ensure the namespace indices in the call match the expected namespace URIs in the *Server*. One option for the caller is to provide the list of namespace URIs used to build the namespace indices. This works best for single *Session-less Service* invocations. The second option is to provide the *UriVersion* to ensure consistency of namespace arrays between *Client* and *Server*. The *UriVersion* is first read from the *Server* together with the *NamespaceArray* and *ServerArray*. This reduces the overhead per call for a sequence of *Session-less Service* invocations.

6.3.2 Parameters

Table 107 defines the parameters for the *Service*.

Table 107 – SessionlessInvoke Service parameters

Name	Type	Description
Request		
urisVersion	VersionTime	The version of the <i>NamespaceArray</i> and the <i>ServerArray</i> used for the <i>Service</i> invocation. The version shall match the value of the <i>UriVersion Property</i> that defines the version for the URI lists in the <i>NamespaceArray</i> and the <i>ServerArray Properties</i> defined in IEC 62541-5. If the <i>urisVersion</i> parameter does not match the <i>Server's UriVersion Property</i> , the <i>Server</i> shall return <i>Bad_VersionTimeInvalid</i> . In this case the <i>Client</i> shall read the <i>UriVersion</i> , <i>NamespaceArray</i> and the <i>ServerArray</i> from the <i>Server Object</i> to repeat the <i>Service</i> invocation with the right version. The <i>VersionTime DataType</i> is defined in 7.38. If the value is 0, the parameter is ignored and the URIs are defined by the <i>namespaceUri</i> and <i>serverUri</i> parameters in request and response. If the value is non-zero, the <i>namespaceUri</i> and <i>serverUri</i> parameters in the request are ignored by the <i>Server</i> and set to null arrays in the response.
namespaceUri []	String	A list of URIs referenced by <i>NodeIds</i> or <i>QualifiedNames</i> in the request. NamespaceIndex 0 shall not be in this list. The first entry in this list is NamespaceIndex 1. The parameter shall be ignored by the <i>Server</i> if the <i>urisVersion</i> is not 0.
serverUri []	String	A list of URIs referenced by <i>ExpandedNodeIds</i> in the request. ServerIndex 0 shall not be in this list. The first entry in this list is ServerIndex 1. The parameter shall be ignored by the <i>Server</i> if the <i>urisVersion</i> is not 0.
localeIds []	LocaleId	List of locale ids in priority order for localized strings. The first <i>LocaleId</i> in the list has the highest priority. If the <i>Server</i> returns a localized string to the <i>Client</i> , the <i>Server</i> shall return the translation with the highest priority that it can. If it does not have a translation for any of the locales identified in this list, then it shall return the string value that it has and include the locale id with the string. See IEC 62541-3 for more detail on locale ids. If <i>localeIds</i> is empty, the returned language variant is <i>Server</i> specific.
serviceId	UInt32	The numeric identifier assigned to the <i>Service</i> request <i>DataType</i> describing the body.
body	*	The body of the request. The body is an embedded structure containing the corresponding <i>Service</i> request for the <i>serviceId</i> .

Name	Type	Description
Response		
namespaceUris []	String	A list of URIs referenced by <i>NodeIds</i> or <i>QualifiedNames</i> in the response. NamespaceIndex 0 shall not be in this list. The first entry in this list is NamespaceIndex 1. An empty array shall be returned if the <i>urisVersion</i> is not 0.
serverUris []	String	A list of URIs referenced by <i>ExpandedNodeIds</i> in the response. ServerIndex 0 shall not be in this list. The first entry in this list is ServerIndex 1. An empty array shall be returned if the <i>urisVersion</i> is not 0.
serviceld	UInt32	The numeric identifier assigned to the <i>Service</i> response <i>DataType</i> describing the body.
body	*	The body of the response. The body is an embedded structure containing the corresponding <i>Service</i> response for the <i>serviceld</i> .

6.3.3 Service results

Table 108 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 108 – SessionlessInvoke Service result codes

Symbolic Id	Description
Bad_VersionTimeInvalid	The provided version time is no longer valid.

6.4 Software Certificates

6.2.1 Overview

All OPC UA Applications may have one or more software *Certificates* that are issued by certification authorities and specify the profiles that the application supports. This clause describes how *SoftwareCertificates* are obtained, installed and validated.

SoftwareCertificates are used to identify a software product that is an OPC UA Application and to describe the capabilities of the product. The capabilities are described through *Profiles* defined in IEC 62541-7. *SoftwareCertificates* are not relevant for security in OPC UA.

6.2.2 Obtaining and Installing a Software Certificate

The software *Certificates* contain the following information:


- The name of the vendor responsible for the product;
- The name of the product;
- The URI for the product;
- The software version and build number;
- The vendor product *Certificate*;
- A list of profiles supported by the application;
- The certification testing status for each supported profile;
- The name of the testing authority that issued the *Certificate*;

- The issue and expiry date for the *Certificate*;
- The public key issued to the application by the certifying authority;
- A digital signature created by the certifying authority.

The product vendor is responsible for completing the certification process and requesting software *Certificates* from the certification authorities. The vendor shall install the software *Certificates* when an application is installed on a machine. When distributing these *Certificates*, the application vendors should take precautions to prevent unauthorized users from acquiring their software *Certificates* and using them for applications that the vendor did not develop. Misused software *Certificates* are not a security risk, but vendors could find that they are blamed for interoperability problems caused by *Certificates* used by unauthorized applications.

Vendors may choose to assign their own *Certificate* to the application (called the Vendor Product *Certificate*). This *Certificate* would then be incorporated in the Software *Certificate*. The private key for the vendor product *Certificate* is managed by the vendor.

The steps involved in acquiring and installing a software *Certificate* from a *Certificate Authority* are shown in Figure 24.



IEC

Figure 24 – Obtaining and Installing a Software Certificate

The figure above illustrates the interactions between the vendor, the tester and the certifying authority. The vendor is the organisation responsible for the *OPC UA Application*. The tester may be the vendor (for self certification) or it may be a third party testing facility. The certifying authority is a PKI *Certificate* authority managed by the organisation that created the *OPC UA Application* profiles and certification programmes.

The certifying authority will issue *Certificates* only to vendors it trusts. For that reason, the tester shall provide proof of identity before the certifying authority will issue a *Certificate*.

~~6.2.3 Validating a Software Certificate~~

~~OPC UA Applications shall validate the SoftwareCertificates provided by the applications that they communicate with.~~

~~A SoftwareCertificate is valid if:~~

- ~~• The signature on the SoftwareCertificate is valid;~~
- ~~• The SoftwareCertificate has passed its issue date and it has not expired;~~
- ~~• The SoftwareCertificate has not been revoked by the issuer;~~
- ~~• The Issuer Certificate is valid and has not been revoked by the CA that issued it.~~

~~The steps used to validate SoftwareCertificates are almost the same as the steps described for Application Instance Certificates in 6.1.3. The only difference is SoftwareCertificates are not checked against the TrustList for the Application.~~

~~SoftwareCertificates always contain an Application Instance Certificates which is owned by the vendor that distributes the application. The application Certificate may also be in the Certificate chain used to issue Application Instance Certificates. For this reason, OPC UA Applications may reject SoftwareCertificates provided by applications if the application Certificate is not part of the Certificate chain for the Application Instance Certificate. OPC UA Applications should allow Administrators to require this behaviour. Profiles defined in IEC 62541-7 may further specify expected Certificate handling behaviours.~~

NOTE Details on SoftwareCertificates will be defined in a future edition of IEC 62541-4.

6.5 Auditing

6.5.1 Overview

Auditing is a requirement in many systems. It provides a means of tracking activities that occur as part of normal operation of the system. It also provides a means of tracking abnormal behaviour. It is also a requirement from a security standpoint. For more information on the security aspects of auditing, see IEC TR 62541-2. Subclause 6.5 describes what is expected of an OPC UA Server and Client with respect to auditing and it details the audit requirements for each service set. Auditing can be accomplished using one or both of the following methods:

- a) The OPC UA Application that generates the audit event can log the audit entry in a log file or other storage location;
- b) The OPC UA Server that generates the audit event can publish the audit event using the OPC UA event mechanism. This allows an external OPC UA Client to subscribe to and log the audit entries to a log file or other storage location.

6.5.2 General audit logs

Each OPC UA Service request contains a string parameter that is used to carry an audit record id. A Client or any Server operating as a Client, such as an aggregating Server, can create a local audit log entry for a request that it submits. This parameter allows this Client to pass the identifier for this entry with the request. If this Server also maintains an audit log, it should include this id in its audit log entry that it writes. When this log is examined and that entry is found, the examiner will be able to relate it directly to the audit log entry created by the Client. This capability allows for traceability across audit logs within a system.

6.5.3 General audit Events

A Server that maintains an audit log shall provide the audit log entries via Event Messages. The AuditEventType and its sub-types are defined in IEC 62541-3. An audit Event Message also includes the audit record Id. The details of the AuditEventType and its subtypes are defined in IEC 62541-5. A Server that is an aggregating Server that supports auditing shall

also subscribe for audit events for all of the *Servers* that it is aggregating (assuming they provide auditing). The combined stream should be available from the aggregating *Server*.

6.5.4 Auditing for Discovery Service Set

This *Service Set* can be separated into two groups: *Services* that are called by OPC UA *Clients* and *Services* that are invoked by OPC UA *Servers*. The *FindServers* and *GetEndpoints* *Services* that are called by OPC UA *Clients* may generate audit entries for failed *Service* invocations. The *RegisterServer* *Service* that is invoked by OPC UA *Servers* shall generate audit entries for all new registrations and for failed *Service* invocations. These audit entries shall include the *Server* URL, *Server* names, *Discovery* URIs and *isOnline* status. Audit entries should not be generated for *RegisterServer* invocation that does not cause changes to the registered *Servers*.

6.5.5 Auditing for SecureChannel Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for failed service invocations and for successful invocation of the *OpenSecureChannel* and *CloseSecureChannel* *Services*. The *Client* generated audit entries should be setup prior to the actual call, allowing the correct audit record *Id* to be provided. The *OpenSecureChannel* *Service* shall generate an audit *Event* of type *AuditOpenSecureChannelEventType* or a subtype of it for the *requestType* *ISSUE_0*. Audit *Events* for the *requestType* *RENEW_1* are only created if the renew fails. The *CloseSecureChannel* *Service* shall generate an audit *Event* of type *AuditChannelEventType* or a subtype of it. Both of these *Event* types are subtypes of the *AuditChannelEventType*. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure cases the *Message* for *Events* of this type should include a description of why the service failed. This description should be more detailed than what was returned to the client. From a security point of view a *Client* only needs to know that it failed, but from an *Auditing* point of view the exact details of the failure need to be known. In the case of *Certificate* validation errors the description should include the audit *EventId* of the specific *AuditCertificateEventType* that was generated to report the *Certificate* error. The *AuditCertificateEvent* *AuditCertificateEventType* shall also contain the detailed *Certificate* validation error. The additional parameters should include the details of the request. It is understood that these events may be generated by the underlying *Communication Stacks* in many cases, but they shall be made available to the *Server* and the *Server* shall report them.

6.5.6 Auditing for Session Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed *Service* invocations. These *Services* shall generate an audit *Event* of type *AuditSessionEventType* or a subtype of it. In particular, they shall generate the base *EventType* or the appropriate subtype, depending on the *Service* that was invoked. The *CreateSession* *Service* shall generate *AuditCreateSessionEventType* events or sub-types of it. The *ActivateSession* *Service* shall generate *AuditActivateSessionEventType* events or subtypes of it. When the *ActivateSession* *Service* is called to change the user identity then the *Server* shall generate *AuditActivateSessionEventType* events or subtypes of it. The *CloseSession* *Service* shall generate ~~the base *EventType* of~~ *AuditSessionEventType* events or subtypes of it. It shall always be generated if a *Session* is terminated like *Session* timeout expiration or *Server* shutdown. The *SourceName* for *Events* of this type shall be "Session/Timeout" for a *Session* timeout, "Session/CloseSession" for a *CloseSession* *Service* call and "Session/Terminated" for all other cases. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case the *Message* for *Events* of this type should include a description of why the *Service* failed. The additional parameters should include the details of the request.

This *Service Set* shall also generate additional audit events in the cases when *Certificate* validation errors occur. These audit *Events* are generated in addition to the *AuditSessionEventTypes* *AuditSessionEventType* *Events*. See IEC 62541-3 for the definition of *AuditCertificateEventType* and its subtypes.

For *Clients*, that support auditing, accessing the services in the *Session Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. These audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.7 Auditing for NodeManagement Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed *Service* invocations. These *Services* shall generate an audit *Event* of type *AuditNodeManagementEventType* or subtypes of it. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case, the *Message* for *Events* of this type should include a description of why the service failed. The additional parameters should include the details of the request.

For *Clients* that support auditing, accessing the *Services* in the *NodeManagement Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.8 Auditing for Attribute Service Set

The *Write* or *HistoryUpdate* *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed *Service* invocations. These *Services* shall generate an audit *Event* of type *AuditUpdateEventType* or subtypes of it. In particular, the *Write Service* shall generate an audit event of type *AuditWriteUpdateEventType* or a subtype of it. The *HistoryUpdate Service* shall generate an audit *Event* of type ~~*AuditHistoryUpdateEvent*~~ *AuditHistoryUpdateEventType* or a subtype of it. Three subtypes of ~~*AuditHistoryUpdateEvent*~~ *AuditHistoryUpdateEventType* are defined as *AuditHistoryEventUpdateEventType*, *AuditHistoryValueUpdateEventType* and *AuditHistoryDeleteEventType*. The subtype depends on the type of operation being performed, historical event update, historical data value update or a historical delete. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case the *Message* for *Events* of this type should include a description of why the *Service* failed. The additional parameters should include the details of the request.

The *Read* and *HistoryRead* *Services* may generate audit entries and audit *Events* for failed *Service* invocations. These *Services* should generate an audit *Event* of type *AuditEventType* or a subtype of it. See IEC 62541-5 for the detailed assignment of the *SourceNode*, *SourceName* and additional parameters. The *Message* for *Events* of this type should include a description of why the *Service* failed.

For *Clients* that support auditing, accessing the *Write* or *HistoryUpdate* services in the *Attribute Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. Invocations of the other *Services* in this *Service Set* may generate audit entries. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.9 Auditing for Method Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed service invocations if the invocation modifies the *AddressSpace*, writes a value or modifies the state of the system (alarm acknowledge, batch sequencing or other system changes). These method calls shall generate an audit *Event* of type *AuditUpdateMethodEventType* or subtypes of it. Methods that do not modify the *AddressSpace*, write values or modify the state of the system may generate events. See IEC 62541-5 for the detailed assignment of the *SourceNode*, *SourceName* and additional parameters.

For *Clients* that support auditing, accessing the *Method Service Set* shall generate audit entries for both successful and failed invocations of the *Service*, if the invocation modifies the *AddressSpace*, writes a value or modifies the state of the system (alarm acknowledge, batch sequencing or other system changes). Invocations of the other *Methods* may generate audit entries. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.10 Auditing for View, Query, MonitoredItem and Subscription Service Set

All of the *Services* in these four *Service Sets* only provide the *Client* with information, with the exception of the *TransferSubscriptions Service* in the *Subscription Service Set*. In general, these services will not generate audit entries or audit *Event Messages*. The *TransferSubscriptions Service* shall generate an audit *Event* of type *AuditSessionEventType* or subtypes of it for both successful and failed *Service* invocations. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case, the *Message* for *Events* of this type should include a description of why the service failed.

For *Clients* that support auditing, accessing the *TransferSubscriptions Service* in the *Subscription Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. Invocations of the other *Services* in this *Service Set* do not require audit entries. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.6 Redundancy

6.4.1 Redundancy overview

~~Redundancy in OPC UA ensures that both Clients and Server can be redundant. OPC UA does not provide redundancy; it provides the data structures and services by which redundancy may be achieved in a standardized manner.~~

6.4.2 Server redundancy overview

6.4.2.1 General

~~Server redundancy comes in two modes, transparent and non-transparent. By definition, in transparent redundancy the failover of Server responsibilities from one Server to another is transparent to the Client; the Client does not care or even know that failover has occurred; the Client does not need to do anything at all to keep data flowing. In contrast, non-transparent failover requires some activity on the part of the Client.~~

~~The ServerRedundancy Object defined in IEC 62541-5 indicates the mode supported by the Server. The ServerRedundancyType ObjectType and its subtypes TransparentRedundancyType and NonTransparentRedundancyType defined in IEC 62541-5 specify information for the supported redundancy mode.~~

~~The two areas where redundancy creates specific needs are in keeping the Server and Client information synchronised across Servers, and in controlling the failover of data flow from one Server to another.~~

~~Independent of the used redundancy mode it is expected that all Servers in the redundant set have an identical address space including identical NodeIds and the identical logic for setting the service level.~~

6.4.2.2 Transparent redundancy

~~For transparent redundancy, all that OPC UA provides is the data structures to allow the Client to identify which Servers are in the redundant set, what the service level of each Server~~

is and which *Server* is currently responsible for the *Client Session*. This information is specified in *TransparentRedundancyType ObjectType* defined in IEC 62541-5.

All OPC UA interactions within a given session shall be supported by one *Server* and the *Client* is able to identify which *Server* that is, allowing a complete audit trail for the data. It is the responsibility of the *Servers* to ensure that information is synchronised between the *Servers*. A functional *Server* will take over the *Session* and *Subscriptions* from the failed *Server*. Failover may require a transport layer reconnect of the *Client* but the *Endpoint URL* of the *Server* shall not change. See 6.5 for more details on re-establishing connections.

Figure 25 shows a typical transparent redundancy setup.




Figure 25 — Transparent Redundancy setup

6.4.2.3 Non-transparent redundancy

For non-transparent redundancy, OPC UA provides the data structures to allow the *Client* to identify what *Servers* are available in the redundant set and also *Server* information which tells the *Client* what modes of failover the *Server* supports. This information allows the *Client* to determine what actions it may need to take in order to accomplish failover. This information is specified in *NonTransparentRedundancyType ObjectType* defined in IEC 62541-5.

Figure 26 shows a typical non-transparent redundancy setup.




Figure 26 — Non-Transparent Redundancy setup

For non-transparent redundancy the Server has additional concepts of *Cold*, *Warm*, *Hot* and *HotAndMirrored* failover modes. The failover mode provides information about the failover capabilities the Server supports and allows a Client to determine the available failover actions described in Table 102.

Cold failovers are for Servers where only one Server can be active at a time. The Client can only connect to one of the Servers in the redundant set.

Warm failovers are for Servers where the backup Servers can be active, but cannot connect to actual data points (typically, a system where the underlying devices are limited to a single connection). The *ServiceLevel Variable* defined in IEC 62541-5 indicates the ability of the Server to provide its data to the Client. The Client shall use the Server with the highest service level. The designation of the active Server is vendor specific.

Hot failovers are for Servers where more than one Server can be active and fully operational. The *ServiceLevel Variable* defined in IEC 62541-5 should be used by the Client to find the Servers with the highest service level to achieve load balancing.

HotAndMirrored failovers are for Servers that are mirroring their internal states to all Servers in the redundant set and more than one Server can be active and fully operational. Mirroring state minimally includes *Sessions*, *Subscriptions*, registered *Nodes*, sequence numbers and sent *Notifications*. This allows Clients to fail over without creating a new context for communication. The *ServiceLevel Variable* defined in IEC 62541-5 should be used by the Client to find the Servers with the highest service level to achieve load balancing. This failover mode is similar to the transparent redundancy. The advantage is that the Client has full control over selecting the Server. The disadvantage is that the Client needs to be able to handle failovers.

Each Server maintains a list of *ServerUris* for all redundant Servers in the redundant set. The list is provided together with the failover mode in the *ServerRedundancy Object* defined in IEC 62541-5. To enable clients to connect to all Servers in the list, each Server in the list shall provide the *ApplicationDescription* for all Servers in the redundant set through the *FindServers Service*. This information is needed by the Client to translate the *ServerUri* into information needed to connect to the other Servers in the redundant set. Therefore a Client needs to know only one of the redundant Servers to find the other Servers based on the provided information.

Table 102 defines the list of failover actions.

Table 102 — Redundancy failover actions

Failover mode and Client options	Cold	Warm	Hot (1)	Hot (2)	HotAndMirrored
On initial connection in addition to actions on active server:					
— Connect to more than one OPC UA Server.		X	X	X	Optional for status check
— Creating Subscriptions and adding monitored items.		X	X	X	
— Activating sampling on the Subscriptions.			X	X	
— Activate publishing.				X	
At Failover:					
— CreateSecureChannel to backup OPC UA Server	X				X
— CreateSession on backup OPC UA Server	X				
— ActivateSession on backup OPC UA Server	X				X
— Creating Subscriptions and adding monitored items.	X				
— Activating sampling on the Subscriptions.	X	X			
— Activate publishing.	X	X	X		

~~A Client connected to a Cold failover redundant set can only connect to one Server. For failover the Client connects to a backup Server, creates Subscriptions and MonitoredItems and activates publishing. There is a loss of data from the time the connection to the active Server is interrupted until the time the Client gets Publish Responses from the backup Server.~~

~~A Client connected to a Warm failover redundant set can connect and create Subscriptions and MonitoredItems on more than one Server. Sampling and publishing can only be activated on one Server. This one Server can be found by reading the ServiceLevel Variable from all Servers. The Server with the highest ServiceLevel is used. For failover the Client activates sampling and publishing on the Server with the highest ServiceLevel. There is a loss of data from the time the connection to the active Server is interrupted until the time the Client gets Publish Responses from the backup Server.~~

~~A Client connected to a Hot failover redundant set has two options to ensure that there is no loss of data in a failover scenario.~~

- ~~• Hot (1)
The Client creates MonitoredItems and is sampling on more than one Server but activates publishing only on the Server with the highest ServiceLevel. The Client should setup the queue size for the MonitoredItems such that it can buffer all changes during the failover time. The failover time is the time between the connection interruption and the time the Client gets Publish Responses from the backup Server.~~
- ~~• Hot (2)
The Client receives the same information from more than one Server by sampling and publishing on more than one Server. This mode ensures that there is no interruption in the data stream and no data loss but the Client should ensure that it detects duplicate data and it shall process two data streams.~~

~~A Client connected to a HotAndMirrored failover redundant set can assume that the Servers are mirroring their state, including connection information and queues. Therefore the Client just connects to one Server and in case of a failover re-establishes its connection using a different Server. In order to validate the capability to connect to other redundant Servers it is allowed to create Sessions with other Servers and maintain the open connections by periodically reading the ServiceLevel. A Client shall not create Subscriptions on the backup Servers for status monitoring. Failover works as described in 6.5. The only difference is that the new SecureChannel is created on another Server with the highest ServiceLevel. Sessions and Subscriptions can be reused since they are mirrored to all Servers.~~

~~A Client can always use a lesser mode than the server supports. For example, the Server supports Hot failover mode and the Client can use the Warm actions. In the case of failover mode HotAndMirrored, the Client shall not use a lesser mode as this would generate unnecessary load on the Servers.~~

6.4.2.4 ~~Hiding Failover with a Server Proxy (Informative)~~

~~A vendor can use the non-transparent redundancy features to create a Server proxy running on the Client machine that makes redundancy transparent to Clients. This reduces the amount of functionality that shall be designed into the Client and to enable simpler Clients to take advantage of non-transparent redundancy. The Server proxy simply duplicates Subscriptions and modifications to Subscriptions, by passing the calls on to both Servers, but only enabling publishing and sampling on one Server. When the proxy detects a failure, it enables publishing and/or sampling on the backup Server, just as the Client would if it were a redundancy-aware Client. The proxy shall run on the Client machine to avoid another single point of failure.~~

~~Figure 27 shows the Server proxy used to provide redundancy to non-redundancy-aware Clients.~~




Figure 27 — Server proxy for redundancy

6.4.3 ~~Client redundancy~~

~~Client redundancy is supported in OPC UA by the TransferSubscriptions call and by exposing Client information in the Server information structures. Since Subscription lifetime is not tied to the Session in which it was created, backup Clients can monitor the active Client's Session with the Server, just as they would monitor any other data variable. If the active Client ceases to be active, the Server shall send a data update to any Client which has that variable monitored. Upon receiving such notification, a backup Client would then instruct the Server to transfer the Subscriptions to its own session. If the Subscription is crafted carefully, with sufficient resources to buffer data during the change-over, there need be no data loss from a Client failover.~~

~~OPC UA does not provide a standardized mechanism for conveying the SessionId and SubscriptionIds from the active Client to the backup Clients, but as long as the backup Clients know the Client name of the active Client, this information is readily available using the SessionDiagnostics and SubscriptionDiagnostics portions of the ServerDiagnostics data.~~

6.4.4 ~~Network redundancy~~

~~Redundant networks can be used with OPC UA either transparent or non-transparent.~~

~~In the transparent case a single *Server Endpoint* can be reached through different network paths. This case is completely handled by the network infrastructure. The selected network path and failover are transparent to the *Client*.~~

~~In the non-transparent case the *Server* provides different *Endpoints* for the different network paths. This requires both the *Server* and the *Client* to support multiple network connections. In this case the *Client* is responsible for selecting the *Endpoint* and for failover. For failover the normal reconnect scenario described in 6.5 can be used. Only the *SecureChannel* is created with another *Endpoint*. *Sessions* and *Subscriptions* can be reused.~~

~~The information about the different network paths is specified in *NonTransparentRedundancyType ObjectType* defined in IEC 62541-5.~~

~~Network redundancy can be combined with server redundancy.~~

6.5 Re-establishing connections

~~After a *Client* establishes a connection to a *Server* and creates a *Subscription*, the *Client* monitors the connection status. Figure 28 shows the steps to connect a *Client* to a *Server* and the general logic for reconnect handling. Not all possible error scenarios are covered.~~

~~The preferred mechanism for a *Client* to monitor the connection status is through the keep-alive of the *Subscription*. A *Client* should subscribe for the *State Variable* in the *ServerStatus* to detect shutdown or other failure states. If no *Subscription* is created or the *Server* does not support *Subscriptions*, the connection can be monitored by periodically reading the *State Variable*.~~

6.6.1 Redundancy overview

OPC UA enables *Servers*, *Clients* and networks to be redundant. OPC UA provides the data structures and *Services* by which *Redundancy* may be achieved in a standardized manner.

Server Redundancy allows *Clients* to have multiple sources from which to obtain the same data. *Server Redundancy* can be achieved in multiple manners, some of which require *Client* interaction, others that require no interaction from a *Client*. Redundant *Servers* could exist in systems without redundant networks or *Clients*. Redundant *Servers* could also coexist in systems with network and *Client Redundancy*. *Server Redundancy* is formally defined in 6.6.2.

Client Redundancy allows identically configured *Clients* to behave as if they were single *Clients*, but not all *Clients* are obtaining data at a given time. Ideally there should be no loss of information when a *Client Failover* occurs. Redundant *Clients* could exist in systems without redundant networks or *Servers*. Redundant *Clients* could also coexist in systems with network and *Server Redundancy*. *Client Redundancy* is formally defined in 6.6.3.

Network Redundancy allows a *Client* and *Server* to have multiple communication paths to obtain the same data. Redundant networks could exist in systems without redundant *Servers* or *Clients*. Redundant networks could also coexist in systems with *Client* and *Server Redundancy*. *Network Redundancy* is formally defined in 6.6.4.

6.6.2 Server Redundancy

6.6.2.1 General

There are two general modes of *Server Redundancy*, transparent and non-transparent.

In transparent *Redundancy* the *Failover* of *Server* responsibilities from one *Server* to another is transparent to the *Client*. The *Client* is unaware that a *Failover* has occurred and the *Client*

has no control over the *Failover* behaviour. Furthermore, the *Client* does not need to perform any actions to continue to send or receive data.

In non-transparent *Redundancy* the *Failover* from one *Server* to another and actions to continue to send or receive data are performed by the *Client*. The *Client* shall be aware of the *Redundant Server Set* and shall perform the required actions to benefit from the *Server Redundancy*.

The *ServerRedundancy Object* defined in IEC 62541-5 indicates the mode supported by the *Server*. The *ServerRedundancyType ObjectType* and its subtypes *TransparentRedundancyType* and *NonTransparentRedundancyType* defined in IEC 62541-5 specify information for the supported *Redundancy* mode.

6.6.2.2 Redundant Server Set requirements

OPC UA *Servers* that are part of a *Redundant Server Set* have certain *AddressSpace* requirements. These requirements allow a *Client* to consistently access information from *Servers* in a *Redundant Server Set* and to make intelligent choices related to the health and availability of *Servers* in the *Redundant Server Set*.

Servers in the *Redundant Server Set* shall have an identical *AddressSpace* including:

- identical *NodeIds*
- identical browse paths and structure of the *AddressSpace*
- identical logic for setting the *ServiceLevel*

The only *Nodes* that can differ between *Servers* in a *Redundant Server Set* are the *Nodes* that are in the local *Server* namespace like the *Server* diagnostic *Nodes*. A *Client* that fails over shall not be required to translate browse paths or otherwise resolve *NodeIds*. *Servers* are allowed to add and delete *Nodes* as long as all *Servers* in the *Redundant Server Set* will be updated with the same *Node* changes.

All *Servers* in a *Redundant Server Set* shall be synchronized with respect to time. This may mean installing a NTP service or a PTP service.

There are other important considerations for a redundant system regarding synchronization:

- EventIds:** Each UA *Server* in a *Transparent and HotAndMirrored Redundant Server Set* shall synchronize *EventIds* to prevent a *Client* from mistakenly processing the same event multiple times simply because the *EventIds* are different. This is very important for Alarms & Conditions. For Cold, Warm, and Hot *Redundant Server Sets*, *Clients* shall be able to handle *EventIds* that are not synchronized. Following any *Failover* the *Client* shall call *ConditionRefresh* defined in IEC 62541-9.
- Timestamp (Source/Server):** If a *Server* is exposing data from a downstream device (PLC, DCS, etc.), then the *SourceTimestamp* and *ServerTimestamp* reported by all redundant *Servers* should match as closely as possible. *Clients* should favour the use of the *SourceTimestamp*.
- ContinuationPoints:** Behaviour of continuation points does not change, in that *Clients* shall be prepared for lost continuation points. *Servers* in *Transparent and HotAndMirrored Redundancy* sets shall synchronize continuation points and they may do so in other modes.

6.6.2.3 Transparent Redundancy

6.6.2.3.1.1 Client behaviour

To a *Client* the transparent *Redundant Server Set* appears as if it is just a single *Server* and the *Client* has no *Failover* actions to perform. All *Servers* in the *Redundant Server Set* have an identical *ServerUri* and an identical *EndpointUrl*.

Figure 26 shows a typical transparent *Redundancy* setup.




Figure 26 – Transparent Redundancy setup example

For transparent *Redundancy*, OPC UA provides data structures to allow *Clients* to identify which *Servers* are in the *Redundant Server Set*, the *ServiceLevel* of each *Server*, and which *Server* is currently responsible for the *Client Session*. This information is specified in *TransparentRedundancyType ObjectType* defined in IEC 62541-5. Since the *ServerUri* is identical for all *Servers* in the *Redundant Server Set*, the *Servers* are identified with a *ServerId* contained in the information provided in the *TransparentRedundancyType Object*.

In transparent *Redundancy*, a *Client* is not able to control which physical *Server* it actually connects to. *Failover* is controlled by the *Redundant Server Set* and a *Client* is also not able to actively *Failover* to another *Server* in the *Redundant Server Set*.

6.6.2.3.1.2 Server requirements

All OPC UA interactions within a given *Session* shall be supported by one *Server* and the *Client* is able to identify which *Server* that is, allowing a complete audit trail for the data. It is the responsibility of the *Servers* to ensure that information is synchronised between the *Servers*. A functional *Server* will take over the *Session* and *Subscriptions* from the *Failed Server*. *Failover* may require a reconnection of the *Client's SecureChannel* but the *EndpointUrl* of the *Server* and the *ServerUri* shall not change. The *Client* shall be able to continue communication with the *Sessions* and *Subscriptions* created on the previously used *Server*.

Figure 26 provides an abstract view of a transparent *Redundant Server Set*. The two or more *Servers* in the *Redundant Server Set* share a virtual network address and therefore all *Servers* have the identical *EndpointUrl*. How this virtual network address is created and managed is vendor specific. There may be special hardware that mediates the network address displayed to the rest of the network. There may be custom hardware, where all components are redundant and *Failover* at a hardware level automatically. There may even be software based systems where all the transparency is governed completely by software.

6.6.2.4 Non-transparent Redundancy

6.6.2.4.1 Overview

For non-transparent *Redundancy*, OPC UA provides the data structures to allow the *Client* to identify what *Servers* are available in the *Redundant Server Set* and also *Server* information which tells the *Client* what modes of *Failover* the *Server* supports. This information allows the *Client* to determine what actions it may need to take in order to accomplish *Failover*. This information is specified in *NonTransparentRedundancyType ObjectType* defined in IEC 62541-5.

Figure 27 shows a typical non-transparent *Redundancy* setup.




Figure 27 – Non-Transparent Redundancy setup

For non-transparent *Redundancy*, the *Servers* will have unique IP addresses. The *Server* also has additional *Failover* modes of *Cold*, *Warm*, *Hot* and *HotAndMirrored*. The *Client* shall be aware of the *Redundant Server Set* and shall be required to perform some actions depending on the *Failover* mode. These actions are described in Table 111 and additional examples and

explanations are provided in 6.6.2.4.5.2 for *Cold*, 6.6.2.4.5.3 for *Warm*, 6.6.2.4.5.4 for *Hot* and 6.6.2.4.5.5 for *HotAndMirrored*.

A *Client* needs to be able to expect that the *SourceTimestamp* associated with a value is approximately the same from all *Servers* in the *Redundant Server Set* for the same value.

6.6.2.4.2 ServiceLevel

The *ServiceLevel* provides information to a *Client* regarding the health of a *Server* and its ability to provide data. See IEC 62541-5 for a formal definition for *ServiceLevel*. The *ServiceLevel* is a byte with a range of 0 to 255, where the values fall into the sub-ranges defined in Table 109.

The algorithm used by a *Server* to determine its *ServiceLevel* within each sub-range is *Server* specific. However, all *Servers* in a *Redundant Server Set* shall use the same algorithm to determine the *ServiceLevel*. All *Servers*, regardless of *Redundant Server Set* membership, shall adhere to the sub-ranges defined in Table 109.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 109 – ServiceLevel ranges

Sub-range	Name	Description
0-0	Maintenance	The <i>Failed Server</i> is in maintenance sub-range. Therefore, new <i>Clients</i> shall not connect and currently connected <i>Clients</i> shall disconnect. The <i>Server</i> should expose a target time at which the <i>Clients</i> are able to reconnect. See <i>EstimatedReturnTime</i> defined in IEC 62541-5 for additional information. A <i>Server</i> that has been set to <i>Maintenance</i> is typically undergoing some maintenance or updates. The main goal for the <i>Maintenance ServiceLevel</i> is to ensure that <i>Clients</i> do not generate load on the <i>Server</i> and allow time for the <i>Server</i> to complete any actions that are required. This load includes even simple connections attempts or monitoring of the <i>ServiceLevel</i>. The <i>EstimatedReturnTime</i> indicates when the <i>Client</i> should check to see if the <i>Server</i> is available. If updates or patches are taking longer than expected the <i>Client</i> may discover that the <i>EstimatedReturnTime</i> has been extended further into the future. If the <i>Server</i> does not provide the <i>EstimatedReturnTime</i>, or if the time has lapsed, the <i>Client</i> should use a much longer interval between reconnects to a <i>Server</i> in the <i>Maintenance</i> sub-range than its normal reconnect interval.
1-1	NoData	The <i>Failed Server</i> is not operational. Therefore, a <i>Client</i> is not able to exchange any information with it. The <i>Server</i> most likely has no data other than <i>ServiceLevel</i>, <i>ServerStatus</i> and diagnostic information available. A <i>Failed Server</i> in this sub-range has no data available. <i>Clients</i> may connect to it to obtain <i>ServiceLevel</i>, <i>ServerStatus</i> and other diagnostic information. If the underlying system has failed, typically the <i>ServerStatus</i> would indicate COMMUNICATION_FAULT_6. The <i>Client</i> may monitor this <i>Server</i> for a <i>ServerStatus</i> and <i>ServiceLevel</i> change, which would indicate that normal communication could be resumed.
2-199	Degraded	The <i>Server</i> is partially operational, but is experiencing problems such that portions of the <i>AddressSpace</i> are out of service or unavailable. To understand <i>Client</i> options, see Degraded <i>Servers</i> discussion in this section. An example usage of this <i>ServiceLevel</i> sub-range would be if 3 of 10 devices connected to a <i>Server</i> are unavailable. <i>Servers</i> that report a <i>ServiceLevel</i> in the <i>Degraded</i> sub-range are partially able to service <i>Client</i> requests. The degradation could be caused by loss of connection to underlying systems. Alternatively, it could be that the <i>Server</i> is overloaded to the point that it is unable to reliably deliver data to <i>Clients</i> in a timely manner. If <i>Clients</i> are experiencing difficulties obtaining required data, they shall switch to another <i>Server</i> if any <i>Servers</i> in the <i>Healthy</i> range are available. If no <i>Servers</i> are available in the <i>Healthy</i> range, then <i>Clients</i> may switch to a <i>Server</i> with a higher <i>ServiceLevel</i> or one that provides the required data. Some <i>Clients</i> may also be configured for higher priority data and may check all <i>Degraded Servers</i>, to see if any of the <i>Servers</i> are able to report as good quality the high priority data, but this functionality would be <i>Client</i> specific. In some cases a <i>Client</i> may connect to multiple <i>Degraded Servers</i> to maximize the available information.
200-255	Healthy	The <i>Server</i> is fully operational. Therefore, a <i>Client</i> can obtain all information from this <i>Server</i>. The sub-range allows a <i>Server</i> to provide information that can be used by <i>Clients</i> to load balance. An example usage of this <i>ServiceLevel</i> sub-range would be to reflect the <i>Server's</i> CPU load where data is delivered as expected. <i>Servers</i> in the <i>Healthy ServiceLevel</i> sub-range are able to deliver information in a timely manner. This <i>ServiceLevel</i> may change for an internal <i>Server</i> reason or it may be used for load balancing described in 6.6.2.4.3. <i>Clients</i> shall connect to the <i>Server</i> with the highest <i>ServiceLevel</i>. Once connected, the <i>ServiceLevel</i> may change, but a <i>Client</i> shall not <i>Failover</i> to a different <i>Server</i> as long as the <i>ServiceLevel</i> of the <i>Server</i> is accessible and in the <i>Healthy</i> sub-range.

6.6.2.4.3 Load balancing

In systems where multiple Hot *Servers* (see 6.6.2.4.5.4) are available, the *Servers* in the *Redundant Server Set* can share the load generated by *Clients* by setting the *ServiceLevel* in the *Healthy* sub-range based on the current load. *Clients* are expected to connect to the *Server* with the highest *ServiceLevel*. *Clients* shall not *Failover* to a different *Server* in the *Redundant Server Set* of *Servers* as long as the *Server* is in the *Healthy* sub-range. This is the normal behaviour for all *Clients*, when communicating with redundant *Servers*. *Servers* can adjust their *ServiceLevel* based on the number of *Clients* that are connected, CPU loading, memory utilization, or any other *Server* specific criteria.

For example in a system with 3 *Servers*, all *Servers* are initially at *ServiceLevel* 255, but when a *Client* connects, the *Server* with the *Client* connection sets its level to 254. The next *Client* would connect to a different *Server* since both of the other *Servers* are still at 255.

It is up to the *Server* vendor to define the logic for spreading the load and the number of expected *Clients*, CPU load or other criteria on each *Server* before the *ServiceLevel* is decremented. It is envisioned that some *Servers* would be able to accomplish this without any communication between the *Servers*.

6.6.2.4.4 Server Failover modes

The *Failover* mode of a *Server* is provided in the *ServerRedundancy Object* defined in IEC 62541-5. The different Failover modes for non-transparent Redundancy are described in Table 110.

Table 110 – Server Failover modes

Name	Description
Cold	<i>Cold Failover</i> mode is where only one <i>Server</i> can be active at a time. This may mean that redundant <i>Servers</i> are unavailable (not powered up) or are available but not running (PC is running, but application is not started)
Warm	<i>Warm Failover</i> mode is where the backup <i>Server(s)</i> can be active, but cannot connect to actual data points (typically, a system where the underlying devices are limited to a single connection). Underlying devices, such as PLCs, may have limited resources that permit a single <i>Server</i> connection. Therefore, only a single <i>Server</i> will be able to consume data. The <i>ServiceLevel Variable</i> defined in IEC 62541-5 indicates the ability of the <i>Server</i> to provide its data to the <i>Client</i> .
Hot	<i>Hot Failover</i> mode is where all <i>Servers</i> are powered-on, and are up and running. In scenarios where <i>Servers</i> acquire data from a downstream device, such as a PLC, then one or more <i>Servers</i> are actively connected to the downstream device(s) in parallel. These <i>Servers</i> have minimal knowledge of the other <i>Servers</i> in their group and are independently functioning. When a <i>Server</i> fails or encounters a serious problem then its <i>ServiceLevel</i> drops. On recovery, the <i>Server</i> returns to the <i>Redundant Server Set</i> with an appropriate <i>ServiceLevel</i> to indicate that it is available.
HotAndMirrored	<i>HotAndMirrored Failover</i> mode is where <i>Failovers</i> are for <i>Servers</i> that are mirroring their internal states to all <i>Servers</i> in the <i>Redundant Server Set</i> and more than one <i>Server</i> can be active and fully operational. Mirroring state minimally includes <i>Sessions</i> , <i>Subscriptions</i> , registered <i>Nodes</i> , <i>ContinuationPoints</i> , sequence numbers, and sent <i>Notifications</i> . The <i>ServiceLevel Variable</i> defined in IEC 62541-5 should be used by the <i>Client</i> to find the <i>Servers</i> with the highest <i>ServiceLevel</i> to achieve load balancing.

6.6.2.4.5 Client Failover behaviour

6.6.2.4.5.1 General


Each *Server* maintains a list of *ServerUris* for all redundant *Servers* in the *Redundant Server Set*. The list is provided together with the *Failover* mode in the *ServerRedundancy Object* defined in IEC 62541-5. To enable *Clients* to connect to all *Servers* in the list, each *Server* in the list shall provide the *ApplicationDescription* for all *Servers* in the *Redundant Server Set* through the *FindServers Service*. This information is needed by the *Client* to translate the *ServerUri* into information needed to connect to the other *Servers* in the *Redundant Server Set*. Therefore a *Client* needs to connect to only one of the redundant *Servers* to find the other *Servers* based on the provided information. A *Client* should persist information about other *Servers* in the *Redundant Server Set*.

Table 111 defines a list of *Client* actions for initial connections and *Failovers*.

Table 111 – Redundancy Failover actions

Failover mode and <i>Client</i> options	Cold	Warm	Hot (a)	Hot (b)	HotAndMirrored
On initial connection in addition to actions on Active Server:					
Connect to more than one OPC UA Server.		X	X	X	Optional for status check
Create <i>Subscriptions</i> and add monitored items.		X	X	X	
Activate sampling on the <i>Subscriptions</i> .			X	X	
Activate publishing.				X	
At Failover:					
OpenSecureChannel to backup OPC UA Server	X				X
CreateSession on backup OPC UA Server	X				
ActivateSession on backup OPC UA Server	X				X
Create <i>Subscriptions</i> and add monitored items.	X				
Activate sampling on the <i>Subscriptions</i> .	X	X			
Activate publishing.	X	X	X		

Clients communicating with a non-transparent *Redundant Server Set* of *Servers* require some additional logic to be able to handle *Server* failures and to *Failover* to another *Server* in the *Redundant Server Set*. Figure 28 provides an overview of the steps a *Client* typically performs when it is first connecting to a *Redundant Server Set*. The figure does not cover all possible error scenarios.

**Figure 28 – Client Start-up steps**

The initial *Server* may be obtained via standard discovery or from a persisted list of *Servers* in the *Redundant Server Set*. But in any case the *Client* needs to check which *Server* in the *Server* set it should connect to. Individual actions will depend on the *Server Failover* mode the *Server* provides and the *Failover* mode the *Client* will use.

Clients once connected to a redundant *Server* shall be aware of the modes of *Failover* supported by a *Server* since this support affects the available options related to *Client* behaviour. A *Client* may always treat a *Server* using a lesser *Failover* mode, i.e. for a *Server* that provides *Hot Redundancy*, a *Client* might connect and choose to treat it as if the *Server* was running in *Warm Redundancy* or *Cold Redundancy*. This choice is up to the client. In the case of *Failover* mode *HotAndMirrored*, the *Client* shall not use *Failover* mode *Hot* or *Warm* as it would generate unnecessary load on the *Servers*.

6.6.2.4.5.2 Cold

A *Cold Failover* mode is where the *Client* can only connect to one *Server* at a time. When the *Client* loses connectivity with the *Active Server* it will attempt a connection to the redundant *Server(s)* which may or may not be available. In this situation the *Client* may need to wait for the redundant *Server* to become available and then create *Subscriptions* and *MonitoredItems* and activate publishing. The *Client* shall cache any information that is required related to the list of available *Servers* in the *Redundant Server Set*. Figure 29 illustrates the action a *Client* would take if it is talking to a *Server* using *Cold Failover* mode.




Figure 29 – Cold Failover

NOTE. There may be a loss of data from the time the connection to the *Active Server* is interrupted until the time the *Client* gets *Publish Responses* from the backup *Server*.

6.6.2.4.5.3 Warm

A *Warm Failover* mode is where the *Client* should connect to one or more *Servers* in the *Redundant Server Set* primarily to monitor the *ServiceLevel*. A *Client* can connect and create *Subscriptions* and *MonitoredItems* on more than one *Server*, but sampling and publishing can only be active on one *Server*. However, the active *Server* will return actual data, whereas the other *Servers* in the *Redundant Server Set* will return an appropriate error for the *MonitoredItems* in the *Publish* response such as *Bad_NoCommunication*. The one *Active Server* can be found by reading the *ServiceLevel Variable* from all *Servers*. The *Server* with the highest *ServiceLevel* is the *Active Server*. For *Failover* the *Client* activates sampling and publishing on the *Server* with the highest *ServiceLevel*. Figure 30 illustrates the steps a *Client* would perform when communicating with a *Server* using *Warm Failover* mode.




Figure 30 – Warm Failover

NOTE There may be a temporary loss of data from the time the connection to the *Active Server* is interrupted until the time the *Client* gets *Publish Responses* from the backup *Server*.

6.6.2.4.5.4 Hot

A *Hot Failover* mode is where the *Client* should connect to two or more *Servers* in the *Redundant Server Set* and to subscribe to the *ServiceLevel* variable defined in IEC 62541-5 to find the highest *ServiceLevel* to achieve load balancing; this means that *Clients* should issue *Service* requests such as *Browse*, *Read*, *Write* to the *Server* with the highest *ServiceLevel*. *Subscription* related activities will need to be invoked for each connected *Server*. *Clients* have the following choices for implementing subscription behaviour in a *Hot Failover* mode:

- The *Client* connects to multiple *Servers* and establishes subscription(s) in each where only one is *Reporting*; the others are *Sampling* only. The *Client* should setup the queue size for the *MonitoredItems* such that it can buffer all changes during the *Failover* time. The *Failover* time is the time between the connection interruption and the time the *Client* gets *Publish Responses* from the backup *Server*. On a *Failover* the *Client* shall enable *Reporting* on the *Server* with the next highest availability.
- The *Client* connects to multiple *Servers* and establishes subscription(s) in each where all subscriptions are *Reporting*. The *Client* is responsible for handling/processing multiple subscription streams concurrently.

Figure 31 illustrate the functionality a *Client* would perform when communicating with a *Server* using *Hot Failover* mode (Figure 31 includes both (a) and (b) options).





Figure 31 – Hot Failover

Clients are not expected to automatically switch over to a *Server* that has recovered from a failure, but the *Client* should establish a connection to it.

6.6.2.4.5.5 HotAndMirrored

A *HotAndMirrored Failover* mode is where a *Client* only connects to one *Server* in the *Redundant Server Set* because the *Server* will share this session/state information with the other *Servers*. In order to validate the capability to connect to other redundant *Servers* it is allowed to create *Sessions* with other *Servers* and maintain the open connections by periodically reading the *ServiceLevel*. A *Client* shall not create *Subscriptions* on the backup *Servers* for status monitoring (to prevent excessive load on the *Servers*). This mode allows *Clients* to fail over without creating a new context for communication. On a *Failover* the *Client* will simply create a new *SecureChannel* on an alternate *Server* and then call *ActivateSession*; all *Client* activities (browsing, subscriptions, history reads, etc.) will then resume. Figure 32

illustrates the behaviour a *Client* would perform when communicating to a *Server* in *HotAndMirrored Failover* mode.



IEC

Figure 32 – HotAndMirrored Failover

This *Failover* mode is similar to the transparent *Redundancy*. The advantage is that the *Client* has full control over selecting the *Server*. The disadvantage is that the *Client* needs to be able to handle *Failovers*.

6.6.2.5 Hiding Failover with a Server Proxy

A vendor can use the non-transparent *Redundancy* features to create a *Server* proxy running on the *Client* machine to provide transparent *Redundancy* to the client. This reduces the amount of functionality that needs to be designed into the *Client* and to enable simpler *Clients* to take advantage of non-transparent *Redundancy*. The *Server* proxy simply duplicates *Subscriptions* and modifications to *Subscriptions*, by passing the calls on to both *Servers*, but only enabling publishing and sampling on one *Server*. When the proxy detects a failure, it enables publishing and/or sampling on the backup *Server*, just as the *Client* would if it were a *Redundancy* aware *Client*.

Figure 33 shows the *Server* proxy used to provide transparent *Redundancy*.




Figure 33 – Server proxy for Redundancy

6.6.3 Client Redundancy

Client Redundancy is supported in OPC UA by the *TransferSubscriptions Service* and by exposing *Client* information in the *Server* diagnostic information. Since *Subscription* lifetime is not tied to the *Session* in which it was created, backup *Clients* may use standard diagnostic information available to monitor the active *Client's Session* with the *Server*. Upon detection of an active *Client* failure, a backup *Client* would then instruct the *Server* to transfer the *Subscriptions* to its own session. If the *Subscription* is crafted carefully, with sufficient resources to buffer data during the change-over, data loss from a *Client Failover* can be prevented.

OPC UA does not provide a standardized mechanism for conveying the *SessionId* and *SubscriptionIds* from the active *Client* to the backup *Clients*, but as long as the backup *Clients* know the *Client* name of the active *Client*, this information is readily available using the *SessionDiagnostics* and *SubscriptionDiagnostics* portions of the *ServerDiagnostics* data. This information is available for authorized users and for the user active on the *Session*. *TransferSubscriptions* requires the same user on all redundant *Clients* to succeed.

6.6.4 Network Redundancy

6.6.4.1 Overview

Redundant networks can be used with OPC UA in either transparent or non-transparent *Redundancy*.

Network Redundancy can be combined with *Server* and *Client Redundancy*.

6.6.4.2 Transparent

In the transparent network use-case a single *Server Endpoint* can be reached through different network paths. This case is completely handled by the network infrastructure. The selected network path and *Failover* are transparent to the *Client* and the *Server*. Transparent network *Redundancy* is illustrated in Figure 34.





Figure 34 – Transparent network Redundancy

Examples:

- A physical appliance/device such as a router or gateway which automatically changes the network routing to maintain communications.
- A virtual adapter which automatically changes the network adapter to maintain communications.

6.6.4.3 Non-transparent

In the non-transparent network use-case the *Server* provides different *Endpoints* for the different network paths. This requires both the *Server* and the *Client* to support multiple network connections. In this case the *Client* is responsible for selecting the *Endpoint* and for *Failover*. For *Failover* the normal reconnect scenario described in 6.7 can be used. Only the *SecureChannel* is created with another *Endpoint*. *Sessions* and *Subscriptions* can be reused. Non-transparent network *Redundancy* is illustrated in Figure 35.



IEC

Figure 35 – Non-transparent network Redundancy

The information about the different network paths is specified in *NonTransparentRedundancyType* *ObjectType* defined in IEC 62541-5.

6.6.5 Manually forcing Failover

In redundant systems, it is common to require that a particular *Server* in the *Redundant Server Set* be taken out of the *Redundant Server Set* for a period of time. Some items that could cause this may include:

- certificate update;
- security reconfiguration;
- rebooting or restarting of the machine for
 - software updates and patches,
 - installation of new software;
- reconfiguration of the *AddressSpace*.

The removal from the *Redundant Server Set* can be done through a complete shutdown or by setting the *ServiceLevel* of the *Server* to *Maintenance* sub-range. This can be done through a *Server* specific configuration tool or through the *Method RequestServerStateChange* on the *ServerRedundancyType*. The *Method* is formally defined in IEC 62541-5.

This *Method* requires that the *Client* provide credentials with administrative rights on the *Server*.

6.7 Re-establishing connections

After a *Client* establishes a connection to a *Server* and creates a *Subscription*, the *Client* monitors the connection status. Figure 36 shows the steps to connect a *Client* to a *Server* and the general logic for reconnect handling. Not all possible error scenarios are covered.

The preferred mechanism for a *Client* to monitor the connection status is through the keep-alive of the *Subscription*. A *Client* should subscribe for the *State Variable* in the *ServerStatus* to detect shutdown or other failure states. If no *Subscription* is created or the *Server* does not support *Subscriptions*, the connection can be monitored by periodically reading the *State Variable*.




Figure 36 – Reconnect sequence

When a *Client* loses the connection to the *Server*, the goal is to reconnect without losing information. To do this the *Client* shall re-establish the connection by creating a new *SecureChannel* and activating the *Session* with the *Service ActivateSession*. This assigns the new *SecureChannel* to the existing *Session* and allows the *Client* to reuse the *Session* and *Subscriptions* in the *Server*. To re-establish the *SecureChannel* and activate the *Session*, the *Client* shall use the same security policy, application instance certificate and the same user credential used to create the original *SecureChannel*. This will result in the *Client* receiving data and event *Notifications* without losing information provided the queues in the *MonitoredItems* do not overflow.

The *Client* shall only create a new *Session* if *ActivateSession* fails. *TransferSubscriptions* is used to transfer the *Subscription* to the new *Session*. If *TransferSubscriptions* fails, the *Client* needs to create a new *Subscription*.

When the connection is lost, *Publish* responses may have been sent but not received by the *Client*.

After re-establishing the connection the *Client* shall call *Republish* in a loop, starting with the next expected sequence number and incrementing the sequence number until the *Server* returns the status *Bad_MessageNotAvailable*. After receiving this status, the *Client* shall start sending *Publish* requests with the normal *Publish* handling. This sequence ensures that the lost *NotificationMessages* queued in the *Server* are not overwritten by new *Publish* responses.

If the *Client* detects missing sequence numbers in the *Publish* and is not able to get the lost *NotificationMessages* through *Republish*, the *Client* should use the *Method ResendData* or should read the values of all data *MonitoredItems* to make sure the *Client* has the latest values for all *MonitoredItems*.

The *Server Object* provides a *Method ResendData* that initiates resending of all data monitored items in a *Subscription*. This *Method* is defined in IEC 62541-5. If this *Method* is called, subsequent *Publish* responses shall contain the current values of all data *MonitoredItems* in the *Subscription* where the *MonitoringMode* is set to *Reporting*. If a value is queued for a data *MonitoredItem*, the next value in the queue is sent in the *Publish* response. If no value is queued for a data *MonitoredItem*, the last value sent is repeated in the *Publish* response. The *Server* shall verify that the *Method* is called within the *Session* context of the *Session* that owns the *Subscription*.

Independent of the detailed recovery strategy, the *Client* should make sure that it does not overwrite newer data in the *Client* with older values provided through *Republish*.

If the *Republish* returns *Bad_SubscriptionIdInvalid*, then the *Client* needs to create a new *Subscription*.

Re-establishing the connection by creating a new *SecureChannel* may be rejected, because of a new *Server Application Instance Certificate* or other security errors. In case of security failures, the *Client* shall use the *GetEndpoints Service* to fetch the most up to date security information from the *Server*.

IEC 62541-6 defines a reverse connect mechanism where the *Server* initiates the logical connection. All subsequent steps like creating a *SecureChannel* are initiated by the *Client*. In this scenario the *Client* is only able to initiate a reconnect if the *Server* initiates a new logical connection after a connection interruption. The *Client* side reconnect handling described in Figure 36 applies also to the reverse connect case. A *Server* is not able to actively check the connection status; therefore the *Server* shall initiate a new connection in a configurable interval, even if a connection to the *Client* is established. This ensures that an initiated connection is available for the reconnect handling in addition to other scenarios where the *Client* needs more than one connection.

6.8 Durable Subscriptions

MonitoredItems are used to monitor *Variable Values* for data changes and event notifier *Objects* for new *Events*. Subscriptions are used to combine data changes and events of the assigned *MonitoredItems* to an optimized stream of network messages. A reliable delivery is ensured as long as the lifetime of the *Subscription* and the queues in the *MonitoredItems* are long enough for a network interruption between OPC UA *Client* and *Server*. All queues that ensure reliable delivery are normally kept in memory and a *Server* restart would delete them.

There are use cases where OPC UA *Clients* have no permanent network connection to the OPC UA *Server* or where reliable delivery of data changes and events is necessary even if the OPC UA *Server* is restarted or the network connection is interrupted for a longer time.

To ensure this reliable delivery, the OPC UA *Server* shall store collected data and events in non-volatile memory until the OPC UA *Client* has confirmed reception. It is possible that there will be data lost if the *Server* is not shut down gracefully or in case of power failure. But the OPC UA *Server* should store the queues frequently even if the *Server* is not shut down.

The *Method SetSubscriptionDurable* defined in IEC 62541-5 is used to set a *Subscription* into this durable mode and to allow much longer lifetimes and queue sizes than for normal *Subscriptions*. The *Method* shall be called before the *MonitoredItems* are created in the durable *Subscription*. The *Server* shall verify that the *Method* is called within the *Session* context of the *Session* that owns the *Subscription*.

A value of 0 for the parameter *lifetimeInHours* requests the highest lifetime supported by the *Server*.

An OPC UA *Server* providing durable *Subscriptions* shall:

- support the *SetSubscriptionDurable Method* defined in IEC 62541-5;
- support *Service TransferSubscriptions*;
- support long *Subscription* lifetimes, minimum requirement is defined in IEC 62541-7;
- support large *MonitoredItem* queues, minimum requirement is defined in IEC 62541-7;
- store *Subscriptions* settings and sent notification messages with sequence numbers;
- store *MonitoredItem* settings and queues.

An OPC UA *Client* using durable *Subscriptions* shall:

- use the *SetSubscriptionDurable Method* defined in IEC 62541-5 to create a durable *Subscription*;
- close *Sessions* for planned communication interruptions;
- use the *Service TransferSubscriptions* to assign the durable *Subscription* to a new *Session* for data transfer;
- store *SubscriptionId*, *MonitoredItem* client and server handles and the last confirmed sequence number.

7 Common parameter type definitions

7.1 ApplicationDescription

The components of this parameter are defined in Table 112.

Table 112 – ApplicationDescription

Name	Type	Description
ApplicationDescription	structure	Specifies an application that is available.
applicationUri	String	The globally unique identifier for the application instance. This URI is used as <i>ServerUri</i> in <i>Services</i> if the application is a <i>Server</i> .
productUri	String	The globally unique identifier for the product.
applicationName	LocalizedText	A localized descriptive name for the application.
applicationType	Enum ApplicationType	The type of application. This value is an enumeration with one of the following values: SERVER_0 The application is a <i>Server</i>. CLIENT_1 The application is a <i>Client</i>. CLIENTANDSERVER_2 The application is a <i>Client</i> and a <i>Server</i>. DISCOVERYSERVER_3 The application is a <i>DiscoveryServer</i>.
gatewayServerUri	String	A URI that identifies the <i>Gateway Server</i> associated with the <i>discoveryUrls</i> . This value is not specified if the <i>Server</i> can be accessed directly. This field is not used if the <i>applicationType</i> is CLIENT_1.
discoveryProfileUri	String	A URI that identifies the discovery profile supported by the URLs provided. This field is not used if the <i>applicationType</i> is CLIENT_1. If this value is not specified then the Endpoints shall support the Discovery Services defined in 5.4. Alternate discovery profiles are defined in IEC 62541-7.
discoveryUrls []	String	A list of URLs for the <i>DiscoveryEndpoints</i> provided by the application. If the <i>applicationType</i> is CLIENT_1, this field shall contain an empty list.

7.2 ApplicationInstanceCertificate

An *ApplicationInstanceCertificate* is a *ByteString* containing an encoded *Certificate*. The encoding of an *ApplicationInstanceCertificate* depends on the security technology mapping and is defined completely in IEC 62541-6. Table 113 specifies the information that shall be contained in an *ApplicationInstanceCertificate*.

Table 113 – ApplicationInstanceCertificate

Name	Type	Description
ApplicationInstanceCertificate	structure	<i>ApplicationInstanceCertificate</i> with signature created by a <i>Certificate Authority</i> .
version	String	An identifier for the version of the <i>Certificate</i> encoding.
serialNumber	ByteString	A unique identifier for the <i>Certificate</i> assigned by the Issuer.
signatureAlgorithm	String	The algorithm used to sign the <i>Certificate</i> . The syntax of this field depends on the <i>Certificate</i> encoding.
signature	ByteString	The signature created by the Issuer.
issuer	Structure	A name that identifies the <i>Issuer Certificate</i> used to create the signature.
validFrom	UtcTime	When the <i>Certificate</i> becomes valid.
validTo	UtcTime	When the <i>Certificate</i> expires.
subject	Structure	A name that identifies the application instance that the <i>Certificate</i> describes. This field shall contain the <i>productName</i> and the name of the organization responsible for the application instance.
applicationUri	String	The <i>applicationUri</i> specified in the <i>ApplicationDescription</i> . The <i>ApplicationDescription</i> is described in 7.1.
hostnames []	String	The name of the machine where the application instance runs. A machine may have multiple names if it is accessible via multiple networks. The hostname may be a numeric network address or a descriptive name. <i>Server Certificates</i> shall have at least one hostname defined.
publicKey	ByteString	The public key associated with the <i>Certificate</i> .
keyUsage []	String	Specifies how the <i>Certificate</i> key may be used. <i>ApplicationInstanceCertificates</i> shall support Digital Signature, Non-Repudiation Key Encryption, Data Encryption and Client/Server Authorization. The contents of this field depend on the <i>Certificate</i> encoding.

7.3 BrowseResult

The components of this parameter are defined in Table 114.

Table 114 – BrowseResult

Name	Type	Description
BrowseResult	structure	The results of a Browse operation.
statusCode	StatusCode	The status for the <i>BrowseDescription</i> . This value is set to <i>Good</i> if there are still references to return for the <i>BrowseDescription</i> .
continuationPoint	ContinuationPoint	A <i>Server</i> defined opaque value that identifies the continuation point. The <i>ContinuationPoint</i> type is defined in 7.6.
References []	ReferenceDescription	The set of references that meet the criteria specified in the <i>BrowseDescription</i> . Empty, if no <i>References</i> met the criteria. The Reference Description type is defined in 7.25.

7.4 ContentFilter

7.4.1 ContentFilter structure

The ContentFilter structure defines a collection of elements that define filtering criteria. Each element in the collection describes an operator and an array of operands to be used by the operator. The operators that can be used in a ContentFilter are described in Table 119. The filter is evaluated by evaluating the first entry in the element array starting with the first operand in the operand array. The operands of an element may contain References to sub-elements resulting in the evaluation continuing to the referenced elements in the element array. The evaluation shall not introduce loops. For example evaluation starting from element "A" shall never be able to return to element "A". However there may be more than one path leading to another element "B". If an element cannot be traced back to the starting element it is ignored. Extra operands for any operator shall result in an error. Annex B provides examples using the ContentFilter structure.

Table 115 defines the *ContentFilter* structure.

Table 115 – ContentFilter structure

Name	Type	Description
ContentFilter	structure	
elements []	ContentFilterElement	List of operators and their operands that compose the filter criteria. The filter is evaluated by starting with the first entry in this array. This structure is defined in-line with the following indented items.
filterOperator	Enum FilterOperator	Filter operator to be evaluated. The <i>FilterOperator</i> enumeration is defined in Table 119.
filterOperands []	Extensible Parameter FilterOperand	Operands used by the selected operator. The number and use depend on the operators defined in Table 119. This array needs at least one entry. This extensible parameter type is the <i>FilterOperand</i> parameter type specified in 7.4.4. It specifies the list of valid <i>FilterOperand</i> values.

7.4.2 ContentFilterResult

The components of this data type are defined in Table 116.

Table 116 – ContentFilterResult structure

Name	Type	Description
ContentFilterResult	structure	A structure that contains any errors associated with the filter.
elementResults []	ContentFilter ElementResult	A list of results for individual elements in the filter. The size and order of the list matches the size and order of the elements in the <i>ContentFilter</i> parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	The status code for a single element.
operandStatusCodes []	StatusCode	A list of status codes for the operands in an element. The size and order of the list matches the size and order of the operands in the <i>ContentFilterElement</i> . This list is empty if no operand errors occurred.
operandDiagnosticInfos []	DiagnosticInfo	A list of diagnostic information for the operands in an element. The size and order of the list matches the size and order of the operands in the <i>ContentFilterElement</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the operands.
elementDiagnosticInfos []	DiagnosticInfo	A list of diagnostic information for individual elements in the filter. The size and order of the list matches the size and order of the elements in the <i>filter</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the elements.

Table 117 defines values for the *StatusCode* parameter that are specific to this structure. Common *StatusCodes* are defined in Table 178.

Table 117 – ContentFilterResult result codes

Symbolic Id	Description
Bad_FilterOperandCountMismatch	The number of operands provided for the filter operator was less than expected for the operand provided.
Bad_FilterOperatorInvalid	An unrecognized operator was provided in a filter.
Bad_FilterOperatorUnsupported	A valid operator was provided, but the <i>Server</i> does not provide support for this filter operator.

Table 118 defines values for the *operandStatusCodes* parameter that are specific to this structure. Common *StatusCodes* are defined in Table 178.

Table 118 – ContentFilterResult Operand result codes

Symbolic Id	Description
Bad_FilterOperandInvalid	See Table 178 for the description of this result code.
Bad_FilterElementInvalid	The referenced element is not a valid element in the content filter.
Bad_FilterLiteralInvalid	The referenced literal is not a valid <i>BaseDataType</i> .
Bad_AttributeIdInvalid	The attribute id is not a valid attribute id in the system.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_NotTypeDefinition	The provided <i>NodeId</i> was not a type definition <i>NodeId</i> .

7.4.3 FilterOperator

Table 119 defines the basic operators that can be used in a *ContentFilter*. See Table 120 for a description of advanced operators. See 7.4.4 for a definition of operands.

Table 119 – Basic FilterOperator definition

Operator	Number of Operands	Description
Equals_0	2	TRUE if operand[0] is equal to operand[1]. If the operands are of different types, the system shall perform any implicit conversion to a common type. This operator resolves to FALSE if no implicit conversion is available and the operands are of different types. This operator returns FALSE if the implicit conversion fails. See the discussion on data type precedence in Table 123 for more information how to convert operands of different types.
IsNull_1	1	TRUE if operand[0] is a null value.
GreaterThan_2	2	TRUE if operand[0] is greater than operand[1]. The following restrictions apply to the operands: [0]: Any operand that resolves to an ordered value. [1]: Any operand that resolves to an ordered value. The same conversion rules as defined for <i>Equals</i> apply.

Operator	Number of Operands	Description
LessThan_3	2	TRUE if operand[0] is less than operand[1]. The same conversion rules and restrictions as defined for <i>GreaterThan</i> apply.
GreaterThanOrEqual_4	2	TRUE if operand[0] is greater than or equal to operand[1]. The same conversion rules and restrictions as defined for <i>GreaterThan</i> apply.
LessThanOrEqual_5	2	TRUE if operand[0] is less than or equal to operand[1]. The same conversion rules and restrictions as defined for <i>GreaterThan</i> apply.
Like_6	2	TRUE if operand[0] matches a pattern defined by operand[1]. See Table 121 for the definition of the pattern syntax. The following restrictions apply to the operands: [0]: Any operand that resolves to a String. [1]: Any operand that resolves to a String. This operator resolves to FALSE if no operand can be resolved to a string.
Not_7	1	TRUE if operand[0] is FALSE. The following restrictions apply to the operands: [0]: Any operand that resolves to a Boolean. If the operand cannot be resolved to a Boolean, the result is a NULL. See below for a discussion on the handling of NULL.
Between_8	3	TRUE if operand[0] is greater or equal to operand[1] and less than or equal to operand[2]. The following restrictions apply to the operands: [0]: Any operand that resolves to an ordered value. [1]: Any operand that resolves to an ordered value. [2]: Any operand that resolves to an ordered value. If the operands are of different types, the system shall perform any implicit conversion to match all operands to a common type. If no implicit conversion is available and the operands are of different types, the particular result is FALSE. See the discussion on data type precedence in Table 123 for more information how to convert operands of different types.
InList_9	2..n	TRUE if operand[0] is equal to one or more of the remaining operands. The Equals Operator is evaluated for operand[0] and each remaining operand in the list. If any Equals evaluation is TRUE, InList returns TRUE.
And_10	2	TRUE if operand[0] and operand[1] are TRUE. The following restrictions apply to the operands: [0]: Any operand that resolves to a Boolean. [1]: Any operand that resolves to a Boolean. If any operand cannot be resolved to a Boolean it is considered a NULL. See below for a discussion on the handling of NULL.
Or_11	2	TRUE if operand[0] or operand[1] are TRUE. The following restrictions apply to the operands: [0]: Any operand that resolves to a Boolean. [1]: Any operand that resolves to a Boolean. If any operand cannot be resolved to a Boolean it is considered a NULL. See below for a discussion on the handling of NULL.

Operator	Number of Operands	Description
Cast_12	2	Converts operand[0] to a value with a data type with a NodeId identified by operand[1]. The following restrictions apply to the operands: [0]: Any operand. [1]: Any operand that resolves to a NodeId or ExpandedNodeId where the <i>Node</i> is of the <i>NodeClass DataType</i>. If there is any error in conversion or in any of the parameters then the Cast Operation evaluates to a NULL. See below for a discussion on the handling of NULL.
BitwiseAnd_16	2	The result is an integer which matches the size of the largest operand and contains a bitwise And operation of the two operands where both have been converted to the same size (largest of the two operands). The following restrictions apply to the operands: [0]: Any operand that resolves to an integer. [1]: Any operand that resolves to an integer. If any operand cannot be resolved to an integer it is considered a NULL. See below for a discussion on the handling of NULL.
BitwiseOr_17	2	The result is an integer which matches the size of the largest operand and contains a bitwise Or operation of the two operands where both have been converted to the same size (largest of the two operands). The following restrictions apply to the operands: [0]: Any operand that resolves to an integer. [1]: Any operand that resolves to an integer. If any operand cannot be resolved to an integer it is considered a NULL. See below for a discussion on the handling of NULL.

Many operands have restrictions on their type. This requires the operand to be evaluated to determine what the type is. In some cases the type is specified in the operand (i.e. a *LiteralOperand*). In other cases the type requires that the value of an attribute be read. An *ElementOperand* evaluates to a Boolean value unless the operator is a Cast or a nested *RelatedTo* operator.

Table 120 defines complex operators that require a target node (i.e. row) to evaluate. These operators shall be re-evaluated for each possible target node in the result set.

Table 120 – Complex FilterOperator definition

Operator	Number of Operands	Description
InView_13	1	TRUE if the target <i>Node</i> is contained in the <i>View</i> defined by operand[0]. The following restrictions apply to the operands: [0]: Any operand that resolves to a <i>NodeId</i> that identifies a <i>View Node</i>. If operand[0] does not resolve to a <i>NodeId</i> that identifies a <i>View Node</i>, this operation shall always be False.
OfType_14	1	TRUE if the target <i>Node</i> is of type operand[0] or of a subtype of operand[0]. The following restrictions apply to the operands: [0]: Any operand that resolves to a <i>NodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i>. If operand[0] does not resolve to a <i>NodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i>, this operation shall always be False.
RelatedTo_15	6	TRUE if the target <i>Node</i> is of type operand[0] and is related to a <i>NodeId</i> of the type defined in operand[1] by the <i>Reference</i> type defined in operand[2]. operand[0] or operand[1] can also point to an element <i>Reference</i> where the referred to element is another <i>RelatedTo</i> operator. This allows chaining of relationships (e.g. A is related to B is related to C), where the relationship is defined by the <i>ReferenceType</i> defined in operand[2]. In this case, the referred to element returns a list of <i>NodeIds</i> instead of TRUE or FALSE. In this case if any errors occur or any of the operands cannot be resolved to an appropriate value, the result of the chained relationship is an empty list of nodes. Operand[3] defines the number of hops for which the relationship should be followed. If operand[3] is 1, then objects shall be directly related. If a hop is greater than 1, then a <i>NodeId</i> of the type described in operand[1] is checked for at the depth specified by the hop. In this case, the type of the intermediate <i>Node</i> is undefined, and only the <i>Reference</i> type used to reach the end <i>Node</i> is defined. If the requested number of hops cannot be followed, then the result is FALSE, i.e., an empty <i>Node</i> list. If operand[3] is 0, the relationship is followed to its logical end in a forward direction and each <i>Node</i> is checked to be of the type specified in operand[1]. If any <i>Node</i> satisfies this criterion, then the result is TRUE, i.e., the <i>NodeId</i> is included in the sub-list. Operand [4] defines if operands [0] and [1] should include support for subtypes of the types defined by these operands. A TRUE indicates support for subtypes operand [5] defines if operand [2] should include support for subtypes of the reference type. A TRUE indicates support for subtypes. The following restrictions apply to the operands: [0]: Any operand that resolves to a <i>NodeId</i> or <i>ExpandedNodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i> or a reference to another element which is a <i>RelatedTo</i> operator. [1]: Any operand that resolves to a <i>NodeId</i> or <i>ExpandedNodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i> or a reference to another element which is a <i>RelatedTo</i> operator. [2]: Any operand that resolves to a <i>NodeId</i> that identifies a <i>ReferenceType Node</i>. [3]: Any operand that resolves to a value implicitly convertible to UInt32. [4]: Any operand that resolves to a value implicitly convertible to a Boolean; if this operand does not resolve to a Boolean, then a value of FALSE is used. [5]: Any operand that resolves to a value implicitly convertible to a Boolean; if this operand does not resolve to a Boolean, then a value of FALSE is used. If none of the operands [0],[1],[2],[3] resolves to an appropriate value then the result of this operation shall always be False (or an Empty set in the case of a nested <i>RelatedTo</i> operand). See examples for <i>RelatedTo</i> in B.2.

The *RelatedTo* operator can be used to identify if a given type, set as operand[1], is a subtype of another type set as operand[0] by setting operand[2] to the *HasSubtype ReferenceType* and operand[3] to 0.

The *Like* operator can be used to perform wildcard comparisons. Several special characters can be included in the second operand of the *Like* operator. The valid characters are defined in Table 121. The wildcard characters can be combined in a single string (i.e. 'Th[ia][ts]%' would match 'That is fine', 'This is fine', 'That as one', 'This it is', 'Then at any', etc.). The *Like* operator is case sensitive.

Table 121 – Wildcard characters

Special Character	Description
%	Match any string of zero or more characters (i.e. 'main%' would match any string that starts with 'main', '%en%' would match any string that contains the letters 'en' such as 'entail', 'green' and 'content'.) If a '%' sign is intend in a string the list operand can be used (i.e. 5[%] would match '5%').
_	Match any single character (i.e. '_ould' would match 'would', 'could'). If the '_' is intended in a string then the list operand can be used (i.e. 5[_] would match '5_').
\	Escape character allows literal interpretation (i.e. \\ is \, \% is %, _ is _)
[]	Match any single character in a list (i.e. 'abc[13-68]' would match 'abc1', 'abc3', 'abc4', 'abc5', 'abc6', and 'abc8'. 'xyz[c-f]' would match 'xyzc', 'xyzd', 'xyze', 'xyzf').
[^]	Not Matching any single character in a list. The ^ shall be the first character inside on the []. (i.e. 'ABC[^13-5]' would NOT match 'ABC1', 'ABC3', 'ABC4', and 'ABC5'. xyz[^dgh] would NOT match 'xyzd', 'xyzg', 'xyzh'.)

Table 122 defines the conversion rules for the operand values. The types are automatically converted if an implicit conversion exists (I). If an explicit conversion exists (E) then type can be converted with the cast operator. If no conversion exists (X) the then types cannot be converted, however, some servers may support application specific explicit conversions. The types used in the table are defined in IEC 62541-3. A data type that is not in the table does not have any defined conversions.

Table 122 – Conversion rules

Source type (from)	Target type (to)																				
	Boolean	Byte	ByteString	DateTime	Double	ExpandedNodeId	Float	Guid	Int16	Int32	Int64	NodeId	SByte	StatusCode	String	LocalizedText	QualifiedName	UInt16	UInt32	UInt64	XmlElement
Boolean	-	I	X	X	I	X	I	X	I	I	I	X	I	X	E	X	X	I	I	I	X
Byte	E	-	X	X	I	X	I	X	I	I	I	X	I	X	E	X	X	I	I	I	X
ByteString	X	X	-	X	X	X	X	E	X	X	X	X	X	X	X	X	X	X	X	X	X
DateTime	X	X	X	-	X	X	X	X	X	X	X	X	X	X	E	X	X	X	X	X	X
Double	E	E	X	X	-	X	E	X	E	E	E	X	E	X	E	X	X	E	E	E	X
ExpandedNodeId	X	X	X	X	X	-	X	X	X	X	X	E	X	X	I	X	X	X	X	X	X
Float	E	E	X	X	I	X	-	X	E	E	E	X	E	X	E	X	X	E	E	E	X
Guid	X	X	E	X	X	X	X	-	X	X	X	X	X	X	E	X	X	X	X	X	X
Int16	E	E	X	X	I	X	I	X	-	I	I	X	E	X	E	X	X	E	I	I	X
Int32	E	E	X	X	I	X	I	X	E	-	I	X	E	E	E	X	X	E	E	I	X
Int64	E	E	X	X	I	X	I	X	E	E	E	X	E	E	E	X	X	E	E	E	X
NodeId	X	X	X	X	X	I	X	X	X	X	X	-	X	X	I	X	X	X	X	X	X
SByte	E	E	X	X	I	X	I	X	I	I	I	X	-	X	E	X	X	I	I	I	X
StatusCode	X	X	X	X	X	X	X	X	X	I	I	X	X	-	X	X	X	E	I	I	X
String	I	I	X	E	I	E	I	I	I	I	I	E	I	X	-	E	E	I	I	I	X
LocalizedText	X	X	X	X	X	X	X	X	X	X	X	X	X	X	I	-	X	X	X	X	X
QualifiedName	X	X	X	X	X	X	X	X	X	X	X	X	X	X	I	I	-	X	X	X	X
UInt16	E	E	X	X	I	X	I	X	I	I	I	X	E	I	E	X	X	-	I	I	X
UInt32	E	E	X	X	I	X	I	X	E	I	I	X	E	E	E	X	X	E	-	I	X
UInt64	E	E	X	X	I	X	I	X	E	E	I	X	E	E	E	X	X	E	E	-	X
XmlElement	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	-

Arrays of a source type can be converted to arrays of the target type by converting each element. A conversion error for any element causes the entire conversion to fail.

Arrays of length 1 can be implicitly converted to a scalar value of the same type.

Guid, *NodeId* and *ExpandedNodeId* are converted to and from *String* using the syntax defined in IEC 62541-6.

Floating point values are rounded by adding 0,5 and truncating when they are converted to integer values.

Converting a negative value to an unsigned type causes a conversion error. If the conversion fails the result is a null value.

Converting a value that is outside the range of the target type causes a conversion error. If the conversion fails the result is a null value.

ByteString is converted to *String* by formatting the bytes as a sequence of hexadecimal digits.

LocalizedText values are converted to *Strings* by dropping the *Locale*. *Strings* are converted to *LocalizedText* values by setting the *Locale* to "".

QualifiedName values are converted to *Strings* by dropping the *NamespaceIndex*. *Strings* are converted to *QualifiedName* values by setting the *NamespaceIndex* to 0.

A *StatusCode* can be converted to and from a *UInt32* and *Int32* by copying the bits. Only the top 16-bits of the *StatusCode* are copied when it is converted to and from a *UInt16* or *Int16* value.

Boolean values are converted to '1' when true and '0' when false. Non zero numeric values are converted to true *Boolean* values. Numeric values of 0 are converted to false *Boolean* values. *String* values containing "true", "false", "1" or "0" can be converted to *Boolean* values. Other string values cause a conversion error. In this case *Strings* are case-insensitive.

It is sometimes possible to use implicit casts when operands with different data types are used in an operation. In this situation the precedence rules defined in Table 123 are used to determine which implicit conversion to use. The first data type in the list (top down) has the most precedence. If a data type is not in this table then it cannot be converted implicitly while evaluating an operation.

For example, assume that A = 1,1 (*Float*) and B = 1 (*Int32*) and that these values are used with an *Equals* operator. This operation would be evaluated by casting the *Int32* value to a *Float* since the *Float* data type has more precedence.

Table 123 – Data Precedence rules

Rank	Data Type
1	Double
2	Float
3	Int64
4	UInt64
5	Int32
6	UInt32
7	StatusCode
8	Int16
9	UInt16
10	SByte
11	Byte
12	Boolean
13	Guid
14	String
15	ExpandedNodeId
16	NodeId
17	LocalizedText
18	QualifiedName

Operands may contain null values (i.e. values which do not exist). When this happens, the element always evaluates to NULL (unless the *IsNull_1* operator has been specified). Table 124 defines how to combine elements that evaluate to NULL with other elements in a logical AND operation.

Table 124 – Logical AND Truth table

	TRUE	FALSE	NULL
TRUE	TRUE	FALSE	NULL
FALSE	FALSE	FALSE	FALSE
NULL	NULL	FALSE	NULL

Table 125 defines how to combine elements that evaluate to NULL with other elements in a logical OR operation.

Table 125 – Logical OR Truth table

	TRUE	FALSE	NULL
TRUE	TRUE	TRUE	TRUE
FALSE	TRUE	FALSE	NULL
NULL	TRUE	NULL	NULL

The NOT operator always evaluates to NULL if applied to a NULL operand.

A *ContentFilter* which evaluates to NULL after all elements are evaluated is evaluated as false.

7.4.4 FilterOperand parameters

7.4.4.1 Overview

The *ContentFilter* structure specified in 7.4 defines a collection of elements that makes up filter criteria and contains different types of *FilterOperands*. The *FilterOperand* parameter is an extensible parameter. This parameter is defined in Table 126. The *ExtensibleParameter* type is defined in 7.12.

Table 126 – FilterOperand parameter Typelds

Symbolic Id	Description
Element	Specifies an index into the array of elements. This type is used to build a logic tree of sub-elements by linking the operand of one element to a sub-element.
Literal	Specifies a literal value.
Attribute	Specifies any <i>Attribute</i> of an <i>Object</i> or <i>Variable Node</i> using a <i>Node</i> in the type system and relative path constructed from <i>ReferenceTypes</i> and <i>BrowseNames</i> .
SimpleAttribute	Specifies any <i>Attribute</i> of an <i>Object</i> or <i>Variable Node</i> using a <i>TypeDefinition</i> and a relative path constructed from <i>BrowseNames</i> .

7.4.4.2 ElementOperand

The *ElementOperand* provides the linking to sub-elements within a *ContentFilter*. The link is in the form of an integer that is used to index into the array of elements contained in the *ContentFilter*. An index is considered valid if its value is greater than the element index it is part of and it does not *Reference* a non-existent element. *Clients* shall construct filters in this way to avoid circular and invalid *References*. *Servers* should protect against invalid indexes by verifying the index prior to using it.

Table 127 defines the *ElementOperand* type.

Table 127 – ElementOperand

Name	Type	Description
ElementOperand	structure	ElementOperand value.
index	UInt32	Index into the element array.

7.4.4.3 LiteralOperand

Table 128 defines the *LiteralOperand* type.

Table 128 – LiteralOperand

Name	Type	Description
LiteralOperand	structure	LiteralOperand value.
value	BaseDataType	A literal value.

7.4.4.4 AttributeOperand

Table 129 defines the *AttributeOperand* type.

Table 129 – AttributeOperand

Name	Type	Description
AttributeOperand	structure	Attribute of a <i>Node</i> in the <i>AddressSpace</i> .
nodeId	NodeId	<i>NodeId</i> of a <i>Node</i> from the type system.
alias	String	An optional parameter used to identify or refer to an alias. An alias is a symbolic name that can be used to alias this operand and use it in other locations in the filter structure.
browsePath	RelativePath	Browse path relative to the <i>Node</i> identified by the <i>nodeId</i> parameter. See 7.26 for the definition of <i>RelativePath</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>AttributeId</i> . The <i>IntegerId</i> is defined in 7.14. The <i>IntegerIds</i> for the <i>Attributes</i> are defined in IEC 62541-6.
indexRange	NumericRange	This parameter is used to identify a single element of an array or a single range of indexes for an array. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is not used if the specified <i>Attribute</i> is not an array. However, if the specified <i>Attribute</i> is an array and this parameter is not used, then all elements are to be included in the range. The parameter is null if not used.

7.4.4.5 SimpleAttributeOperand

The *SimpleAttributeOperand* is a simplified form of the *AttributeOperand* and all of the rules that apply to the *AttributeOperand* also apply to the *SimpleAttributeOperand*. The examples provided in B.1 only use *AttributeOperand*, however, the *AttributeOperand* can be replaced by a *SimpleAttributeOperand* whenever all *ReferenceTypes* in the *RelativePath* are subtypes of *HierarchicalReferences* and the targets are *Object* or *Variable Nodes* and an *Alias* is not required.

Table 130 defines the *SimpleAttributeOperand* type.

Table 130 – SimpleAttributeOperand

Name	Type	Description
SimpleAttributeOperand	structure	Attribute of a <i>Node</i> in the <i>AddressSpace</i> .
typeId	NodeId	<i>NodeId</i> of a <i>TypeDefinitionNode</i> . This parameter restricts the operand to instances of the <i>TypeDefinitionNode</i> or one of its subtypes.
browsePath []	QualifiedName	A relative path to a <i>Node</i> . This parameter specifies a relative path using a list of <i>BrowseNames</i> instead of the <i>RelativePath</i> structure used in the <i>AttributeOperand</i> . The list of <i>BrowseNames</i> is equivalent to a <i>RelativePath</i> that specifies forward references which are subtypes of the <i>HierarchicalReferences ReferenceType</i> . All <i>Nodes</i> followed by the <i>browsePath</i> shall be of the <i>NodeClass Object</i> or <i>Variable</i> . If this list is empty the <i>Node</i> is the instance of the <i>TypeDefinition</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . The <i>IntegerId</i> is defined in 7.14. The <i>Value Attribute</i> shall be supported by all <i>Servers</i> . The support of other <i>Attributes</i> depends on requirements set in Profiles or other parts of this specification.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for an array. The first element is identified by index 0 (zero). This parameter is ignored if the selected <i>Node</i> is not a <i>Variable</i> or the <i>Value</i> of a <i>Variable</i> is not an array. The parameter is null if not specified. All values in the array are used if this parameter is not specified. The <i>NumericRange</i> type is defined in 7.22.

7.5 Counter

This primitive data type is a UInt32 that represents the value of a counter. The initial value of a counter is specified by its use. Modulus arithmetic is used for all calculations, where the modulus is max value + 1. Therefore,

$$x + y = (x + y) \bmod (\text{max value} + 1)$$

For example:

$$\text{max value} + 1 = 0$$

$$\text{max value} + 2 = 1$$

7.6 ContinuationPoint

A *ContinuationPoint* is used to pause a *Browse*, *QueryFirst* or *HistoryRead* operation and allow it to be restarted later by calling *BrowseNext*, *QueryNext* or *HistoryRead*. Operations are paused when the number of results found exceeds the limits set by either the *Client* or the *Server*.

The *Client* specifies the maximum number of results per operation in the request message. A *Server* shall not return more than this number of results but it may return fewer results. The *Server* allocates a *ContinuationPoint* if there are more results to return. ~~The Server shall always return at least one result if it returns a ContinuationPoint.~~

Servers shall support at least one *ContinuationPoint* per *Session*. *Servers* specify a maximum number of *ContinuationPoints* per *Session* in the *ServerCapabilities Object* defined in IEC 62541-5. *ContinuationPoints* remain active until the *Client* retrieves the remaining results, the *Client* releases the *ContinuationPoint* or the *Session* is closed. A *Server* shall

automatically free *ContinuationPoints* from prior requests from a *Session* if they are needed to process a new request from this *Session*. The *Server* returns a *Bad_ContinuationPointInvalid* error if a *Client* tries to use a *ContinuationPoint* that has been released. A *Client* can avoid this situation by completing paused operations before starting new operations.

Requests will often specify multiple operations that may or may not require a *ContinuationPoint*. A *Server* shall process the operations until it ~~runs out of~~ *ContinuationPoints* uses the maximum number of continuation points in this response. Once that happens the *Server* shall return a *Bad_NoContinuationPoints* error for any remaining operations. A *Client* can avoid this situation by sending requests with a number of operations that do not exceed the maximum number of *ContinuationPoints* per *Session* defined for the *Service* in the *ServerCapabilities Object* defined in IEC 62541-5.

A *Client* restarts an operation by passing the *ContinuationPoint* back to the *Server*. *Server* should always be able to reuse the *ContinuationPoint* provided so *Servers* shall never return *Bad_NoContinuationPoints* error when continuing a previously halted operation.

A *ContinuationPoint* is a subtype of the *ByteString* data type.

7.7 DataValue

7.7.1 General

The components of this parameter are defined in Table 131.

Table 131 – DataValue

Name	Type	Description
DataValue	structure	The value and associated information.
value	BaseDataType	The data value. If the <i>StatusCode</i> indicates an error then the value is to be ignored and the <i>Server</i> shall set it to null.
statusCode	StatusCode	The <i>StatusCode</i> that defines with the <i>Server's</i> ability to access/provide the value. The <i>StatusCode</i> type is defined in 7.34
sourceTimestamp	UtcTime	The source timestamp for the value.
sourcePicoSeconds	UInteger	Specifies the number of 10 picoseconds (1,0 e-11 seconds) intervals which shall be added to the sourceTimestamp.
serverTimestamp	UtcTime	The <i>Server</i> timestamp for the value.
serverPicoSeconds	UInteger	Specifies the number of 10 picoseconds (1,0 e-11 seconds) intervals which shall be added to the serverTimestamp.

7.7.2 PicoSeconds

Some applications require high resolution timestamps. The *PicoSeconds* fields allow applications to specify timestamps with a resolution of 10 picoseconds. The actual size of the *PicoSeconds* field depends on the resolution of the *UtcTime DataType*. For example, if the *UtcTime DataType* has a resolution of 100 nanoseconds then the *PicoSeconds* field would ~~have~~ need to store values up to 10 000 in order to provide the resolution of 10 picoseconds. The resolution of the *UtcTime DataType* depends on the *Mappings* defined in IEC 62541-6.

7.7.3 SourceTimestamp

The *sourceTimestamp* is used to reflect the timestamp that was applied to a *Variable* value by the data source. Once a value has been assigned a source timestamp, the source timestamp for that value instance never changes. In this context, "value instance" refers to the value received, independent of its actual value.

The *sourceTimestamp* shall be UTC time and should indicate the time of the last change of the *value* or *statusCode*.

The *sourceTimestamp* should be generated as close as possible to the source of the value but the timestamp needs to be set always by the same physical clock. In the case of redundant sources, the clocks of the sources should be synchronised.

If the OPC UA *Server* receives the *Variable* value from another OPC UA *Server*, then the OPC UA *Server* shall always pass the source timestamp without changes. If the source that applies the timestamp is not available, the source timestamp is set to null. For example, if a value could not be read because of some error during processing like invalid arguments passed in the request then the *sourceTimestamp* shall be null.

In the case of a bad or uncertain status *sourceTimestamp* is used to reflect the time that the source recognized the non-good status or the time the *Server* last tried to recover from the bad or uncertain status.

The *sourceTimestamp* is only returned with a *Value Attribute*. For all other *Attributes* the returned *sourceTimestamp* is set to null.

7.7.4 ServerTimestamp

The *serverTimestamp* is used to reflect the time that the *Server* received a *Variable* value or knew it to be accurate.

In the case of a bad or uncertain status, *serverTimestamp* is used to reflect the time that the *Server* received the status or that the *Server* last tried to recover from the bad or uncertain status.

In the case where the OPC UA *Server* subscribes to a value from another OPC UA *Server*, each *Server* applies its own *serverTimestamp*. This is in contrast to the *sourceTimestamp* in which only the originator of the data is allowed to apply the *sourceTimestamp*.

If the *Server* subscribes to the value from another *Server* every ten seconds and the value changes, then the *serverTimestamp* is updated each time a new value is received. If the value does not change, then new values will not be received on the *Subscription*. However, in the absence of errors, the receiving *Server* applies a new *serverTimestamp* every ten seconds because not receiving a value means that the value has not changed. Thus, the *serverTimestamp* reflects the time at which the *Server* knew the value to be accurate.

This concept also applies to OPC UA *Servers* that receive values from exception-based data sources. For example, suppose that a *Server* is receiving values from an exception-based device, and that

- a) the device is checking values every 0,5 seconds,
- b) the connection to the device is good,
- c) the device sent an update 3 minutes ago with a value of 1,234.

In this case, the *Server* value would be 1,234 and the *serverTimestamp* would be updated every 0,5 seconds after the receipt of the value.

7.7.5 StatusCode assigned to a value

The *StatusCode* is used to indicate the conditions under which a *Variable* value was generated, and thereby can be used as an indicator of the usability of the value. The *StatusCode* is defined in 7.34.

Overall condition (severity)

- A *StatusCode* with severity Good means that the value is of good quality.
- A *StatusCode* with severity Uncertain means that the quality of the value is uncertain for reasons indicated by the *SubCode*.
- A *StatusCode* with severity Bad means that the value is not usable for reasons indicated by the *SubCode*.

Rules

- The *StatusCode* indicates the usability of the value. Therefore, It is required that *Clients* minimally check the *StatusCode Severity* of all results, even if they do not check the other fields, before accessing and using the value.
- A *Server*, which does not support status information, shall return a severity code of Good. It is also acceptable for a *Server* to simply return a severity and a non-specific (0) ~~Substatus~~ *SubCode*.
- If the *Server* has no known value - in particular when *Severity* is BAD, it shall return a NULL value.

7.8 DiagnosticInfo

The components of this parameter are defined in Table 132.

Table 132 – DiagnosticInfo

Name	Type	Description
DiagnosticInfo	structure	Vendor-specific diagnostic information.
namespaceUri	Int32	The <i>symbolicId</i> is defined within the context of a namespace. This namespace is represented as a string and is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>namespaceIndex</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified. The <i>namespaceUri</i> shall not be the standard OPC UA namespace. There are no <i>symbolicIds</i> provided for standard <i>StatusCodes</i>.
symbolicId	Int32	The <i>symbolicId</i> shall be used to identify a vendor-specific error or condition; typically the result of some <i>Server</i> internal operation. The maximum length of this string is 32 characters. <i>Servers</i> wishing to return a numeric return code should convert the return code into a string and use this string as <i>symbolicId</i> (e.g., "0xC0040007" or "-4"). This symbolic identifier string is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>symbolicId</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified. The <i>symbolicId</i> shall not contain <i>StatusCodes</i>. If the <i>localizedText</i> contains a translation for the description of a <i>StatusCode</i>, the <i>symbolicId</i> is -1.
locale	Int32	The locale part of the vendor-specific localized text describing the symbolic id. This localized text string is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>localizedTextIndex</i> <i>locale</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified.
localizedText	Int32	A vendor-specific localized text string describes the symbolic id. The maximum length of this text string is 256 characters. This localized text string is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>localizedTextIndex</i> <i>localizedText</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified. The <i>localizedText</i> refers to the <i>symbolicId</i> if present or the string that describes the standard <i>StatusCode</i> if the <i>Server</i> provides translations. If the index is -1, the <i>Server</i> has no translation to return and the <i>Client</i> should use the invariant <i>StatusCode</i> description from the specification.
additionalInfo	String	Vendor-specific diagnostic information.
innerStatusCode	StatusCode	The <i>StatusCode</i> from the inner operation. Many applications will make calls into underlying systems during OPC UA request processing. An OPC UA <i>Server</i> has the option of reporting the status from the underlying system in the diagnostic info.
innerDiagnosticInfo	DiagnosticInfo	The diagnostic info associated with the inner <i>StatusCode</i> .

7.9 DiscoveryConfiguration parameters

7.9.1 Overview

The *DiscoveryConfiguration* structure used in the *RegisterServer2 Service* allows *Servers* to provide additional configuration parameters to *Discovery Servers* for registration. Table 133 defines the current set of discovery configuration options. The *ExtensibleParameter* type is defined in 7.12.

Table 133 – DiscoveryConfiguration parameterTypeIds

Symbolic Id	Description
MdnsDiscoveryConfiguration	Configuration parameters for mDNS discovery.

7.9.2 MdnsDiscoveryConfiguration

Table 134 defines the *MdnsDiscoveryConfiguration* parameter.

Table 134 – MdnsDiscoveryConfiguration

Name	Type	Description
MdnsDiscoveryConfiguration	structure	mDNS discovery configuration.
mdnsServerName	String	The name of the <i>Server</i> when it is announced via mDNS. See IEC 62541-12 for the details about mDNS. This string shall be less than 64 bytes. If not specified the first element of the <i>serverNames</i> array is used (truncated to 63 bytes if necessary).
serverCapabilities []	String	The set of <i>Server</i> capabilities supported by the <i>Server</i> . A <i>Server</i> capability is a short identifier for a feature. The set of allowed <i>Server</i> capabilities are defined in IEC 62541-12.

7.10 EndpointDescription

The components of this parameter are defined in Table 135.

Table 135 – EndpointDescription

Name	Type	Description
EndpointDescription	structure	Describes an <i>Endpoint</i> for a <i>Server</i> .
endpointUrl	String	The URL for the <i>Endpoint</i> described.
server	ApplicationDescription	The description for the <i>Server</i> that the <i>Endpoint</i> belongs to. The <i>ApplicationDescription</i> type is defined in 7.1.
serverCertificate	ApplicationInstance Certificate	The <i>Application Instance Certificate</i> issued to the <i>Server</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2.
securityMode	Enum MessageSecurityMode	The type of security to apply to the messages. The type <i>MessageSecurityMode</i> type is defined in 7.15. A <i>SecureChannel</i> may ~~have~~ need to be created even if the <i>securityMode</i> is NONE. The exact behaviour depends on the mapping used and is described in the IEC 62541-6.
securityPolicyUri	String	The URI for <i>SecurityPolicy</i> to use when securing messages. The set of known URIs and the <i>SecurityPolicies</i> associated with them are defined in IEC 62541-7.
userIdentityTokens []	UserTokenPolicy	The user identity tokens that the <i>Server</i> will accept. The <i>Client</i> shall pass one of the <i>UserIdentityTokens</i> in the <i>ActivateSession</i> request. The <i>UserTokenPolicy</i> type is described in 7.37.
transportProfileUri	String	The URI of the <i>Transport Profile</i> supported by the <i>Endpoint</i> . IEC 62541-7 defines URIs for the <i>Transport Profiles</i> .
securityLevel	Byte	A numeric value that indicates how secure the <i>EndpointDescription</i> is compared to other <i>EndpointDescriptions</i> for the same <i>Server</i> . A value of 0 indicates that the <i>EndpointDescription</i> is not recommended and is only supported for backward compatibility. A higher value indicates better security.

7.11 ExpandedNodeId

The components of this parameter are defined in Table 136. *ExpandedNodeId* allows the namespace to be specified explicitly as a string or with an index in the *Server*'s namespace table.

Table 136 – ExpandedNodeId

Name	Type	Description
ExpandedNodeId	structure	The <i>NodeId</i> with the namespace expanded to its string representation.
serverIndex	Index	Index that identifies the <i>Server</i> that contains the <i>TargetNode</i> . This <i>Server</i> may be the local <i>Server</i> or a remote <i>Server</i> . This index is the index of that <i>Server</i> in the local <i>Server</i> 's <i>Server</i> table. The index of the local <i>Server</i> in the <i>Server</i> table is always 0. All remote <i>Servers</i> have indexes greater than 0. The <i>Server</i> table is contained in the <i>Server Object</i> in the <i>AddressSpace</i> (see IEC 62541-3 and IEC 62541-5). The <i>Client</i> may read the <i>Server</i> table <i>Variable</i> to access the description of the target <i>Server</i> .
namespaceUri	String	The URI of the namespace. If this parameter is specified then the namespace index is ignored. 5.4 and IEC 62541-12 ⁴ describe discovery mechanism that can be used to resolve URIs into URLs.
namespaceIndex	Index	The index in the <i>Server</i> 's namespace table. This parameter shall be 0 and is ignored in the <i>Server</i> if the namespace URI is specified.
identifierType	IdType	Type of the identifier element of the <i>NodeId</i> .
identifier	*	The identifier for a <i>Node</i> in the <i>AddressSpace</i> of an OPC UA <i>Server</i> (see <i>NodeId</i> definition in IEC 62541-3).

7.12 ExtensibleParameter

The extensible parameter types can only be extended by additional parts of IEC 62541.

The *ExtensibleParameter* defines a data structure with two elements. The *parameterTypeId* specifies the data type encoding of the second element. Therefore the second element is specified as "--". The *ExtensibleParameter* base type is defined in Table 137.

Concrete extensible parameters that are common to OPC UA are defined in Clause 7. Additional parts of IEC 62541 can define additional extensible parameter types.

Table 137 – ExtensibleParameter base type

Name	Type	Description
ExtensibleParameter	structure	Specifies the details of an extensible parameter type.
parameterTypeId	NodeId	Identifies the data type of the parameter that follows.
parameterData	--	The details for the extensible parameter type.

7.13 Index

This primitive data type is a UInt32 that identifies an element of an array.

⁴—Under consideration.

7.14 IntegerId

This primitive data type is a UInt32 that is used as an identifier, such as a handle. All values, except for 0, are valid.

7.15 MessageSecurityMode

The *MessageSecurityMode* is an enumeration that specifies what security should be applied to messages exchanges during a Session. The possible values are described in Table 138.

Table 138 – MessageSecurityMode values

Value	Description
INVALID_0	The MessageSecurityMode is invalid. This value is the default value to avoid an accidental choice of no security is applied. This choice will always be rejected.
NONE_1	No security is applied.
SIGN_2	All messages are signed but not encrypted.
SIGNANDENCRYPT_3	All messages are signed and encrypted.

7.16 MonitoringParameters

The components of this parameter are defined in Table 139.

Table 139 – MonitoringParameters

Name	Type	Description
MonitoringParameters	structure	Parameters that define the monitoring characteristics of a <i>MonitoredItem</i> .
clientHandle	IntegerId	<i>Client</i> -supplied id of the <i>MonitoredItem</i> . This id is used in <i>Notifications</i> generated for the list <i>Node</i> . The <i>IntegerId</i> type is defined in 7.14.
samplingInterval	Duration	The interval that defines the fastest rate at which the <i>MonitoredItem</i> (s) should be accessed and evaluated. This interval is defined in milliseconds. The value 0 indicates that the <i>Server</i> should use the fastest practical rate. The value -1 indicates that the default sampling interval defined by the publishing interval of the <i>Subscription</i> is requested. A different sampling interval is used if the publishing interval is not a supported sampling interval. Any negative number is interpreted as -1. The sampling interval is not changed if the publishing interval is changed by a subsequent call to the <i>ModifySubscription</i> Service. The <i>Server</i> uses this parameter to assign the <i>MonitoredItems</i> to a sampling interval that it supports. The assigned interval is provided in the <i>revisedSamplingInterval</i> parameter. The <i>Server</i> shall always return a <i>revisedSamplingInterval</i> that is equal or higher than the requested <i>samplingInterval</i> . If the requested <i>samplingInterval</i> is higher than the maximum sampling interval supported by the <i>Server</i> , the maximum sampling interval is returned.
filter	Extensible Parameter MonitoringFilter	A filter used by the <i>Server</i> to determine if the <i>MonitoredItem</i> should generate a <i>Notification</i> . If not used, this parameter is null. The <i>MonitoringFilter</i> parameter type is an extensible parameter type specified in 7.17. It specifies the types of filters that can be used.

Name	Type	Description												
queueSize	Counter	The requested size of the <i>MonitoredItem</i> queue. The following values have special meaning for data monitored items: <table><tr><th>Value</th><th>Meaning</th></tr><tr><td>0 or 1</td><td>the <i>Server</i> returns the default queue size which shall be 1 as <i>revisedQueueSize</i> for data monitored items. The queue has a single entry, effectively disabling queuing.</td></tr></table> For values larger than one a first-in-first-out queue is to be used. The <i>Server</i> may limit the size in <i>revisedQueueSize</i>. In the case of a queue overflow, the <i>Overflow</i> bit (flag) in the <i>InfoBits</i> portion of the <i>DataValue statusCode</i> is set in the new value. The following values have special meaning for event monitored items: <table><tr><th>Value</th><th>Meaning</th></tr><tr><td>0</td><td>the <i>Server</i> returns the default queue size for <i>Event Notifications</i> as <i>revisedQueueSize</i> for event monitored items.</td></tr><tr><td>1</td><td>the <i>Server</i> returns the minimum queue size the <i>Server</i> requires for <i>Event Notifications</i> as <i>revisedQueueSize</i>.</td></tr><tr><td>MaxUInt32</td><td>the <i>Server</i> returns the maximum queue size that the <i>Server</i> can support for <i>Event Notifications</i> as <i>revisedQueueSize</i>.</td></tr></table> If a <i>Client</i> chooses a value between the minimum and maximum settings of the <i>Server</i> the value shall be returned in the <i>revisedQueueSize</i>. If the requested <i>queueSize</i> is outside the minimum or maximum, the <i>Server</i> shall return the corresponding bounding value. In the case of a queue overflow, an <i>Event</i> of the type <i>EventQueueOverflowEventType</i> is generated.	Value	Meaning	0 or 1	the <i>Server</i> returns the default queue size which shall be 1 as <i>revisedQueueSize</i> for data monitored items. The queue has a single entry, effectively disabling queuing.	Value	Meaning	0	the <i>Server</i> returns the default queue size for <i>Event Notifications</i> as <i>revisedQueueSize</i> for event monitored items.	1	the <i>Server</i> returns the minimum queue size the <i>Server</i> requires for <i>Event Notifications</i> as <i>revisedQueueSize</i> .	MaxUInt32	the <i>Server</i> returns the maximum queue size that the <i>Server</i> can support for <i>Event Notifications</i> as <i>revisedQueueSize</i> .
Value	Meaning													
0 or 1	the <i>Server</i> returns the default queue size which shall be 1 as <i>revisedQueueSize</i> for data monitored items. The queue has a single entry, effectively disabling queuing.													
Value	Meaning													
0	the <i>Server</i> returns the default queue size for <i>Event Notifications</i> as <i>revisedQueueSize</i> for event monitored items.													
1	the <i>Server</i> returns the minimum queue size the <i>Server</i> requires for <i>Event Notifications</i> as <i>revisedQueueSize</i> .													
MaxUInt32	the <i>Server</i> returns the maximum queue size that the <i>Server</i> can support for <i>Event Notifications</i> as <i>revisedQueueSize</i> .													
discardOldest	Boolean	A boolean parameter that Specifies the discard policy when the queue is full and a new <i>Notification</i> is to be queued. It has the following values: <table><tr><td>TRUE</td><td>the oldest (first) <i>Notification</i> in the queue is discarded. The new <i>Notification</i> is added to the end of the queue.</td></tr><tr><td>FALSE</td><td>the last <i>Notification</i> added to the queue gets replaced with the new <i>Notification</i>.</td></tr></table>	TRUE	the oldest (first) <i>Notification</i> in the queue is discarded. The new <i>Notification</i> is added to the end of the queue.	FALSE	the last <i>Notification</i> added to the queue gets replaced with the new <i>Notification</i> .								
TRUE	the oldest (first) <i>Notification</i> in the queue is discarded. The new <i>Notification</i> is added to the end of the queue.													
FALSE	the last <i>Notification</i> added to the queue gets replaced with the new <i>Notification</i> .													

7.17 MonitoringFilter parameters

7.17.1 Overview

The *CreateMonitoredItem Service* allows specifying a filter for each *MonitoredItem*. The *MonitoringFilter* is an extensible parameter whose structure depends on the type of item being monitored. The *parameterTypeIds* are defined in Table 140. Other types can be defined by additional parts of this multi-part specification or other specifications based on OPC UA. The *ExtensibleParameter* type is defined in 7.12.

Each *MonitoringFilter* may have an associated *MonitoringFilterResult* structure which returns revised parameters and/or error information to clients in the response. The result structures, when they exist, are described in the section that defines the *MonitoringFilter*.

Table 140 – MonitoringFilter parameterTypeIds

Symbolic Id	Description
DataChangeFilter	The change in a data value that shall cause a <i>Notification</i> to be generated.
EventFilter	If a <i>Notification</i> conforms to the <i>EventFilter</i> , the <i>Notification</i> is sent to the <i>Client</i> .
AggregateFilter	The <i>Aggregate</i> and its intervals when it will be calculated and a <i>Notification</i> is generated.

7.17.2 DataChangeFilter

The *DataChangeFilter* defines the conditions under which a *DataChange Notification* should be reported and, optionally, a range or band for value changes where no *DataChange Notification* is generated. This range is called *Deadband*. The *DataChangeFilter* is defined in Table 141.

Table 141 – DataChangeFilter

Name	Type	Description								
DataChangeFilter	structure									
trigger	Enum DataChangeTrigger	Specifies the conditions under which a data change notification should be reported. It has the following values: STATUS_0 Report a notification ONLY if the <i>StatusCode</i> associated with the value changes. See Table 178 for <i>StatusCodes</i> defined in this document. IEC 62541-8 specifies additional <i>StatusCodes</i> that are valid in particular for device data. STATUS_VALUE_1 Report a notification if either the <i>StatusCode</i> or the value change. The <i>Deadband</i> filter can be used in addition for filtering value changes. For floating point values a <i>Server</i> shall check for NaN and only report a single notification with NaN when the value enters the NaN state. This is the default setting if no filter is set. STATUS_VALUE_TIMESTAMP_2 Report a notification if either <i>StatusCode</i>, value or the <i>SourceTimestamp</i> change. If a <i>Deadband</i> filter is specified, this trigger has the same behaviour as STATUS_VALUE_1. If the <i>DataChangeFilter</i> is not applied to the monitored item, STATUS_VALUE_1 is the default reporting behaviour.								
deadbandType	UInt32	A value that defines the <i>Deadband</i> type and behaviour. <table><tr><td><u>Value</u></td><td><u>deadbandType</u></td></tr><tr><td>None_0</td><td>No <i>Deadband</i> calculation should be applied.</td></tr><tr><td>Absolute_1</td><td>AbsoluteDeadband (see below)</td></tr><tr><td>Percent_2</td><td>PercentDeadband (This type is specified in IEC 62541-8).</td></tr></table>	<u>Value</u>	<u>deadbandType</u>	None_0	No <i>Deadband</i> calculation should be applied.	Absolute_1	AbsoluteDeadband (see below)	Percent_2	PercentDeadband (This type is specified in IEC 62541-8).
<u>Value</u>	<u>deadbandType</u>									
None_0	No <i>Deadband</i> calculation should be applied.									
Absolute_1	AbsoluteDeadband (see below)									
Percent_2	PercentDeadband (This type is specified in IEC 62541-8).									
deadbandValue	Double	The <i>Deadband</i> is applied only if * the <i>trigger</i> includes value changes and * the <i>deadbandType</i> is set appropriately. Deadband is ignored if the status of the data item changes. DeadbandType = AbsoluteDeadband: For this type the <i>deadbandValue</i> contains the absolute change in a data value that shall cause a <i>Notification</i> to be generated. This parameter applies only to <i>Variables</i> with any <i>Number</i> data type. An exception that causes a <i>DataChange Notification</i> based on an AbsoluteDeadband is determined as follows: Generate a Notification if (absolute value of (last cached value - current value) > AbsoluteDeadband) The last cached value is defined as the last value pushed to the queue. If the item is an array of values, the entire array is returned if any array element exceeds the AbsoluteDeadband, or the size or dimension of the array changes. DeadbandType = PercentDeadband: This type is specified in IEC 62541-8								

The *DataChangeFilter* does not have an associated result structure.

7.17.3 EventFilter

The *EventFilter* provides for the filtering and content selection of *Event Subscriptions*.

If an *Event Notification* conforms to the filter defined by the *where* parameter of the *EventFilter*, then the *Notification* is sent to the *Client*.

Each *Event Notification* shall include the fields defined by the *selectClauses* parameter of the *EventFilter*. The defined *EventTypes* are specified in IEC 62541-5.

The *selectClauses* and *whereClause* parameters are specified with the *SimpleAttributeOperand* structure (see 7.4.4.5). This structure requires the *NodeId* of an *EventType* supported by the *Server* and a path to an *InstanceDeclaration*. An *InstanceDeclaration* is a *Node* which can be found by following forward hierarchical references from the fully inherited *EventType* where the *Node* is also the source of a *HasModellingRule* reference. *EventTypes*, *InstanceDeclarations* and *Modelling Rules* are described completely in IEC 62541-3.

In some cases the same *BrowsePath* will apply to multiple *EventTypes*. If the *Client* specifies the *BaseEventType* in the *SimpleAttributeOperand* then the *Server* shall evaluate the *BrowsePath* without considering the *Type*.

Each *InstanceDeclaration* in the path shall be *Object* or *Variable Node*. The final *Node* in the path may be an *Object Node*; however, *Object Nodes* are only available for *Events* which are visible in the *Server's AddressSpace*.

The *SimpleAttributeOperand* structure allows the *Client* to specify any *Attribute*; however, the *Server* is only required to support the *Value Attribute* for *Variable Nodes* and the *NodeId Attribute* for *Object Nodes*. That said, profiles defined in IEC 62541-7 may make support for additional *Attributes* mandatory.

The *SimpleAttributeOperand* structure is used in the *selectClauses* to select the value to return if an *Event* meets the criteria specified by the *whereClause*. A null value is returned in the corresponding event field in the Publish response if the selected *field* is not part of the *Event* or an error was returned in the *selectClauseResults* of the *EventFilterResult*. If the selected *field* is supported but not available at the time of the event notification, the event field shall contain a *StatusCode* that indicates the reason for the unavailability. For example, the *Server* shall set the event field to *Bad_UserAccessDenied* if the value is not accessible to the user associated with the *Session*. If a *Value Attribute* has an uncertain or bad *StatusCode* associated with it then the *Server* shall provide the *StatusCode* instead of the *Value Attribute*. The *Server* shall set the event field to *Bad_EncodingLimitsExceeded* if a value exceeds the *maxResponseMessageSize*. The *EventId*, *EventType* and *ReceiveTime* cannot contain a *StatusCode* or a null value.

The *Server* shall validate the *selectClauses* when a *Client* creates or updates the *EventFilter*. Any errors which are true for all possible *Events* are returned in the *selectClauseResults* parameter described in Table 143. Some *Servers*, like aggregating *Servers*, may not know all possible *EventTypes* at the time the *EventFilter* is set. These *Servers* do not return errors for unknown *EventTypes* or *BrowsePaths*. The *Server* shall not report errors that might occur depending on the state or the *Server* or type of *Event*. For example, a *selectClauses* that requests a single element in an array would always produce an error if the *DataType* of the *Attribute* is a scalar. However, even if the *DataType* is an array an error could occur if the requested index does not exist for a particular *Event*, the *Server* would not report an error in the *selectClauseResults* parameter if the latter situation existed.

The *SimpleAttributeOperand* is used in the *whereClause* to select a value which forms part of a logical expression. These logical expressions are then used to determine whether a particular *Event* should be reported to the *Client*. The *Server* shall use a null value if any error occurs when a *whereClause* is evaluated for a particular *Event*. If a *Value Attribute* has an

uncertain or bad *StatusCode* associated with it, then the *Server* shall use a null value instead of the *Value*.

Any basic *FilterOperator* in Table 119 may be used in the *whereClause*, however, only the *OfType_14 FilterOperator* from Table 120 is permitted.

The *Server* shall validate the *whereClause* when a *Client* creates or updates the *EventFilter*. Any structural errors in the construction of the filter and any errors which are true for all possible *Events* are returned in the *whereClauseResult* parameter described in Table 143. Errors that could occur depending on the state of the *Server* or the *Event* are not reported. Some *Servers*, like aggregating *Servers*, may not know all possible *EventTypes* at the time the *EventFilter* is set. These *Servers* do not return errors for unknown *EventTypes* or *BrowsePaths*.

~~SubscriptionControlEvents~~ *EventQueueOverflowEventType Events* are special *Events* which are used to provide control information to the *Client*. These *Events* are only published to the *MonitoredItems* in the *Subscription* that produced the ~~SubscriptionControlEvent~~ *EventQueueOverflowEventType Event*. These *Events* bypass the *whereClause*.

Table 142 defines the *EventFilter* structure.

Table 142 – EventFilter structure

Name	Type	Description
EventFilter	structure	
selectClauses []	SimpleAttribute Operand	List of the values to return with each <i>Event</i> in a <i>Notification</i> . At least one valid clause shall be specified. See 7.4.4.5 for the definition of <i>SimpleAttributeOperand</i> .
whereClause	ContentFilter	Limit the <i>Notifications</i> to those <i>Events</i> that match the criteria defined by this <i>ContentFilter</i> . The <i>ContentFilter</i> structure is described in 7.4. The <i>AttributeOperand</i> structure may not be used in an <i>EventFilter</i> .

Table 143 defines the *EventFilterResult* structure. This is the *MonitoringFilterResult* associated with the *EventFilter* *MonitoringFilter*.

Table 143 – EventFilterResult structure

Name	Type	Description
EventFilterResult	structure	
selectClauseResults []	StatusCode	List of status codes for the elements in the select clause. The size and order of the list matches the size and order of the elements in the <i>selectClauses</i> request parameter. The <i>Server</i> returns null for unavailable or rejected <i>Event</i> fields.
selectClauseDiagnosticInfos []	DiagnosticInfo	A list of diagnostic information for individual elements in the select clause. The size and order of the list matches the size and order of the elements in the <i>selectClauses</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the select clauses.
whereClauseResult	ContentFilter Result	Any results associated with the <i>whereClause</i> request parameter. The <i>ContentFilterResult</i> type is defined in 7.4.2.

Table 144 defines values for the *selectClauseResults* parameter. Common *StatusCodes* are defined in Table 178.

Table 144 – EventFilterResult result codes

Symbolic Id	Description
Bad_TypeDefinitionInvalid	See Table 178 for the description of this result code. The typeId is not the NodeId for BaseEventType or a subtype of it.
Bad_NodeIdUnknown	See Table 178 for the description of this result code. The browsePath is specified but it will never exist in any Event.
Bad_BrowseNameInvalid	See Table 178 for the description of this result code. The browsePath is specified and contains a null element.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code. The node specified by the browse path will never allow the given ~~attribute id~~ <i>AttributeId</i> to be returned.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_TypeMismatch	See Table 178 for the description of this result code. The indexRange is valid but the value of the Attribute is never an array.

7.17.4 AggregateFilter

The *AggregateFilter* defines the *Aggregate* function that should be used to calculate the values to be returned. See IEC 62541-13 for details on possible *Aggregate* functions. It specifies a *startTime* of the first *Aggregate* to be calculated. The *samplingInterval* of the ~~MonitoringAttributes~~ *MonitoringParameters* (see 7.16) defines how the *Server* should internally sample the underlying data source. The *processingInterval* specifies the size of a time-period where the *Aggregate* is calculated. The *queueSize* from the *MonitoringAttributes* specifies the number of processed values that should be kept.

The intention of the *AggregateFilter* is not to read historical data, the *HistoryRead* service should be used for this purpose. However, it is allowed that the *startTime* is set to a time that is in the past when received from the server. The number of *Aggregates* to be calculated in the past should not exceed the *queueSize* defined in the *MonitoringAttributes* since the values exceeding the *queueSize* would directly be discharged and never returned to the client.

The *startTime* and the *processingInterval* can be revised by the server, but the *startTime* should remain in the same boundary ($\text{startTime} + \text{revisedProcessingInterval} * n = \text{revisedStartTime}$). That behaviour simplifies accessing historical values of the *Aggregates* using the same boundaries by calling the *HistoryRead* service. The extensible Parameter *AggregateFilterResult* is used to return the revised values for the *AggregateFilter*.

Some underlying systems may poll data and produce multiple samples with the same value. Other systems may only report changes to the values. The definition for each *Aggregate* type explains how to handle the two different scenarios.

The *MonitoredItem* only reports values for intervals that have completed when the publish timer expires. Unused data is carried over and used to calculate a value returned in the next publish.

The *ServerTimestamp* for each interval shall be the time of the end of the processing interval.

The *AggregateFilter* is defined in Table 145.

Table 145 – AggregateFilter structure

Name	Type	Description
AggregateFilter	structure	
startTime	UtcTime	Beginning of period to calculate the <i>Aggregate</i> the first time. The size of each period used to calculate the <i>Aggregate</i> is defined by the <i>samplingInterval</i> of the <i>MonitoringAttributes</i> <i>MonitoringParameters</i> (see 7.16).
aggregateType	NodeId	The <i>NodeId</i> of the <i>AggregateFunctionType</i> <i>Object</i> that indicates the <i>Aggregate</i> to be used when retrieving processed data. See IEC 62541-13 for details.
processingInterval	Duration	The period be used to compute the <i>Aggregate</i> .
aggregateConfiguration	Aggregate Configuration	This parameter allows <i>Clients</i> to override the <i>Aggregate</i> configuration settings supplied by the <i>AggregateConfiguration</i> <i>Object</i> on a per monitored item basis. See IEC 62541-13 for more information on <i>Aggregate</i> configurations. If the <i>Server</i> does not support the ability to override the <i>Aggregate</i> configuration settings it shall return a <i>StatusCode</i> of <i>Bad_AggregateListMismatch</i> . This structure is defined in-line with the following indented items.
useServerCapabilities Defaults	Boolean	If value = TRUE use <i>Aggregate</i> configuration settings as outlined by the <i>AggregateConfiguration</i> object. If value=FALSE use configuration settings as outlined in the following <i>aggregateConfiguration</i> parameters. Default is TRUE.
treatUncertainAsBad	Boolean	As described in IEC 62541-13.
percentDataBad	Byte	As described in IEC 62541-13.
percentDataGood	Byte	As described in IEC 62541-13.
useSloped Extrapolation	Boolean	As described in IEC 62541-13.

The *AggregateFilterResult* defines the revised *AggregateFilter* the *Server* can return when an *AggregateFilter* is defined for a *MonitoredItem* in the *CreateMonitoredItems* or *ModifyMonitoredItems* *Services*. The *AggregateFilterResult* is defined in Table 146. This is the *MonitoringFilterResult* associated with the *AggregateFilter* *MonitoringFilter*.

Table 146 – AggregateFilterResult structure

Name	Type	Description
AggregateFilterResult	structure	
revisedStartTime	UtcTime	The actual <i>StartTime</i> interval that the <i>Server</i> shall use. This value is based on a number of factors, including capabilities of the <i>Server</i> to access historical data. The <i>revisedStartTime</i> should remain in the same boundary as the <i>startTime</i> ($startTime + samplingInterval * n = revisedStartTime$).
revisedProcessingInterval	Duration	The actual <i>processingInterval</i> that the <i>Server</i> shall use. The <i>revisedProcessingInterval</i> shall be at least twice the <i>revisedSamplingInterval</i> for the <i>MonitoredItem</i> .

7.18 MonitoringMode

The *MonitoringMode* is an enumeration that specifies whether sampling and reporting are enabled or disabled for a *MonitoredItem*. The value of the publishing enabled parameter for a *Subscription* does not affect the value of the monitoring mode for a *MonitoredItem* of the *Subscription*. The values of this parameter are defined in Table 147.

Table 147 – MonitoringMode values

Value	Description
DISABLED_0	The item being monitored is not sampled or evaluated, and <i>Notifications</i> are not generated or queued. <i>Notification</i> reporting is disabled.
SAMPLING_1	The item being monitored is sampled and evaluated, and <i>Notifications</i> are generated and queued. <i>Notification</i> reporting is disabled.
REPORTING_2	The item being monitored is sampled and evaluated, and <i>Notifications</i> are generated and queued. <i>Notification</i> reporting is enabled.

7.19 NodeAttributes parameters

7.19.1 Overview

The *AddNodes Service* allows specifying the *Attributes* for the *Nodes* to add. The *NodeAttributes* is an extensible parameter whose structure depends on the type of the *Attribute NodeClass* being added. It identifies the *NodeClass* that defines the structure of the *Attributes* that follow. The *parameterTypeIds* are defined in Table 148. The *ExtensibleParameter* type is defined in 7.12.

Table 148 – NodeAttributes parameterTypeIds

Symbolic Id	Description
ObjectAttributes	Defines the <i>Attributes</i> for the <i>Object NodeClass</i> .
VariableAttributes	Defines the <i>Attributes</i> for the <i>Variable NodeClass</i> .
MethodAttributes	Defines the <i>Attributes</i> for the <i>Method NodeClass</i> .
ObjectTypeAttributes	Defines the <i>Attributes</i> for the <i>ObjectType NodeClass</i> .
VariableTypeAttributes	Defines the <i>Attributes</i> for the <i>VariableType NodeClass</i> .
ReferenceTypeAttributes	Defines the <i>Attributes</i> for the <i>ReferenceType NodeClass</i> .
DataTypeAttributes	Defines the <i>Attributes</i> for the <i>DataType NodeClass</i> .
ViewAttributes	Defines the <i>Attributes</i> for the <i>View NodeClass</i> .
GenericAttributes	Defines an id and value list for passing in any number of <i>Attribute</i> values. It should be used instead of the <i>NodeClass</i> specific structures since it allows the handling of additional <i>Attributes</i> defined in future specification versions.

Table 149 defines the bit mask used in the *NodeAttributes* parameters to specify which *Attributes* are set by the *Client*.

Table 149 – Bit mask for specified Attributes

Field	Bit	Description
AccessLevel	0	Indicates if the AccessLevel Attribute is set.
ArrayDimensions	1	Indicates if the ArrayDimensions Attribute is set.
Reserved	2	Reserved to be consistent with WriteMask defined in IEC 62541-3.
ContainsNoLoops	3	Indicates if the ContainsNoLoops Attribute is set.
DataType	4	Indicates if the DataType Attribute is set.
Description	5	Indicates if the Description Attribute is set.
DisplayName	6	Indicates if the DisplayName Attribute is set.
EventNotifier	7	Indicates if the EventNotifier Attribute is set.
Executable	8	Indicates if the Executable Attribute is set.
Historizing	9	Indicates if the Historizing Attribute is set.
InverseName	10	Indicates if the InverseName Attribute is set.
IsAbstract	11	Indicates if the IsAbstract Attribute is set.
MinimumSamplingInterval	12	Indicates if the MinimumSamplingInterval Attribute is set.
Reserved	13	Reserved to be consistent with WriteMask defined in IEC 62541-3.
Reserved	14	Reserved to be consistent with WriteMask defined in IEC 62541-3.
Symmetric	15	Indicates if the Symmetric Attribute is set.
UserAccessLevel	16	Indicates if the UserAccessLevel Attribute is set.
UserExecutable	17	Indicates if the UserExecutable Attribute is set.
UserWriteMask	18	Indicates if the UserWriteMask Attribute is set.
ValueRank	19	Indicates if the ValueRank Attribute is set.
WriteMask	20	Indicates if the WriteMask Attribute is set.
Value	21	Indicates if the Value Attribute is set.
Reserved	22:32	Reserved for future use. Shall always be zero.

7.19.2 ObjectAttributes parameter

Table 150 defines the *ObjectAttributes* parameter.

Table 150 – ObjectAttributes

Name	Type	Description
ObjectAttributes	structure	Defines the <i>Attributes</i> for the <i>Object NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
eventNotifier	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.3 VariableAttributes parameter

Table 151 defines the *VariableAttributes* parameter.

Table 151 – VariableAttributes

Name	Type	Description
VariableAttributes	structure	Defines the <i>Attributes</i> for the <i>Variable NodeClass</i>
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
value	Defined by the <i>DataType Attribute</i>	See IEC 62541-3 for the description of this <i>Attribute</i> .
dataType	NodeId	See IEC 62541-3 for the description of this <i>Attribute</i> .
valueRank	Int32	See IEC 62541-3 for the description of this <i>Attribute</i> .
arrayDimensions	UInt32 []	See IEC 62541-3 for the description of this <i>Attribute</i> .
accessLevel	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
userAccessLevel	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
minimumSamplingInterval	Duration	See IEC 62541-3 for the description of this <i>Attribute</i> .
historizing	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.4 MethodAttributes parameter

Table 152 defines the *MethodAttributes* parameter.

Table 152 – MethodAttributes

Name	Type	Description
BaseAttributes	structure	Defines the <i>Attributes</i> for the <i>Method NodeClass</i>
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
executable	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
userExecutable	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.5 ObjectTypeAttributes parameter

Table 153 defines the *ObjectTypeAttributes* parameter.

Table 153 – ObjectTypeAttributes

Name	Type	Description
ObjectTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>ObjectType NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.6 VariableTypeAttributes parameter

Table 154 defines the *VariableTypeAttributes* parameter.

Table 154 – VariableTypeAttributes

Name	Type	Description
VariableTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>VariableType NodeClass</i>
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
value	Defined by the <i>Data Type Attribute</i>	See IEC 62541-3 for the description of this <i>Attribute</i> .
dataType	NodeId	See IEC 62541-3 for the description of this <i>Attribute</i> .
valueRank	Int32	See IEC 62541-3 for the description of this <i>Attribute</i> .
arrayDimensions	UInt32 []	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.7 ReferenceTypeAttributes parameter

Table 155 defines the *ReferenceTypeAttributes* parameter.

Table 155 – ReferenceTypeAttributes

Name	Type	Description
ReferenceTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>ReferenceType NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
symmetric	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
inverseName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.8 DataTypeAttributes parameter

Table 156 defines the *DataTypeAttributes* parameter.

Table 156 – DataTypeAttributes

Name	Type	Description
DataTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>DataType NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.9 ViewAttributes parameter

Table 157 defines the *ViewAttributes* parameter.

Table 157 – ViewAttributes

Name	Type	Description
ViewAttributes	structure	Defines the <i>Attributes</i> for the <i>View NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
containsNoLoops	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
eventNotifier	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.10 GenericAttributes parameter

This structure should be used instead of the *NodeClass* specific structures defined in the other subclauses of 7.19 since it allows the handling of additional *Attributes* defined in future specification versions.

Table 158 defines the *GenericAttributes* parameter.

Table 158 – GenericAttributes

Name	Type	Description
GenericAttributes	structure	Defines a generic structure for passing in any number of <i>Attributes</i> .
attributeValues	GenericAttributeValue []	Defines one <i>attributeId</i> and <i>value</i> combination.
attributeId	IntegerId	Id of the <i>Attribute</i> specified.
value	BaseDataType	Value of the <i>Attribute</i> specified.

7.20 NotificationData parameters

7.20.1 Overview

The *NotificationMessage* structure used in the *Subscription Service* set allows specifying different types of *NotificationData*. The *NotificationData* parameter is an extensible parameter whose structure depends on the type of *Notification* being sent. This parameter is defined in Table 159. Other types can be defined by additional parts of IEC 62541 or other specifications based on OPC UA. The *ExtensibleParameter* type is defined in 7.12.

There may be multiple notifications for a single *MonitoredItem* in a single *NotificationData* structure. When that happens the *Server* shall ensure the notifications appear in the same order that they were queued in the *MonitoredItem*. These notifications do not need to appear as a contiguous block.

Table 159 – NotificationData parameterTypeIds

Symbolic Id	Description
DataChange	<i>Notification</i> data parameter used for data change <i>Notifications</i> .
Event	<i>Notification</i> data parameter used for <i>Event Notifications</i> .
StatusChange	<i>Notification</i> data parameter used for Subscription status change <i>Notifications</i> .

7.20.2 DataChangeNotification parameter

Table 160 defines the *NotificationData* parameter used for data change notifications. This structure contains the monitored data items that are to be reported. Monitored data items are reported under two conditions:

- if the *MonitoringMode* is set to REPORTING and a change in value or its status (represented by its *StatusCode*) is detected;
- if the *MonitoringMode* is set to SAMPLING, the *MonitoredItem* is linked to a triggering item and the triggering item triggers.

See 5.12 for a description of the *MonitoredItem Service* set, and in particular the *MonitoringItemModel*, *MonitoredItem model* and the *Triggering model*.

After creating a *MonitoredItem*, the current value or status of the monitored Attribute shall be queued without applying the filter. If the current value is not available after the first sampling interval the first *Notification* shall be queued after getting the initial value or status from the data source.

Table 160 – DataChangeNotification

Name	Type	Description
DataChangeNotification	structure	Data change <i>Notification</i> data.
monitoredItems []	MonitoredItem Notification	The list of <i>MonitoredItems</i> for which a change has been detected. This structure is defined in-line with the following indented items.
clientHandle	IntegerId	<i>Client</i> -supplied handle for the <i>MonitoredItem</i> . The <i>IntegerId</i> type is defined in 7.14
Value	DataValue	The <i>StatusCode</i> , value and timestamp(s) of the monitored <i>Attribute</i> depending on the sampling and queuing configuration. If the <i>StatusCode</i> indicates an error then the value is to be ignored. If not every detected change has been returned since the <i>Server's</i> queue buffer for the <i>MonitoredItem</i> reached its limit and had to purge out data and the size of the queue is larger than one, the <i>Overflow</i> bit in the <i>DataValue InfoBits</i> of the <i>statusCode</i> is set. <i>DataValue</i> is a common type defined in 7.7.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information. The size and order of this list matches the size and order of the <i>monitoredItems</i> parameter. There is one entry in this list for each <i>Node</i> contained in the <i>monitoredItems</i> parameter. This list is empty if diagnostics information was not requested or is not available for any of the <i>MonitoredItems</i> . <i>DiagnosticInfo</i> is a common type defined in 7.8.

7.20.3 EventNotificationList parameter

Table 161 defines the *NotificationData* parameter used for *Event* notifications.

The EventNotificationList defines a table structure that is used to return *Event* fields to a *Client Subscription*. The structure is in the form of a table consisting of one or more *Events*, each containing an array of one or more fields. The selection and order of the fields returned for each *Event* is identical to the selected parameter of the *EventFilter*.

Table 161 – EventNotificationList

Name	Type	Description
EventNotificationList	structure	Event <i>Notification</i> data.
events []	EventFieldList	The list of <i>Events</i> being delivered. This structure is defined in-line with the following indented items.
clientHandle	IntegerId	<i>Client</i> -supplied handle for the <i>MonitoredItem</i> . The <i>IntegerId</i> type is defined in 7.14.
eventFields []	BaseDataType	List of selected <i>Event</i> fields. This shall be a one to one match with the fields selected in the <i>EventFilter</i> . 7.17.3 specifies how the <i>Server</i> shall deal with error conditions.

7.20.4 StatusChangeNotification parameter

Table 162 defines the *NotificationData* parameter used for a *StatusChangeNotification*.

The *StatusChangeNotification* informs the *Client* about a change in the status of a *Subscription*.

Table 162 – StatusChangeNotification

Name	Type	Description
StatusChangeNotification	structure	Event <i>Notification</i> data
status	StatusCode	The <i>StatusCode</i> that indicates the status change.
diagnosticInfo	DiagnosticInfo	Diagnostic information for the status change

7.21 NotificationMessage

The components of this parameter are defined in Table 163.

Table 163 – NotificationMessage

Name	Type	Description
NotificationMessage	structure	The <i>Message</i> that contains one or more <i>Notifications</i> .
sequenceNumber	Counter	The sequence number of the <i>NotificationMessage</i> .
publishTime	UtcTime	The time that this <i>Message</i> was sent to the <i>Client</i> . If this <i>Message</i> is retransmitted to the <i>Client</i> , this parameter contains the time it was first transmitted to the <i>Client</i> .
notificationData []	Extensible Parameter NotificationData	The list of <i>NotificationData</i> structures. The <i>NotificationData</i> parameter type is an extensible parameter type specified in 7.20. It specifies the types of <i>Notifications</i> that can be sent. The <i>ExtensibleParameter</i> type is specified in 7.12. Notifications of the same type should be grouped into one <i>NotificationData</i> element. If a <i>Subscription</i> contains <i>MonitoredItems</i> for events and data, this array should have not more than 2 elements. If the <i>Subscription</i> contains <i>MonitoredItems</i> only for data or only for events, the array size should always be one for this <i>Subscription</i> .

7.22 NumericRange

This parameter is defined in Table 164. A formal BNF definition of the numeric range can be found in Clause A.3.

The syntax for the string contains one of the following two constructs. The first construct is the string representation of an individual integer. For example, "6" is valid, but "6,0" and "3,2" are not. The minimum and maximum values that can be expressed are defined by the use of this parameter and not by this parameter type definition. The second construct is a range represented by two integers separated by the colon (":") character. The first integer shall always have a lower value than the second. For example, "5:7" is valid, while "7:5" and "5:5" are not. The minimum and maximum values that can be expressed by these integers are defined by the use of this parameter, and not by this parameter type definition. No other characters, including white-space characters, are permitted.

Multi-dimensional arrays can be indexed by specifying a range for each dimension separated by a ','. For example, a 2x2 block in a 4x4 matrix could be selected with the range "1:2,0:1". A single element in a multi-dimensional array can be selected by specifying a single number instead of a range. For example, "1,1" ~~specifies~~ selects the [1,1] element in a two dimensional array.

Dimensions are specified in the order that they appear in the *ArrayDimensions* Attribute. All dimensions shall be specified for a *NumericRange* to be valid.

All indexes start with 0. The maximum value for any index is one less than the length of the dimension.

~~When reading a value, the indexes may not specify a range that is within the bounds of the array. The Server shall return a partial result if some elements exist within the range. the Server shall return a *Bad_IndexRangeNoData* if no elements exist within the range.~~

When reading a value and any of the lower bounds of the indexes is out of range, the Server shall return a *Bad_IndexRangeNoData*. If any of the upper bounds of the indexes is out of range, the Server shall return partial results

Bad_IndexRangeInvalid is only used for invalid syntax of the *NumericRange*. All other invalid requests with a valid syntax shall result in *Bad_IndexRangeNoData*.

When writing a value, the size of the array shall match the size specified by the *NumericRange*. The Server shall return an error if it cannot write all elements specified by the Client.

The *NumericRange* can also be used to specify substrings for *ByteString* and *String* values. Arrays of *ByteString* and *String* values are treated as two dimensional arrays where the final index specifies the substring range within the *ByteString* or *String* value. The entire *ByteString* or *String* value is selected if the final index is omitted.

Table 164 – NumericRange

Name	Type	Description
NumericRange	String	A number or a numeric range. A null string indicates that this parameter is not used.

7.23 QueryDataSet

The components of this parameter are defined in Table 165.

Table 165 – QueryDataSet

Name	Type	Description
QueryDataSet	structure	Data related to a <i>Node</i> returned in a Query response.
nodeId	ExpandedNodeId	The <i>NodeId</i> for this <i>Node</i> description.
typeDefinitionNode	ExpandedNodeId	The <i>NodeId</i> for the type definition for this <i>Node</i> description.
values []	BaseDataType	Values for the selected <i>Attributes</i>. The order of returned items matches the order of the requested items. There is an entry for each requested item for the given <i>TypeDefinitionNode</i> that matches the selected instance, this includes any related nodes that were specified using a relative path from the selected instance's <i>TypeDefinitionNode</i>. If no values were found for a given requested item a null value is returned for that item. If a value has a bad status, the <i>StatusCode</i> is returned instead of the value. If multiple values exist for a requested item then an array of values is returned. If the requested item is a reference then a <i>ReferenceDescription</i> or array of <i>ReferenceDescription</i> is returned for that item. If the <i>QueryDataSet</i> is returned in a <i>QueryNext</i> to continue a list of <i>ReferenceDescription</i>, the <i>values</i> array will have the same size but the other values already returned are null.

7.24 ReadValueId

The components of this parameter are defined in Table 166.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 166 – ReadValueId

Name	Type	Description
ReadValueId	structure	Identifier for an item to read or to monitor.
nodeId	NodeId	<i>NodeId</i> of a <i>Node</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>Attribute</i> id. The <i>IntegerId</i> is defined in 7.14. The <i>IntegerIds</i> for the <i>Attributes</i> are defined in IEC 62541-6.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for arrays. If a range of elements is specified, the values are returned as a composite. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is null if the specified <i>Attribute</i> is not an array. However, if the specified <i>Attribute</i> is an array, and this parameter is null, then all elements are to be included in the range.
dataEncoding	QualifiedName	This parameter specifies the <i>BrowseName</i> of the <i>DataTypeEncoding</i> that the <i>Server</i> should use when returning the Value <i>Attribute</i> of a <i>Variable</i>. It is an error to specify this parameter for other <i>Attributes</i>. This parameter only applies if the <i>DataType</i> of the <i>Variable</i> is a subtype of <i>Structure</i>. It is an error to specify this parameter if the <i>DataType</i> of the <i>Variable</i> is not a subtype of <i>Structure</i>. A <i>Client</i> can discover what <i>DataTypeEncodings</i> are available by following the <i>HasEncoding Reference</i> from the <i>DataType Node</i> for a <i>Variable</i>. OPC UA defines <i>BrowseNames</i> which <i>Servers</i> shall recognize even if the <i>DataType Nodes</i> are not visible in the <i>Server AddressSpace</i>. These <i>BrowseNames</i> are: Default Binary The default or native binary (or non-XML) encoding. Default XML The default XML encoding. Default JSON The default JSON encoding Each <i>DataType</i> shall support at least one of these encodings. <i>DataTypes</i> that do not have a true binary encoding (e.g. they only have a non-XML text encoding) should use the Default Binary name to identify the encoding that is considered to be the default non-XML encoding. <i>DataTypes</i> that support at least one XML-based encoding shall identify one of the encodings as the Default XML encoding. Other standards bodies may define other well-known data encodings that could be supported. If this parameter is not specified then the <i>Server</i> shall choose ~~either the default Binary or Default XML encoding~~ according to what <i>Message</i> encoding (see IEC 62541-6) is used for the <i>Session</i>. If the <i>Server</i> does not support the encoding that matches the <i>Message</i> encoding then the <i>Server</i> shall choose the default encoding that it does support. ~~If this parameter is specified for a <i>MonitoredItem</i>, the <i>Server</i> shall set the <i>StructureChanged</i> bit in the <i>StatusCode</i> (see 7.33) if the <i>DataTypeEncoding</i> changes. The <i>DataTypeEncoding</i> changes if the <i>DataTypeVersion</i> of the <i>DataTypeDescription</i> or the <i>DataTypeDictionary</i> associated with the <i>DataTypeEncoding</i> changes.~~

7.25 ReferenceDescription

The components of this parameter are defined in Table 167.

Table 167 – ReferenceDescription

Name	Type	Description
ReferenceDescription	structure	Reference parameters returned for the <i>Browse Service</i> .
referenceTypeid	Nodeid	<i>Nodeid</i> of the <i>ReferenceType</i> that defines the <i>Reference</i> .
isForward	Boolean	If the value is TRUE, the <i>Server</i> followed a forward <i>Reference</i> . If the value is FALSE, the <i>Server</i> followed an inverse <i>Reference</i> .
nodeid	Expanded Nodeid	<i>Nodeid</i> of the <i>TargetNode</i> as assigned by the <i>Server</i> identified by the <i>Server</i> index. The <i>ExpandedNodeid</i> type is defined in 7.11. If the <i>serverIndex</i> indicates that the <i>TargetNode</i> is a remote <i>Node</i> , then the <i>nodeid</i> shall contain the absolute namespace URI. If the <i>TargetNode</i> is a local <i>Node</i> the <i>nodeid</i> shall contain the namespace index.
browseName ¹⁾	QualifiedName	The <i>BrowseName</i> of the <i>TargetNode</i> .
displayName	LocalizedText	The <i>DisplayName</i> of the <i>TargetNode</i> .
nodeClass ¹⁾	NodeClass	<i>NodeClass</i> of the <i>TargetNode</i> .
typeDefinition ¹⁾	Expanded Nodeid	Type definition <i>Nodeid</i> of the <i>TargetNode</i> . Type definitions are only available for the <i>NodeClasses Object</i> and <i>Variable</i> . For all other <i>NodeClasses</i> a null <i>Nodeid</i> shall be returned.
¹⁾ If the <i>Server</i> index indicates that the <i>TargetNode</i> is a remote <i>Node</i> , then the <i>browseName</i> , <i>nodeClass</i> and <i>typeDefinition</i> may be null or empty. If they are not, they might not be up to date because the local <i>Server</i> might not continuously monitor the remote <i>Server</i> for changes. The <i>displayName</i> shall be provided for remote <i>Nodes</i> .		

7.26 RelativePath

The components of this parameter are defined in Table 168.

Table 168 – RelativePath

Name	Type	Description
RelativePath	structure	Defines a sequence of <i>References</i> and <i>BrowseNames</i> to follow.
elements []	RelativePath Element	A sequence of <i>References</i> and <i>BrowseNames</i> to follow. This structure is defined in-line with the following indented items. Each element in the sequence is processed by finding the targets and then using those targets as the starting nodes for the next element. The targets of the final element are the target of the <i>RelativePath</i> .
referenceTypeid	Nodeid	The type of reference to follow from the current node. The current path cannot be followed any further if the <i>referenceTypeid</i> is not available on the <i>Node</i> instance. If not specified then all <i>References</i> are included and the parameter <i>includeSubtypes</i> is ignored.
isInverse	Boolean	~~Indicates whether the~~ Only inverse ~~Reference should~~ references shall be followed. ~~The inverse reference~~ if this value is TRUE. Only forward references shall be followed if this value is TRUEFALSE.
includeSubtypes	Boolean	Indicates whether subtypes of the <i>ReferenceType</i> should be followed. Subtypes are included if this value is TRUE.
targetName	QualifiedName	The <i>BrowseName</i> of the target node. The final element may have an empty <i>targetName</i> . In this situation all targets of the references identified by the <i>referenceTypeid</i> are the targets of the <i>RelativePath</i> . The <i>targetName</i> shall be specified for all other elements. The current path cannot be followed any further if no targets with the specified <i>BrowseName</i> exist.

A *RelativePath* can be applied to any starting *Node*. The targets of the *RelativePath* are the set of *Nodes* that are found by sequentially following the elements in *RelativePath*.

A text format for the *RelativePath* can be found in Clause A.2. This format is used in examples that explain the *Services* that make use of the *RelativePath* structure.

7.27 RegisteredServer

The components of this parameter are defined in Table 169.

Table 169 – RegisteredServer

Name	Type	Description
RegisteredServer	structure	The <i>Server</i> to register.
serverUri	String	The globally unique identifier for the <i>Server</i> instance. The <i>serverUri</i> matches the <i>applicationUri</i> from the <i>ApplicationDescription</i> defined in 7.1.
productUri	String	The globally unique identifier for the <i>Server</i> product.
serverNames []	LocalizedText	A list of localized descriptive names for the <i>Server</i> . The list shall have at least one valid entry.
serverType	Enum ApplicationType	The type of application. The enumeration values are defined in Table 112. The value "CLIENT_1" (the application is a <i>Client</i>) is not allowed. The <i>Service</i> result shall be <i>Bad_InvalidArgument</i> in this case.
gatewayServerUri	String	The URI of the <i>Gateway Server</i> associated with the <i>discoveryUrls</i> . This value is only specified by <i>Gateway Servers</i> that wish to register the <i>Servers</i> that they provide access to. For <i>Servers</i> that do not act as a <i>Gateway Server</i> this parameter shall be null.
discoveryUrls []	String	A list of <i>DiscoveryEndpoints</i> for the <i>Server</i> . The list shall have at least one valid entry.
semaphoreFilePath	String	The path to the semaphore file used to identify an automatically-launched <i>Server</i> instance; manually-launched servers will not use this parameter. If a Semaphore file is provided, the <i>isOnline</i> flag is ignored. If a Semaphore file is provided and exists, the <i>LocalDiscoveryServer</i> shall save the registration information in a persistent data store that it reads whenever the <i>LocalDiscoveryServer</i> starts. If a Semaphore file is specified but does not exist, the <i>Discovery Server</i> shall remove the registration from any persistent data store. If the <i>Server</i> has registered with a <i>semaphoreFilePath</i> , the <i>Discovery Server</i> shall check that this file exists before returning the <i>ApplicationDescription</i> to the client. If the <i>Server</i> did not register with a <i>semaphoreFilePath</i> (it is null or empty), then the <i>Discovery Server</i> does not attempt to verify the existence of the file before returning the <i>ApplicationDescription</i> to the client.
isOnline	Boolean	True if the <i>Server</i> is currently able to accept connections from <i>Clients</i> . The <i>Discovery Server</i> shall return <i>ApplicationDescriptions</i> to the <i>Client</i> . The <i>Server</i> is expected to periodically re-register with the <i>Discovery Server</i> . False if the <i>Server</i> is currently unable to accept connections from <i>Clients</i> . The <i>Discovery Server</i> shall NOT return <i>ApplicationDescriptions</i> to the <i>Client</i> . This parameter is ignored if a <i>semaphoreFilePath</i> is provided.

7.28 RequestHeader

The components of this parameter are defined in Table 170.

Table 170 – RequestHeader

Name	Type	Description																																				
RequestHeader	structure	Common parameters for all requests submitted on a <i>Session</i> .																																				
authenticationToken	Session AuthenticationToken	The secret <i>Session</i> identifier used to verify that the request is associated with the <i>Session</i> . The <i>SessionAuthenticationToken</i> type is defined in 7.31.																																				
timestamp	UtcTime	The time the <i>Client</i> sent the request. The parameter is only used for diagnostic and logging purposes in the server.																																				
requestHandle	IntegerId	A <i>requestHandle</i> associated with the request. This <i>Client</i> defined handle can be used to cancel the request. It is also returned in the response.																																				
returnDiagnostics	UInt32	A bit mask that identifies the types of vendor-specific diagnostics to be returned in <i>diagnosticInfo</i> response parameters. The value of this parameter may consist of zero, one or more of the following values. No value indicates that diagnostics are not to be returned. <table><thead><tr><th>Bit Value</th><th>Diagnostics to return</th></tr></thead><tbody><tr><td>0x0000 0001</td><td>ServiceLevel / SymbolicId</td></tr><tr><td>0x0000 0002</td><td>ServiceLevel / LocalizedText</td></tr><tr><td>0x0000 0004</td><td>ServiceLevel / AdditionalInfo</td></tr><tr><td>0x0000 0008</td><td>ServiceLevel / Inner <i>StatusCode</i></td></tr><tr><td>0x0000 0010</td><td>ServiceLevel / Inner Diagnostics</td></tr><tr><td>0x0000 0020</td><td>OperationLevel / SymbolicId</td></tr><tr><td>0x0000 0040</td><td>OperationLevel / LocalizedText</td></tr><tr><td>0x0000 0080</td><td>OperationLevel / AdditionalInfo</td></tr><tr><td>0x0000 0100</td><td>OperationLevel / Inner <i>StatusCode</i></td></tr><tr><td>0x0000 0200</td><td>OperationLevel / Inner Diagnostics</td></tr></tbody></table> Each of these values is composed of two components, <i>level</i> and <i>type</i>, as described below. If none are requested, as indicated by a 0 value, or if no diagnostic information was encountered in processing of the request, then diagnostics information is not returned. Level: <table><tbody><tr><td>ServiceLevel</td><td>return diagnostics in the <i>diagnosticInfo</i> of the <i>Service</i>.</td></tr><tr><td>OperationLevel</td><td>return diagnostics in the <i>diagnosticInfo</i> defined for individual operations requested in the <i>Service</i>.</td></tr></tbody></table> Type: <table><tbody><tr><td>SymbolicId</td><td>return a namespace-qualified, symbolic identifier for an error or condition. The maximum length of this identifier is 32 characters.</td></tr><tr><td>LocalizedText</td><td>return up to 256 bytes of localized text that describes the symbolic id.</td></tr><tr><td>AdditionalInfo</td><td>return a byte string that contains additional diagnostic information, such as a memory image. The format of this byte string is vendor-specific, and may depend on the type of error or condition encountered.</td></tr><tr><td>InnerStatusCode</td><td>return the inner <i>StatusCode</i> associated with the operation or <i>Service</i>.</td></tr><tr><td>InnerDiagnostics</td><td>return the inner diagnostic info associated with the operation or <i>Service</i>. The contents of the inner diagnostic info structure are determined by other bits in the mask. Note that setting this bit could cause multiple levels of nested diagnostic info structures to be returned.</td></tr></tbody></table>	Bit Value	Diagnostics to return	0x0000 0001	ServiceLevel / SymbolicId	0x0000 0002	ServiceLevel / LocalizedText	0x0000 0004	ServiceLevel / AdditionalInfo	0x0000 0008	ServiceLevel / Inner <i>StatusCode</i>	0x0000 0010	ServiceLevel / Inner Diagnostics	0x0000 0020	OperationLevel / SymbolicId	0x0000 0040	OperationLevel / LocalizedText	0x0000 0080	OperationLevel / AdditionalInfo	0x0000 0100	OperationLevel / Inner <i>StatusCode</i>	0x0000 0200	OperationLevel / Inner Diagnostics	ServiceLevel	return diagnostics in the <i>diagnosticInfo</i> of the <i>Service</i> .	OperationLevel	return diagnostics in the <i>diagnosticInfo</i> defined for individual operations requested in the <i>Service</i> .	SymbolicId	return a namespace-qualified, symbolic identifier for an error or condition. The maximum length of this identifier is 32 characters.	LocalizedText	return up to 256 bytes of localized text that describes the symbolic id.	AdditionalInfo	return a byte string that contains additional diagnostic information, such as a memory image. The format of this byte string is vendor-specific, and may depend on the type of error or condition encountered.	InnerStatusCode	return the inner <i>StatusCode</i> associated with the operation or <i>Service</i> .	InnerDiagnostics	return the inner diagnostic info associated with the operation or <i>Service</i> . The contents of the inner diagnostic info structure are determined by other bits in the mask. Note that setting this bit could cause multiple levels of nested diagnostic info structures to be returned.
Bit Value	Diagnostics to return																																					
0x0000 0001	ServiceLevel / SymbolicId																																					
0x0000 0002	ServiceLevel / LocalizedText																																					
0x0000 0004	ServiceLevel / AdditionalInfo																																					
0x0000 0008	ServiceLevel / Inner <i>StatusCode</i>																																					
0x0000 0010	ServiceLevel / Inner Diagnostics																																					
0x0000 0020	OperationLevel / SymbolicId																																					
0x0000 0040	OperationLevel / LocalizedText																																					
0x0000 0080	OperationLevel / AdditionalInfo																																					
0x0000 0100	OperationLevel / Inner <i>StatusCode</i>																																					
0x0000 0200	OperationLevel / Inner Diagnostics																																					
ServiceLevel	return diagnostics in the <i>diagnosticInfo</i> of the <i>Service</i> .																																					
OperationLevel	return diagnostics in the <i>diagnosticInfo</i> defined for individual operations requested in the <i>Service</i> .																																					
SymbolicId	return a namespace-qualified, symbolic identifier for an error or condition. The maximum length of this identifier is 32 characters.																																					
LocalizedText	return up to 256 bytes of localized text that describes the symbolic id.																																					
AdditionalInfo	return a byte string that contains additional diagnostic information, such as a memory image. The format of this byte string is vendor-specific, and may depend on the type of error or condition encountered.																																					
InnerStatusCode	return the inner <i>StatusCode</i> associated with the operation or <i>Service</i> .																																					
InnerDiagnostics	return the inner diagnostic info associated with the operation or <i>Service</i> . The contents of the inner diagnostic info structure are determined by other bits in the mask. Note that setting this bit could cause multiple levels of nested diagnostic info structures to be returned.																																					

Name	Type	Description
auditEntryId	String	An identifier that identifies the <i>Client</i> 's security audit log entry associated with this request. An empty string value means that this parameter is not used. The <i>auditEntryId</i> typically contains who initiated the action and from where it was initiated. The <i>AuditEventId</i> <i>auditEntryId</i> is included in the <i>AuditEvent</i> to allow human readers to correlate an <i>Event</i> with the initiating action. More details of the <i>Audit</i> mechanisms are defined in 6.5 and in IEC 62541-3.
timeoutHint	UInt32	This timeout in milliseconds is used in the <i>Client</i> side <i>Communication Stack</i> to set the timeout on a per-call base. For a <i>Server</i> this timeout is only a hint and can be used to cancel long running operations to free resources. If the <i>Server</i> detects a timeout, he can cancel the operation by sending the <i>Service</i> result <i>Bad_Timeout</i> . The <i>Server</i> should wait at minimum the timeout after he received the request before cancelling the operation. The <i>Server</i> shall check the <i>timeoutHint</i> parameter of a <i>Publish</i> request before processing a <i>Publish</i> response. If the request timed out, a <i>Bad_Timeout Service</i> result is sent and another <i>Publish</i> request is used. The value of 0 indicates no timeout.
additionalHeader	Extensible Parameter AdditionalHeader	Reserved for future use. Applications that do not understand the header should ignore it.

7.29 ResponseHeader

The components of this parameter are defined in Table 171.

Table 171 – ResponseHeader

Name	Type	Description
ResponseHeader	structure	Common parameters for all responses.
timestamp	UtcTime	The time the <i>Server</i> sent the response.
requestHandle	IntegerId	The requestHandle given by the <i>Client</i> to the request.
serviceResult	StatusCode	OPC UA-defined result of the <i>Service</i> invocation. The <i>StatusCode</i> type is defined in 7.34.
serviceDiagnostics	DiagnosticInfo	Diagnostic information for the <i>Service</i> invocation. This parameter is empty if diagnostics information was not requested in the request header. The <i>DiagnosticInfo</i> type is defined in 7.8.
stringTable []	String	There is one string in this list for each unique namespace, symbolic identifier, and localized text string contained in all of the diagnostics information parameters contained in the response (see 7.8). Each is identified within this table by its zero-based index.
additionalHeader	Extensible Parameter AdditionalHeader	Reserved for future use. Applications that do not understand the header should ignore it.

7.30 ServiceFault

The components of this parameter are defined in Table 172.

The *ServiceFault* parameter is returned instead of the *Service* response message when a service level error occurs. The *requestHandle* in the *ResponseHeader* should be set to what was provided in the *RequestHeader* even if these values were not valid. The level of diagnostics returned in the *ResponseHeader* is specified by the *returnDiagnostics* parameter in the *RequestHeader*.

The exact use of this parameter depends on the mappings defined in IEC 62541-6.

Table 172 – ServiceFault

Name	Type	Description
ServiceFault	structure	An error response sent when a service level error occurs.
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

7.31 SessionAuthenticationToken

The *SessionAuthenticationToken* type is an opaque identifier that is used to identify requests associated with a particular *Session*. This identifier is used in conjunction with the *SecureChannelId* or *Client Certificate* to authenticate incoming messages. It is the secret form of the *sessionId* for internal use in the *Client* and *Server Applications*.

A *Server* returns a *SessionAuthenticationToken* in the *CreateSession* response. The *Client* then sends this value with every request which allows the *Server* to verify that the sender of the request is the same as the sender of the original *CreateSession* request.

For the purposes of this discussion, a *Server* consists of application (code) and a *Communication Stack* as shown in Figure 37. The security provided by the *SessionAuthenticationToken* depends on a trust relationship between the *Server* application and the *Communication Stack*. The *Communication Stack* shall be able to verify the sender of the message and it uses the *SecureChannelId* or the *Client Certificate* to identify the sender to the *Server*. In these cases, the *SessionAuthenticationToken* is a UInt32 identifier that allows the *Server* to distinguish between different *Sessions* created by the same sender.



Figure 37 – Logical layers of a Server

In some cases, the application and the *Communication Stack* cannot exchange information at runtime which means the application will not have access to the *SecureChannelId* or the *Certificate* used to create the *SecureChannel*. In these cases the application shall create a random *ByteString* value that is at least 32 bytes long. This value shall be kept secret and shall always be exchanged over a *SecureChannel* with encryption enabled. The Administrator is responsible for ensuring that encryption is enabled. The *Profiles* in IEC 62541-7 may define additional requirements for a *ByteString SessionAuthenticationToken*.

Client and *Server* applications should be written to be independent of the *SecureChannel* implementation. Therefore, they should always treat the *SessionAuthenticationToken* as secret information even if it is not required when using some *SecureChannel* implementations.

Figure 38 illustrates the information exchanged between the *Client*, the *Server* and the *Server Communication Stack* when the *Client* obtains a *SessionAuthenticationToken*. In this figure the *GetSecureChannelInfo* step represents an API that depends on the *Communication Stack* implementation.




Figure 38 – Obtaining a SessionAuthenticationToken

The *SessionAuthenticationToken* is a subtype of the *NodeId* data type; however, it is never used to identify a *Node* in the *AddressSpace*. Servers may assign a value to the *NamespaceIndex*; however, its meaning is *Server* specific.

7.32 SignatureData

The components of this parameter are defined in Table 173.

Table 173 – SignatureData

Name	Type	Description
SignatureData	structure	Contains a digital signature created with a <i>Certificate</i> .
algorithm	String	A string containing the URI of the <i>algorithm</i> . The URI string values are defined as part of the security profiles specified in IEC 62541-7.
signature	ByteString	This is a signature generated with the private key associated with a <i>Certificate</i> .

~~7.32 SoftwareCertificate~~

~~The components of this parameter are defined in Table 162.~~

Table 162 — SoftwareCertificate

Name	Type	Description
SoftwareCertificate	Structure	A <i>Certificate</i> describing a product.
— productName	String	The name of the product that is certified. This field shall be specified.
— productUri	String	A globally unique identifier for the product that is certified. This field shall be specified.
— vendorName	String	The name of the vendor responsible for the product. This field shall be specified.
— vendorProductCertificate	ByteString	The DER-encoded form of the X509 <i>Certificate</i> which is assigned to the product by the vendor. This field may be omitted.
— softwareVersion	String	Software version. This field shall be specified.
— buildNumber	String	Build number. This field shall be specified.
— buildDate	UtcTime	Date and time of the build. This field shall be specified.
— issuedBy	String	URI of the certifying authority. This field shall be specified.
— issuedDate	UtcTime	Specifies when the <i>Certificate</i> was issued by the certifying authority. This field shall be specified.
— supportedProfiles []	structure	List of supported <i>Profiles</i> . This structure is defined in-line with the following indented items:
—— organizationUri	String	A URI that identifies the organization that defined the profile.
—— profileId	String	A string that identifies the <i>Profile</i> .
—— complianceTool	String	A string that identifies the tool or certification method used for compliance testing.
—— complianceDate	UtcTime	Date and time of the compliance test.
—— complianceLevel	Enum Compliance Level	An enumeration that specifies the compliance level of the <i>Profile</i> . It has the following values: — UNTESTED_0 — the profiled capability has not been tested successfully — PARTIAL_1 — the profiled capability has been partially tested and has passed critical tests, as defined by the certifying authority. — SELFTESTED_2 — the profiled capability has been successfully tested using a self-test system authorized by the certifying authority. — CERTIFIED_3 — the profiled capability has been successfully tested by a testing organisation authorized by the certifying authority.
—— unsupportedUnitIds []	String	The identifiers for the optional conformance units that were not tested. See IEC 62541-7 for a detailed explanation.

7.33 SignedSoftwareCertificate

A *SignedSoftwareCertificate* is a *ByteString* containing an encoded *Certificate*. The encoding of a *SignedSoftwareCertificate* depends on the security technology mapping and is defined completely in IEC 62541-6. Table 161 specifies the information that shall be contained in a *SignedSoftwareCertificate*.

NOTE Details on SoftwareCertificates will be defined in a future edition of IEC 62541-4.

Table 174 specifies *SignedSoftwareCertificate Structure*.

Table 174 – SignedSoftwareCertificate

Name	Type	Description
SignedSoftwareCertificate	structure	A <i>SoftwareCertificate</i> with a signature created by Certifying Authority.
—version	String	An identifier for the version of the <i>Certificate</i> encoding.
—serialNumber	ByteString	A unique identifier for the <i>Certificate</i> assigned by the Issuer.
—signatureAlgorithm	String	The algorithm used to sign the <i>Certificate</i> . The syntax of this field depends on the <i>Certificate</i> encoding.
—signature	ByteString	The signature created by the Issuer.
—issuer	Structure	A name that identifies the <i>Issuer Certificate</i> used to create the signature.
—validFrom	UtcTime	When the <i>Certificate</i> becomes valid.
—validTo	UtcTime	When the <i>Certificate</i> expires.
—subject	Structure	A name that identifies the product which the <i>Certificate</i> describes. This field shall contain the <i>productName</i> and <i>vendorName</i> from the <i>SoftwareCertificate</i> .
—subjectAltName []	Structure	A list of alternate names for the product. This list shall include the <i>productUri</i> specified in the <i>SoftwareCertificate</i> .
—publicKey	ByteString	The public key associated with the <i>Certificate</i> .
—keyUsage []	String	Specifies how the <i>Certificate</i> key may be used. <i>SignedSoftwareCertificate</i> may only be used for creating Digital Signatures and Non-Repudiation. The contents of this field depends on the <i>Certificate</i> encoding.
—softwareCertificate	ByteString	The XML encoded form of the <i>SoftwareCertificate</i> stored as UTF8 text. IEC 62541-6 describes the XML representation for the <i>SoftwareCertificate</i> .

Name	Type	Description
SignedSoftwareCertificate	structure	
certificateData	ByteString	The certificate data serialized as a ByteString.
signature	ByteString	The signature for the certificateData.

7.34 StatusCode

7.34.1 General

A *StatusCode* in OPC UA is numerical value that is used to report the outcome of an operation performed by an OPC UA Server. This code may have associated diagnostic information that describes the status in more detail; however, the code by itself is intended to provide *Client* applications with enough information to make decisions on how to process the results of an OPC UA Service.

The *StatusCode* is a 32-bit unsigned integer. The top 16 bits represent the numeric value of the code that shall be used for detecting specific errors or conditions. The bottom 16 bits are bit flags that contain additional information but do not affect the meaning of the *StatusCode*.

All OPC UA *Clients* shall always check the *StatusCode* associated with a result before using it. Results that have an uncertain/warning status associated with them shall be used with care since these results might not be valid in all situations. Results with a bad/failed status shall never be used.

OPC UA Servers should return good/success *StatusCodes* if the operation completed normally and the result is always valid. Different *StatusCode* values can provide additional information to the *Client*.

OPC UA Servers should use uncertain/warning *StatusCodes* if they could not complete the operation in the manner requested by the *Client*, however, the operation did not fail entirely.

The list of *StatusCodes* is managed by OPC UA. The complete list of *StatusCodes* is defined in IEC 62541-6. Servers shall not define their own *StatusCodes*. OPC UA companion working

groups may request additional *StatusCodes* from the OPC Foundation to be added to the list in IEC 62541-6.

The exact bit assignments are shown in Table 175.

Table 175 – StatusCode bit assignments

Field	Bit Range	Description												
Severity	30:31	Indicates whether the <i>StatusCode</i> represents a good, bad or uncertain condition. These bits have the following meanings: <table> <tr> <td>Good Success</td><td>00</td><td>Indicates that the operation was successful and the associated results may be used.</td></tr> <tr> <td>Uncertain Warning</td><td>01</td><td>Indicates that the operation was partially successful and that associated results might not be suitable for some purposes.</td></tr> <tr> <td>Bad Failure</td><td>10</td><td>Indicates that the operation failed and any associated results cannot be used.</td></tr> <tr> <td>Reserved</td><td>11</td><td>Reserved for future use. All <i>Clients</i> should treat a <i>StatusCode</i> with this severity as "Bad".</td></tr> </table>	Good Success	00	Indicates that the operation was successful and the associated results may be used.	Uncertain Warning	01	Indicates that the operation was partially successful and that associated results might not be suitable for some purposes.	Bad Failure	10	Indicates that the operation failed and any associated results cannot be used.	Reserved	11	Reserved for future use. All <i>Clients</i> should treat a <i>StatusCode</i> with this severity as "Bad".
Good Success	00	Indicates that the operation was successful and the associated results may be used.												
Uncertain Warning	01	Indicates that the operation was partially successful and that associated results might not be suitable for some purposes.												
Bad Failure	10	Indicates that the operation failed and any associated results cannot be used.												
Reserved	11	Reserved for future use. All <i>Clients</i> should treat a <i>StatusCode</i> with this severity as "Bad".												
Reserved	29:29	Reserved for use in OPC UA application specific APIs. This bit shall always be zero on the wire but may be used by OPC UA application specific APIs for API specific status codes.												
Reserved	28:28	Reserved for future use. Shall always be zero.												
SubCode	16:27	The code is a numeric value assigned to represent different conditions. Each code has a symbolic name and a numeric value. All descriptions in the OPC UA specification refer to the symbolic name. IEC 62541-6 maps the symbolic names onto a numeric value.												
StructureChanged	15:15	Indicates that the structure of the associated data value has changed since the last <i>Notification</i> . <i>Clients</i> should not process the data value unless they re-read the metadata. <i>Servers</i> shall set this bit if the <i>DataTypeEncoding</i> used for a <i>Variable</i> changes. 7.24 describes how the <i>DataTypeEncoding</i> is specified for a <i>Variable</i> . <i>Servers</i> shall also set this bit if the <i>EnumStrings Property</i> of the <i>DataType</i> of the <i>Variable</i> changes. This bit is provided to warn <i>Clients</i> that parse complex data values that their parsing routines could fail because the serialized form of the data value has changed. This bit has meaning only for <i>StatusCodes</i> returned as part of a data change <i>Notification</i> or the <i>HistoryRead</i> . <i>StatusCodes</i> used in other contexts shall always set this bit to zero.												
SemanticsChanged	14:14	Indicates that the semantics of the associated data value have changed. <i>Clients</i> should not process the data value until they re-read the metadata associated with the <i>Variable</i> . <i>Servers</i> should set this bit if the metadata has changed in way that could cause application errors if the <i>Client</i> does not re-read the metadata. For example, a change to the engineering units could create problems if the <i>Client</i> uses the value to perform calculations. IEC 62541-8 defines the conditions where a <i>Server</i> shall set this bit for a DA <i>Variable</i> . Other specifications may define additional conditions. A <i>Server</i> may define other conditions that cause this bit to be set. This bit has meaning only for <i>StatusCodes</i> returned as part of a data change <i>Notification</i> or the <i>HistoryRead</i> . <i>StatusCodes</i> used in other contexts shall always set this bit to zero.												
Reserved	12:13	Reserved for future use. Shall always be zero.												
InfoType	10:11	The type of information contained in the info bits. These bits have the following meanings: <table> <tr> <td>NotUsed</td><td>00</td><td>The info bits are not used and shall be set to zero.</td></tr> <tr> <td>DataValue</td><td>01</td><td>The <i>StatusCode</i> and its info bits are associated with a data value returned from the <i>Server</i>. This flag is only used in combination with <i>StatusCodes</i> defined in IEC 62541-8. The info bits are defined in Table 176.</td></tr> <tr> <td>Reserved</td><td>1X</td><td>Reserved for future use. The info bits shall be ignored.</td></tr> </table>	NotUsed	00	The info bits are not used and shall be set to zero.	DataValue	01	The <i>StatusCode</i> and its info bits are associated with a data value returned from the <i>Server</i> . This flag is only used in combination with <i>StatusCodes</i> defined in IEC 62541-8. The info bits are defined in Table 176.	Reserved	1X	Reserved for future use. The info bits shall be ignored.			
NotUsed	00	The info bits are not used and shall be set to zero.												
DataValue	01	The <i>StatusCode</i> and its info bits are associated with a data value returned from the <i>Server</i> . This flag is only used in combination with <i>StatusCodes</i> defined in IEC 62541-8. The info bits are defined in Table 176.												
Reserved	1X	Reserved for future use. The info bits shall be ignored.												
InfoBits	0:9	Additional information bits that qualify the <i>StatusCode</i> . The structure of these bits depends on the Info Type field.												

Table 176 describes the structure of the *InfoBits* when the Info Type is set to *DataValue* (01).

Table 176 – DataValue InfoBits

Info Type	Bit Range	Description																					
LimitBits	8:9	The limit bits associated with the data value. The limits bits have the following meanings: <table> <tr> <th>Limit</th><th>Bits</th><th>Description</th></tr> <tr> <td>None</td><td>00</td><td>The value is free to change.</td></tr> <tr> <td>Low</td><td>01</td><td>The value is at the lower limit for the data source.</td></tr> <tr> <td>High</td><td>10</td><td>The value is at the higher limit for the data source.</td></tr> <tr> <td>Constant</td><td>11</td><td>The value is constant and cannot change.</td></tr> </table>	Limit	Bits	Description	None	00	The value is free to change.	Low	01	The value is at the lower limit for the data source.	High	10	The value is at the higher limit for the data source.	Constant	11	The value is constant and cannot change.						
Limit	Bits	Description																					
None	00	The value is free to change.																					
Low	01	The value is at the lower limit for the data source.																					
High	10	The value is at the higher limit for the data source.																					
Constant	11	The value is constant and cannot change.																					
Overflow	7	This bit shall only be set if the <i>MonitoredItem</i> queue size is greater than 1. If this bit is set, not every detected change has been returned since the <i>Server's</i> queue buffer for the <i>MonitoredItem</i> reached its limit and had to purge out data.																					
Reserved	5:6	Reserved for future use. Shall always be zero.																					
HistorianBits	0:4	These bits are set only when reading historical data. They indicate where the data value came from and provide information that affects how the <i>Client</i> uses the data value. The historian bits have the following meaning: <table> <tr> <td>Raw</td><td>XXX00</td><td>A raw data value.</td></tr> <tr> <td>Calculated</td><td>XXX01</td><td>A data value which was calculated.</td></tr> <tr> <td>Interpolated</td><td>XXX10</td><td>A data value which was interpolated.</td></tr> <tr> <td>Reserved</td><td>XXX11</td><td>Undefined.</td></tr> <tr> <td>Partial</td><td>XX1XX</td><td>A data value which was calculated with an incomplete interval.</td></tr> <tr> <td>Extra Data</td><td>X1XXX</td><td>A raw data value that hides other data at the same timestamp.</td></tr> <tr> <td>Multi Value</td><td>1XXXX</td><td>Multiple values match the <i>Aggregate</i> criteria (i.e. multiple minimum values at different timestamps within the same interval).</td></tr> </table> IEC 62541-11 describes how these bits are used in more detail.	Raw	XXX00	A raw data value.	Calculated	XXX01	A data value which was calculated.	Interpolated	XXX10	A data value which was interpolated.	Reserved	XXX11	Undefined.	Partial	XX1XX	A data value which was calculated with an incomplete interval.	Extra Data	X1XXX	A raw data value that hides other data at the same timestamp.	Multi Value	1XXXX	Multiple values match the <i>Aggregate</i> criteria (i.e. multiple minimum values at different timestamps within the same interval).
Raw	XXX00	A raw data value.																					
Calculated	XXX01	A data value which was calculated.																					
Interpolated	XXX10	A data value which was interpolated.																					
Reserved	XXX11	Undefined.																					
Partial	XX1XX	A data value which was calculated with an incomplete interval.																					
Extra Data	X1XXX	A raw data value that hides other data at the same timestamp.																					
Multi Value	1XXXX	Multiple values match the <i>Aggregate</i> criteria (i.e. multiple minimum values at different timestamps within the same interval).																					

7.34.2 Common StatusCodes

Table 177 defines the common *StatusCodes* for all *Service* results used in more than one service. It does not provide a complete list. These *StatusCodes* may also be used as operation level result code. IEC 62541-6 maps the symbolic names to a numeric value and provides a complete list of *StatusCodes* including codes defines in other parts.

Table 177 – Common Service result codes

Symbolic Id	Description
Good	The operation was successful.
Good_CompletesAsynchronously	The processing will complete asynchronously.
Good_SubscriptionTransferred	The subscription was transferred to another session.
Bad_CertificateHostNameInvalid	The <i>HostName</i> used to connect to a <i>Server</i> does not match a <i>HostName</i> in the <i>Certificate</i> .
Bad_CertificateChainIncomplete	The <i>Certificate</i> chain is incomplete.
Bad_CertificateIssuerRevocationUnknown	It was not possible to determine if the <i>Issuer Certificate</i> has been revoked.
Bad_CertificateIssuerUseNotAllowed	The <i>Issuer Certificate</i> may not be used for the requested operation.
Bad_CertificateIssuerTimeInvalid	An <i>Issuer Certificate</i> has expired or is not yet valid.
Bad_CertificateIssuerRevoked	The <i>Issuer Certificate</i> has been revoked.
Bad_CertificateInvalid	The <i>Certificate</i> provided as a parameter is not valid.
Bad_CertificateRevocationUnknown	It was not possible to determine if the <i>Certificate</i> has been revoked.
Bad_CertificateRevoked	The <i>Certificate</i> has been revoked.
Bad_CertificateTimeInvalid	The <i>Certificate</i> has expired or is not yet valid.
Bad_CertificateUriInvalid	The URI specified in the <i>ApplicationDescription</i> does not match the URI in the <i>Certificate</i> .
Bad_CertificateUntrusted	The <i>Certificate</i> is not trusted.
Bad_CertificateUseNotAllowed	The <i>Certificate</i> may not be used for the requested operation.
Bad_CommunicationError	A low level communication error occurred.
Bad_DataTypeIdUnknown	The <i>ExtensionObject</i> cannot be (de)serialized because the data type id is not recognized.
Bad_DecodingError	Decoding halted because of invalid data in the stream.
Bad_EncodingError	Encoding halted because of invalid data in the objects being serialized.
Bad_EncodingLimitsExceeded	The message encoding/decoding limits imposed by the <i>Communication Stack</i> have been exceeded.
Bad_IdentityTokenInvalid	The user identity token is not valid.
Bad_IdentityTokenRejected	The user identity token is valid but the <i>Server</i> has rejected it.
Bad_InternalError	An internal error occurred as a result of a programming or configuration error.
Bad_InvalidArgument	One or more arguments are invalid. Each service defines parameter-specific <i>StatusCodes</i> and these <i>StatusCodes</i> shall be used instead of this general error code. This error code shall be used only by the <i>Communication Stack</i> and in services where it is defined in the list of valid <i>StatusCodes</i> for the service.
Bad_InvalidState	The operation cannot be completed because the object is closed, uninitialized or in some other invalid state.
Bad_InvalidTimestamp	The timestamp is outside the range allowed by the <i>Server</i> .
Bad_LicenseExpired	The UA Server requires a license to operate in general or to perform a service or operation, but the existing license is expired.
Bad_LicenseLimitsExceeded	The UA Server has limits on number of allowed operations / objects, based on installed licenses, and these limits were exceeded.
Bad_LicenseNotAvailable	The UA Server does not have a license which is required to operate in general or to perform a service or operation.
Bad_NonceInvalid	The nonce does appear to be not a random value or it is not the correct length.
Bad_NothingToDo	There was nothing to do because the <i>Client</i> passed a list of operations with no elements.

Symbolic Id	Description
Bad_OutOfMemory	Not enough memory to complete the operation.
Bad_RequestCancelledByClient	The request was cancelled by the client.
Bad_RequestTooLarge	The request message size exceeds limits set by the <i>Server</i> .
Bad_ResponseTooLarge	The response message size exceeds limits set by the client.
Bad_RequestHeaderInvalid	The header for the request is missing or invalid.
Bad_ResourceUnavailable	An operating system resource is not available.
Bad_SecureChannelIdInvalid	The specified secure channel is no longer valid.
Bad_SecurityChecksFailed	An error occurred while verifying security.
Bad_ServerHalted	The <i>Server</i> has stopped and cannot process any requests.
Bad_ServerNotConnected	The operation could not complete because the <i>Client</i> is not connected to the <i>Server</i> .
Bad_ServerUriInvalid	The <i>Server</i> URI is not valid.
Bad_ServiceUnsupported	The <i>Server</i> does not support the requested service.
Bad_SessionIdInvalid	The <i>Session</i> id is not valid.
Bad_SessionClosed	The <i>Session</i> was closed by the client.
Bad_SessionNotActivated	The <i>Session</i> cannot be used because <i>ActivateSession</i> has not been called.
Bad_Shutdown	The operation was cancelled because the application is shutting down.
Bad_SubscriptionIdInvalid	The subscription id is not valid.
Bad_Timeout	The operation timed out.
Bad_TimestampsToReturnInvalid	The timestamps to return parameter is invalid.
Bad_TooManyOperations	The request could not be processed because it specified too many operations.
Bad_UnexpectedError	An unexpected error occurred.
Bad_UnknownResponse	An unrecognized response was received from the <i>Server</i> .
Bad_UserAccessDenied	User does not have permission to perform the requested operation.
Bad_ViewIdUnknown	The view id does not refer to a valid view Node.
Bad_ViewTimestampInvalid	The view timestamp is not available or not supported.
Bad_ViewParameterMismatchInvalid	The view parameters are not consistent with each other.
Bad_ViewVersionInvalid	The view version is not available or not supported.

Table 178 defines the common *StatusCodes* for all operation level results used in more than one service. It does not provide a complete list. IEC 62541-6 maps the symbolic names to a numeric value and provides a complete list of *StatusCodes* including codes defines in other parts. The common *Service* result codes can be also contained in the operation level.

Table 178 – Common Operation Level result codes

Symbolic Id	Description
Good_Clamped	The value written was accepted but was clamped.
Good_Overload	Sampling has slowed down due to resource limitations.
Uncertain	The value is uncertain but no specific reason is known.
Bad	The value is bad but no specific reason is known.
Bad_AttributeIdInvalid	The attribute is not supported for the specified node.
Bad_BrowseDirectionInvalid	The browse direction is not valid.
Bad_BrowseNameInvalid	The browse name is invalid.
Bad_ContentFilterInvalid	The content filter is not valid.
Bad_ContinuationPointInvalid	The continuation point provided is no longer valid. This status is returned if the continuation point was deleted or the address space was changed between the browse calls.
Bad_DataEncodingInvalid	The data encoding is invalid. This result is used if no <i>dataEncoding</i> can be applied because an <i>Attribute</i> other than <i>Value</i> was requested or the <i>DataType</i> of the <i>Value Attribute</i> is not a subtype of the <i>Structure DataType</i> .
Bad_DataEncodingUnsupported	The <i>Server</i> does not support the requested data encoding for the node. This result is used if a <i>dataEncoding</i> can be applied but the passed data encoding is not known to the <i>Server</i> .
Bad_EventFilterInvalid	The event filter is not valid.
Bad_FilterNotAllowed	A monitoring filter cannot be used in combination with the attribute specified.
Bad_FilterOperandInvalid	The operand used in a content filter is not valid.
Bad_HistoryOperationInvalid	The history details parameter is not valid.
Bad_HistoryOperationUnsupported	The <i>Server</i> does not support the requested operation.
Bad_IndexRangeInvalid	The syntax of the index range parameter is invalid.
Bad_IndexRangeNoData	No data exists within the range of indexes specified.
Bad_MonitoredItemFilterInvalid	The monitored item filter parameter is not valid.
Bad_MonitoredItemFilterUnsupported	The <i>Server</i> does not support the requested monitored item filter.
Bad_MonitoredItemIdInvalid	The monitoring item id does not refer to a valid monitored item.
Bad_MonitoringModeInvalid	The monitoring mode is invalid.
Bad_NoCommunication	Communication with the data source is defined, but not established, and there is no last known value available. This status/sub-status is used for cached values before the first value is received or for Write and Call if the communication is not established.
Bad_NoContinuationPoints	The operation could not be processed because all continuation points have been allocated.
Bad_NodeClassInvalid	The node class is not valid.
Bad_NodeIdInvalid	The syntax of the node id is not valid.
Bad_NodeIdUnknown	The node id refers to a node that does not exist in the <i>Server</i> address space.
Bad_NoDeleteRights	The <i>Server</i> will not allow the node to be deleted.
Bad_NodeNotInView	The <i>nodesToBrowse</i> is not part of the view.
Bad_NotFound	A requested item was not found or a search operation ended without success.
Bad_NotImplemented	Requested operation is not implemented.
Bad_NotReadable	The access level does not allow reading or subscribing to the <i>Node</i> .

Symbolic Id	Description
Bad_NotSupported	The requested operation is not supported.
Bad_NotWritable	The access level does not allow writing to the <i>Node</i> .
Bad_ObjectDeleted	The <i>Object</i> cannot be used because it has been deleted.
Bad_OutOfRange	The value was out of range.
Bad_ReferenceTypeIdInvalid	The reference type id does not refer to a valid reference type node.
Bad_SecurityModeInsufficient	The <i>SecurityPolicy</i> and/or <i>MessageSecurityMode</i> do not match the <i>Server</i> requirements to complete the operation. For example, a user may have the right to receive the data but the data can only be transferred through an encrypted channel with an appropriate <i>SecurityPolicy</i> .
Bad_SourceNodeIdInvalid	The source node id does not refer to a valid node.
Bad_StructureMissing	A mandatory structured parameter was missing or null.
Bad_TargetNodeIdInvalid	The target node id does not refer to a valid node.
Bad_TypeDefinitionInvalid	The type definition node id does not reference an appropriate type node.
Bad_TypeMismatch	The value supplied for the attribute is not of the same type as the attribute's value.
Bad_WaitingForInitialData	Waiting for the <i>Server</i> to obtain values from the underlying data source. After creating a <i>MonitoredItem</i> or after setting the <i>MonitoringMode</i> from <i>DISABLED</i> to <i>REPORTING</i> or <i>SAMPLING</i> , it may take some time for the <i>Server</i> to actually obtain values for these items. In such cases the <i>Server</i> can send a <i>Notification</i> with this status prior to the <i>Notification</i> with the first value or status from the data source.

7.35 TimestampsToReturn

The *TimestampsToReturn* is an enumeration that specifies the *Timestamp Attributes* to be transmitted for *MonitoredItems* or *Nodes* in *Read* and *HistoryRead*. The values of this parameter are defined in Table 179.

Table 179 – TimestampsToReturn values

Value	Description
SOURCE_0	Return the source timestamp.
SERVER_1	Return the <i>Server</i> timestamp.
BOTH_2	Return both the source and <i>Server</i> timestamps.
NEITHER_3	Return neither timestamp. This is the default value for <i>MonitoredItems</i> if a <i>Variable</i> value is not being accessed. For <i>HistoryRead</i> this is not a valid setting.

7.36 UserIdentityToken parameters

7.36.1 Overview

The *UserIdentityToken* structure used in the *Server Service Set* allows *Clients* to specify the identity of the user they are acting on behalf of. The exact mechanism used to identify users depends on the system configuration. The different types of identity tokens are based on the most common mechanisms that are used in systems today. Table 180 defines the current set of user identity tokens. The *ExtensibleParameter* type is defined in 7.12.

Table 180 – UserIdentityToken parameterTypeIds

Symbolic Id	Description
AnonymousIdentityToken	No user information is available.
UserNameIdentityToken	A user identified by user name and password.
X509IdentityToken	A user identified by an X.509 v3 Certificate.
IssuedIdentityToken	A user identified by a ~~WS-SecurityToken~~ token issued by an external Authorization Service.

7.36.2 Token Encryption and Proof of Possession

7.36.2.1 Overview

The *Client* shall always prove possession of a *UserIdentityToken* when it passes it to the *Server*. Some tokens include a secret such as a password which the *Server* will accept as proof. In order to protect these secrets the *Token* may be encrypted before it is passed to the *Server*. ~~The profiles defined in IEC 62541-7 specify where encryption shall be applied.~~ Other types of tokens allow the *Client* to create a signature with the secret associated with the *Token*. In these cases, the *Client* proves possession of a *UserIdentityToken* by ~~appending the last ServerNonce to the ServerCertificate and uses the secret to produce a signature which is passed to the Server~~ creating a signature with the secret and passing it to the *Server*.

Each *UserIdentityToken* allowed by an *Endpoint* shall have a *UserTokenPolicy* specified in the *EndpointDescription*. The *UserTokenPolicy* specifies what *SecurityPolicy* to use when encrypting or signing. If this *SecurityPolicy* is omitted then the *Client* uses the *SecurityPolicy* in the *EndpointDescription*. If the matching *SecurityPolicy* is set to *None* then no encryption or signature is required. The possible *SecurityPolicies* are defined in IEC 62541-7.

It is recommended that applications never set the *SecurityPolicy* to *None* for ~~UserTokens~~ *UserIdentityTokens* that include a secret because these secrets could be used by an attacker to gain access to the system.

Clients shall validate the *Server Certificate* and ensure it is trusted before sending a *UserIdentityToken* encrypted with the *Certificate*.

The encrypted secret and *Signature* are embedded in a *ByteString* which is part of the *UserIdentityToken*. The format of this *ByteString* depends on the type of *UserIdentityToken* and the *SecurityPolicy*.

The legacy token secret format defined in 7.36.2.2 is not extensible and provides only encryption but the encrypted data is not signed. It is used together with the *USERNAME_1 UserIdentityToken*. The password secret exchanged with this format shall not exceed 64 bytes.

The *EncryptedSecret* format defined in 7.36.2.3 provides an extensible secret format together with the definition how the secret is signed and encrypted. It allows for the layout to be updated as new token types are defined or new *SecurityPolicies* are added.

The *UserIdentityToken* types and the token formats supported by the *Endpoint* are identified by the *UserTokenPolicy* defined in 7.37.

To prevent the leakage of information useful to attackers, *Servers* shall ensure that the process of validating *UserIdentityTokens* completes in a fixed interval independent of whether an error occurs or not. The process of validation includes decrypting, check for padding and checking for a valid nonce. If any errors occur the return code is *Bad_IdentityTokenInvalid*.

Servers shall log details of any failure to validate a *UserIdentityToken* and shall lock out *Client* applications after five failures.

7.36.2.2 Legacy Encrypted Token Secret Format

When encrypting a *UserIdentityToken*, the Client appends the last *ServerNonce* to the secret. The data is then encrypted with the public key from the *Server's Certificate*.

A *Client* should not add any padding after the secret. If a *Client* adds padding then all bytes shall be zero. A *Server* shall check for padding added by *Clients* and ensure that all padding bytes are zeros. *Servers* shall reject *UserIdentityTokens* with invalid padding. Administrators shall be able to configure *Servers* to accept *UserIdentityTokens* with invalid padding.

If no encryption is applied, the structure is not used and only the secret without any *Nonce* is passed to the *Server*.


Table 181 describes how to serialize *UserIdentityTokens* before applying encryption.

Table 181 – Legacy UserIdentityToken Encrypted Token Secret Format

Name	Type	Description
Length	Byte [4]	The length of the ~~data to be encrypted~~ ~~data~~ including the <i>ServerNonce</i> but excluding the <i>length</i> field. This field is a 4-byte unsigned integer encoded with the least significant bytes appearing first.
tokenData	Byte [*]	The token data.
serverNonce	Byte [*]	The last <i>ServerNonce</i> returned by the <i>Server</i> in the <i>CreateSession</i> or <i>ActivateSession</i> response.

7.36.2.3 EncryptedSecret Format

The *EncryptedSecret* uses an extensible format which has the *TypeId* of a *DataType Node* as a prefix as defined for the *ExtensionObject* encoding in IEC 62541-6. The general layout of the *EncryptedSecret* is shown in Figure 39.



IEC

Figure 39 – EncryptedSecret layout

The *TypeId* specifies how the *EncryptedSecret* is serialized and secured. For example, *RsaEncryptedSecrets* require that the policy header be encrypted with the public key associated with the *EncryptingCertificate* before it is serialized. Other *TypeIds* may or may not require the policy header to be encrypted.

The *SecurityPolicyUri* is used to determine what algorithms were used to encrypt and sign the data.

The payload is always encrypted using the symmetric encryption algorithm specified by the *SecurityPolicyUri*. The *EncryptingKey* and *InitializationVector* are used to initialize this algorithm. The mechanisms used to initialize the symmetric encryption algorithm depend on the *TypeId*. The lengths of the fields are specified by the *SecurityPolicyUri*.

The *Signature* is used for the symmetric signature algorithm specified by the *SecurityPolicyUri*. The *SymmetricKey* and *InitializationVector* are used to initialize this algorithm. The mechanisms used to initialize the symmetric signature algorithm depend on the *TypeId*. The length of the *SymmetricKey* is specified by the *SecurityPolicyUri*.

The *EncryptedSecret* is secured and serialized as follows:

- serialize the common header;
- serialize the policy header;
- encrypt the policy header and append the result to the common header;
- update the *PolicyHeaderLength* with the length of the encrypted header;
- append the *Nonce* and the *Secret* to the encrypted policy header;
- calculate padding required on the payload and append after the *Secret*;
- encrypt the payload;
- Calculate a *Signature* on {common header | encrypted policy header | encrypted payload};
- append the *Signature*.

Individual fields are serialized using the UA Binary encoding (see IEC 62541-6) for the *DataType* specified in Table 182. The *Padding* is used to ensure there is enough data to fill an integer multiple of encryption blocks. The size of the encryption block depends on the encryption algorithm. Two separate padding operations are needed because two different encryption algorithms may be used. The total length of the *Padding*, not including the *PaddingSize*, is encoded as a *UInt16*. The individual bytes of the *Padding* are set to the least significant byte of the *PaddingSize*.

The *EncryptedSecret* is deserialized and validated as follows:

- deserialize the common header;
- decrypt the policy header using the private key associated with the *EncryptingCertificate*;
- verify the padding in the policy header;
- verify the *Signature* with the *SigningKey*;
- decrypt the payload;
- verify the padding on the payload;
- extract the *Secret*;

If the *TypeId* does not require the policy header to be encrypted then the padding on the policy header is omitted and the *PolicyHeaderLength* specifies the length of the unencrypted data.

The fields in the *EncryptedSecret* are described in Table 182. The first three fields *TypeId*, *EncodingMask* and *Length* belong to the *ExtensionObject* encoding defined in IEC 62541-6.

Table 182 – EncryptedSecret layout

Name	Type	Description
TypeId	NodeId	The <i>NodeId</i> of the <i>DataType Node</i> .
EncodingMask	Byte	This value is always 1.
Length	Int32	The length of the data that follows including the <i>Signature</i> .
SecurityPolicyUri	String	The URI for the <i>SecurityPolicy</i> used to apply security.
EncryptingCertificate	ByteString	The SHA1 thumbprint of the DER form of the encrypting <i>Certificate</i> .
SigningTime	DateTime	When the <i>Signature</i> was created.
PolicyHeaderLength	UInt16	The length of the policy header that follows. If the policy header is encrypted this is the length of the encrypted data. Otherwise, it is the length of the unencrypted data.
SigningKey	ByteString	The key data used to create the <i>Signature</i> . The <i>TypeId</i> specifies how this data is used.
EncryptingKey	ByteString	The key data used to encrypt the payload. The <i>TypeId</i> specifies how this data is used.
InitializationVector	ByteString	The data used to initialize the algorithm used to encrypt the payload. The <i>TypeId</i> specifies how this data is used.
Nonce	ByteString	This is the last <i>serverNonce</i> returned in the <i>CreateSession</i> or <i>ActivateSession Response</i> when a <i>UserIdentityToken</i> is passed with the <i>ActivateSession Request</i> . If used outside of an <i>ActivateSession</i> call, the source of the <i>Nonce</i> is defined by the context in which this <i>EncryptedSecret</i> is used. The length of the <i>Nonce</i> shall equal the <i>SecureChannelNonceLength</i> specified by the <i>SecurityPolicy</i> .
Secret	ByteString	The <i>tokenData</i> that depends on the <i>IssuedIdentityToken</i> . If the <i>tokenData</i> is a <i>String</i> it is encoded using UTF-8 first.
PayloadPadding	Byte[*]	Additional padding added to ensure the size of the encrypted payload is an integer multiple of the input block size for the symmetric encryption algorithm specified by the <i>SecurityPolicyUri</i> . The value of each byte is the least significant byte of the <i>PayloadPaddingSize</i> .
PayloadPaddingSize	UInt16	The size of the padding added to the payload.
Signature	Byte[*]	The signature calculated using the symmetric signing algorithm specified by the <i>SecurityPolicyUri</i> . The length of the signature is specified by the <i>SecurityPolicyUri</i> .

The currently available *EncryptedSecret DataTypes* are defined in Table 183.

Table 183 – EncryptedSecret DataTypes

Type Name	When to Use
RsaEncryptedSecret	Used when the <i>SecurityPolicy</i> requires the use of RSA based asymmetric encryption. It is described in 7.36.2.4.

7.36.2.4 RsaEncryptedSecret DataType

The *RsaEncryptedSecret* uses RSA based asymmetric encryption to encrypt the policy header.

Additional semantics for the fields in the *EncryptedSecret* layout for the *RsaEncryptedSecret Structure* are described in Table 184.

Table 184 – RsaEncryptedSecret structure

Name	Type	Description
TypeId	NodeId	The <i>NodeId</i> of the <i>RsaEncryptedSecret DataType Node</i> .
EncodingMask	Byte	See Table 182.
Length	UInt32	See Table 182.
SecurityPolicyUri	String	See Table 182.
EncryptingCertificate	ByteString	See Table 182.
SigningTime	DateTime	See Table 182.
PolicyHeaderLength	UInt16	See Table 182.
SigningKey	ByteString	The key used to compute the <i>Signature</i> . See Table 182 for additional details.
EncryptingKey	ByteString	The key used to encrypt payload. See Table 182 for additional details.
InitializationVector	ByteString	The initialization vector used with the <i>EncryptingKey</i> . See Table 182 for additional details.
Nonce	ByteString	See Table 182.
Secret	ByteString	See Table 182.
PayloadPadding	Byte[*]	See Table 182.
PayloadPaddingSize	UInt16	See Table 182.
Signature	Byte[*]	See Table 182.

7.36.3 AnonymousIdentityToken

The *AnonymousIdentityToken* is used to indicate that the Client has no user credentials.

Table 185 defines the *AnonymousIdentityToken* parameter.

Table 185 – AnonymousIdentityToken

Name	Type	Description
AnonymousIdentityToken	Structure	An anonymous user identity.
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.

7.36.4 UserNameIdentityToken

The *UserNameIdentityToken* is used to pass simple username/password credentials to the *Server*.

This token shall be encrypted by the *Client* if required by the ~~*Server in the*~~ *SecurityPolicy* of the *UserTokenPolicy* ~~*in the EndpointDescription*~~. The *Server* should specify a *SecurityPolicy* for the *UserTokenPolicy* if the *SecureChannel* has a *SecurityPolicy* of *None* and no transport layer encryption is available. If *None* is specified for the *UserTokenPolicy* and *SecurityPolicy* is *None* then the password only contains the UTF-8 encoded password. The *SecurityPolicy* of the *SecureChannel* is used if no *SecurityPolicy* is specified in the *UserTokenPolicy*.

If the token is to be encrypted the password shall be converted to a UTF-8 *ByteString*, encrypted and then serialized as shown in Table 181.

The *Server* shall decrypt the password and verify the *ServerNonce*.

If the *SecurityPolicy* is *None* then the password only contains the UTF-8 encoded password. This configuration should not be used unless the network is encrypted in some other manner such as a VPN. The use of this configuration without network encryption would result in a serious security fault, in that it would cause the appearance of a secure user access, but it would make the password visible in clear text.

~~If the token is encrypted the password shall be converted to a UTF-8 *ByteString* and then serialized as shown in Table 169.~~

~~The *Server* shall decrypt the password and verify the *ServerNonce*.~~

~~If the *SecurityPolicy* is *None* then the password only contains the UTF-8 encoded password.~~

Table 186 defines the *UserNameIdentityToken* parameter.

Table 186 – UserNameIdentityToken

Name	Type	Description
UserNameIdentityToken	Structure	UserName value.
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.
userName	String	A string that identifies the user.
password	ByteString	The password for the user. The password can be an empty string. This parameter shall be encrypted with the ~~Server's Certificate~~ Server's public key using the algorithm specified by the <i>SecurityPolicy</i> . The format used for the encrypted data is described in 7.36.2.2.
encryptionAlgorithm	String	A string containing the URI of the <i>AsymmetricEncryptionAlgorithm</i> . The URI string values are defined names that may be used as part of the security profiles specified in IEC 62541-7. This parameter is null if the password is not encrypted.

Table 187 describes the dependencies for selecting the *AsymmetricEncryptionAlgorithm* for the *UserNameIdentityToken*. The *SecureChannel SecurityPolicy* URI is specified in the *EndpointDescription* and used in subsequent *OpenSecureChannel* requests. The *UserTokenPolicy SecurityPolicy* URI is specified in the *EndpointDescription*. The *encryptionAlgorithm* is specified in the *UserNameIdentityToken* or *IssuedIdentityToken* provided by the *Client* in the *ActivateSession* call. The *SecurityPolicy* Other in Table 187 refers to any *SecurityPolicy* other than *None*. The selection of the *EncryptionAlgorithm* is based on the *UserTokenPolicy*. The *SecureChannel SecurityPolicy* is used if the *UserTokenPolicy* is null or empty.

Table 187 – EncryptionAlgorithm selection

SecureChannel SecurityPolicy	UserTokenPolicy SecurityPolicy	UserNamelIdentityToken UserIdentityToken EncryptionAlgorithm
Security Policy - None	Null or empty	Null No encryption
Security Policy - None	Security Policy - None	Null No encryption
Security Policy - None	Security Policy - Other	Asymmetric algorithm for "Other"
Security Policy - Other	Null or empty	Asymmetric algorithm for "Other"
Security Policy - Other	Security Policy - Yet another	Asymmetric algorithm for "Yet another"
Security Policy - Other	Security Policy - Other	Asymmetric algorithm for "Other"
Security Policy - Other	Security Policy - None	Null No encryption

7.36.5 X509IdentityTokens

The *X509IdentityToken* is used to pass an X.509 v3 *Certificate* which is issued by the user.

This token shall always be accompanied by a *Signature* in the *userTokenSignature* parameter of *ActivateSession* if required by the *SecurityPolicy*. The *Server* should specify a *SecurityPolicy* for the *UserTokenPolicy* if the *SecureChannel* has a *SecurityPolicy* of None.

Table 188 defines the *X509IdentityToken* parameter.

Table 188 – X.509 v3 Identity Token

Name	Type	Description
X509IdentityToken	structure	X.509 v3 value.
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.
certificateData	ByteString	The X.509 v3 <i>Certificate</i> in DER format.

7.36.6 IssuedIdentityToken

~~The *IssuedIdentityToken* is used to pass WS Security compliant *SecurityTokens* to the *Server*.~~

~~WS Security defines a number of token profiles that may be used to represent different types of *SecurityTokens*. For example, Kerberos and SAML tokens have WSS token profiles and shall be exchanged in OPC UA as XML Security Tokens.~~

~~The WSS X509 and UserName tokens should not be exchanged as XML security tokens. OPC UA applications should use the appropriate OPC UA identity tokens to pass the information contained in these types of WSS *SecurityTokens*.~~

~~These tokens may be encrypted or require a signature. IEC 62541-7 defines profiles that include user related security; they also include any requirements for encryption and signatures. Additional security profiles specify encryption and signature algorithms.~~

~~If the token is encrypted then the XML shall be converted to an UTF8 *ByteString* and then serialized as shown in Table 169.~~

~~The *Server* shall decrypt the *tokenData* and verify the *ServerNonce*.~~

~~If the *SecurityPolicy* is *None* or if the token only requires signing then the *tokenData* contains the UTF-8 encoded XML representation of the token.~~

The *IssuedIdentityToken* is used to pass *SecurityTokens* issued by an external *Authorization Service* to the *Server*. These tokens may be text or binary.

OAuth2 defines a standard for *Authorization Services* that produce JSON Web Tokens (JWT). These JWTs are passed as an *Issued Token* to an OPC UA Server which uses the signature contained in the JWT to validate the token. IEC 62541-6 describes OAuth2 and JWTs in more detail. If the token is encrypted, it shall use the *EncryptedSecret* format defined in 7.36.2.3.

This token shall be encrypted by the *Client* if required by the *SecurityPolicy* of the *UserTokenPolicy*. The *Server* should specify a *SecurityPolicy* for the *UserTokenPolicy* if the *SecureChannel* has a *SecurityPolicy* of *None* and no transport layer encryption is available. The *SecurityPolicy* of the *SecureChannel* is used If no *SecurityPolicy* is specified in the *UserTokenPolicy*.

If the *SecurityPolicy* is not *None*, the *tokenData* shall be encoded in UTF-8 (if it is not already binary), signed and encrypted according the rules specified for the *tokenType* of the associated *UserTokenPolicy* (see 7.37).

If the *SecurityPolicy* is *None* then the *tokenData* only contains the UTF-8 encoded *tokenData*. This configuration should not be used unless the network is encrypted in some other manner such as a VPN. The use of this configuration without network encryption would result in a serious security fault, in that it would cause the appearance of a secure user access, but it would make the token visible in clear text.

Table 189 defines the *IssuedIdentityToken* parameter.

Table 189 – IssuedIdentityToken

Name	Type	Description
IssuedIdentityToken	structure	~~WSS value.~~ The token provided by an <i>Authorization Service</i> .
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.
tokenData	ByteString	The ~~XML~~ text or binary representation of the token. ~~This parameter may be encrypted with the Server's Certificate.~~ The format of the data depends on the associated <i>UserTokenPolicy</i> .
encryptionAlgorithm	String	~~A string containing~~ The URI of the <i>AsymmetricEncryptionAlgorithm</i> . The list of OPC UA-defined names that may be used is specified in IEC 62541-7. See Table 187 for details on picking the correct URI. ~~This parameter is null if the token is not encrypted.~~ ~~The encryptionAlgorithm is AsymmetricEncryptionAlgorithm for the security-policy identified by the securityPolicyUri for the UserTokenPolicy.~~ This parameter is null if the <i>tokenData</i> is not encrypted or if the <i>EncryptedSecret</i> format is used.

7.37 UserTokenPolicy

The components of this parameter are defined in Table 190.

Table 190 – UserTokenPolicy

Name	Type	Description
UserTokenPolicy	structure	Specifies a <i>UserIdentityToken</i> that a <i>Server</i> will accept.
policyId	String	An identifier for the UserTokenPolicy assigned by the <i>Server</i> . The <i>Client</i> specifies this value when it constructs a <i>UserIdentityToken</i> that conforms to the policy. This value is only unique within the context of a single <i>Server</i> .
tokenType	Enum <i>UserIdentityTokenType</i>	The type of user identity token required. This value is an enumeration with one of the following values: ANONYMOUS_0 No token is required. USERNAME_1 A username/password token. CERTIFICATE_2 An X.509 v3 <i>Certificate</i> token. ISSUEDTOKEN_3 Any ~~WS-Security defined~~ token issued by an <i>Authorization Service</i> . A <i>tokenType</i> of ANONYMOUS indicates that the <i>Server</i> does not require any user identification. In this case, the <i>Client Application Instance Certificate</i> is used as the user identification.
issuedTokenType	String	A URI for the type of token. IEC 62541-6 defines URIs for common issued token types. Vendors may specify their own token types. This field may only be specified if <i>TokenType</i> is ISSUEDTOKEN_3.
issuerEndpointUrl	String	An optional ~~URL for the token issuing service~~ string which depends on the <i>Authorization Service</i> . The meaning of this value depends on the <i>issuedTokenType</i> . Further details for the different token types are defined in IEC 62541-6. For Kerberos this string is the name of the Service Principal Name (SPN). For JWTs this is a JSON object with fields defined in IEC 62541-6.
securityPolicyUri	String	The security policy to use when encrypting or signing the <i>UserIdentityToken</i> when it is passed to the <i>Server</i> in the <i>ActivateSession</i> request. Clause 7.36 describes how this parameter is used. The security policy for the <i>SecureChannel</i> is used if this value is ~~omitted~~ null or empty.

7.38 VersionTime

This primitive data type is a UInt32 that represents the time in seconds since the year 2000. The epoch date is midnight UTC (00:00) on January 1, 2000.

It is used as version number based on the last change time. If the version is updated, the new value shall be greater than the previous value.

If a *Variable* is initialized with a *VersionTime* value, the value shall be either loaded from persisted configuration or time synchronization shall be available to ensure a unique version is applied.

The value 0 is used to indicate that no version information is available.

7.39 ViewDescription

The components of this parameter are defined in Table 191.

Table 191 – ViewDescription

Name	Type	Description
ViewDescription	structure	Specifies a <i>View</i> .
viewId	NodeId	<i>NodeId</i> of the <i>View</i> to <i>Query</i> . A null value indicates the entire <i>AddressSpace</i> .
timestamp	UtcTime	The time date desired. The corresponding version is the one with the closest previous creation timestamp. Either the <i>Timestamp</i> or the <i>viewVersion</i> parameter may be set by a <i>Client</i> , but not both. If <i>ViewVersion</i> is set this parameter shall be null.
viewVersion	UInt32	The version number for the <i>View</i> desired. When <i>Nodes</i> are added to or removed from a <i>View</i> , the value of a <i>View</i> 's <i>ViewVersion Property</i> is updated. Either the <i>Timestamp</i> or the <i>viewVersion</i> parameter may be set by a <i>Client</i> , but not both. The <i>ViewVersion Property</i> is defined in IEC 62541-3. If <i>timestamp</i> is set this parameter shall be 0. The current view is used if <i>timestamp</i> is null and <i>viewVersion</i> is 0.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Annex A (informative)

BNF definitions

A.1 Overview over BNF

The BNF (Backus-Naur form) used in this annex uses '<' and '>' to mark symbols, '[' and ']' to identify optional paths and '|' to identify alternatives. If the '(' and ')' symbols are used, it indicates sets.

A.2 BNF of RelativePath

A *RelativePath* is a structure that describes a sequence of *References* and *Nodes* to follow. This annex describes a text format for a *RelativePath* that can be used in documentation or in files used to store configuration information.

The components of a *RelativePath* text format are specified in Table A.1.

Table A.1 – RelativePath

Symbol	Meaning										
/	The forward slash character indicates that the <i>Server</i> is to follow any subtype of <i>HierarchicalReferences</i> .										
.	The period (dot) character indicates that the <i>Server</i> is to follow any subtype of a <i>Aggregates ReferenceType</i> .										
<[#!ns:]ReferenceType>	A string delimited by the '<' and '>' symbols specifies the <i>BrowseName</i> of a <i>ReferenceType</i> to follow. By default, any <i>References</i> of the subtypes the <i>ReferenceType</i> are followed as well. A '#' placed in front of the <i>BrowseName</i> indicates that subtypes should not be followed. A '!' in front of the <i>BrowseName</i> is used to indicate that the inverse <i>Reference</i> should be followed. The <i>BrowseName</i> may be qualified with a namespace index (indicated by a numeric prefix followed by a colon). This namespace index is used specify the namespace component of the <i>BrowseName</i> for the <i>ReferenceType</i>. If the namespace prefix is omitted then namespace index 0 is used.										
[ns:]BrowseName	A string that follows a '/', '.', or '>' symbol specifies the <i>BrowseName</i> of a target <i>Node</i> to return or follow. This <i>BrowseName</i> may be prefixed by its namespace index. If the namespace prefix is omitted then namespace index 0 is used. Omitting the final <i>BrowseName</i> from a path is equivalent to a wildcard operation that matches all <i>Nodes</i> which are the target of the <i>Reference</i> specified by the path.										
&	The & sign character is the escape character. It is used to specify reserved characters that appear within a <i>BrowseName</i>. A reserved character is escaped by inserting the '&' in front of it. Examples of <i>BrowseNames</i> with escaped characters are: <table> <tr> <th>Received browse path name</th><th>Resolves to</th></tr> <tr> <td>"&/Name_1"</td><td>"/Name_1"</td></tr> <tr> <td>"&.Name_2"</td><td>".Name_2"</td></tr> <tr> <td>"&:Name_3"</td><td>".Name_3"</td></tr> <tr> <td>"&&Name_4"</td><td>"&Name_4"</td></tr> </table>	Received browse path name	Resolves to	"&/Name_1"	"/Name_1"	"&.Name_2"	".Name_2"	"&:Name_3"	".Name_3"	"&&Name_4"	"&Name_4"
Received browse path name	Resolves to										
"&/Name_1"	"/Name_1"										
"&.Name_2"	".Name_2"										
"&:Name_3"	".Name_3"										
"&&Name_4"	"&Name_4"										

Table A.2 provides *RelativePaths* examples in text format.

Table A.2 – RelativePath Examples

Browse Path	Description
"/2:Block&.Output"	Follows any forward hierarchical <i>Reference</i> with target <i>BrowseName</i> = "2:Block.Output".
"/3:Truck.0:NodeVersion"	Follows any forward hierarchical <i>Reference</i> with target <i>BrowseName</i> = "3:Truck" and from there a forward <i>Aggregates Reference</i> to a target with <i>BrowseName</i> "0:NodeVersion".
"<1:ConnectedTo>1:Boiler/1:HeatSensor"	Follows any forward <i>Reference</i> with a <i>BrowseName</i> = '1:ConnectedTo' and finds targets with <i>BrowseName</i> = '1:Boiler'. From there follows any hierarchical <i>Reference</i> and find targets with <i>BrowseName</i> = '1:HeatSensor'.
"<1:ConnectedTo>1:Boiler/"	Follows any forward <i>Reference</i> with a <i>BrowseName</i> = '1:ConnectedTo' and finds targets with <i>BrowseName</i> = '1:Boiler'. From there it finds all targets of hierarchical <i>References</i> .
"<0:HasChild>2:Wheel"	Follows any forward <i>Reference</i> with a <i>BrowseName</i> = 'HasChild' and qualified with the default OPC UA namespace. Then find targets with <i>BrowseName</i> = 'Wheel' qualified with namespace index '2'.
"<!HasChild>Truck"	Follows any inverse <i>Reference</i> with a <i>BrowseName</i> = 'HasChild'. Then find targets with <i>BrowseName</i> = 'Truck'. In both cases, the namespace component of the <i>BrowseName</i> is assumed to be 0.
"<0:HasChild>"	Finds all targets of forward <i>References</i> with a <i>BrowseName</i> = 'HasChild' and qualified with the default OPC UA namespace.

The following BNF describes the syntax of the *RelativePath* text format.

```

<relative-path> ::= <reference-type> <browse-name> [relative-path]
<reference-type> ::= '/' | '.' | '<' | '#' | ['!'] <browse-name> '>'
<browse-name> ::= [<namespace-index> ':' ] <name>
<namespace-index> ::= <digit> [<digit>]
<digit> ::= '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9'
<name> ::= (<name-char> | '&' <reserved-char>) [<name>]
<reserved-char> ::= '/' | '.' | '<' | '>' | ':' | '#' | '!' | '&'
<name-char> ::= All valid characters for a String (see IEC 62541-3) excluding reserved-chars.

```

A.3 BNF of NumericRange

The following BNF describes the syntax of the *NumericRange* parameter type.

```

<numeric-range> ::= <dimension> [',' <dimension>]
<dimension> ::= <index> [':' <index>]
<index> ::= <digit> [<digit>]
<digit> ::= '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9'

```

Annex B (informative)

ContentFilter and Query examples

B.1 Simple ContentFilter examples

B.1.1 Overview

These examples provide fairly simple *ContentFilters*. Filter similar to these examples may be used in processing events.

The following conventions apply to these examples with regard to how *Attribute* operands are used (for a definition of this operand see 7.4.4):

- **AttributeOperand:** Refers to a *Node*, an *Attribute* of a *Node* or the *Value Attribute* of a *Property* associated with a *Node*. In the examples, the character names of NodeIds are used instead of an actual nodeId, this also applies to Attribute Ids.
- The string representation of relative paths is used instead of the actual structure.
- The NamespaceIndex used in all examples is 12 (it could just as easily have been 4 or 23 or any value). For more information about NamespaceIndex, see IEC 62541-3. The use of the NamespaceIndex illustrates that the information model being used in the examples is not a model defined by this document, but one created for the examples.

B.1.2 Example 1

For example the logic described by '(((AType.A = 5) or InList(BType.B, 3,5,7)) and BaseObjectType.displayName LIKE "Main%")' would result in a logic tree as shown in Figure B.1 and a *ContentFilter* as shown in Table B.1. For this example to return anything AType and BType both shall be subtypes of BaseObjectType, or the resulting "And" operation would always be false.




Figure B.1 – Filter logic tree example

Table B.1 describes the elements, operators and operands used in the example.

Table B.1 – ContentFilter example

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	And	ElementOperand = 1	Element Operand = 4		
1	Or	ElementOperand = 2	Element Operand = 3		
2	Equals	AttributeOperand = NodeId: AType, BrowsePath: ".12:A", Attribute:value	LiteralOperand = '5'		
3	InList	AttributeOperand = NodeId: BType, BrowsePath: ".12:B", Attribute:value	LiteralOperand = '3'	LiteralOperand = '5'	LiteralOperand = '7'
4	Like	AttributeOperand = NodeId: BaseObjectType, BrowsePath: ".", Attribute: displayName	LiteralOperand = "Main%"		

B.1.3 Example 2

As another example a filter to select all *SystemEvents* (including derived types) that are contained in the *Area1 View* or the *Area2 View* would result in a logic tree as shown in Figure B.2 and a ContentFilter as shown in Table B.2.


**Figure B.2 – Filter logic tree example**

Table B.2 describes the elements, operators and operands used in the example.

Table B.2 – ContentFilter example

Element[]	Operator	Operand[0]	Operand[1]
0	And	ElementOperand = 1	ElementOperand = 4
1	Or	ElementOperand = 2	ElementOperand = 3
2	InView	AttributeOperand = NodeId: Area1, BrowsePath: ".", Attribute: NodeId	
3	InView	AttributeOperand = NodeId: Area2, BrowsePath: ".", Attribute: NodeId	
4	OfType	AttributeOperand = NodeId: SystemEventType, BrowsePath: ".", Attribute: NodeId	

B.2 Complex examples of Query filters

B.2.1 Overview

These Query examples illustrate complex filters. The following conventions apply to these examples with regard to Attribute operands (for a definition of these operands, see 7.4.4).

- **AttributeOperand:** Refers to a *Node*, an *Attribute* of a *Node* or the *Value Attribute* of a *Property* associated with a *Node*. In the examples character names of *ExpandedNodeId* are used instead of an actual *ExpandedNodeId*, this also applies to *Attribute* Ids.
- The string representation of relative paths is used instead of the actual structure.
- The *NamespaceIndex* used in all examples is 12 (it could just as easily have been 4 or 23 or any value). For more information about *NamespaceIndex*, see IEC 62541-3. The use of the *NamespaceIndex* illustrates that the information model being used in the examples is not a model defined by this document, but one created for the examples.

B.2.2 Used type model

The following examples use the type model described below. All *Property* values are assumed to be string unless otherwise noted.

New Reference types:

- "HasChild" derived from HierarchicalReference.
- "HasAnimal" derived from HierarchicalReference.
- "HasPet" derived from HasAnimal.
- "HasFarmAnimal" derived from HasAnimal.
- "HasSchedule" derived from HierarchicalReference.

PersonType derived from BaseObjectType adds:

- HasProperty "LastName".
- HasProperty "FirstName".
- HasProperty "StreetAddress".
- HasProperty "City".
- HasProperty "ZipCode".
- May have HasChild reference to a node of type PersonType.
- May have HasAnimal reference to a node of type AnimalType (or a subtype of this Reference type).

AnimalType derived from BaseObjectType adds:

- May have HasSchedule reference to a node of type FeedingScheduleType.
- HasProperty "Name".

DogType derived from AnimalType adds:

- HasProperty "NickName".
- HasProperty "DogBreed".
- HasProperty "License".

CatType derived from AnimalType adds:

HasProperty "NickName".

HasProperty "CatBreed".

PigType derived from AnimalType adds:

HasProperty "PigBreed".

ScheduleType derived from BaseObjectType adds:

HasProperty "Period".

FeedingScheduleType derived from ScheduleType adds:

HasProperty "Food".

HasProperty "Amount" (Stored as an *Int32*).

AreaType derived from *BaseObjectType* is just a simple *Folder* and contains no *Properties*.

This example type system is shown in Figure B.3. In Figure B.3, the OPC UA notation is used for all *References* to *ObjectTypes*, *Variables*, *Properties* and subtypes. Additionally, supported *References* are contained in an inner box. The actual references only exist in the instances, thus, no connections to other *Objects* are shown in Figure B.3 and they are subtypes of the listed *Reference*.





Figure B.3 – Example Type Nodes

A corresponding example set of instances is shown in Figure B.4. These instances include a type *Reference* for *Objects*. Properties also have type *References*, but the *References* are omitted for simplicity. The name of the *Object* is provided in the box and a numeric instance *NodeId* in brackets. Standard *ReferenceTypes* use the OPC UA notation, custom *ReferenceTypes* are listed as a named *Reference*. For *Properties*, the *BrowseName*, *NodeId*, and *Value* are shown. The *Nodes* that are included in a *View* (View1) are enclosed in the coloured box. Two *Area* nodes are included for grouping of the existing person nodes. All custom nodes are defined in namespace 12 which is not included in Figure B.4.



IEC

Figure B.4 – Example Instance Nodes

B.2.3 Example Notes

For all of the examples in 7.4.4, the type definition *Node* is listed in its symbolic form, in the actual call it would be the *ExpandedNodeId* assigned to the *Node*. The *Attribute* is the symbolic name of the *Attribute*, in the actual call they would be translated to the *IntegerId* of the *Attribute*. Also in all of the examples the *BrowseName* is included in the result table for clarity; normally this would not be returned.

All of the examples include the following items:

- an English description of the object of the query;
- an SQL like description of the query;
- a table that has a *NodeTypeDescription* of the items that are to be returned;
- a figure illustrating the query filter;
- a table describing the content filter;
- a table describing the resulting dataset.

The examples assume namespace 12 is the namespace for all of the custom definitions described for the examples.

B.2.4 Example 1

This example requests a simple layered filter, a person has a pet and the pet has a schedule.

Example 1: Get PersonType.LastName, AnimalType.Name, ScheduleType.Period where the Person Has a Pet and that Pet Has a Schedule.

The *NodeTypeDescription* parameters used in the example are described in Table B.3.

Table B.3 – Example 1 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	value	N/A
		"<12:HasPet>12:AnimalType. 12:Name"	value	N/A
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:Schedule. 12:Period"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.5.




Figure B.5 – Example 1 Filter

Table B.4 describes the *ContentFilter* elements, operators and operands used in the example.

Table B.4 – Example 1 ContentFilter

Element[]	filterOperator Operator	filterOperand Operand[0]	filterOperand Operand[1]	filterOperand Operand[2]	filterOperand Operand[3]
1	RelatedTo_15	AttributeOperand = NodeId: PersonType, BrowsePath ".", Attribute: NodeId	ElementOperand = 2	AttributeOperand = NodeId: HasPet, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	RelatedTo_15	AttributeOperand = NodeId: AnimalType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: ScheduleType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

Table B.5 describes the *QueryDataSet* that results from this query if it were executed against the instances described in Figure B.4

Table B.5 – Example 1 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (JFamily1)	PersonType	".12:LastName"	Jones
		"<12:HasPet>12:AnimalType. 12:Name"	Rosemary
			Basil
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:Schedule.12:Period"	Hourly
			Daily
12:42(HFamily1)	PersonType	".12:LastName"	Hervey
		"<12:HasPet>12:AnimalType. 12:Name"	Oliver
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:Schedule.12:Period"	Daily

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

The Value column is returned as an array for each *Node* description, where the order of the items in the array would correspond to the order of the items that were requested for the given Node Type. In Addition, if a single *Attribute* has multiple values then it would be returned as an array within the larger array, for example in this table Rosemary and Basil would be returned in a array for the .<HasPet>.AnimalType.Name item. They are show as separate rows for ease of viewing. The actual value array for JFamily1 would be ("Jones", {"RoseMary", "Basil"}, {"Hourly", "Daily"})

B.2.5 Example 2

The second example illustrates receiving a list of disjoint *Nodes* and also illustrates that an array of results can be received.

Example 2: Get PersonType.LastName, AnimalType.Name where a person has a child or (a pet is of type cat and has a feeding schedule).

The NodeTypeDescription parameters used in the example are described in Table B.6.

Table B.6 – Example 2 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	Value	N/A
AnimalType	TRUE	".12:Name"	Value	N/A

The corresponding ContentFilter is illustrated in Figure B.6.




Figure B.6 – Example 2 Filter logic tree

Table B.7 describes the elements, operators and operands used in the example. It is worth noting that a CatType is a subtype of AnimalType.

Table B.7 – Example 2 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	Or	ElementOperand=1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodId: PersonType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: PersonType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: HasChild, BrowsePath ".", Attribute: NodId	LiteralOperand = '1'
2	RelatedTo	AttributeOperand = NodId: CatType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: FeedingScheduleType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: HasSchedule, BrowsePath ".", Attribute: NodId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.8.

Table B.8 – Example 2 QueryDataSets

NodId	TypeDefinition NodId	RelativePath	Value
12:30 (Jfamily1)	PersonType	. 12:LastName	Jones
12:42 (HFamily1)	PersonType	. 12:LastName	Hervey
12:48 (HFamily2)	PersonType	. 12:LastName	Hervey
12:70 (Cat1)	CatType	. 12:Name	Rosemary
12:74 (Cat2)	CatType	. 12:Name	Basil

NOTE The relative path column and browse name (in parentheses in the *NodId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodId* would be a *NodId* not the symbolic name that is included in the table.

B.2.6 Example 3

The third example provides a more complex *Query* in which the results are filtered on multiple criteria.

Example 3: Get PersonType.LastName, AnimalType.Name, ScheduleType.Period where a person has a pet and the animal has a feeding schedule and the person has a Zipcode = '02138' and (the Schedule.Period is Daily or Hourly) and Amount to feed is > 10.

Table B.9 describes the NodeDescription parameters used in the example.

Table B.9 – Example 3 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		RelativePath	Attribute	Index Range
PersonType	FALSE	"12:LastName"	Value	N/A
		"<12:HasPet>12:AnimalType. 12:Name"	Value	N/A
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:FeedingSchedule.Period"	Value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.7.


**Figure B.7 – Example 3 Filter logic tree**

Table B.10 describes the elements, operators and operands used in the example.

Table B.10 – Example 3 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	And	Element Operand = 1	ElementOperand = 2		
1	And	ElementOperand = 4	ElementOperand = 6		
2	And	ElementOperand = 3	ElementOperand = 9		
3	Or	ElementOperand = 7	ElementOperand = 8		
4	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	ElementOperand = 5	AttributeOperand = NodeId: 12:HasPet, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
5	RelatedTo	AttributeOperand = Node: 12:AnimalType, BrowsePath ".", Attribute: NodeId Alias: AT	AttributeOperand = NodeId: 12:FeedingScheduleType, BrowsePath ".", Attribute: NodeId Alias: FST	AttributeOperand = NodeId: 12:HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
6	Equals	AttributeOperand = NodeId: 12:PersonType BrowsePath 12:Zipcode ".", Attribute: Value	LiteralOperand = '02138'		
7	Equals	AttributeOperand = NodeId: 12:PersonType BrowsePath "12:HasPet>12:AnimalType<12: HasSchedule>12: FeedingSchedule/12:Period", Attribute: Value Alias: FST	LiteralOperand = 'Daily'		
8	Equals	AttributeOperand = NodeId: 12:PersonType BrowsePath "12:HasPet>12:AnimalType<12: HasSchedule>12: FeedingSchedule/12:Period", Attribute: Value Alias: FST	LiteralOperand = 'Hourly'		
9	Greater Than	AttributeOperand = NodeId: 12:PersonType BrowsePath "12:HasPet>12:AnimalType<12: HasSchedule>12: FeedingSchedule/12:Amount", Attribute: Value Alias: FST	ElementOperand = 10		
10	Cast	LiteralOperand = 10	AttributeOperand = NodeId: Int32, BrowsePath ".", Attribute: NodeId		

The results from this query would contain the *QueryDataSets* shown in Table B.11.

Table B.11 – Example 3 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (JFamily1)	PersonType	".12:LastName"	Jones
		"<12:HasPet>12:PersonType. 12:Name"	Rosemary
			Basil
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:FeedingSchedule. 12:Period"	Hourly
			Daily

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

B.2.7 Example 4

The fourth example provides an illustration of the Hop parameter that is part of the RelatedTo Operator.

Example 4: Get PersonType.LastName where a person has a child who has a child who has a pet.

Table B.12 describes the NodeTypeDescription parameters used in the example.

Table B.12 – Example 4 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.8.


**Figure B.8 – Example 4 Filter logic tree**

Table B.13 describes the elements, operators and operands used in the example.

Table B.13 – Example 4 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	Element Operand = 1	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '2'
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:AnimalType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasPet, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.14. It is worth noting that the pig "Pig1" is referenced as a pet by Sara, but is referenced as a farm animal by Sara's parent Paul.

Table B.14 – Example 4 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	".12:LastName"	Hervey

NOTE The RelativePath column and browse name (in parentheses in the NodeId column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

B.2.8 Example 5

The fifth example provides an illustration of the use of alias.

Example 5: Get the last names of children that have the same first name as a parent of theirs

Table B.15 describes the *NodeTypeDescription* parameters used in the example.

Table B.15 – Example 5 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	"<12:HasChild>12:PersonType. 12:LastName"	Value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.9.




Figure B.9 – Example 5 Filter logic tree

In this example, one *Reference* to *PersonType* is aliased to "Parent" and another *Reference* to *PersonType* is aliased to "Child". The value of *Parent.firstName* and *Child.firstName* are then compared. Table B.16 describes the elements, operators and operands used in the example.

Table B.16 – Example 5 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	And	ElementOperand = 1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodId: 12:PersonType, BrowsePath ".", Attribute: NodId, Alias: "Parent"	AttributeOperand = NodId: 12:PersonType, BrowsePath ".", Attribute: NodId, Alias: "Child"	AttributeOperand = NodId: 12:HasChild, Attribute: NodId	LiteralOperand = "1"
2	Equals	AttributeOperand = NodId: 12:PersonType, BrowsePath ""/12:FirstName", Attribute: Value, Alias: "Parent"	AttributeOperand = NodId: 12:PersonType, BrowsePath ""/12:FirstName", Attribute: Value, Alias: "Child"		

The results from this query would contain the *QueryDataSets* shown in Table B.17.

Table B.17 – Example 5 QueryDataSets

NodId	TypeDefinition NodId	RelativePath	Value
12:42 (HFamily1)	PersonType	"<12:HasChild>12:PersonType.12:LastName"	Hervey

NOTE The *RelativePath* column and browse name (in parentheses in the *NodId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodId* would be an integer not the symbolic name that is included in the table.

B.2.9 Example 6

The sixth example provides an illustration a different type of request, one in which the *Client* is interested in displaying part of the *AddressSpace* of the *Server*. This request includes listing a *Reference* as something that is to be returned.

Example 6: Get *PersonType.NodId*, *AnimalType.NodId*, *PersonType.HasChild* Reference, *PersonType.HasAnimal* Reference where a person has a child who has a *Animal*.

Table B.18 describes the *NodeTypeDescription* parameters used in the example.

Table B.18 – Example 6 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:NodeId"	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>12:AnimalType.NodeId	value	N/A
		<12:HasChild>	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.10.




Figure B.10 – Example 6 Filter logic tree

Table B.19 describes the elements, operators and operands used in the example.

Table B.19 – Example 6 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	ElementOperand = 1	AttributeOperand = Node: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:AnimalType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasAnimal, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.20.

Table B.20 – Example 6 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	".NodeId"	12:42 (HFamily1)
		<12:HasChild>12:PersonType<12:HasAnimal> 12:AnimalType.NodeId	12:91 (Pig1)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasChild>12:PersonType<12:HasAnimal>	HasFarmAnimal <i>ReferenceDescription</i>
12:48 (HFamily2)	PersonType	".NodeId"	12:48 (HFamily2)
		<12:HasChild>12:PersonType<12:HasAnimal> 12:AnimalType.NodeId	12:91 (Pig1)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasChild>12:PersonType<12:HasAnimal>	HasPet <i>ReferenceDescription</i>

NOTE The *RelativePath* and browse name (in parentheses) is not in the *QueryDataSet* and is only shown here for clarity and the *TypeDefinition NodeId* would be an integer, not the symbolic name that is included in the table. The value field would in this case be the *NodeId* where it was requested, but for the example the browse name is provided in parentheses and in the case of *Reference* types on the browse name is provided. For the *References* listed in Table B.20, the value would be a *ReferenceDescription* which are described in 7.25.

Table B.21 provides an example of the same *QueryDataSet* as shown in Table B.20 without any additional fields and minimal symbolic ids. There is an entry for each requested Attribute, in the cases where an Attribute would return multiple entries the entries are separated by comas. If a structure is being returned then the structure is enclosed in square brackets. In the case of a *ReferenceDescription* the structure contains a structure and *DisplayName* and *BrowseName* are assumed to be the same and defined in Figure B.4.

Table B.21 – Example 6 QueryDataSets without additional information

NodeId	TypeDefinition NodeId	Value
12:42	PersonType	12:42
		12:91
		[HasChild,TRUE,[48,HFamily2,HFamily2,PersonType]],
		[HasFarmAnimal,TRUE[91,Pig1,Pig1,PigType]
12:48	PersonType	12:54
		12:91
		[HasChild,TRUE,[54,HFamily3,HFamily3,PersonType]]
		[HasPet, TRUE,[91,Pig1,Pig1,PigType]]

The PersonType, HasChild, PigType, HasPet, HasFarmAnimal identifiers used in the above table would be translated to actual ~~ExtendedNodeIds~~ *ExpandedNodeIds*.

B.2.10 Example 7

The seventh example provides an illustration a request in which a *Client* wants to display part of the *AddressSpace* based on a starting point that was obtained via browsing. This request

includes listing *References* as something that is to be returned. In this case the Person Browsed to Area2 and wanted to *Query* for information below this starting point.

Example 7: Get PersonType.NodeId, AnimalType.NodeId, PersonType.HasChild Reference, PersonType.HasAnimal Reference where the person is in Area2 (Cleveland nodes) and the person has a child.

Table B.22 describes the NodeTypeDescription parameters used in the example.

Table B.22 – Example 7 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".NodeId"	Value	N/A
		<12:HasChild>	Value	N/A
		<12:HasAnimal>NodeId	Value	N/A
		<12:HasAnimal>	Value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.11. Note that the *Browse* call would typically return a *NodeId*, thus the first filter is for the *BaseObjectType* with a *NodeId* of 95 where 95 is the *NodeId* associated with the Area2 node, all *Nodes* descend from *BaseObjectType*, and *NodeId* is a base *Property* so this filter will work for all *Queries* of this nature.




Figure B.11 – Example 7 Filter logic tree

Table B.23 describes the elements, operators and operands used in the example.

Table B.23 – Example 7 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	RelatedTo	ElementOperand = 2	ElementOperand = 1	AttributeOperand = Node:HierarchicalReference, BrowsePath ".", Attribute:NodeId	LiteralOperand = '1'
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	Equals	AttributeOperand = NodeId: BaseObjectType, BrowsePath ".", Attribute: NodeId,	LiteralOperand = '95		

The results from this *Query* would contain the *QueryDataSets* shown in Table B.24.

Table B.24 – Example 7 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	".NodeId"	12:42 (HFamily1)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasAnimal>12:AnimalType.NodeId	NULL
		<12:HasAnimal>	HasFarmAnimal <i>ReferenceDescription</i>
12:48 (HFamily2)	PersonType	".NodeId"	12:48 (HFamily2)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasAnimal>12:AnimalType.NodeId	12:91 (Pig1)
		<12:HasAnimal>	HasFarmAnimal <i>ReferenceDescription</i>

NOTE The *RelativePath* and browse name (in parentheses) is not in the *QueryDataSet* and is only shown here for clarity and the *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table. The value field would in this case be the *NodeId* where it was requested, but for the example the browse name is provided in parentheses and in the case of *Reference* types on the browse name is provided. For the *References* listed in Table B.24, the value would be a *ReferenceDescription* which are described in 7.25.

B.2.11 Example 8

The eighth example provides an illustration of a request in which the *AddressSpace* is restricted by a *Server* defined *View*. This request is the same as in the second example which illustrates receiving a list of disjoint *Nodes* and also illustrates that an array of results can be received. It is important to note that all of the parameters and the *ContentFilter* are the same, only the *View* description would be specified as "View1".

Example 8: Get PersonType.LastName, AnimalType.Name where a person has a child or (a pet is of type cat and has a feeding schedule) limited by the *AddressSpace* in View1.

The *NodeTypeDescription* parameters used in the example are described in Table B.25

Table B.25 – Example 8 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	value	N/A
AnimalType	TRUE	".12.Name"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.12.


**Figure B.12 – Example 8 Filter logic tree**

Table B.26 describes the elements, operators and operands used in the example. It is worth noting that a CatType is a subtype of AnimalType.

Table B.26 – Example 8 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	Or	ElementOperand=1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	RelatedTo	AttributeOperand = NodeId: 12:CatType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:FeedingScheduleType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.27. If this is compared to the result set from example 2, the only difference is the omission of the Cat Nodes. These Nodes are not in the View and thus are not included in the result set.

Table B.27 – Example 8 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (Jfamily1)	Persontype	.12:LastName	Jones

NOTE The RelativePath column and browse name (in parentheses in the NodeId column) are not in the QueryDataSet and are only shown here for clarity. The TypeDefinition NodeId would be an integer not the symbolic name that is included in the table.

B.2.12 Example 9

The ninth example provides a further illustration for a request in which the *AddressSpace* is restricted by a *Server* defined *View*. This request is similar to the second example except that some of the requested nodes are expressed in terms of a relative path. It is important to note that the *ContentFilter* is the same, only the *View* description would be specified as "View1".

Example 9: Get PersonType.LastName, AnimalType.Name where a person has a child or (a pet is of type cat and has a feeding schedule) limited by the *AddressSpace* in View1.

Table B.28 describes the *NodeTypeDescription* parameters used in the example.

Table B.28 – Example 9 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".NodeId"	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>12:AnimalType.NodeId	value	N/A
		<12:HasChild>	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>	value	N/A
PersonType	FALSE	".12:LastName"	value	N/A
		<12:HasAnimal>12:AnimalType. 12:Name	value	N/A
AnimalType	TRUE	".12:name"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.13.




Figure B.13 – Example 9 Filter logic tree

Table B.29 describes the elements, operators and operands used in the example.

Table B.29 – Example 9 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	Or	ElementOperand=1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	RelatedTo	AttributeOperand = NodeId: 12:CatType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:FeedingScheduleType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this *Query* would contain the *QueryDataSets* shown in Table B.30. If this is compared to the result set from example 2, the Pet *Nodes* are included in the list, even though they are outside of the *View*. This is possible since the name referenced via the relative path and the root *Node* is in the *View*.

Table B.30 – Example 9 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (Jfamily1)	PersonType	. 12:LastName	Jones
		<12:HasAnimal>12:AnimalType. 12:Name	Rosemary
		<12:HasAnimal>12:AnimalType. 12:Name	Basil

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the QueryDataSet and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

INTERNATIONAL STANDARD

NORME INTERNATIONALE



**OPC unified architecture –
Part 4: Services**

**Architecture unifiée OPC –
Partie 4: Services**

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

CONTENTS

FOREWORD	13
1 Scope	15
2 Normative references	15
3 Terms, definitions, abbreviated terms and conventions	16
3.1 Terms and definitions	16
3.2 Abbreviated terms	17
3.3 Conventions for Service definitions	17
4 Overview	19
4.1 Service Set model	19
4.2 Request/response Service procedures	22
5 Service Sets	22
5.1 General	22
5.2 Service request and response header	23
5.3 Service results	23
5.4 Discovery Service Set	24
5.4.1 Overview	24
5.4.2 FindServers	26
5.4.3 FindServersOnNetwork	27
5.4.4 GetEndpoints	29
5.4.5 RegisterServer	31
5.4.6 RegisterServer2	34
5.5 SecureChannel Service Set	35
5.5.1 Overview	35
5.5.2 OpenSecureChannel	36
5.5.3 CloseSecureChannel	40
5.6 Session Service Set	41
5.6.1 Overview	41
5.6.2 CreateSession	41
5.6.3 ActivateSession	46
5.6.4 CloseSession	49
5.6.5 Cancel	50
5.7 NodeManagement Service Set	50
5.7.1 Overview	50
5.7.2 AddNodes	50
5.7.3 AddReferences	52
5.7.4 DeleteNodes	54
5.7.5 DeleteReferences	56
5.8 View Service Set	57
5.8.1 Overview	57
5.8.2 Browse	57
5.8.3 BrowseNext	60
5.8.4 TranslateBrowsePathsToNodeIds	62
5.8.5 RegisterNodes	64
5.8.6 UnregisterNodes	65
5.9 Query Service Set	66
5.9.1 Overview	66

5.9.2	Querying Views	66
5.9.3	QueryFirst	67
5.9.4	QueryNext	70
5.10	Attribute Service Set	71
5.10.1	Overview	71
5.10.2	Read	72
5.10.3	HistoryRead	73
5.10.4	Write	76
5.10.5	HistoryUpdate	79
5.11	Method Service Set	81
5.11.1	Overview	81
5.11.2	Call	81
5.12	MonitoredItem Service Set	84
5.12.1	MonitoredItem model	84
5.12.2	CreateMonitoredItems	89
5.12.3	ModifyMonitoredItems	92
5.12.4	SetMonitoringMode	94
5.12.5	SetTriggering	95
5.12.6	DeleteMonitoredItems	97
5.13	Subscription Service Set	98
5.13.1	Subscription model	98
5.13.2	CreateSubscription	107
5.13.3	ModifySubscription	108
5.13.4	SetPublishingMode	110
5.13.5	Publish	111
5.13.6	Republish	113
5.13.7	TransferSubscriptions	114
5.13.8	DeleteSubscriptions	116
6	Service behaviours	117
6.1	Security	117
6.1.1	Overview	117
6.1.2	Obtaining and installing an Application Instance Certificate	117
6.1.3	Determining if a Certificate is trusted	118
6.1.4	Creating a SecureChannel	121
6.1.5	Creating a Session	123
6.1.6	Impersonating a User	124
6.2	Authorization Services	124
6.2.1	Overview	124
6.2.2	Indirect handshake with an Identity Provider	124
6.2.3	Direct handshake with an Identity Provider	125
6.3	Session-less Service invocation	126
6.3.1	Description	126
6.3.2	Parameters	127
6.3.3	Service results	128
6.4	Software Certificates	128
6.5	Auditing	128
6.5.1	Overview	128
6.5.2	General audit logs	128
6.5.3	General audit Events	129

6.5.4	Auditing for Discovery Service Set	129
6.5.5	Auditing for SecureChannel Service Set	129
6.5.6	Auditing for Session Service Set	129
6.5.7	Auditing for NodeManagement Service Set	130
6.5.8	Auditing for Attribute Service Set	130
6.5.9	Auditing for Method Service Set	131
6.5.10	Auditing for View, Query, MonitoredItem and Subscription Service Set	131
6.6	Redundancy	131
6.6.1	Redundancy overview	131
6.6.2	Server Redundancy	132
6.6.3	Client Redundancy	143
6.6.4	Network Redundancy	143
6.6.5	Manually forcing Failover	145
6.7	Re-establishing connections	145
6.8	Durable Subscriptions	147
7	Common parameter type definitions	148
7.1	ApplicationDescription	148
7.2	ApplicationInstanceCertificate	149
7.3	BrowseResult	150
7.4	ContentFilter	151
7.4.1	ContentFilter structure	151
7.4.2	ContentFilterResult	151
7.4.3	FilterOperator	152
7.4.4	FilterOperand parameters	159
7.5	Counter	161
7.6	ContinuationPoint	161
7.7	DataValue	162
7.7.1	General	162
7.7.2	PicoSeconds	162
7.7.3	SourceTimestamp	162
7.7.4	ServerTimestamp	163
7.7.5	StatusCodes assigned to a value	163
7.8	DiagnosticInfo	164
7.9	DiscoveryConfiguration parameters	165
7.9.1	Overview	165
7.9.2	MdnsDiscoveryConfiguration	166
7.10	EndpointDescription	166
7.11	ExpandedNodeId	167
7.12	ExtensibleParameter	167
7.13	Index	167
7.14	IntegerId	167
7.15	MessageSecurityMode	168
7.16	MonitoringParameters	168
7.17	MonitoringFilter parameters	169
7.17.1	Overview	169
7.17.2	DataChangeFilter	170
7.17.3	EventFilter	171
7.17.4	AggregateFilter	173
7.18	MonitoringMode	174

7.19	NodeAttributes parameters	175
7.19.1	Overview	175
7.19.2	ObjectAttributes parameter	176
7.19.3	VariableAttributes parameter	176
7.19.4	MethodAttributes parameter	177
7.19.5	ObjectTypeAttributes parameter	177
7.19.6	VariableTypeAttributes parameter	178
7.19.7	ReferenceTypeAttributes parameter	178
7.19.8	DataTypeAttributes parameter	179
7.19.9	ViewAttributes parameter	179
7.19.10	GenericAttributes parameter	180
7.20	NotificationData parameters	180
7.20.1	Overview	180
7.20.2	DataChangeNotification parameter	181
7.20.3	EventNotificationList parameter	181
7.20.4	StatusChangeNotification parameter	182
7.21	NotificationMessage	182
7.22	NumericRange	182
7.23	QueryDataSet	183
7.24	ReadValueId	184
7.25	ReferenceDescription	185
7.26	RelativePath	186
7.27	RegisteredServer	187
7.28	RequestHeader	187
7.29	ResponseHeader	189
7.30	ServiceFault	189
7.31	SessionAuthenticationToken	190
7.32	SignatureData	191
7.33	SignedSoftwareCertificate	191
7.34	StatusCode	192
7.34.1	General	192
7.34.2	Common StatusCodes	194
7.35	TimestampsToReturn	198
7.36	UserIdentityToken parameters	198
7.36.1	Overview	198
7.36.2	Token Encryption and Proof of Possession	199
7.36.3	AnonymousIdentityToken	203
7.36.4	UserNameIdentityToken	203
7.36.5	X509IdentityTokens	205
7.36.6	IssuedIdentityToken	205
7.37	UserTokenPolicy	206
7.38	VersionTime	207
7.39	ViewDescription	207
Annex A	(informative) BNF definitions	208
A.1	Overview over BNF	208
A.2	BNF of RelativePath	208
A.3	BNF of NumericRange	209
Annex B	(informative) ContentFilter and Query examples	210
B.1	Simple ContentFilter examples	210

B.1.1	Overview	210
B.1.2	Example 1	210
B.1.3	Example 2	211
B.2	Complex examples of Query filters	212
B.2.1	Overview	212
B.2.2	Used type model	212
B.2.3	Example Notes	215
B.2.4	Example 1	216
B.2.5	Example 2	217
B.2.6	Example 3	218
B.2.7	Example 4	221
B.2.8	Example 5	222
B.2.9	Example 6	223
B.2.10	Example 7	225
B.2.11	Example 8	227
B.2.12	Example 9	228
Figure 1	– Discovery Service Set	19
Figure 2	– SecureChannel Service Set	19
Figure 3	– Session Service Set	20
Figure 4	– NodeManagement Service Set	20
Figure 5	– View Service Set	20
Figure 6	– Attribute Service Set	21
Figure 7	– Method Service Set	21
Figure 8	– MonitoredItem and Subscription Service Sets	22
Figure 9	– Discovery process	25
Figure 10	– Using a Gateway Server	30
Figure 11	– The registration process – Manually launched servers	32
Figure 12	– The registration process – Automatically launched servers	32
Figure 13	– SecureChannel and Session Services	36
Figure 14	– Multiplexing users on a Session	43
Figure 15	– MonitoredItem model	84
Figure 16	– Typical delay in change detection	86
Figure 17	– Queue overflow handling	87
Figure 18	– Triggering model	88
Figure 19	– Obtaining and installing an Application Instance Certificate	118
Figure 20	– Determining if an Application Instance Certificate is trusted	121
Figure 21	– Establishing a SecureChannel	122
Figure 22	– Establishing a Session	123
Figure 23	– Impersonating a User	124
Figure 24	– Indirect handshake with an Identity Provider	125
Figure 25	– Direct handshake with an Identity Provider	126
Figure 26	– Transparent Redundancy setup example	133
Figure 27	– Non-Transparent Redundancy setup	134
Figure 28	– Client Start-up steps	138

Figure 29 – Cold Failover.....	139
Figure 30 – Warm Failover.....	140
Figure 31 – Hot Failover	141
Figure 32 – HotAndMirrored Failover	142
Figure 33 – Server proxy for Redundancy	143
Figure 34 – Transparent network Redundancy	144
Figure 35 – Non-transparent network Redundancy.....	145
Figure 36 – Reconnect sequence.....	146
Figure 37 – Logical layers of a <i>Server</i>	190
Figure 38 – Obtaining a SessionAuthenticationToken	191
Figure 39 – EncryptedSecret layout	200
Figure B.1 – Filter logic tree example.....	210
Figure B.2 – Filter logic tree example.....	211
Figure B.3 – Example Type Nodes	214
Figure B.4 – Example Instance Nodes	215
Figure B.5 – Example 1 Filter.....	216
Figure B.6 – Example 2 Filter logic tree	218
Figure B.7 – Example 3 Filter logic tree	219
Figure B.8 – Example 4 Filter logic tree	221
Figure B.9 – Example 5 Filter logic tree	222
Figure B.10 – Example 6 Filter logic tree	224
Figure B.11 – Example 7 Filter logic tree	226
Figure B.12 – Example 8 Filter logic tree	227
Figure B.13 – Example 9 Filter logic tree	228
Table 1 – Service definition table.....	18
Table 2 – Parameter Types defined in IEC 62541-3	18
Table 3 – FindServers Service parameters	27
Table 4 – FindServersOnNetwork Service parameters	28
Table 5 – GetEndpoints Service parameters	31
Table 6 – RegisterServer Service parameters	33
Table 7 – RegisterServer Service result codes.....	33
Table 8 – RegisterServer2	34
Table 9 – RegisterServer2 Service result codes	35
Table 10 – RegisterServer2 Operation Level result codes	35
Table 11 – OpenSecureChannel Service parameters	38
Table 12 – OpenSecureChannel Service result codes	40
Table 13 – CloseSecureChannel Service parameters.....	41
Table 14 – CloseSecureChannel Service result codes	41
Table 15 – CreateSession Service parameters.....	44
Table 16 – CreateSession Service result codes	46
Table 17 – ActivateSession Service parameters.....	48
Table 18 – ActivateSession Service result codes	49

Table 19 – CloseSession Service parameters	49
Table 20 – CloseSession Service result codes	50
Table 21 – Cancel Service parameters	50
Table 22 – AddNodes Service parameters	51
Table 23 – AddNodes Service result codes	52
Table 24 – AddNodes Operation Level result codes	52
Table 25 – AddReferences Service parameters	53
Table 26 – AddReferences Service result codes	53
Table 27 – AddReferences Operation Level result codes	54
Table 28 – DeleteNodes Service parameters	55
Table 29 – DeleteNodes Service result codes	55
Table 30 – DeleteNodes Operation Level result codes	56
Table 31 – DeleteReferences Service parameters	56
Table 32 – DeleteReferences Service result codes	57
Table 33 – DeleteReferences Operation Level result codes	57
Table 34 – Browse Service parameters	58
Table 35 – Browse Service result codes	59
Table 36 – Browse Operation Level result codes	60
Table 37 – BrowseNext Service parameters	61
Table 38 – BrowseNext Service result codes	61
Table 39 – BrowseNext Operation Level result codes	62
Table 40 – TranslateBrowsePathsToNodeIds Service parameters	63
Table 41 – TranslateBrowsePathsToNodeIds Service result codes	63
Table 42 – TranslateBrowsePathsToNodeIds Operation Level result codes	64
Table 43 – RegisterNodes Service parameters	65
Table 44 – RegisterNodes Service result codes	65
Table 45 – UnregisterNodes Service parameters	66
Table 46 – UnregisterNodes Service result codes	66
Table 47 – QueryFirst Request parameters	68
Table 48 – QueryFirst Response parameters	69
Table 49 – QueryFirst Service result codes	70
Table 50 – QueryFirst Operation Level result codes	70
Table 51 – QueryNext Service parameters	71
Table 52 – QueryNext Service result codes	71
Table 53 – Read Service parameters	72
Table 54 – Read Service result codes	73
Table 55 – Read Operation Level result codes	73
Table 56 – HistoryRead Service parameters	74
Table 57 – HistoryRead Service result codes	76
Table 58 – HistoryRead Operation Level result codes	76
Table 59 – Write Service parameters	78
Table 60 – Write Service result codes	79
Table 61 – Write Operation Level result codes	79

Table 62 – HistoryUpdate Service parameters	80
Table 63 – HistoryUpdate Service result codes	80
Table 64 – HistoryUpdate Operation Level result codes	81
Table 65 – Call Service parameters	82
Table 66 – Call Service result codes	83
Table 67 – Call Operation Level result codes	83
Table 68 – Call Input Argument Result Codes	84
Table 69 – CreateMonitoredItems Service parameters	91
Table 70 – CreateMonitoredItems Service result codes	92
Table 71 – CreateMonitoredItems Operation Level result codes	92
Table 72 – ModifyMonitoredItems Service parameters	93
Table 73 – ModifyMonitoredItems Service result codes	94
Table 74 – ModifyMonitoredItems Operation Level result codes	94
Table 75 – SetMonitoringMode service parameters	95
Table 76 – SetMonitoringMode Service result codes	95
Table 77 – SetMonitoringMode Operation Level result codes	95
Table 78 – SetTriggering Service parameters	96
Table 79 – SetTriggering Service result codes	96
Table 80 – SetTriggering Operation Level result codes	97
Table 81 – DeleteMonitoredItems Service parameters	97
Table 82 – DeleteMonitoredItems Service result codes	98
Table 83 – DeleteMonitoredItems Operation Level result codes	98
Table 84 – Subscription states	101
Table 85 – Subscription state table	102
Table 86 – State variables and parameters	105
Table 87 – Functions	106
Table 88 – CreateSubscription Service parameters	107
Table 89 – CreateSubscription Service result codes	108
Table 90 – ModifySubscription Service parameters	109
Table 91 – ModifySubscription Service result codes	110
Table 92 – SetPublishingMode Service parameters	110
Table 93 – SetPublishingMode Service result codes	110
Table 94 – SetPublishingMode Operation Level result codes	111
Table 95 – Publish Service parameters	112
Table 96 – Publish Service result codes	112
Table 97 – Publish Operation Level Result Codes	113
Table 98 – Republish Service parameters	113
Table 99 – Republish Service result codes	113
Table 100 – TransferSubscriptions Service parameters	115
Table 101 – TransferSubscriptions Service result codes	115
Table 102 – TransferSubscriptions Operation Level result codes	116
Table 103 – DeleteSubscriptions Service parameters	116
Table 104 – DeleteSubscriptions Service result codes	117

Table 105 – DeleteSubscriptions Operation Level result codes	117
Table 106 – Certificate validation steps	119
Table 107 – SessionlessInvoke Service parameters.....	127
Table 108 – SessionlessInvoke Service result codes	128
Table 109 – ServiceLevel ranges	136
Table 110 – Server Failover modes.....	137
Table 111 – Redundancy Failover actions.....	138
Table 112 – ApplicationDescription	149
Table 113 – ApplicationInstanceCertificate	150
Table 114 – BrowseResult	150
Table 115 – ContentFilter structure.....	151
Table 116 – ContentFilterResult structure	151
Table 117 – ContentFilterResult result codes.....	152
Table 118 – ContentFilterResult Operand result codes	152
Table 119 – Basic FilterOperator definition	152
Table 120 – Complex FilterOperator definition	155
Table 121 – Wildcard characters.....	156
Table 122 – Conversion rules	157
Table 123 – Data Precedence rules	158
Table 124 – Logical AND Truth table	159
Table 125 – Logical OR Truth table.....	159
Table 126 – FilterOperand parameter Typelds	159
Table 127 – ElementOperand	160
Table 128 – LiteralOperand	160
Table 129 – AttributeOperand	160
Table 130 – SimpleAttributeOperand	161
Table 131 – DataValue	162
Table 132 – DiagnosticInfo	165
Table 133 – DiscoveryConfiguration parameterTypelds	165
Table 134 – MdnsDiscoveryConfiguration	166
Table 135 – EndpointDescription	166
Table 136 – ExpandedNodeId	167
Table 137 – ExtensibleParameter base type	167
Table 138 – MessageSecurityMode values	168
Table 139 – MonitoringParameters	168
Table 140 – MonitoringFilter parameterTypelds	169
Table 141 – DataChangeFilter	170
Table 142 – EventFilter structure	172
Table 143 – EventFilterResult structure	172
Table 144 – EventFilterResult result codes	173
Table 145 – AggregateFilter structure	174
Table 146 – AggregateFilterResult structure	174
Table 147 – MonitoringMode values.....	175

Table 148 – NodeAttributes parameterTypelds	175
Table 149 – Bit mask for specified Attributes	176
Table 150 – ObjectAttributes	176
Table 151 – VariableAttributes	177
Table 152 – MethodAttributes	177
Table 153 – ObjectTypeAttributes	178
Table 154 – VariableTypeAttributes	178
Table 155 – ReferenceTypeAttributes	179
Table 156 – DataTypeAttributes	179
Table 157 – ViewAttributes	180
Table 158 – GenericAttributes	180
Table 159 – NotificationData parameterTypelds	181
Table 160 – DataChangeNotification	181
Table 161 – EventNotificationList	182
Table 162 – StatusChangeNotification	182
Table 163 – NotificationMessage	182
Table 164 – NumericRange	183
Table 165 – QueryDataSet	184
Table 166 – ReadValueId	185
Table 167 – ReferenceDescription	186
Table 168 – RelativePath	186
Table 169 – RegisteredServer	187
Table 170 – RequestHeader	188
Table 171 – ResponseHeader	189
Table 172 – ServiceFault	190
Table 173 – SignatureData	191
Table 174 – SignedSoftwareCertificate	192
Table 175 – StatusCode bit assignments	193
Table 176 – DataValue InfoBits	194
Table 177 – Common Service result codes	195
Table 178 – Common Operation Level result codes	197
Table 179 – TimestampsToReturn values	198
Table 180 – UserIdentityToken parameterTypelds	199
Table 181 – Legacy UserIdentityToken Encrypted Token Secret Format	200
Table 182 – EncryptedSecret layout	202
Table 183 – EncryptedSecret DataTypes	202
Table 184 – RsaEncryptedSecret structure	203
Table 185 – AnonymousIdentityToken	203
Table 186 – UserNameIdentityToken	204
Table 187 – EncryptionAlgorithm selection	204
Table 188 – X.509 v3 Identity Token	205
Table 189 – IssuedIdentityToken	206
Table 190 – UserTokenPolicy	206

Table 191 – ViewDescription.....	207
Table A.1 – RelativePath	208
Table A.2 – <i>RelativePath</i> Examples	209
Table B.1 – ContentFilter example	211
Table B.2 – ContentFilter example	211
Table B.3 – Example 1 NodeTypeDescription	216
Table B.4 – Example 1 ContentFilter.....	216
Table B.5 – Example 1 QueryDataSets	217
Table B.6 – Example 2 NodeTypeDescription	217
Table B.7 – Example 2 ContentFilter.....	218
Table B.8 – Example 2 QueryDataSets	218
Table B.9 – Example 3 NodeTypeDescription	219
Table B.10 – Example 3 ContentFilter.....	220
Table B.11 – Example 3 QueryDataSets	221
Table B.12 – Example 4 NodeTypeDescription.....	221
Table B.13 – Example 4 ContentFilter	222
Table B.14 – Example 4 QueryDataSets	222
Table B.15 – Example 5 NodeTypeDescription.....	222
Table B.16 – Example 5 ContentFilter.....	223
Table B.17 – Example 5 QueryDataSets	223
Table B.18 – Example 6 NodeTypeDescription.....	223
Table B.19 – Example 6 ContentFilter.....	224
Table B.20 – Example 6 QueryDataSets.....	224
Table B.21 – Example 6 QueryDataSets without additional information.....	225
Table B.22 – Example 7 NodeTypeDescription.....	225
Table B.23 – Example 7 ContentFilter.....	226
Table B.24 – Example 7 QueryDataSets	226
Table B.25 – Example 8 NodeTypeDescription.....	227
Table B.26 – Example 8 ContentFilter.....	227
Table B.27 – Example 8 QueryDataSets	228
Table B.28 – Example 9 NodeTypeDescription.....	228
Table B.29 – Example 9 ContentFilter.....	229
Table B.30 – Example 9 QueryDataSets	229

INTERNATIONAL ELECTROTECHNICAL COMMISSION

OPC UNIFIED ARCHITECTURE –

Part 4: Services

FOREWORD

- 1) The International Electrotechnical Commission (IEC) is a worldwide organization for standardization comprising all national electrotechnical committees (IEC National Committees). The object of IEC is to promote international co-operation on all questions concerning standardization in the electrical and electronic fields. To this end and in addition to other activities, IEC publishes International Standards, Technical Specifications, Technical Reports, Publicly Available Specifications (PAS) and Guides (hereafter referred to as "IEC Publication(s)"). Their preparation is entrusted to technical committees; any IEC National Committee interested in the subject dealt with may participate in this preparatory work. International, governmental and non-governmental organizations liaising with the IEC also participate in this preparation. IEC collaborates closely with the International Organization for Standardization (ISO) in accordance with conditions determined by agreement between the two organizations.
- 2) The formal decisions or agreements of IEC on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC National Committees.
- 3) IEC Publications have the form of recommendations for international use and are accepted by IEC National Committees in that sense. While all reasonable efforts are made to ensure that the technical content of IEC Publications is accurate, IEC cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 4) In order to promote international uniformity, IEC National Committees undertake to apply IEC Publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any IEC Publication and the corresponding national or regional publication shall be clearly indicated in the latter.
- 5) IEC itself does not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. IEC is not responsible for any services carried out by independent certification bodies.
- 6) All users should ensure that they have the latest edition of this publication.
- 7) No liability shall attach to IEC or its directors, employees, servants or agents including individual experts and members of its technical committees and IEC National Committees for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication, use of, or reliance upon, this IEC Publication or any other IEC Publications.
- 8) Attention is drawn to the Normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.
- 9) Attention is drawn to the possibility that some of the elements of this IEC Publication may be the subject of patent rights. IEC shall not be held responsible for identifying any or all such patent rights.

International Standard IEC 62541-4 has been prepared by subcommittee 65E: Devices and integration in enterprise systems, of IEC technical committee 65: Industrial-process measurement, control and automation.

This third edition cancels and replaces the second edition published in 2015. This edition constitutes a technical revision.

This edition includes the following significant technical changes with respect to the previous edition:

- a) Added ability to resend all data of monitored items in a Subscription using the ResendData Method.
- b) Added support for durable Subscriptions (lifetime of hours or days).
- c) Added Register2 and FindServersOnNetwork Services to support network-wide discovery using capability filters.
- d) Removed definition of software certificates. Will be defined in a future edition.

- e) Extended and partially revised the redundancy definition. Added sub-range definitions for ServiceLevel and added more terms for redundancy.
- f) Added a section on how to use Authorization Services to request user access tokens.
- g) Added JSON Web Tokens (JWTs) as a new user token.
- h) Added the concept of session-less service invocation.
- i) Added a generic structure that allows passing any number of attributes to the AddNodes Service.
- j) Added requirement to protect against user identity token attacks.
- k) Added new EncryptedSecret format for user identity tokens.

The text of this standard is based on the following documents:

FDIS	Report on voting
65E/716/FDIS	65E/732/RVD

Full information on the voting for the approval of this International Standard can be found in the report on voting indicated in the above table.

This document has been drafted in accordance with the ISO/IEC Directives, Part 2.

Throughout this document and the other parts of the IEC 62541 series, certain document conventions are used:

Italics are used to denote a defined term or definition that appears in Clause 3 in one of the parts of the series.

Italics are also used to denote the name of a service input or output parameter or the name of a structure or element of a structure that are usually defined in tables.

The *italicized terms and names* are also, with a few exceptions, written in camel-case (the practice of writing compound words or phrases in which the elements are joined without spaces, with each element's initial letter capitalized within the compound). For example the defined term is *AddressSpace* instead of Address Space. This makes it easier to understand that there is a single definition for *AddressSpace*, not separate definitions for Address and Space.

A list of all parts of the IEC 62541 series, published under the general title *OPC Unified Architecture*, can be found on the IEC website.

The committee has decided that the contents of this document will remain unchanged until the stability date indicated on the IEC website under "<http://webstore.iec.ch>" in the data related to the specific document. At this date, the document will be

- reconfirmed,
- withdrawn,
- replaced by a revised edition, or
- amended.

IMPORTANT – The 'colour inside' logo on the cover page of this publication indicates that it contains colours which are considered to be useful for the correct understanding of its contents. Users should therefore print this document using a colour printer.

OPC UNIFIED ARCHITECTURE –

Part 4: Services

1 Scope

This part of IEC 62541 defines the OPC Unified Architecture (OPC UA) *Services*. The *Services* defined are the collection of abstract Remote Procedure Calls (RPC) that are implemented by OPC UA *Servers* and called by OPC UA *Clients*. All interactions between OPC UA *Clients* and *Servers* occur via these *Services*. The defined *Services* are considered abstract because no particular RPC mechanism for implementation is defined in this document. IEC 62541-6 specifies one or more concrete mappings supported for implementation. For example, one mapping in IEC 62541-6 is to XML Web Services. In that case the *Services* described in this document appear as the Web service methods in the WSDL contract.

Not all OPC UA *Servers* will need to implement all of the defined *Services*. IEC 62541-7 defines the *Profiles* that dictate which *Services* need to be implemented in order to be compliant with a particular *Profile*.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

IEC TR 62541-1, *OPC Unified Architecture – Part 1: Overview and Concepts*

IEC TR 62541-2, *OPC Unified Architecture – Part 2: Security Model*

IEC 62541-3, *OPC Unified Architecture – Part 3: Address Space Model*

IEC 62541-5, *OPC Unified Architecture – Part 5: Information Model*

IEC 62541-6, *OPC Unified Architecture – Part 6: Mappings*

IEC 62541-7, *OPC Unified Architecture – Part 7: Profiles*

IEC 62541-8, *OPC Unified Architecture – Part 8: Data Access*

IEC 62541-11, *OPC Unified Architecture – Part 11: Historical Access*

IEC 62541-12¹, *OPC Unified Architecture – Part 12: Discovery and Global Services*

IEC 62541-13, *OPC Unified Architecture – Part 13: Aggregates*

¹ Under preparation. Stage at the time of publication: IEC CDV 62541-12:2018.

3 Terms, definitions, abbreviated terms and conventions

3.1 Terms and definitions

For the purposes of this document, the terms and definitions given in IEC TR 62541-1, IEC TR 62541-2, IEC 62541-3 and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1.1

Active Server

Server which is currently sourcing information

Note 1 to entry: In OPC UA redundant systems, an *Active Server* is the *Server* that a *Client* is using as the source of data.

3.1.2

Deadband

permitted range for value changes that will not trigger a data change *Notification*

Note 1 to entry: *Deadband* can be applied as a filter when subscribing to *Variables* and is used to keep noisy signals from updating the *Client* unnecessarily. This document defines *AbsoluteDeadband* as a common filter. IEC 62541-8 defines an additional *Deadband* filter.

3.1.3

DiscoveryEndpoint

Endpoint that allows *Clients* access to *Discovery Services* without security

Note 1 to entry: A *DiscoveryEndpoint* allows access to *Discovery Services* without a *Session* and without message security.

3.1.4

Endpoint

physical address available on a network that allows *Clients* to access one or more *Services* provided by a *Server*

Note 1 to entry: Each *Server* may have multiple *Endpoints*. Each *Endpoint* includes a *HostName*.

3.1.5

Failed Server

Server that is not operational.

Note 1 to entry: In OPC UA redundant systems, a *Failed Server* is a *Server* that is unavailable or is not able to serve data.

3.1.6

Failover

act of switching the source or target of information

Note 1 to entry: In OPC UA redundant systems, a *Failover* is the act of a *Client* switching away from a failed or degraded *Server* to another *Server* in the redundant set (*Server Failover*). In some cases a *Client* may have no knowledge of a *Failover* action occurring (transparent redundancy). A *Client Failover* is the act of an alternate *Client* replacing an existing failed or degraded *Client* connection to a *Server*.

3.1.7

Gateway Server

Server that acts as an intermediary for one or more *Servers*

Note 1 to entry: *Gateway Servers* may be deployed to limit external access, provide protocol conversion or to provide features that the underlying *Servers* do not support.

3.1.8

HostName

unique identifier for a machine on a network

Note 1 to entry: This identifier is unique within a local network; however, it may also be globally unique. The identifier can be an IP address.

3.1.9

Redundancy

presence of duplicate components enabling the continued operation after a failure of an OPC UA component

Note 1 to entry: This may apply to *Servers*, *Clients* or networks.

3.1.10

Redundant server set

two or more *Servers* that are redundant with each other

Note 1 to entry: A *Redundant Server Set* is a group of *Servers* that are configured to provide *Redundancy*. These *Servers* have requirements related to the address space and provide *Failovers*.

3.1.11

SecurityToken

identifier for a cryptographic key set

Note 1 to entry: All *SecurityTokens* belong to a security context. For OPC UA the security context is the *SecureChannel*.

3.1.12

SoftwareCertificate

digital certificate for a software product that can be installed on several hosts to describe the capabilities of the software product

Note 1 to entry: Different installations of one software product could have the same software certificate. Software certificates are not relevant for security. They are used to identify a software product and its supported features. *SoftwareCertificates* are described in 6.4.

3.2 Abbreviated terms

API	Application Programming Interface
BNF	Backus-Naur Form
CA	Certificate Authority
CRL	Certificate Revocation List
CTL	Certificate Trust List
DA	Data Access
NAT	Network Address Translation
UA	Unified Architecture
URI	Uniform Resource Identifier
URL	Uniform Resource Locator

3.3 Conventions for Service definitions

OPC UA *Services* contain parameters that are conveyed between the *Client* and the *Server*. The OPC UA *Service* specifications use tables to describe *Service* parameters, as shown in Table 1. Parameters are organized in this table into request parameters and response parameters.

Table 1 – Service definition table

Name	Type	Description
Request		Defines the request parameters of the <i>Service</i>
Simple Parameter Name		Description of this parameter
Constructed Parameter Name		Description of the constructed parameter
Component Parameter Name		Description of the component parameter
Response		Defines the response parameters of the <i>Service</i>

The Name, Type and Description columns contain the name, data type and description of each parameter. All parameters are mandatory, although some may be unused under certain circumstances. The Description column specifies the value to be supplied when a parameter is unused.

Two types of parameters are defined in these tables, simple and constructed. Simple parameters have a simple data type, such as *Boolean* or *String*.

Constructed parameters are composed of two or more component parameters, which can be simple or constructed. Component parameter names are indented below the constructed parameter name.

The data types used in these tables may be base types, common types to multiple *Services* or *Service*-specific types. Base data types are defined in IEC 62541-3. The base types used in *Services* are listed in Table 2. Data types that are common to multiple *Services* are defined in Clause 7. Data types that are *Service*-specific are defined in the parameter table of the *Service*.

Table 2 – Parameter Types defined in IEC 62541-3

Parameter Type
BaseDataType
Boolean
ByteString
Double
Duration
Guid
Int32
LocaleId
NodeId
QualifiedName
String
UInt16
UInt32
UInteger
UtcTime
XmlElement

The parameters of the Request and Indication service primitives are represented in Table 1 as Request parameters. Likewise, the parameters of the Response and Confirmation service primitives are represented in Table 1 as Response parameters. All request and response parameters are conveyed between the sender and receiver without change. Therefore, separate columns for request, indication, response and confirmation parameter values are not needed and have been intentionally omitted to improve readability.

4 Overview

4.1 Service Set model

This clause specifies the OPC UA *Services*. The OPC UA *Service* definitions are abstract descriptions and do not represent a specification for implementation. The mapping between the abstract descriptions and the *Communication Stack* derived from these *Services* are defined in IEC 62541-6. In the case of an implementation as web services, the OPC UA *Services* correspond to the web service and an OPC UA *Service* corresponds to an operation of the web service.

These *Services* are organized into *Service Sets*. Each *Service Set* defines a set of related *Services*. The organization in *Service Sets* is a logical grouping used in this document and is not used in the implementation.

The *Discovery Service Set*, illustrated in Figure 1, defines *Services* that allow a *Client* to discover the *Endpoints* implemented by a *Server* and to read the security configuration for each of those *Endpoints*.




Figure 1 – Discovery Service Set

The *SecureChannel Service Set*, illustrated in Figure 2, defines *Services* that allow a *Client* to establish a communication channel to ensure the *Confidentiality* and *Integrity* of *Messages* exchanged with the *Server*.




Figure 2 – SecureChannel Service Set

The *Session Service Set*, illustrated in Figure 3, defines *Services* that allow the *Client* to authenticate the user on whose behalf it is acting and to manage *Sessions*.




Figure 3 – Session Service Set

The *NodeManagement Service Set*, illustrated in Figure 4, defines *Services* that allow the *Client* to add, modify and delete *Nodes* in the *AddressSpace*.




Figure 4 – NodeManagement Service Set

The *View Service Set*, illustrated in Figure 5, defines *Services* that allow *Clients* to browse through the *AddressSpace* or subsets of the *AddressSpace* called *Views*. The *Query Service Set* allows *Clients* to get a subset of data from the *AddressSpace* or the *View*.




Figure 5 – View Service Set

The *Attribute Service Set* is illustrated in Figure 6. It defines *Services* that allow *Clients* to read and write *Attributes* of *Nodes*, including their historical values. Since the value of a *Variable* is modelled as an *Attribute*, these *Services* allow *Clients* to read and write the values of *Variables*.




Figure 6 – Attribute Service Set

The *Method Service Set* is illustrated in Figure 7. It defines *Services* that allow *Clients* to call methods. Methods run to completion when called. They may be called with method-specific input parameters and may return method-specific output parameters.




Figure 7 – Method Service Set

The *MonitoredItem Service Set* and the *Subscription Service Set*, illustrated in Figure 8, are used together to subscribe to *Nodes* in the OPC UA AddressSpace.

The *MonitoredItem Service Set* defines *Services* that allow *Clients* to create, modify, and delete *MonitoredItems* used to monitor *Attributes* for value changes and *Objects* for *Events*.

These *Notifications* are queued for transfer to the *Client* by *Subscriptions*.

The *Subscription Service Set* defines *Services* that allow *Clients* to create, modify and delete *Subscriptions*. *Subscriptions* send *Notifications* generated by *MonitoredItems* to the *Client*. *Subscription Services* also provide for *Client* recovery from missed *Messages* and communication failures.




Figure 8 – MonitoredItem and Subscription Service Sets

4.2 Request/response Service procedures

Request/response *Service* procedures describe the processing of requests received by the *Server*, and the subsequent return of responses. The procedures begin with the requesting *Client* submitting a *Service* request *Message* to the *Server*.

Upon receipt of the request, the *Server* processes the *Message* in two steps. In the first step, it attempts to decode and locate the *Service* to execute. The error handling for this step is specific to the communication technology used and is described in IEC 62541-6.

If it succeeds, then it attempts to access each operation identified in the request and perform the requested operation. For each operation in the request, it generates a separate success/failure code that it includes in a positive response *Message* along with any data that is to be returned.

To perform these operations, both the *Client* and the *Server* may make use of the API of a *Communication Stack* to construct and interpret *Messages* and to access the requested operation.

The implementation of each service request or response handling shall check that each service parameter lies within the specified range for that parameter.

5 Service Sets

5.1 General

This clause defines the OPC UA *Service Sets* and their *Services*. Clause 7 contains the definitions of common parameters used by these *Services*. Subclause 6.5 describes auditing requirements for all services.

Whether or not a *Server* supports a *Service Set*, or a *Service* within a *Service Set*, is defined by its *Profile*. *Profiles* are described in IEC 62541-7.

5.2 Service request and response header

Each *Service* request has a *RequestHeader* and each *Service* response has a *ResponseHeader*.

The *RequestHeader* structure is defined in 7.28 and contains common request parameters such as *authenticationToken*, *timestamp* and *requestHandle*.

The *ResponseHeader* structure is defined in 7.29 and contains common response parameters such as *serviceResult* and *diagnosticInfo*.

5.3 Service results

Service results are returned at two levels in OPC UA responses, one that indicates the status of the *Service* call, and the other that indicates the status of each operation requested by the *Service*.

Service results are defined via the *StatusCode* (see 7.34).

The status of the *Service* call is represented by the *serviceResult* contained in the *ResponseHeader* (see 7.29). The mechanism for returning this parameter is specific to the communication technology used to convey the *Service* response and is defined in IEC 62541-6.

The status of individual operations in a request is represented by individual *StatusCodes*.

The following cases define the use of these parameters.

- a) A bad code is returned in *serviceResult* if the *Service* itself failed. In this case, a *ServiceFault* is returned. The *ServiceFault* is defined in 7.30.
- b) The good code is returned in *serviceResult* if the *Service* fully or partially succeeded. In this case, other response parameters are returned. The *Client* shall always check the response parameters, especially all *StatusCodes* associated with each operation. These *StatusCodes* may indicate bad or uncertain results for one or more operations requested in the *Service* call.

All *Services* with arrays of operations in the request shall return a bad code in the *serviceResult* if the array is empty.

The *Services* define various specific *StatusCodes* and a *Server* shall use these specific *StatusCodes* as described in the *Service*. A *Client* should be able to handle these *Service* specific *StatusCodes*. In addition, a *Client* shall expect other common *StatusCodes* defined in Table 177 and Table 178. Additional details for *Client* handling of specific *StatusCodes* may be defined in IEC 62541-7.

If the *Server* discovers, through some out-of-band mechanism that the application or user credentials used to create a *Session* or *SecureChannel* have been compromised, then the *Server* should immediately terminate all sessions and channels that use those credentials. In this case, the *Service* result code should be either *Bad_IdentityTokenRejected* or *Bad_CertificateUntrusted*.

Message parsing can fail due to syntax errors or if data contained within the message exceeds ranges supported by the receiver. When this happens messages shall be rejected by the receiver. If the receiver is a *Server* then it shall return a *ServiceFault* with result code of

Bad_DecodingError or *Bad_EncodingLimitsExceeded*. If the receiver is the *Client* then the *Communication Stack* should report these errors to the *Client* application.

Many applications will place limits on the size of messages and/or data elements contained within these messages. For example, a *Server* may reject requests containing string values longer than a certain length. These limits are typically set by administrators and apply to all connections between a *Client* and a *Server*.

Clients that receive *Bad_EncodingLimitsExceeded* faults from the *Server* will likely need to reformulate their requests. The administrator may need to increase the limits for the *Client* if it receives a response from the *Server* with this fault.

In some cases, parsing errors are fatal and it is not possible to return a fault. For example, the incoming message could exceed the buffer capacity of the receiver. In these cases, these errors may be treated as a communication fault which requires the *SecureChannel* to be re-established (see 5.5).

The *Client* and *Server* reduce the chances of a fatal error by exchanging their message size limits in the *CreateSession* service. This will allow either party to avoid sending a message that causes a communication fault. The *Server* should return a *Bad_ResponseTooLarge* fault if a serialized response message exceeds the message size specified by the *Client*. Similarly, the *Client Communication Stack* should report a *Bad_RequestTooLarge* error to the application before sending a message that exceeds the *Server's* limit.

Note that the message size limits only apply to the raw message body and do not include headers or the effect of applying any security. This means that a message body that is smaller than the specified maximum could still cause a fatal error.

5.4 Discovery Service Set

5.4.1 Overview

This *Service Set* defines *Services* used to discover the *Endpoints* implemented by a *Server* and to read the security configuration for those *Endpoints*. The *Discovery Services* are implemented by individual *Servers* and by dedicated *Discovery Servers*. IEC 62541-12 describes how to use the *Discovery Services* with dedicated *Discovery Servers*.

Every *Server* shall have a *DiscoveryEndpoint* that *Clients* can access without establishing a *Session*. This *Endpoint* may or may not be the same *Session Endpoint* that *Clients* use to establish a *SecureChannel*. *Clients* read the security information necessary to establish a *SecureChannel* by calling the *GetEndpoints* Service on the *DiscoveryEndpoint*.

In addition, *Servers* may register themselves with a well-known *Discovery Server* using the *RegisterServer* Service. *Clients* can later discover any registered *Servers* by calling the *FindServers* Service on the *Discovery Server*.

The discovery process using *FindServers* is illustrated in Figure 9. The establishment of a *SecureChannel* (with *MessageSecurityMode* NONE) for *FindServers* and *GetEndpoints* is omitted from the figure for clarity.




Figure 9 – Discovery process

The URL for a *DiscoveryEndpoint* shall provide all of the information that the *Client* needs to connect to the *DiscoveryEndpoint*.

Once a *Client* retrieves the *Endpoints*, the *Client* can save this information and use it to connect directly to the *Server* again without going through the discovery process. If the *Client* finds that it cannot connect then the *Server* configuration may have changed and the *Client* needs to go through the discovery process again.

DiscoveryEndpoints shall not require any message security, but it may require transport layer security. In production systems, Administrators may disable discovery for security reasons and *Clients* shall rely on cached *EndpointDescriptions*. To provide support for systems with disabled *Discovery Services* *Clients* shall allow *Administrators* to manually update the *EndpointDescriptions* used to connect to a *Server*. *Servers* shall allow *Administrators* to disable the *DiscoveryEndpoint*.

A *Client* shall be careful when using the information returned from a *DiscoveryEndpoint* since it has no security. A *Client* does this by comparing the information returned from the *DiscoveryEndpoint* to the information returned in the *CreateSession* response. A *Client* shall verify that:

- The *ApplicationUri* specified in the *Server Certificate* is the same as the *ApplicationUri* provided in the *EndpointDescription*.
- The *Server Certificate* returned in *CreateSession* response is the same as the *Certificate* used to create the *SecureChannel*.
- The *EndpointDescriptions* returned from the *DiscoveryEndpoint* are the same as the *EndpointDescriptions* returned in the *CreateSession* response.

If the *Client* detects that one of the above requirements is not fulfilled, then the *Client* shall close the *SecureChannel* and report an error.

A *Client* shall verify the *HostName* specified in the *Server Certificate* is the same as the *HostName* contained in the *endpointUrl* provided in the *EndpointDescription* returned by *CreateSession*. If there is a difference then the *Client* shall report the difference and may close the *SecureChannel*. *Servers* shall add all possible *HostNames* like *MyHost* and *MyHost.local* into the *Server Certificate*. This includes IP addresses of the host or the *HostName* exposed by a NAT router used to connect to the *Server*.

5.4.2 FindServers

5.4.2.1 Description

This *Service* returns the *Servers* known to a *Server* or *Discovery Server*. The behaviour of *Discovery Servers* is described in detail in IEC 62541-12.

The *Client* may reduce the number of results returned by specifying filter criteria. A *Discovery Server* returns an empty list if no *Servers* match the criteria specified by the client. The filter criteria supported by this *Service* are described in 5.4.2.2.

Every *Server* shall provide a *DiscoveryEndpoint* that supports this *Service*. The *Server* shall always return a record that describes itself, however in some cases more than one record may be returned. *Gateway Servers* shall return a record for each *Server* that they provide access to plus (optionally) a record that allows the *Gateway Server* to be accessed as an ordinary OPC UA *Server*. Non-transparent redundant *Servers* shall provide a record for each *Server* in the *Redundant Server Set*.

Every *Server* shall have a globally unique identifier called the *ServerUri*. This identifier should be a fully qualified domain name; however, it may be a GUID or similar construct that ensures global uniqueness. The *ServerUri* returned by this *Service* shall be the same value that appears in index 0 of the *ServerArray* property (see IEC 62541-5). The *ServerUri* is returned as the *applicationUri* field in the *ApplicationDescription* (see 7.1)

Every *Server* shall also have a human readable identifier called the *ServerName* which is not necessarily globally unique. This identifier may be available in multiple locales.

A *Server* may have multiple *HostNames*. For this reason, the *Client* shall pass the URL it used to connect to the *Endpoint* to this *Service*. The implementation of this *Service* shall use this information to return responses that are accessible to the *Client* via the provided URL.

This *Service* shall not require message security but it may require transport layer security.

Some *Servers* may be accessed via a *Gateway Server* and shall have a value specified for *gatewayServerUri* in their *ApplicationDescription* (see 7.1). The *discoveryUrls* provided in *ApplicationDescription* shall belong to the *Gateway Server*. Some *Discovery Servers* may return multiple records for the same *Server* if that *Server* can be accessed via multiple paths.

This *Service* can be used without security and it is therefore vulnerable to Denial of Service (DOS) attacks. A *Server* should minimize the amount of processing required to send the response for this *Service*. This can be achieved by preparing the result in advance. The *Server* should also add a short delay before starting processing of a request during high traffic conditions.

5.4.2.2 Parameters

Table 3 defines the parameters for the *Service*.

Table 3 – FindServers Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The <i>authenticationToken</i> shall be ignored if it is provided. The type <i>RequestHeader</i> is defined in 7.28.
endpointUrl	String	The network address that the <i>Client</i> used to access the <i>DiscoveryEndpoint</i> . The <i>Server</i> uses this information for diagnostics and to determine what URLs to return in the response. The <i>Server</i> should return a suitable default URL if it does not recognize the <i>HostName</i> in the URL.
localeIds []	LocaleId	List of locales to use. The <i>Server</i> should return the <i>applicationName</i> in the <i>ApplicationDescription</i> defined in 7.1 using one of locales specified. If the <i>Server</i> supports more than one of the requested locales then the <i>Server</i> shall use the locale that appears first in this list. If the <i>Server</i> does not support any of the requested locales it chooses an appropriate default locale. The <i>Server</i> chooses an appropriate default locale if this list is empty.
serverUris []	String	List of <i>Servers</i> to return. All known <i>Servers</i> are returned if the list is empty. A <i>serverUri</i> matches the <i>applicationUri</i> from the <i>ApplicationDescription</i> defined in 7.1.
Response		
responseHeader	ResponseHeader	Common response parameters. The <i>ResponseHeader</i> type is defined in 7.29.
servers []	ApplicationDescription	List of <i>Servers</i> that meet criteria specified in the request. This list is empty if no <i>Servers</i> meet the criteria. The <i>ApplicationDescription</i> type is defined in 7.1.

5.4.2.3 Service results

Common *StatusCodes* are defined in Table 177.

5.4.3 FindServersOnNetwork

5.4.3.1 Description

This *Service* returns the *Servers* known to a *Discovery Server*. Unlike *FindServers*, this *Service* is only implemented by *Discovery Servers*.

The *Client* may reduce the number of results returned by specifying filter criteria. An empty list is returned if no *Server* matches the criteria specified by the *Client*.

This *Service* shall not require message security but it may require transport layer security.

Each time the *Discovery Server* creates or updates a record in its cache, it shall assign a monotonically increasing identifier to the record. This allows *Clients* to request records in batches by specifying the identifier for the last record received in the last call to *FindServersOnNetwork*. To support this the *Discovery Server* shall return records in numerical order starting from the lowest record identifier. The *Discovery Server* shall also return the last time the counter was reset, for example due to a restart of the *Discovery Server*. If a *Client*

detects that this time is more recent than the last time the *Client* called the *Service*, it shall call the *Service* again with a startingRecordId of 0.

This *Service* can be used without security and it is therefore vulnerable to denial of service (DOS) attacks. A *Server* should minimize the amount of processing required to send the response for this *Service*. This can be achieved by preparing the result in advance.

5.4.3.2 Parameters

Table 4 defines the parameters for the *Service*.

Table 4 – FindServersOnNetwork Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The <i>authenticationToken</i> shall be ignored if it is provided. The type <i>RequestHeader</i> is defined in 7.28.
startingRecordId	Counter	Only records with an identifier greater than this number will be returned. Specify 0 to start with the first record in the cache.
maxRecordsToReturn	UInt32	The maximum number of records to return in the response. 0 indicates that there is no limit.
serverCapabilityFilter[]	String	List of <i>Server</i> capability filters. The set of allowed <i>Server</i> capabilities are defined in IEC 62541-12. Only records with all of the specified <i>Server</i> capabilities are returned. The comparison is case insensitive. If this list is empty then no filtering is performed.
Response		
responseHeader	ResponseHeader	Common response parameters. The <i>ResponseHeader</i> type is defined in 7.29.
lastCounterResetTime	UtcTime	The last time the counters were reset.
servers[]	ServerOnNetwork	List of DNS service records that meet criteria specified in the request. This list is empty if no <i>Servers</i> meet the criteria.
recordId	UInt32	A unique identifier for the record. This can be used to fetch the next batch of <i>Servers</i> in a subsequent call to <i>FindServersOnNetwork</i> .
serverName	String	The name of the <i>Server</i> specified in the mDNS announcement (see IEC 62541-12). This may be the same as the <i>ApplicationName</i> for the <i>Server</i> .
discoveryUrl	String	The URL of the <i>DiscoveryEndpoint</i> .
serverCapabilities	String[]	The set of <i>Server</i> capabilities supported by the <i>Server</i> . The set of allowed <i>Server</i> capabilities are defined in IEC 62541-12.

5.4.3.3 Service results

Common *StatusCodes* are defined in Table 177.

5.4.4 GetEndpoints

5.4.4.1 Description

This *Service* returns the *Endpoints* supported by a *Server* and all of the configuration information required to establish a *SecureChannel* and a *Session*.

This *Service* shall not require message security but it may require transport layer security.

A *Client* may reduce the number of results returned by specifying filter criteria based on *LocaleIds* and *Transport Profile* URIs. The *Server* returns an empty list if no *Endpoints* match the criteria specified by the client. The filter criteria supported by this *Service* are described in 5.4.4.2.

A *Server* may support multiple security configurations for the same *Endpoint*. In this situation, the *Server* shall return separate *EndpointDescription* records for each available configuration. *Clients* should treat each of these configurations as distinct *Endpoints* even if the physical URL happens to be the same.

The security configuration for an *Endpoint* has four components:

- Server Application Instance Certificate
- Message Security Mode
- Security Policy
- Supported User Identity Tokens

The *ApplicationInstanceCertificate* is used to secure the *OpenSecureChannel* request (see 5.5.2). The *MessageSecurityMode* and the *SecurityPolicy* tell the *Client* how to secure messages sent via the *SecureChannel*. The *UserIdentityTokens* tell the *Client* which type of user credentials shall be passed to the *Server* in the *ActivateSession* request (see 5.6.3).

If the *securityPolicyUri* is NONE and none of the *UserTokenPolicies* requires encryption, the *Client* shall ignore the *ApplicationInstanceCertificate*.

Each *EndpointDescription* also specifies a URI for the *Transport Profile* that the *Endpoint* supports. The *Transport Profiles* specify information such as message encoding format and protocol version and are defined in IEC 62541-7.

Messages are secured by applying standard cryptography algorithms to the messages before they are sent over the network. The exact set of algorithms used depends on the *SecurityPolicy* for the *Endpoint*. IEC 62541-7 defines *Profiles* for common *SecurityPolicies* and assigns a unique URI to them. It is expected that applications have built in knowledge of the *SecurityPolicies* that they support, as a result, only the Profile URI for the *SecurityPolicy* is specified in the *EndpointDescription*. A *Client* cannot connect to an *Endpoint* that does not support a *SecurityPolicy* that it recognizes.

An *EndpointDescription* may specify that the message security mode is NONE. This configuration is not recommended unless the applications are communicating on a physically isolated network where the risk of intrusion is extremely small. If the message security is NONE then it is possible for *Clients* to deliberately or accidentally hijack *Sessions* created by other *Clients*.

A *Server* may have multiple *HostNames*. For this reason, the *Client* shall pass the URL it used to connect to the *Endpoint* to this *Service*. The implementation of this *Service* shall use this information to return responses that are accessible to the *Client* via the provided URL.

This *Service* can be used without security and it is therefore vulnerable to Denial of Service (DOS) attacks. A *Server* should minimize the amount of processing required to send the

response for this *Service*. This can be achieved by preparing the result in advance. The *Server* should also add a short delay before starting processing of a request during high traffic conditions.

Some of the *EndpointDescriptions* returned in a response shall specify the *Endpoint* information for a *Gateway Server* that can be used to access another *Server*. In these situations, the *gatewayServerUri* is specified in the *EndpointDescription* and all security checks used to verify *Certificates* shall use the *gatewayServerUri* (see 6.1.3) instead of the *serverUri*.

To connect to a *Server* via the gateway the *Client* shall first establish a *SecureChannel* with the *Gateway Server*. Then the *Client* shall call the *CreateSession* service and pass the *serverUri* specified in the *EndpointDescription* to the *Gateway Server*. The *Gateway Server* shall then connect to the underlying *Server* on behalf of the *Client*. The process of connecting to a *Server* via a *Gateway Server* is illustrated in Figure 10.




Figure 10 – Using a Gateway Server

5.4.4.2 Parameters

Table 5 defines the parameters for the *Service*.

Table 5 – GetEndpoints Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The <i>authenticationToken</i> shall be ignored if it is provided. The type <i>RequestHeader</i> is defined in 7.28.
endpointUrl	String	The network address that the <i>Client</i> used to access the <i>DiscoveryEndpoint</i> . The <i>Server</i> uses this information for diagnostics and to determine what URLs to return in the response. The <i>Server</i> should return a suitable default URL if it does not recognize the <i>HostName</i> in the URL.
localeIds []	LocaleId	List of locales to use. Specifies the locale to use when returning human readable strings. This parameter is described in 5.4.2.2.
profileUris []	String	List of <i>Transport Profile</i> that the returned <i>Endpoints</i> shall support. IEC 62541-7 defines URIs for the <i>Transport Profiles</i> . All <i>Endpoints</i> are returned if the list is empty. If the URI is a URL, this URL may have a query string appended. The <i>Transport Profiles</i> that support query strings are defined in IEC 62541-7.
Response		
responseHeader	ResponseHeader	Common response parameters. The <i>ResponseHeader</i> type is defined in 7.29.
Endpoints []	EndpointDescription	List of <i>Endpoints</i> that meet criteria specified in the request. This list is empty if no <i>Endpoints</i> meet the criteria. The <i>EndpointDescription</i> type is defined in 7.10.

5.4.4.3 Service results

Common *StatusCodes* are defined in Table 177.

5.4.5 RegisterServer

5.4.5.1 Description

This Service is implemented by *Discovery Servers*.

This Service registers a *Server* with a *Discovery Server*. This Service will be called by a *Server* or a separate configuration utility. *Clients* will not use this Service.

A *Server* shall establish a *SecureChannel* with the *Discovery Server* before calling this Service. The *SecureChannel* is described in 5.5. The *Administrator* of the *Server* shall provide the *Server* with an *EndpointDescription* for the *Discovery Server* as part of the configuration process. *Discovery Servers* shall reject registrations if the *serverUri* provided does not match the *applicationUri* in *Server Certificate* used to create the *SecureChannel*.

This Service can only be invoked via *SecureChannels* that support *Client* authentication (i.e. HTTPS cannot be used to call this Service).


A *Server* only provides its *serverUri* and the URLs of the *DiscoveryEndpoints* to the *Discovery Server*. *Clients* shall use the *GetEndpoints Service* to fetch the most up to date configuration information directly from the *Server*.

The *Server* shall provide a localized name for itself in all locales that it supports.

Servers shall be able to register themselves with a *Discovery Server* running on the same machine. The exact mechanisms depend on the *Discovery Server* implementation and are described in IEC 62541-6.

There are two types of *Server* applications: those which are manually launched including a start by the operating system at boot and those that are automatically launched when a *Client* attempts to connect. The registration process that a *Server* shall use depends on which category it falls into.


The registration process for manually launched *Servers* is illustrated in Figure 11.



IEC

Figure 11 – The registration process – Manually launched servers

The registration process for automatically launched *Servers* is illustrated in Figure 12.



IEC

Figure 12 – The registration process – Automatically launched servers

The registration process is designed to be platform independent, robust and able to minimize problems created by configuration errors. For that reason, *Servers* shall register themselves more than once.

Under normal conditions, manually launched *Servers* shall periodically register with the *Discovery Server* as long as they are able to receive connections from *Clients*. If a *Server* goes offline then it shall register itself once more and indicate that it is going offline. The period of the recurring registration should be configurable; however, the maximum is 10 minutes. If an error occurs during registration (e.g. the *Discovery Server* is not running) then the *Server* shall periodically re-attempt registration. The frequency of these attempts should start at 1 second but gradually increase until the registration frequency is the same as what it would be if no errors occurred. The recommended approach would be to double the period of each attempt until reaching the maximum.

When an automatically launched *Server* (or its install program) registers with the *Discovery Server* it shall provide a path to a semaphore file which the *Discovery Server* can use to determine if the *Server* has been uninstalled from the machine. The *Discovery Server* shall have read access to the file system that contains the file.

5.4.5.2 Parameters

Table 6 defines the parameters for the *Service*.

Table 6 – RegisterServer Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
Server	RegisteredServer	The <i>Server</i> to register. The type <i>RegisteredServer</i> is defined in 7.27.
Response		
ResponseHeader	ResponseHeader	Common response parameters. The type <i>ResponseHeader</i> is defined in 7.29.

5.4.5.3 Service results

Table 7 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 7 – RegisterServer Service result codes

Symbolic Id	Description
Bad_InvalidArgument	See Table 177 for the description of this result code.
Bad_ServerUriInvalid	See Table 177 for the description of this result code.
Bad_ServerNameMissing	No <i>ServerName</i> was specified.
Bad_DiscoveryUriMissing	No discovery URL was specified.
Bad_SemaphoreFileMissing	The semaphore file specified is not valid.

5.4.6 RegisterServer2

5.4.6.1 Description

This Service is implemented by *Discovery Servers*.

This *Service* allows a *Server* to register its *DiscoveryUrls* and capabilities with a *Discovery Server*. It extends the registration information from *RegisterServer* with information necessary for *FindServersOnNetwork*. This *Service* will be called by a *Server* or a separate configuration utility. *Clients* will not use this *Service*.

Servers that support *RegisterServer2* shall try to register with the *Discovery Server* using this Service and shall fall back to *RegisterServer* if *RegisterServer2* fails with the status *Bad_ServiceUnsupported*.

A *Discovery Server* that implements this *Service* needs to assign unique record ids each time this *Service* is called. See 5.4.3 for more details.

This *Service* can only be invoked via *SecureChannels* that support *Client* authentication (i.e. HTTPS cannot be used to call this *Service*).

5.4.6.2 Parameters

Table 8 defines the parameters for the *Service*.

Table 8 – RegisterServer2

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
Server	RegisteredServer	The <i>Server</i> to register. The type <i>RegisteredServer</i> is defined in 7.27.
discoveryConfiguration []	ExtensibleParameter DiscoveryConfiguration	Additional configuration settings for the <i>Server</i> to register. The <i>discoveryConfiguration</i> is an extensible parameter type defined in 7.9. Discovery Servers that do not understand a configuration shall return <i>Bad_NotSupported</i> for this configuration.
Response		
ResponseHeader	ResponseHeader	Common response parameters. The type <i>ResponseHeader</i> is defined in 7.29.
configurationResults []	StatusCode	List of results for the <i>discoveryConfiguration</i> parameters.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>discoveryConfiguration</i> parameters. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.4.6.3 Service results

Table 9 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 9 – RegisterServer2 Service result codes

Symbolic Id	Description
Bad_InvalidArgument	See Table 177 for the description of this result code.
Bad_ServerUriInvalid	See Table 177 for the description of this result code.
Bad_ServerNameMissing	No <i>ServerName</i> was specified.
Bad_DiscoveryUriMissing	No discovery URL was specified.
Bad_SemaphoreFileMissing	The semaphore file specified is not valid.
Bad_ServiceUnsupported	See Table 177 for the description of this result code.

5.4.6.4 StatusCodes

Table 10 defines values for the operation level *configurationResults* parameters that are specific to this Service. Common *StatusCodes* are defined in Table 178.

Table 10 – RegisterServer2 Operation Level result codes

Symbolic Id	Description
Bad_NotSupported	See Table 178 for the description of this result code.

5.5 SecureChannel Service Set

5.5.1 Overview

This *Service Set* defines *Services* used to open a communication channel that ensures the confidentiality and *Integrity* of all *Messages* exchanged with the *Server*. The base concepts for OPC UA security are defined in IEC TR 62541-2.

The *SecureChannel Services* are unlike other *Services* because they are not implemented directly by the *OPC UA Application*. Instead, they are provided by the *Communication Stack* on which the *OPC UA Application* is built. For example, an OPC UA *Server* may be built on a stack that allows applications to establish a *SecureChannel* using HTTPS. In these cases, the *OPC UA Application* shall verify that the *Message* it received was in the context of an HTTPS connection. IEC 62541-6 describes how the *SecureChannel Services* are implemented.

A *SecureChannel* is a long-running logical connection between a single *Client* and a single *Server*. This channel maintains a set of keys known only to the *Client* and *Server*, which are used to authenticate and encrypt *Messages* sent across the network. The *SecureChannel Services* allow the *Client* and *Server* to securely negotiate the keys to use.

Logical connections may be initiated by the *Client* or by the *Server* as described in IEC 62541-6. After the connection is initiated, the *SecureChannel* is opened and closed by the *Client* using the *SecureChannel Services*.

An *EndpointDescription* tells a *Client* how to establish a *SecureChannel* with a given *Endpoint*. A *Client* may obtain the *EndpointDescription* from a *Discovery Server*, via some non-UA defined directory server or from its own configuration.

The exact algorithms used to authenticate and encrypt *Messages* are described in the *SecurityPolicy* field of the *EndpointDescription*. A *Client* shall use these algorithms when it creates a *SecureChannel*.

It should be noted that some *SecurityPolicies* defined in IEC 62541-7 will turn off authentication and encryption resulting in a *SecureChannel* that provides no security.

When a *Client* and *Server* are communicating via a *SecureChannel*, they shall verify that all incoming *Messages* have been signed and encrypted according to the requirements specified in the *EndpointDescription*. An *OPC UA Application* shall not process any *Message* that does not conform to these requirements.

The relationship between the *SecureChannel* and the *OPC UA Application* depends on the implementation technology. IEC 62541-6 defines any requirements that depend on the technology used.

The correlation between the *OPC UA Application Session* and the *SecureChannel* is illustrated in Figure 13. The *Communication Stack* is used by the *OPC UA Applications* to exchange *Messages*. In the first step, the *SecureChannel Services* are used to establish a *SecureChannel* between the two *Communication Stacks* which allows the secure exchange of *Messages*. In the second step, the *OPC UA Applications* use the *Session Service Set* to establish an *OPC UA Application Session*.




Figure 13 – SecureChannel and Session Services

Once a *Client* has established a *Session* it may wish to access the *Session* from a different *SecureChannel*. The *Client* can do this by validating the new *SecureChannel* with the *ActivateSession Service* described in 5.6.3.

If a *Server* acts as a *Client* to other *Servers*, which is commonly referred to as *Server chaining*, then the *Server* shall be able to maintain user level security. By this we mean that the user identity should be passed to the underlying *Server* or it should be mapped to an appropriate user identity in the underlying server. It is unacceptable to ignore user level security. This is required to ensure that security is maintained and that a user does not obtain information that they should not have access to. Whenever possible a *Server* should impersonate the original *Client* by passing the original *Client's* user identity to the underlying *Server* when it calls the *ActivateSession Service*. If impersonation is not an option then the *Server* shall map the original *Client's* user identity onto a new user identity which the underlying *Server* does recognize.

5.5.2 OpenSecureChannel

5.5.2.1 Description

This *Service* is used to open or renew a *SecureChannel* that can be used to ensure *Confidentiality* and *Integrity* for *Message* exchange during a *Session*. This *Service* requires the *Communication Stack* to apply the various security algorithms to the *Messages* as they are sent and received. Specific implementations of this *Service* for different *Communication Stacks* are described in IEC 62541-6.

Each *SecureChannel* has a globally-unique identifier and is valid for a specific combination of *Client* and *Server* application instances. Each channel contains one or more *SecurityTokens* that identify a set of cryptography keys that are used to encrypt and authenticate *Messages*. *SecurityTokens* also have globally-unique identifiers which are attached to each *Message* secured with the token. This allows an authorized receiver to know how to decrypt and verify the *Message*.

SecurityTokens have a finite lifetime negotiated with this *Service*. However, differences between the system clocks on different machines and network latencies mean that valid *Messages* could arrive after the token has expired. To prevent valid *Messages* from being discarded, the applications should do the following:

- a) *Clients* should request a new *SecurityToken* after 75 % of its lifetime has elapsed. This should ensure that *Clients* will receive the new *SecurityToken* before the old one actually expires.
- b) *Servers* shall use the existing *SecurityToken* to secure outgoing *Messages* until the *SecurityToken* expires or the *Server* receives a *Message* secured with a new *SecurityToken*. This should ensure that *Clients* do not reject *Messages* secured with the new *SecurityToken* that arrive before the *Client* receives the new *SecurityToken*.
- c) *Clients* should accept *Messages* secured by an expired *SecurityToken* for up to 25 % of the token lifetime. This should ensure that *Messages* sent by the *Server* before the token expired are not rejected because of network delays.

Each *SecureChannel* exists until it is explicitly closed or until the last token has expired and the overlap period has elapsed. A *Server* application should limit the number of *SecureChannels*. To protect against misbehaving *Clients* and denial of service attacks, the *Server* shall close the oldest *SecureChannel* that has no *Session* assigned before reaching the maximum number of supported *SecureChannels*.

The *OpenSecureChannel* request and response *Messages* shall be signed with the sender's private key. These *Messages* shall always be encrypted. If the transport layer does not provide encryption, then these *Messages* shall be encrypted with the receiver's public key. These requirements for *OpenSecureChannel* only apply if the *securityPolicyUri* is not None.

If the protocol defined in IEC 62541-6 requires that *Application Instance Certificates* are used in the *OpenSecureChannel Service*, then *Clients* and *Servers* shall verify that the same *Certificates* are used in the *CreateSession* and *ActivateSession Services*. *Certificates* are not provided and shall not be verified if the *securityPolicyUri* is None.

If the *securityPolicyUri* is not None, a *Client* shall verify the *HostName* specified in the *Server Certificate* is the same as the *HostName* contained in the *endpointUrl*. If there is a difference then the *Client* shall report the difference and may choose to not open the *SecureChannel*. *Servers* shall add all possible *HostNames* like *MyHost* and *MyHost.local* into the *Server Certificate*. This includes IP addresses of the host or the *HostName* exposed by a NAT router used to connect to the *Server*.

Clients should be prepared to replace the *HostName* returned in the *EndpointDescription* with the *HostName* or the IP addresses they used to call *GetEndpoints*.

5.5.2.2 Parameters

Table 11 defines the parameters for the *Service*.

Unlike other *Services*, the parameters for this *Service* provide only an abstract definition. The concrete representation on the network depends on the mappings defined in IEC 62541-6.

Table 11 – OpenSecureChannel Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
clientCertificate	ApplicationInstanceCertificate	A <i>Certificate</i> that identifies the <i>Client</i> . The <i>OpenSecureChannel</i> request shall be signed with the private key for this <i>Certificate</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2. If the <i>securityPolicyUri</i> is None, the <i>Server</i> shall ignore the <i>ApplicationInstanceCertificate</i> .
requestType	Enum SecurityTokenRequestType	The type of <i>SecurityToken</i> request: An enumeration that shall be one of the following: ISSUE_0 creates a new <i>SecurityToken</i> for a new <i>SecureChannel</i> . RENEW_1 creates a new <i>SecurityToken</i> for an existing <i>SecureChannel</i> .
secureChannelId	BaseDataType	The identifier for the <i>SecureChannel</i> that the new token should belong to. This parameter shall be null when creating a new <i>SecureChannel</i> . The concrete security protocol definition in IEC 62541-6 chooses the concrete <i>Data Type</i> .
securityMode	Enum MessageSecurityMode	The type of security to apply to the messages. The type <i>MessageSecurityMode</i> type is defined in 7.15. A <i>SecureChannel</i> may need to be created even if the <i>securityMode</i> is NONE. The exact behaviour depends on the mapping used and is described in the IEC 62541-6.
securityPolicyUri	String	The URI for <i>SecurityPolicy</i> to use when securing messages sent over the <i>SecureChannel</i> . The set of known URIs and the <i>SecurityPolicies</i> associated with them are defined in IEC 62541-7.
clientNonce	ByteString	A random number that shall not be used in any other request. A new <i>clientNonce</i> shall be generated for each time a <i>SecureChannel</i> is renewed. This parameter shall have a length equal to the <i>SecureChannelNonceLength</i> defined for the <i>SecurityPolicy</i> in IEC 62541-7. The <i>SecurityPolicy</i> is identified by the <i>securityPolicyUri</i> .
requestedLifetime	Duration	The requested lifetime, in milliseconds, for the new <i>SecurityToken</i> . It specifies when the <i>Client</i> expects to renew the <i>SecureChannel</i> by calling the <i>OpenSecureChannel Service</i> again. If a <i>SecureChannel</i> is not renewed, then all <i>Messages</i> sent using the current <i>SecurityTokens</i> shall be rejected by the receiver. Several cryptanalytic attacks become easier as more material encrypted with a specific key is available. By limiting the amount of data processed using a particular key, those attacks are made more difficult. Therefore the volume of data exchanged between <i>Client</i> and <i>Server</i> shall be limited by establishing a new <i>SecurityToken</i> after the lifetime. The setting of the requested lifetime depends on the expected number of exchanged messages and their size in the lifetime. A higher volume of data requires shorter lifetime.

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> type definition).
securityToken	ChannelSecurityToken	Describes the new <i>SecurityToken</i> issued by the <i>Server</i> . This structure is defined in-line with the following indented items.
channelId	BaseDataType	A unique identifier for the <i>SecureChannel</i> . This is the identifier that shall be supplied whenever the <i>SecureChannel</i> is renewed. The concrete security protocol definition in IEC 62541-6 chooses the concrete <i>DataType</i> .
tokenId	ByteString	A unique identifier for a single <i>SecurityToken</i> within the channel. This is the identifier that shall be passed with each <i>Message</i> secured with the <i>SecurityToken</i> .
createdAt	UtcTime	The time when the <i>SecurityToken</i> was created.
revisedLifetime	Duration	The lifetime of the <i>SecurityToken</i> in milliseconds. The UTC expiration time for the token may be calculated by adding the lifetime to the <i>createdAt</i> time.
serverNonce	ByteString	A random number that shall not be used in any other request. A new <i>serverNonce</i> shall be generated for each time a <i>SecureChannel</i> is renewed. This parameter shall have a length equal to the <i>SecureChannelNonceLength</i> defined for the <i>SecurityPolicy</i> in IEC 62541-7. The <i>SecurityPolicy</i> is identified by the <i>securityPolicyUri</i> .

5.5.2.3 Service results

Table 12 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 12 – OpenSecureChannel Service result codes

Symbolic Id	Description
Bad_SecurityChecksFailed	See Table 177 for the description of this result code.
Bad_CertificateTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateIssuerTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateHostNameInvalid	See Table 177 for the description of this result code.
Bad_CertificateUriInvalid	See Table 177 for the description of this result code.
Bad_CertificateUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateIssuerUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateUntrusted	See Table 177 for the description of this result code.
Bad_CertificateRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateRevoked	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevoked	See Table 177 for the description of this result code.
Bad_RequestTypeInvalid	The security token request type is not valid.
Bad_SecurityModeRejected	The security mode does not meet the requirements set by the <i>Server</i> .
Bad_SecurityPolicyRejected	The security policy does not meet the requirements set by the <i>Server</i> .
Bad_SecureChannelIdInvalid	See Table 177 for the description of this result code.
Bad_NonceInvalid	See Table 177 for the description of this result code. A <i>Server</i> shall check the minimum length of the <i>Client</i> nonce and return this status if the length is below 32 bytes. A check for duplicated nonce can only be done in <i>OpenSecureChannel</i> calls with the request type RENEW_1.

5.5.3 CloseSecureChannel

5.5.3.1 Description

This *Service* is used to terminate a *SecureChannel*.

The request *Messages* shall be signed with the appropriate key associated with the current token for the *SecureChannel*.

5.5.3.2 Parameters

Table 13 defines the parameters for the *Service*.

Specific protocol mappings defined in IEC 62541-6 may choose to omit the response.

Table 13 – CloseSecureChannel Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
secureChannelId	BaseDataType	The identifier for the <i>SecureChannel</i> to close. The concrete security protocol definition in IEC 62541-6 chooses the concrete <i>DataType</i> .
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

5.5.3.3 Service results

Table 14 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 14 – CloseSecureChannel Service result codes

Symbolic Id	Description
Bad_SecureChannelIdInvalid	See Table 177 for the description of this result code.

5.6 Session Service Set

5.6.1 Overview

This *Service Set* defines *Services* for an application layer connection establishment in the context of a *Session*.

5.6.2 CreateSession

5.6.2.1 Description

This *Service* is used by an OPC UA *Client* to create a *Session* and the *Server* returns two values which uniquely identify the *Session*. The first value is the *sessionId* which is used to identify the *Session* in the audit logs and in the *Server's AddressSpace*. The second is the *authenticationToken* which is used to associate an incoming request with a *Session*.

Before calling this *Service*, the *Client* shall create a *SecureChannel* with the *OpenSecureChannel Service* to ensure the *Integrity* of all *Messages* exchanged during a *Session*. This *SecureChannel* has a unique identifier which the *Server* shall associate with the *authenticationToken*. The *Server* may accept requests with the *authenticationToken* only if they are associated with the same *SecureChannel* that was used to create the *Session*. The *Client* may associate a new *SecureChannel* with the *Session* by calling the *ActivateSession* method.

The *SecureChannel* is always managed by the *Communication Stack* which means it shall provide APIs which the *Server* can use to find out information about the *SecureChannel* used for any given request. The *Communication Stack* shall, at a minimum, provide the *SecurityPolicy* and *SecurityMode* used by the *SecureChannel*. It shall also provide a *SecureChannelId* which uniquely identifies the *SecureChannel* or the *Client Certificate* used to establish the *SecureChannel*. The *Server* uses one of these to identify the *SecureChannel* used to send a request. Clause 7.31 describes how to create the *authenticationToken* for different types of *Communication Stack*.

Depending upon on the *SecurityPolicy* and the *SecurityMode* of the *SecureChannel*, the exchange of *ApplicationInstanceCertificates* and *Nonces* may be optional and the signatures may be empty. See IEC 62541-7 for the definition of *SecurityPolicies* and the handling of these parameters.

The *Server* returns its *EndpointDescriptions* in the response. *Clients* use this information to determine whether the list of *EndpointDescriptions* returned from the *DiscoveryEndpoint* matches the *Endpoints* that the *Server* has. If there is a difference then the *Client* shall close the *Session* and report an error. The *Server* returns all *EndpointDescriptions* for the *serverUri* specified by the *Client* in the request. The *Client* only verifies *EndpointDescriptions* with a *transportProfileUri* that matches the *profileUri* specified in the original *GetEndpoints* request. A *Client* may skip this check if the *EndpointDescriptions* were provided by a trusted source such as the *Administrator*.

The *Session* created with this *Service* shall not be used until the *Client* calls the *ActivateSession Service* and provides its *SoftwareCertificates* and proves possession of its *Application Instance Certificate* and any user identity token that it provided.

A *Server* application should limit the number of *Sessions*. To protect against misbehaving *Clients* and denial of service attacks, the *Server* shall close the oldest *Session* that is not activated before reaching the maximum number of supported *Sessions*.

The *SoftwareCertificates* parameter in the *Server* response is deprecated to reduce the message size for OPC UA Applications with limited resources. The *SoftwareCertificates* are provided in the *Server's AddressSpace* as defined in IEC 62541-5. A *SoftwareCertificate* identifies the capabilities of the *Server* and also contains the list of OPC UA *Profiles* supported by the *Server*. OPC UA *Profiles* are defined in IEC 62541-7.

Additional *Certificates* issued by other organizations may be included to identify additional *Server* capabilities. Examples of these *Profiles* include support for specific information models and support for access to specific types of devices.

When a *Session* is created, the *Server* adds an entry for the *Client* in its *SessionDiagnosticsArray Variable*. See IEC 62541-5 for a description of this *Variable*.

Sessions are created to be independent of the underlying communications connection. Therefore, if a communications connection fails, the *Session* is not immediately affected. The exact mechanism to recover from an underlying communication connection error depends on the *SecureChannel* mapping as described in IEC 62541-6.

Sessions are terminated by the *Server* automatically if the *Client* fails to issue a *Service* request on the *Session* within the timeout period negotiated by the *Server* in the *CreateSession Service* response. This protects the *Server* against *Client* failures and against situations where a failed underlying connection cannot be re-established. *Clients* shall be prepared to submit requests in a timely manner to prevent the *Session* from closing automatically. *Clients* may explicitly terminate *Sessions* using the *CloseSession Service*.

When a *Session* is terminated, all outstanding requests on the *Session* are aborted and *Bad_SessionClosed StatusCodes* are returned to the *Client*. In addition, the *Server* deletes the entry for the *Client* from its *SessionDiagnosticsArray Variable* and notifies any other *Clients* who were subscribed to this entry.

If a *Client* invokes the *CloseSession Service* then all *Subscriptions* associated with the *Session* are also deleted if the *deleteSubscriptions* flag is set to TRUE. If a *Server* terminates a *Session* for any other reason, *Subscriptions* associated with the *Session*, are not deleted. Each *Subscription* has its own lifetime to protect against data loss in the case of a *Session* termination. In these cases, the *Subscription* can be reassigned to another *Client* before its lifetime expires.

Some *Servers*, such as aggregating *Servers*, also act as *Clients* to other *Servers*. These *Servers* typically support more than one system user, acting as their agent to the *Servers* that they represent. Security for these *Servers* is supported at two levels.

First, each OPC UA *Service* request contains a string parameter that is used to carry an audit record id. A *Client*, or any *Server* operating as a *Client*, such as an aggregating *Server*, can create a local audit log entry for a request that it submits. This parameter allows the *Client* to pass the identifier for this entry with the request. If the *Server* also maintains an audit log, then it can include this id in the audit log entry that it writes. When the log is examined and the entry is found, the examiner will be able to relate it directly to the audit log entry created by the *Client*. This capability allows for traceability across audit logs within a system. See IEC TR 62541-2 for additional information on auditing. A *Server* that maintains an audit log shall provide the information in the audit log entries via event *Messages* defined in this document. The *Server* may choose to only provide the *Audit* information via event *Messages*. The *Audit EventType* is defined in IEC 62541-3.

Second, these aggregating *Servers* may open independent *Sessions* to the underlying *Servers* for each *Client* that accesses data from them. Figure 14 illustrates this concept.




Figure 14 – Multiplexing users on a Session

5.6.2.2 Parameters

Table 15 defines the parameters for the *Service*.

Table 15 – CreateSession Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The <i>authenticationToken</i> is always null. The type <i>RequestHeader</i> is defined in 7.28.
clientDescription	Application Description	Information that describes the <i>Client</i> application. The type <i>ApplicationDescription</i> is defined in 7.1.
serverUri	String	This value is only specified if the <i>EndpointDescription</i> has a <i>gatewayServerUri</i> . This value is the <i>applicationUri</i> from the <i>EndpointDescription</i> which is the <i>applicationUri</i> for the underlying <i>Server</i> . The type <i>EndpointDescription</i> is defined in 7.10.
endpointUrl	String	The network address that the <i>Client</i> used to access the <i>Session Endpoint</i> . The <i>HostName</i> portion of the URL should be one of the <i>HostNames</i> for the application that are specified in the <i>Server's ApplicationInstanceCertificate</i> (see 7.2). The <i>Server</i> shall raise an <i>AuditUrlMismatchEventType</i> event if the URL does not match the <i>Server's HostNames</i> . <i>AuditUrlMismatchEventType</i> event type is defined in IEC 62541-5. The <i>Server</i> uses this information for diagnostics and to determine the set of <i>EndpointDescriptions</i> to return in the response.
sessionName	String	Human readable string that identifies the <i>Session</i> . The <i>Server</i> makes this name and the <i>sessionId</i> visible in its <i>AddressSpace</i> for diagnostic purposes. The <i>Client</i> should provide a name that is unique for the instance of the <i>Client</i> . If this parameter is not specified the <i>Server</i> shall assign a value.
clientNonce	ByteString	A random number that should never be used in any other request. This number shall have a minimum length of 32 bytes. Profiles may increase the required length. The <i>Server</i> shall use this value to prove possession of its <i>Application Instance Certificate</i> in the response.
clientCertificate	ApplicationInstance Certificate	The <i>Application Instance Certificate</i> issued to the <i>Client</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2. If the <i>securityPolicyUri</i> is None, the <i>Server</i> shall ignore the <i>ApplicationInstanceCertificate</i> .
Requested SessionTimeout	Duration	Requested maximum number of milliseconds that a <i>Session</i> should remain open without activity. If the <i>Client</i> fails to issue a <i>Service</i> request within this interval, then the <i>Server</i> shall automatically terminate the <i>Client Session</i> .
maxResponse MessageSize	UInt32	The maximum size, in bytes, for the body of any response message. The <i>Server</i> should return a <i>Bad_ResponseTooLarge</i> service fault if a response message exceeds this limit. The value zero indicates that this parameter is not used. The transport protocols defined in IEC 62541-6 may imply minimum message sizes. More information on the use of this parameter is provided in 5.3.

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> type).
sessionId	NodeId	A unique <i>NodeId</i> assigned by the <i>Server</i> to the <i>Session</i> . This identifier is used to access the diagnostics information for the <i>Session</i> in the <i>Server AddressSpace</i> . It is also used in the audit logs and any events that report information related to the <i>Session</i> . The <i>Session</i> diagnostic information is described in IEC 62541-5. Audit logs and their related events are described in 6.5.
authentication Token	Session AuthenticationToken	A unique identifier assigned by the <i>Server</i> to the <i>Session</i> . This identifier shall be passed in the <i>RequestHeader</i> of each request and is used with the <i>SecureChannelId</i> to determine whether a <i>Client</i> has access to the <i>Session</i> . This identifier shall not be reused in a way that the <i>Client</i> or the <i>Server</i> has a chance of confusing them with a previous or existing <i>Session</i> . The <i>SessionAuthenticationToken</i> type is described in 7.31.
revisedSession Timeout	Duration	Actual maximum number of milliseconds that a <i>Session</i> shall remain open without activity. The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.
serverNonce	ByteString	A random number that should never be used in any other request. This number shall have a minimum length of 32 bytes. The <i>Client</i> shall use this value to prove possession of its <i>Application Instance Certificate</i> in the <i>ActivateSession</i> request. This value may also be used to prove possession of the <i>userIdentityToken</i> it specified in the <i>ActivateSession</i> request.
serverCertificate	ApplicationInstance Certificate	The <i>Application Instance Certificate</i> issued to the <i>Server</i> . A <i>Server</i> shall prove possession by using the private key to sign the <i>Nonce</i> provided by the <i>Client</i> in the request. The <i>Client</i> shall verify that this <i>Certificate</i> is the same as the one it used to create the <i>SecureChannel</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2. If the <i>securityPolicyUri</i> is NONE and none of the <i>UserTokenPolicies</i> requires encryption, the <i>Client</i> shall ignore the <i>ApplicationInstanceCertificate</i> .
serverEndpoints []	EndpointDescription	List of <i>Endpoints</i> that the <i>Server</i> supports. The <i>Server</i> shall return a set of <i>EndpointDescriptions</i> available for the <i>serverUri</i> specified in the request. The <i>EndpointDescription</i> type is defined in 7.10. The <i>Client</i> shall verify this list with the list from a <i>DiscoveryEndpoint</i> if it used a <i>DiscoveryEndpoint</i> to fetch the <i>EndpointDescriptions</i> . It is recommended that <i>Servers</i> only include the <i>server.applicationUri</i> , <i>endpointUrl</i> , <i>securityMode</i> , <i>securityPolicyUri</i> , <i>userIdentityTokens</i> , <i>transportProfileUri</i> and <i>securityLevel</i> with all other parameters set to null. Only the recommended parameters shall be verified by the client.
serverSoftware Certificates []	SignedSoftware Certificate	This parameter is deprecated and the array shall be empty. The <i>SoftwareCertificates</i> are provided in the <i>Server AddressSpace</i> as defined in IEC 62541-5.
serverSignature	SignatureData	This is a signature generated with the private key associated with the <i>serverCertificate</i> . This parameter is calculated by appending the <i>clientNonce</i> to the <i>clientCertificate</i> and signing the resulting sequence of bytes. If the <i>clientCertificate</i> contains a chain, the signature calculation shall be done only with the leaf <i>Certificate</i> . For backward compatibility a <i>Client</i> shall check the signature with the full chain if the check with the leaf <i>Certificate</i> fails. The <i>SignatureAlgorithm</i> shall be the <i>AsymmetricSignatureAlgorithm</i> specified in the <i>SecurityPolicy</i> for the <i>Endpoint</i> . The <i>SignatureData</i> type is defined in 7.32.
maxRequest MessageSize	UInt32	The maximum size, in bytes, for the body of any request message. The <i>Client Communication Stack</i> should return a <i>Bad_RequestTooLarge</i> error to the application if a request message exceeds this limit. The value zero indicates that this parameter is not used. See IEC 62541-6 for protocol specific minimum or default values. 5.3 provides more information on the use of this parameter.

5.6.2.3 Service results

Table 16 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 16 – CreateSession Service result codes

Symbolic Id	Description
Bad_SecureChannelIdInvalid	See Table 177 for the description of this result code.
Bad_NonceInvalid	See Table 177 for the description of this result code. A <i>Server</i> shall check the minimum length of the <i>Client</i> nonce and return this status if the length is below 32 bytes. A check for a duplicated nonce is optional and requires access to the nonce used to create the secure channel.
Bad_SecurityChecksFailed	See Table 177 for the description of this result code.
Bad_CertificateTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateIssuerTimeInvalid	See Table 177 for the description of this result code.
Bad_CertificateHostNameInvalid	See Table 177 for the description of this result code.
Bad_CertificateUriInvalid	See Table 177 for the description of this result code.
Bad_CertificateUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateIssuerUseNotAllowed	See Table 177 for the description of this result code.
Bad_CertificateUntrusted	See Table 177 for the description of this result code.
Bad_CertificateRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevocationUnknown	See Table 177 for the description of this result code.
Bad_CertificateRevoked	See Table 177 for the description of this result code.
Bad_CertificateIssuerRevoked	See Table 177 for the description of this result code.
Bad_TooManySessions	The <i>Server</i> has reached its maximum number of <i>Sessions</i> .
Bad_ServerUriInvalid	See Table 177 for the description of this result code.

5.6.3 ActivateSession

5.6.3.1 Description

This *Service* is used by the *Client* to specify the identity of the user associated with the *Session*. This *Service* request shall be issued by the *Client* before it issues any *Service* request other than *CloseSession* after *CreateSession*. Failure to do so shall cause the *Server* to close the *Session*.

Whenever the *Client* calls this *Service* the *Client* shall prove that it is the same application that called the *CreateSession* *Service*. The *Client* does this by creating a signature with the private key associated with the *clientCertificate* specified in the *CreateSession* request. This signature is created by appending the last *serverNonce* provided by the *Server* to the *serverCertificate* and calculating the signature of the resulting sequence of bytes.

Once used, a *serverNonce* cannot be used again. For that reason, the *Server* returns a new *serverNonce* each time the *ActivateSession* *Service* is called.

When the *ActivateSession* *Service* is called for the first time then the *Server* shall reject the request if the *SecureChannel* is not same as the one associated with the *CreateSession* request. Subsequent calls to *ActivateSession* may be associated with different *SecureChannels*. If this is the case then the *Server* shall verify that the *Certificate* the *Client* used to create the new *SecureChannel* is the same as the *Certificate* used to create the original *SecureChannel*. In addition, the *Server* shall verify that the *Client* supplied a

UserIdentityToken that is identical to the token currently associated with the *Session*. Once the *Server* accepts the new *SecureChannel* it shall reject requests sent via the old *SecureChannel*.

The *ActivateSession Service* is used to associate a user identity with a *Session*. When a *Client* provides a user identity then it shall provide proof that it is authorized to use that user identity. The exact mechanism used to provide this proof depends on the type of the *UserIdentityToken*. If the token is a *UserNameIdentityToken* then the proof is the *password* that is included in the token. If the token is an *X509IdentityToken* then the proof is a signature generated with private key associated with the *Certificate*. The data to sign is created by appending the last *serverNonce* to the *serverCertificate* specified in the *CreateSession* response. If a token includes a secret then it should be encrypted using the public key from the *serverCertificate*.

Servers shall take proper measures to protect against attacks on user identity tokens. Such an attack is assumed if repeated connection attempts with invalid user identity tokens happen. One option is to lock out an OPC UA *Client* for a period of time if the user identity token validation fails several times. The OPC UA *Client* is either detected by IP address for unsecured connections or by the *ApplicationInstanceUri* for secured connections. Another option is delaying the *Service* response when the validation of a user identity fails. This delay time could be increased with repeated failures. Sporadic failures shall not delay connections with valid tokens.

Clients can change the identity of a user associated with a *Session* by calling the *ActivateSession Service*. The *Server* validates the signatures provided with the request and then validates the new user identity. If no errors occur the *Server* replaces the user identity for the *Session*. Changing the user identity for a *Session* may cause discontinuities in active *Subscriptions* because the *Server* may need to tear down connections to an underlying system and re-establish them using the new credentials.

When a *Client* supplies a list of locale ids in the request, each locale id is required to contain the language component. It may optionally contain the <country/region> component. When the *Server* returns a *LocalizedText* in the context of the *Session*, it also may return both the language and the country/region or just the language as its default locale id.

When a *Server* returns a string to the *Client*, it first determines if there are available translations for it. If there are, then the *Server* returns the string whose locale id exactly matches the locale id with the highest priority in the *Client*-supplied list.

If there are no exact matches, then the *Server* ignores the <country/region> component of the locale id, and returns the string whose <language> component matches the <language> component of the locale id with the highest priority in the *Client* supplied list.

If there still are no matches, then the *Server* returns the string that it has along with the locale id.

A *Gateway Server* is expected to impersonate the user provided by the *Client* when it connects to the underlying *Server*. This means it shall re-calculate the signatures on the *UserIdentityToken* using the nonce provided by the underlying *Server*. The *Gateway Server* will need to use its own user credentials if the *UserIdentityToken* provided by the *Client* does not support impersonation.

5.6.3.2 Parameters

Table 17 defines the parameters for the *Service*.

Table 17 – ActivateSession Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters. The type <i>RequestHeader</i> is defined in 7.28.
clientSignature	SignatureData	This is a signature generated with the private key associated with the <i>clientCertificate</i>. This parameter is calculated by appending the <i>serverNonce</i> to the <i>serverCertificate</i> and signing the resulting sequence of bytes. If the <i>serverCertificate</i> contains a chain, the signature calculation shall be done only with the leaf <i>Certificate</i>. For backward compatibility a <i>Server</i> shall check the signature with the full chain if the check with the leaf <i>Certificate</i> fails. The <i>SignatureAlgorithm</i> shall be the <i>AsymmetricSignatureAlgorithm</i> specified in the <i>SecurityPolicy</i> for the <i>Endpoint</i>. The <i>SignatureData</i> type is defined in 7.32.
clientSoftwareCertificates []	SignedSoftwareCertificate	Reserved for future use. The <i>SignedSoftwareCertificate</i> type is defined in 7.33.
localeIds []	LocaleId	List of locale ids in priority order for localized strings. The first <i>LocaleId</i> in the list has the highest priority. If the <i>Server</i> returns a localized string to the <i>Client</i>, the <i>Server</i> shall return the translation with the highest priority that it can. If it does not have a translation for any of the locales identified in this list, then it shall return the string value that it has and include the locale id with the string. See IEC 62541-3 for more detail on locale ids. If the <i>Client</i> fails to specify at least one locale id, the <i>Server</i> shall use any that it has. This parameter only needs to be specified during the first call to <i>ActivateSession</i> during a single application <i>Session</i>. If it is not specified the <i>Server</i> shall keep using the current <i>localeIds</i> for the <i>Session</i>.
userIdentityToken	ExtensibleParameter UserIdentityToken	The credentials of the user associated with the <i>Client</i> application. The <i>Server</i> uses these credentials to determine whether the <i>Client</i> should be allowed to activate a <i>Session</i> and what resources the <i>Client</i> has access to during this <i>Session</i>. The <i>UserIdentityToken</i> is an extensible parameter type defined in 7.36. The <i>EndpointDescription</i> specifies what <i>UserIdentityTokens</i> the <i>Server</i> shall accept. Null or empty user token shall always be interpreted as anonymous.
userTokenSignature	SignatureData	If the <i>Client</i> specified a user identity token that supports digital signatures, then it shall create a signature and pass it as this parameter. Otherwise the parameter is null. The <i>SignatureAlgorithm</i> depends on the identity token type. The <i>SignatureData</i> type is defined in 7.32.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
serverNonce	ByteString	A random number that should never be used in any other request. This number shall have a minimum length of 32 bytes. The <i>Client</i> shall use this value to prove possession of its <i>Application Instance Certificate</i> in the next call to <i>ActivateSession</i> request.
results []	StatusCode	List of validation results for the <i>SoftwareCertificates</i> (see 7.34 for <i>StatusCode</i> definition).
diagnosticInfos []	DiagnosticInfo	List of diagnostic information associated with <i>SoftwareCertificate</i> validation errors (see 7.8 for <i>DiagnosticInfo</i> definition). This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.6.3.3 Service results

Table 18 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 18 – ActivateSession Service result codes

Symbolic Id	Description
Bad_IdentityTokenInvalid	See Table 177 for the description of this result code.
Bad_IdentityTokenRejected	See Table 177 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_ApplicationSignatureInvalid	The signature provided by the <i>Client</i> application is missing or invalid.
Bad_UserSignatureInvalid	The user token signature is missing or invalid.
Bad_NoValidCertificates	The <i>Client</i> did not provide at least one <i>Software Certificate</i> that is valid and meets the profile requirements for the <i>Server</i> .
Bad_IdentityChangeNotSupported	The <i>Server</i> does not support changing the user identity assigned to the session.

5.6.4 CloseSession

5.6.4.1 Description

This *Service* is used to terminate a *Session*. The *Server* takes the following actions when it receives a *CloseSession* request:

- It stops accepting requests for the *Session*. All subsequent requests received for the *Session* are discarded.
- It returns negative responses with the *StatusCode* *Bad_SessionClosed* to all requests that are currently outstanding to provide for the timely return of the *CloseSession* response. *Clients* are urged to wait for all outstanding requests to complete before submitting the *CloseSession* request.
- It removes the entry for the *Client* in its *SessionDiagnosticsArray Variable*.

When the *CloseSession Service* is called before the *Session* is successfully activated, the *Server* shall reject the request if the *SecureChannel* is not the same as the one associated with the *CreateSession* request.

5.6.4.2 Parameters

Table 19 defines the parameters for the *Service*.

Table 19 – CloseSession Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
deleteSubscriptions	Boolean	If the value is TRUE, the <i>Server</i> deletes all Subscriptions associated with the <i>Session</i> . If the value is FALSE, the <i>Server</i> keeps the Subscriptions associated with the <i>Session</i> until they timeout based on their own lifetime.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

5.6.4.3 Service results

Table 20 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 20 – CloseSession Service result codes

Symbolic Id	Description
Bad_SessionIdInvalid	See Table 177 for the description of this result code.

5.6.5 Cancel

5.6.5.1 Description

This *Service* is used to cancel outstanding *Service* requests. Successfully cancelled service requests shall respond with *Bad_RequestCancelledByClient*.

5.6.5.2 Parameters

Table 21 defines the parameters for the *Service*.

Table 21 – Cancel Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
requestHandle	IntegerId	The <i>requestHandle</i> assigned to one or more requests that should be cancelled. All outstanding requests with the matching <i>requestHandle</i> shall be cancelled.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
cancelCount	UInt32	Number of cancelled requests.

5.6.5.3 Service results

Common *StatusCodes* are defined in Table 177.

5.7 NodeManagement Service Set

5.7.1 Overview

This *Service Set* defines *Services* to add and delete *AddressSpace Nodes* and *References* between them. All added *Nodes* continue to exist in the *AddressSpace* even if the *Client* that created them disconnects from the *Server*.

5.7.2 AddNodes

5.7.2.1 Description

This *Service* is used to add one or more *Nodes* into the *AddressSpace* hierarchy. Using this *Service*, each *Node* is added as the *TargetNode* of a *HierarchicalReference* to ensure that the *AddressSpace* is fully connected and that the *Node* is added as a child within the *AddressSpace* hierarchy (see IEC 62541-3).

When a *Server* creates an instance of a *TypeDefinitionNode* it shall create the same hierarchy of *Nodes* beneath the new *Object* or *Variable* depending on the *ModellingRule* of each *InstanceDeclaration*. All *Nodes* with a *ModellingRule* of *Mandatory* shall be created or an existing *Node* shall be referenced that conforms to the *InstanceDeclaration*. The creation of *Nodes* with other *ModellingRules* is *Server* specific.

5.7.2.2 Parameters

Table 22 defines the parameters for the *Service*.

Table 22 – AddNodes Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToAdd []	AddNodesItem	List of <i>Nodes</i> to add. All <i>Nodes</i> are added as a <i>Reference</i> to an existing <i>Node</i> using a hierarchical <i>ReferenceType</i> . This structure is defined in-line with the following indented items.
parentNodeid	Expanded Nodeid	<i>ExpandedNodeid</i> of the parent <i>Node</i> for the <i>Reference</i> . The <i>ExpandedNodeid</i> type is defined in 7.11.
referenceTypeid	Nodeid	<i>Nodeid</i> of the hierarchical <i>ReferenceType</i> to use for the <i>Reference</i> from the parent <i>Node</i> to the new <i>Node</i> .
requestedNewNodeid	Expanded Nodeid	<i>Client</i> requested expanded <i>Nodeid</i> of the <i>Node</i> to add. The <i>serverIndex</i> in the expanded <i>Nodeid</i> shall be 0. If the <i>Server</i> cannot use this <i>Nodeid</i> , it rejects this <i>Node</i> and returns the appropriate error code. If the <i>Client</i> does not want to request a <i>Nodeid</i> , then it sets the value of this parameter to the null expanded <i>Nodeid</i> . If the <i>Node</i> to add is a <i>ReferenceType Node</i> , its <i>Nodeid</i> should be a numeric id. See IEC 62541-3 for a description of <i>ReferenceType Nodeids</i> .
browseName	QualifiedName	The browse name of the <i>Node</i> to add.
nodeClass	NodeClass	<i>NodeClass</i> of the <i>Node</i> to add.
nodeAttributes	Extensible Parameter NodeAttributes	The <i>Attributes</i> that are specific to the <i>NodeClass</i> . The <i>NodeAttributes</i> parameter type is an extensible parameter type specified in 7.19. A <i>Client</i> is allowed to omit values for some or all <i>Attributes</i> . If an <i>Attribute</i> value is null, the <i>Server</i> shall use the default values from the <i>TypeDefinitionNode</i> . If a <i>TypeDefinitionNode</i> was not provided the <i>Server</i> shall choose a suitable default value. The <i>Server</i> may still add an optional <i>Attribute</i> to the <i>Node</i> with an appropriate default value even if the <i>Client</i> does not specify a value.
typeDefinition	Expanded Nodeid	<i>Nodeid</i> of the <i>TypeDefinitionNode</i> for the <i>Node</i> to add. This parameter shall be null for all <i>NodeClasses</i> other than <i>Object</i> and <i>Variable</i> in which case it shall be provided.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	AddNodesResult	List of results for the <i>Nodes</i> to add. The size and order of the list matches the size and order of the <i>nodesToAdd</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the <i>Node</i> to add (see 7.34 for <i>StatusCode</i> definition).
addedNodeid	Nodeid	<i>Server</i> assigned <i>Nodeid</i> of the added <i>Node</i> . Null <i>Nodeid</i> if the operation failed.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Nodes</i> to add (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>nodesToAdd</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.2.3 Service results

Table 23 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 23 – AddNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.2.4 StatusCodes

Table 24 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 24 – AddNodes Operation Level result codes

Symbolic Id	Description
Bad_ParentNodeIdInvalid	The parent node id does not refer to a valid node.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_ReferenceNotAllowed	The reference could not be created because it violates constraints imposed by the data model.
Bad_NodeIdRejected	The requested node id was rejected either because it was invalid or because the <i>Server</i> does not allow node ids to be specified by the client.
Bad_NodeIdExists	The requested node id is already used by another node.
Bad_NodeClassInvalid	See Table 178 for the description of this result code.
Bad_BrowseNameInvalid	See Table 178 for the description of this result code.
Bad_BrowseNameDuplicated	The browse name is not unique among nodes that share the same relationship with the parent.
Bad_NodeAttributesInvalid	The node <i>Attributes</i> are not valid for the node class.
Bad_TypeDefinitionInvalid	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.

5.7.3 AddReferences

5.7.3.1 Description

This *Service* is used to add one or more *References* to one or more *Nodes*. The *NodeClass* is an input parameter that is used to validate that the *Reference* to be added matches the *NodeClass* of the *TargetNode*. This parameter is not validated if the *Reference* refers to a *TargetNode* in a remote *Server*.

In certain cases, adding new *References* to the *AddressSpace* shall require that the *Server* add new *Server* ids to the *Server's ServerArray Variable*. For this reason, remote *Servers* are identified by their URI and not by their *ServerArray* index. This allows the *Server* to add the remote *Server* URIs to its *ServerArray*.

5.7.3.2 Parameters

Table 25 defines the parameters for the *Service*.

Table 25 – AddReferences Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
referencesToAdd []	AddReferences Item	List of <i>Reference</i> instances to add to the <i>SourceNode</i> . The <i>targetNodeClass</i> of each <i>Reference</i> in the list shall match the <i>NodeClass</i> of the <i>TargetNode</i> . This structure is defined in-line with the following indented items.
sourceNodeid	Nodeid	<i>Nodeid</i> of the <i>Node</i> to which the <i>Reference</i> is to be added. The source <i>Node</i> shall always exist in the <i>Server</i> to add the <i>Reference</i> . The <i>isForward</i> parameter can be set to FALSE if the target <i>Node</i> is on the local <i>Server</i> and the source <i>Node</i> on the remote <i>Server</i> .
referenceTypeid	Nodeid	<i>Nodeid</i> of the <i>ReferenceType</i> that defines the <i>Reference</i> .
isForward	Boolean	If the value is TRUE, the <i>Server</i> creates a forward <i>Reference</i> . If the value is FALSE, the <i>Server</i> creates an inverse <i>Reference</i> .
targetServerUri	String	URI of the remote <i>Server</i> . If this parameter is not null, it overrides the <i>serverIndex</i> in the <i>targetNodeid</i> .
targetNodeid	Expanded Nodeid	Expanded <i>Nodeid</i> of the <i>TargetNode</i> . The <i>ExpandedNodeid</i> type is defined in 7.11.
targetNodeClass	NodeClass	<i>NodeClass</i> of the <i>TargetNode</i> . The <i>Client</i> shall specify this since the <i>TargetNode</i> might not be accessible directly by the <i>Server</i> .
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	Status Code	List of <i>StatusCodes</i> for the <i>References</i> to add (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>referencesToAdd</i> request parameter.
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>References</i> to add (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>referencesToAdd</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.3.3 Service results

Table 26 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 26 – AddReferences Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.3.4 StatusCodes

Table 27 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 27 – AddReferences Operation Level result codes

Symbolic Id	Description
Bad_SourceNodeIdInvalid	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_ServerUriInvalid	See Table 177 for the description of this result code.
Bad_TargetNodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeClassInvalid	See Table 178 for the description of this result code.
Bad_ReferenceNotAllowed	The reference could not be created because it violates constraints imposed by the data model on this <i>Server</i> .
Bad_ReferenceLocalOnly	The reference type is not valid for a reference to a remote <i>Server</i> .
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_DuplicateReferenceNotAllowed	The reference type between the nodes is already defined.
Bad_InvalidSelfReference	The <i>Server</i> does not allow this type of self reference on this node.

5.7.4 DeleteNodes

5.7.4.1 Description

This *Service* is used to delete one or more *Nodes* from the *AddressSpace*.

When any of the *Nodes* deleted by an invocation of this *Service* is the *TargetNode* of a *Reference*, then those *References* are left unresolved based on the *deleteTargetReferences* parameter.

When any of the *Nodes* deleted by an invocation of this *Service* is being monitored, then a *Notification* containing the status code *Bad_NodeIdUnknown* is sent to the monitoring *Client* indicating that the *Node* has been deleted.

5.7.4.2 Parameters

Table 28 defines the parameters for the *Service*.

Table 28 – DeleteNodes Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToDelete []	DeleteNodes Item	List of <i>Nodes</i> to delete. This structure is defined in-line with the following indented items.
nodeId	NodeId	<i>NodeId</i> of the <i>Node</i> to delete.
deleteTargetReferences	Boolean	A <i>Boolean</i> parameter with the following values: TRUE delete <i>References</i> in <i>TargetNodes</i> that <i>Reference</i> the <i>Node</i> to delete. FALSE delete only the <i>References</i> for which the <i>Node</i> to delete is the source. The <i>Server</i> cannot guarantee that it is able to delete all <i>References</i> from <i>TargetNodes</i> if this parameter is TRUE.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCodes	List of <i>StatusCodes</i> for the <i>Nodes</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the list of the <i>nodesToDelete</i> request parameter.
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>Nodes</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>nodesToDelete</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.4.3 Service results

Table 29 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 29 – DeleteNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.4.4 StatusCodes

Table 30 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 30 – DeleteNodes Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_NoDeleteRights	See Table 178 for the description of this result code.
Uncertain_ReferenceNotDeleted	The <i>Server</i> was not able to delete all target references.

5.7.5 DeleteReferences

5.7.5.1 Description

This *Service* is used to delete one or more *References* of a *Node*.

When any of the *References* deleted by an invocation of this *Service* are contained in a *View*, then the *ViewVersion Property* is updated if this *Property* is supported.

Table 31 defines the parameters for the *Service*.

Table 31 – DeleteReferences Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
referencesToDelete []	DeleteReferences Item	List of <i>References</i> to delete. This structure is defined in-line with the following indented items.
sourceNodeId	NodeId	<i>NodeId</i> of the <i>Node</i> that contains the <i>Reference</i> to delete.
referenceTypeId	NodeId	<i>NodeId</i> of the <i>ReferenceType</i> that defines the <i>Reference</i> to delete.
isForward	Boolean	If the value is TRUE, the <i>Server</i> deletes a forward <i>Reference</i> . If the value is FALSE, the <i>Server</i> deletes an inverse <i>Reference</i> .
targetNodeId	ExpandedNodeId	<i>NodeId</i> of the <i>TargetNode</i> of the <i>Reference</i> . If the <i>Server</i> index indicates that the <i>TargetNode</i> is a remote <i>Node</i> , then the <i>nodeId</i> shall contain the absolute namespace URI. If the <i>TargetNode</i> is a local <i>Node</i> the <i>nodeId</i> shall contain the namespace index.
deleteBidirectional	Boolean	A <i>Boolean</i> parameter with the following values: TRUE delete the specified <i>Reference</i> and the opposite <i>Reference</i> from the <i>TargetNode</i> . If the <i>TargetNode</i> is located in a remote <i>Server</i> , the <i>Server</i> is permitted to delete the specified <i>Reference</i> only. FALSE delete only the specified <i>Reference</i> .
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of <i>StatusCodes</i> for the <i>References</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>referencesToDelete</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>References</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>referencesToDelete</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.7.5.2 Service results

Table 32 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 32 – DeleteReferences Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.7.5.3 StatusCodes

Table 33 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 33 – DeleteReferences Operation Level result codes

Symbolic Id	Description
Bad_SourceNodeIdInvalid	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_ServerIndexInvalid	The <i>Server</i> index is not valid.
Bad_TargetNodeIdInvalid	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_NoDeleteRights	See Table 178 for the description of this result code.

5.8 View Service Set

5.8.1 Overview

Clients use the browse *Services* of the *View Service Set* to navigate through the *AddressSpace* or through a *View* which is a subset of the *AddressSpace*.

A *View* is a subset of the *AddressSpace* created by the *Server*. Future versions of IEC 62541-4 may also define services to create *Client*-defined *Views*. See IEC 62541-5 for a description of the organization of views in the *AddressSpace*.

5.8.2 Browse

5.8.2.1 Description

This *Service* is used to discover the *References* of a specified *Node*. The browse can be further limited by the use of a *View*. This Browse *Service* also supports a primitive filtering capability.

In some cases it may take longer than the *Client* timeout hint to process all nodes to browse. In this case the *Server* may return zero results with a continuation point for the affected nodes before the timeout expires.

5.8.2.2 Parameters

Table 34 defines the parameters for the *Service*.

Table 34 – Browse Service parameters

Name	Type	Description																		
Request																				
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).																		
View	ViewDescription	Description of the <i>View</i> to browse (see 7.39 for <i>ViewDescription</i> definition). An empty <i>ViewDescription</i> value indicates the entire <i>AddressSpace</i> . Use of the empty <i>ViewDescription</i> value causes all <i>References</i> of the <i>nodesToBrowse</i> to be returned. Use of any other <i>View</i> causes only the <i>References</i> of the <i>nodesToBrowse</i> that are defined for that <i>View</i> to be returned.																		
requestedMax ReferencesPerNode	Counter	Indicates the maximum number of references to return for each starting Node specified in the request. The value 0 indicates that the <i>Client</i> is imposing no limitation (see 7.5 for <i>Counter</i> definition).																		
nodesToBrowse []	BrowseDescription	A list of nodes to Browse. This structure is defined in-line with the following indented items.																		
nodeId	NodeId	<i>NodeId</i> of the <i>Node</i> to be browsed. If a <i>view</i> is provided, it shall include this Node.																		
browseDirection	Enum BrowseDirection	An enumeration that specifies the direction of <i>References</i> to follow. It has the following values: FORWARD_0 select only forward <i>References</i>. INVERSE_1 select only inverse <i>References</i>. BOTH_2 select forward and inverse <i>References</i>. The returned <i>References</i> do indicate the direction the <i>Server</i> followed in the <i>isForward</i> parameter of the <i>ReferenceDescription</i>. Symmetric <i>References</i> are always considered to be in forward direction therefore the <i>isForward</i> flag is always set to TRUE and symmetric <i>References</i> are not returned if <i>browseDirection</i> is set to <i>INVERSE_1</i>.																		
referenceTypeId	NodeId	Specifies the <i>NodeId</i> of the <i>ReferenceType</i> to follow. Only instances of this <i>ReferenceType</i> or its subtypes are returned. If not specified then all <i>References</i> are returned and includeSubtypes is ignored.																		
includeSubtypes	Boolean	Indicates whether subtypes of the <i>ReferenceType</i> should be included in the browse. If TRUE, then instances of <i>referenceTypeId</i> and all of its subtypes are returned.																		
nodeClassMask	UInt32	Specifies the <i>NodeClasses</i> of the <i>TargetNodes</i>. Only <i>TargetNodes</i> with the selected <i>NodeClasses</i> are returned. The <i>NodeClasses</i> are assigned the following bits: <table><tr><th>Bit</th><th>NodeClass</th></tr><tr><td>0</td><td>Object</td></tr><tr><td>1</td><td>Variable</td></tr><tr><td>2</td><td>Method</td></tr><tr><td>3</td><td>ObjectType</td></tr><tr><td>4</td><td>VariableType</td></tr><tr><td>5</td><td>ReferenceType</td></tr><tr><td>6</td><td>DataType</td></tr><tr><td>7</td><td>View</td></tr></table> If set to zero, then all <i>NodeClasses</i> are returned. If the <i>NodeClass</i> is unknown for a remote <i>Node</i>, the <i>nodeClassMask</i> is ignored.	Bit	NodeClass	0	Object	1	Variable	2	Method	3	ObjectType	4	VariableType	5	ReferenceType	6	DataType	7	View
Bit	NodeClass																			
0	Object																			
1	Variable																			
2	Method																			
3	ObjectType																			
4	VariableType																			
5	ReferenceType																			
6	DataType																			
7	View																			

Name	Type	Description														
resultMask	UInt32	Specifies the fields in the <i>ReferenceDescription</i> structure that should be returned. The fields are assigned the following bits: <table><tr><th>Bit</th><th>Result</th></tr><tr><td>0</td><td>ReferenceType</td></tr><tr><td>1</td><td>IsForward</td></tr><tr><td>2</td><td>NodeClass</td></tr><tr><td>3</td><td>BrowseName</td></tr><tr><td>4</td><td>DisplayName</td></tr><tr><td>5</td><td>TypeDefinition</td></tr></table> The <i>ReferenceDescription</i> type is defined in 7.25.	Bit	Result	0	ReferenceType	1	IsForward	2	NodeClass	3	BrowseName	4	DisplayName	5	TypeDefinition
Bit	Result															
0	ReferenceType															
1	IsForward															
2	NodeClass															
3	BrowseName															
4	DisplayName															
5	TypeDefinition															
Response																
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).														
results []	BrowseResult	A list of <i>BrowseResults</i>. The size and order of the list matches the size and order of the <i>nodesToBrowse</i> specified in the request. The <i>BrowseResult</i> type is defined in 7.3.														
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>results</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>results</i> response parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.														

5.8.2.3 Service results

Table 35 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 35 – Browse Service result codes

Symbolic Id	Description
Bad_ViewIdUnknown	See Table 177 for the description of this result code.
Bad_ViewTimestampInvalid	See Table 177 for the description of this result code.
Bad_ViewParameterMismatchInvalid	See Table 177 for the description of this result code.
Bad_ViewVersionInvalid	See Table 177 for the description of this result code.
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.8.2.4 StatusCodes

Table 36 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 36 – Browse Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_BrowseDirectionInvalid	See Table 178 for the description of this result code.
Bad_NodeNotInView	See Table 178 for the description of this result code.
Bad_NoContinuationPoints	See Table 178 for the description of this result code.
Uncertain_NotAllNodesAvailable	Browse results may be incomplete because of the unavailability of a subsystem.

5.8.3 BrowseNext

5.8.3.1 Description

This *Service* is used to request the next set of *Browse* or *BrowseNext* response information that is too large to be sent in a single response. "Too large" in this context means that the *Server* is not able to return a larger response or that the number of results to return exceeds the maximum number of results to return that was specified by the *Client* in the original *Browse* request. The *BrowseNext* shall be submitted on the same *Session* that was used to submit the *Browse* or *BrowseNext* that is being continued.

5.8.3.2 Parameters

Table 37 defines the parameters for the *Service*.

Table 37 – BrowseNext Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
releaseContinuationPoints	Boolean	A <i>Boolean</i> parameter with the following values: TRUE passed <i>continuationPoints</i> shall be reset to free resources in the <i>Server</i>. The continuation points are released and the results and diagnosticInfos arrays are empty. FALSE passed <i>continuationPoints</i> shall be used to get the next set of browse information. A <i>Client</i> shall always use the continuation point returned by a <i>Browse</i> or <i>BrowseNext</i> response to free the resources for the continuation point in the <i>Server</i>. If the <i>Client</i> does not want to get the next set of browse information, <i>BrowseNext</i> shall be called with this parameter set to TRUE.
continuationPoints []	Continuation Point	A list of <i>Server</i>-defined opaque values that represent continuation points. The value for a continuation point was returned to the <i>Client</i> in a previous <i>Browse</i> or <i>BrowseNext</i> response. These values are used to identify the previously processed <i>Browse</i> or <i>BrowseNext</i> request that is being continued and the point in the result set from which the browse response is to continue. Clients may mix continuation points from different <i>Browse</i> or <i>BrowseNext</i> responses. The <i>ContinuationPoint</i> type is described in 7.6.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	BrowseResult	A list of references that met the criteria specified in the original <i>Browse</i> request. The size and order of this list matches the size and order of the <i>continuationPoints</i> request parameter. The <i>BrowseResult</i> type is defined in 7.3.
diagnosticInfos []	Diagnostic Info	List of diagnostic information for the <i>results</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>results</i> response parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.8.3.3 Service results

Table 38 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 38 – BrowseNext Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.8.3.4 StatusCodes

Table 39 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 39 – BrowseNext Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_ReferenceTypeIdInvalid	See Table 178 for the description of this result code.
Bad_BrowseDirectionInvalid	See Table 178 for the description of this result code.
Bad_NodeNotInView	See Table 178 for the description of this result code.
Bad_ContinuationPointInvalid	See Table 178 for the description of this result code.

5.8.4 TranslateBrowsePathsToNodeIds

5.8.4.1 Description

This *Service* is used to request that the *Server* translates one or more browse paths to *NodeIds*. Each browse path is constructed of a starting *Node* and a *RelativePath*. The specified starting *Node* identifies the *Node* from which the *RelativePath* is based. The *RelativePath* contains a sequence of *ReferenceTypes* and *BrowseNames*.

One purpose of this *Service* is to allow programming against type definitions. Since *BrowseNames* shall be unique in the context of type definitions, a *Client* may create a browse path that is valid for a type definition and use this path on instances of the type. For example, an *ObjectType* "Boiler" may have a "HeatSensor" *Variable* as *InstanceDeclaration*. A graphical element programmed against the "Boiler" may need to display the *Value* of the "HeatSensor". If the graphical element would be called on "Boiler1", an instance of "Boiler", it would need to call this *Service* specifying the *NodeId* of "Boiler1" as starting *Node* and the *BrowseName* of the "HeatSensor" as browse path. The *Service* would return the *NodeId* of the "HeatSensor" of "Boiler1" and the graphical element could subscribe to its *Value Attribute*.

If a *Node* has multiple targets with the same *BrowseName*, the *Server* shall return a list of *NodeIds*. However, since one of the main purposes of this *Service* is to support programming against type definitions, the *NodeId* of the *Node* based on the type definition of the starting *Node* is returned as the first *NodeId* in the list.

5.8.4.2 Parameters

Table 40 defines the parameters for the *Service*.

Table 40 – TranslateBrowsePathsToNodeIds Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
browsePaths []	BrowsePath	List of browse paths for which <i>NodeIds</i> are being requested. This structure is defined in-line with the following indented items.
startingNode	NodeId	<i>NodeId</i> of the starting <i>Node</i> for the browse path.
relativePath	RelativePath	The path to follow from the <i>startingNode</i> . The last element in the <i>relativePath</i> shall always have a <i>targetName</i> specified. This further restricts the definition of the <i>RelativePath</i> type. The <i>Server</i> shall return <i>Bad_BrowseNameInvalid</i> if the <i>targetName</i> is missing. The <i>RelativePath</i> structure is defined in 7.26.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	BrowsePathResult	List of results for the list of browse paths. The size and order of the list matches the size and order of the <i>browsePaths</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the browse path (see 7.34 for <i>StatusCode</i> definition).
targets []	BrowsePathTarget	List of targets for the <i>relativePath</i> from the <i>startingNode</i> . This structure is defined in-line with the following indented items. A <i>Server</i> may encounter a <i>Reference</i> to a <i>Node</i> in another <i>Server</i> which it cannot follow while it is processing the <i>RelativePath</i> . If this happens the <i>Server</i> returns the <i>NodeId</i> of the external <i>Node</i> and sets the <i>remainingPathIndex</i> parameter to indicate which <i>RelativePath</i> elements still need to be processed. To complete the operation the <i>Client</i> shall connect to the other <i>Server</i> and call this service again using the target as the <i>startingNode</i> and the unprocessed elements as the <i>relativePath</i> .
targetId	ExpandedNodeId	The identifier for a target of the <i>RelativePath</i> .
remainingPathIndex	Index	The index of the first unprocessed element in the <i>RelativePath</i> . This value shall be equal to the maximum value of <i>Index</i> data type if all elements were processed (see 7.13 for <i>Index</i> definition).
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the list of browse paths (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>browsePaths</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.8.4.3 Service results

Table 41 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in 7.34.

Table 41 – TranslateBrowsePathsToNodeIds Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.8.4.4 StatusCodes

Table 42 defines values for the operation level *statusCode* parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 42 – TranslateBrowsePathsToNodeIds Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_NothingToDo	See Table 177 for the description of this result code. This code indicates that the relativePath contained an empty list.
Bad_BrowseNameInvalid	See Table 178 for the description of this result code. This code indicates that a TargetName was missing in a RelativePath.
Uncertain_ReferenceOutOfServer	The path element has targets which are in another server.
Bad_TooManyMatches	The requested operation has too many matches to return. Users should use queries for large result sets. <i>Servers</i> should allow at least 10 matches before returning this error code.
Bad_QueryTooComplex	The requested operation requires too many resources in the server.
Bad_NoMatch	The requested relativePath cannot be resolved to a target to return.

5.8.5 RegisterNodes

5.8.5.1 Description

A *Server* often has no direct access to the information that it manages. Variables or services might be in underlying systems where additional effort is required to establish a connection to these systems. The *RegisterNodes Service* can be used by *Clients* to register the *Nodes* that they know they will access repeatedly (e.g. Write, Call). It allows *Servers* to set up anything needed so that the access operations will be more efficient. Clients can expect performance improvements when using registered *NodeIds*, but the optimization measures are vendor-specific. For *Variable Nodes Servers* shall concentrate their optimization efforts on the *Value Attribute*.

Registered *NodeIds* are only guaranteed to be valid within the current *Session*. *Clients* shall unregister unneeded *Ids* immediately to free up resources.

RegisterNodes does not validate the *NodeIds* from the request. *Servers* will simply copy unknown *NodeIds* in the response. Structural *NodeId* errors (size violations, invalid id types) will cause the complete *Service* to fail.

For the purpose of *Auditing*, *Servers* shall not use the registered *NodeIds* but only the canonical *NodeIds* which is the value of the *NodeId Attribute*.

5.8.5.2 Parameters

Table 43 defines the parameters for the *Service*.

Table 43 – RegisterNodes Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToRegister []	NodeId	List of <i>NodeIds</i> to register that the <i>Client</i> has retrieved through browsing, querying or in some other manner.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
registeredNodeIds []	NodeId	A list of <i>NodeIds</i> which the <i>Client</i> shall use for subsequent access operations. The size and order of this list matches the size and order of the <i>nodesToRegister</i> request parameter. The <i>Server</i> may return the <i>NodeId</i> from the request or a new (an alias) <i>NodeId</i>. It is recommended that the <i>Server</i> return a numeric <i>NodeIds</i> for aliasing. In case no optimization is supported for a <i>Node</i>, the <i>Server</i> shall return the <i>NodeId</i> from the request.

5.8.5.3 Service results

Table 44 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 44 – RegisterNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_NodeIdInvalid	See Table 178 for the description of this result code. <i>Servers</i> shall completely reject the <i>RegisterNodes</i> request if any of the <i>NodeIds</i> in the <i>nodesToRegister</i> parameter are structurally invalid.

5.8.6 UnregisterNodes

5.8.6.1 Description

This *Service* is used to unregister *NodeIds* that have been obtained via the *RegisterNodes* service.

UnregisterNodes does not validate the *NodeIds* from the request. *Servers* shall simply unregister *NodeIds* that are known as registered *NodeIds*. Any *NodeIds* that are in the list, but are not registered *NodeIds* are simply ignored.

5.8.6.2 Parameters

Table 50 defines the parameters for the *Service*.

Table 45 – UnregisterNodes Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToUnregister []	NodeId	A list of <i>NodeIds</i> that have been obtained via the <i>RegisterNodes</i> service.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

5.8.6.3 Service results

Table 51 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 46 – UnregisterNodes Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.9 Query Service Set

5.9.1 Overview

This *Service Set* is used to issue a *Query* to a *Server*. OPC UA *Query* is generic in that it provides an underlying storage mechanism independent *Query* capability that can be used to access a wide variety of OPC UA data stores and information management systems. OPC UA *Query* permits a *Client* to access data maintained by a *Server* without any knowledge of the logical schema used for internal storage of the data. Knowledge of the *AddressSpace* is sufficient.

An *OPC UA Application* is expected to use the OPC UA *Query Services* as part of an initialization process or an occasional information synchronization step. For example, OPC UA *Query* would be used for bulk data access of a persistent store to initialise an analysis application with the current state of a system configuration. A *Query* may also be used to initialise or populate data for a report.

A *Query* defines what instances of one or more *TypeDefinitionNodes* in the *AddressSpace* should supply a set of *Attributes*. Results returned by a *Server* are in the form of an array of *QueryDataSets*. The selected *Attribute* values in each *QueryDataSet* come from the definition of the selected *TypeDefinitionNodes* or related *TypeDefinitionNodes* and appear in results in the same order as the *Attributes* that were passed into the *Query*. *Query* also supports *Node* filtering on the basis of *Attribute* values, as well as relationships between *TypeDefinitionNodes*.

See Annex B for example queries.

5.9.2 Querying Views

A *View* is a subset of the *AddressSpace* available in the *Server*. See IEC 62541-5 for a description of the organization of *Views* in the *AddressSpace*.

For any existing *View*, a *Query* may be used to return a subset of data from the *View*. When an application issues a *Query* against a *View*, only data defined by the *View* is returned. Data not included in the *View* but included in the original *AddressSpace* is not returned.

The *Query Services* supports access to current and historical data. The *Service* supports a *Client* querying a past version of the *AddressSpace*. Clients may specify a *ViewVersion* or a *Timestamp* in a *Query* to access past versions of the *AddressSpace*. OPC UA *Query* is complementary to Historical Access in that the former is used to *Query* an *AddressSpace* that existed at a time and the latter is used to *Query* for the value of *Attributes* over time. In this way, a *Query* can be used to retrieve a portion of a past *AddressSpace* so that *Attribute* value history may be accessed using Historical Access even if the *Node* is no longer in the current *AddressSpace*.

Servers that support *Query* are expected to be able to access the *AddressSpace* that is associated with the local *Server* and any *Views* that are available on the local *Server*. If a *View* or the *AddressSpace* also references a remote *Server*, query may be able to access the *AddressSpace* of the remote *Server*, but it is not required. If a *Server* does access a remote *Server* the access shall be accomplished using the user identity of the *Client* as described in 5.5.1.

5.9.3 QueryFirst

5.9.3.1 Description

This *Service* is used to issue a *Query* request to the *Server*. The complexity of the *Query* can range from very simple to highly sophisticated. The *Query* can simply request data from instances of a *TypeDefinitionNode* or *TypeDefinitionNode* subject to restrictions specified by the filter. On the other hand, the *Query* can request data from instances of related *Node* types by specifying a *RelativePath* from an originating *TypeDefinitionNode*. In the filter, a separate set of paths can be constructed for limiting the instances that supply data. A filtering path can include multiple *RelatedTo* operators to define a multi-hop path between source instances and target instances. For example, one could filter on students that attend a particular school, but return information about students and their families. In this case, the student school relationship is traversed for filtering, but the student family relationship is traversed to select data. For a complete description of *ContentFilter* see 7.4, also see Clause B.1 for simple examples and Clause B.2 for more complex examples of content filter and queries.

The *Client* provides an array of *NodeTypeDescription* which specify the *NodeId* of a *TypeDefinitionNode* and selects what *Attributes* are to be returned in the response. A *Client* can also provide a set of *RelativePaths* through the type system starting from an originating *TypeDefinitionNode*. Using these paths, the *Client* selects a set of *Attributes* from *Nodes* that are related to instances of the originating *TypeDefinitionNode*. Additionally, the *Client* can request the *Server* return instances of subtypes of *TypeDefinitionNodes*. If a selected *Attribute* does not exist in a *TypeDefinitionNode* but does exist in a subtype, it is assumed to have a null value in the *TypeDefinitionNode* in question. Therefore, this does not constitute an error condition and a null value is returned for the *Attribute*.

The *Client* can use the filter parameter to limit the result set by restricting *Attributes* and *Properties* to certain values. Another way the *Client* can use a filter to limit the result set is by specifying how instances should be related, using *RelatedTo* operators. In this case, if an instance at the top of the *RelatedTo* path cannot be followed to the bottom of the path via specified hops, no *QueryDataSets* are returned for the starting instance or any of the intermediate instances.

When querying for related instances in the *RelativePath*, the *Client* can optionally ask for *References*. A *Reference* is requested via a *RelativePath* that only includes a *ReferenceType*. If all *References* are desired then the root *ReferenceType* is listed. These *References* are returned as part of the *QueryDataSets*.

Query Services allow a special handling of the *targetName* field in the *RelativePath*. In several Query use cases a type *NodeId* is necessary in the path instead of a *QualifiedName*. Therefore the Client is allowed to specify a *NodeId* in the *QualifiedName*. This is done by setting the *namespaceIndex* of the *targetName* to zero and the name part of the *targetName* to the XML representation of the *NodeId*. The XML representation is defined in IEC 62541-6. When matching instances are returned as the target node, the target node shall be an instance of the specified type or subtype of the specified type.

Table 47 defines the request parameters and Table 48 the response parameters for the *QueryFirst Service*.

Table 47 – QueryFirst Request parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
View	ViewDescription	Specifies a <i>View</i> and temporal context to a <i>Server</i> (see 7.39 for <i>ViewDescription</i> definition).
nodeTypes []	NodeTypeDescription	This is the <i>Node</i> type description. This structure is defined in-line with the following indented items.
typeDefinitionNode	ExpandedNodeId	<i>NodeId</i> of the originating <i>TypeDefinitionNode</i> of the instances for which data is to be returned.
includeSubtypes	Boolean	A flag that indicates whether the <i>Server</i> should include instances of subtypes of the <i>TypeDefinitionNode</i> in the list of instances of the <i>Node</i> type.
dataToReturn []	QueryDataDescription	Specifies an <i>Attribute</i> or <i>Reference</i> from the originating <i>typeDefinitionNode</i> along a given <i>relativePath</i> for which to return data. This structure is defined in-line with the following indented items.
relativePath	RelativePath	Browse path relative to the originating <i>Node</i> that identifies the <i>Node</i> which contains the data that is being requested, where the originating <i>Node</i> is an instance <i>Node</i> of the type defined by the type definition <i>Node</i>. The instance <i>Nodes</i> are further limited by the filter provided as part of this call. For a definition of <i>relativePath</i> see 7.26. This relative path could end on a <i>Reference</i>, in which case the <i>ReferenceDescription</i> of the <i>Reference</i> would be returned as its value. The <i>targetName</i> field of the <i>relativePath</i> may contain a type <i>NodeId</i>. This is done by setting the <i>namespaceIndex</i> of the <i>targetName</i> to zero and the <i>name</i> part of the <i>targetName</i> to the XML representation of the <i>NodeId</i>. The XML representation is defined in IEC 62541-6. When matching instances are returned as the target node, the target node shall be an instance of the specified type or subtype of the specified type.
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>Attribute</i> Id. The <i>IntegerId</i> is defined in 7.14. The <i>IntegerId</i> for <i>Attributes</i> are defined in IEC 62541-6. If the <i>RelativePath</i> ended in a <i>Reference</i> then this parameter is 0 and ignored by the server.
indexRange	NumericRange	This parameter is used to identify a single element of a structure or an array, or a single range of indexes for arrays. If a range of elements are specified, the values are returned as a composite. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is null if the specified <i>Attribute</i> is not an array or a structure. However, if the specified <i>Attribute</i> is an array or a structure, and this parameter is null, then all elements are to be included in the range.
Filter	ContentFilter	Resulting <i>Nodes</i> shall be limited to the <i>Nodes</i> matching the criteria defined by the filter. <i>ContentFilter</i> is discussed in 7.4. If an empty filter is provided then the entire <i>AddressSpace</i> shall be examined and all <i>Nodes</i> that contain a matching requested <i>Attribute</i> or <i>Reference</i> are returned.

Name	Type	Description
maxDataSetsToReturn	Counter	The number of <i>QueryDataSets</i> that the <i>Client</i> wants the <i>Server</i> to return in the response and on each subsequent continuation call response. The <i>Server</i> is allowed to further limit the response, but shall not exceed this limit. A value of 0 indicates that the <i>Client</i> is imposing no limitation.
maxReferencesToReturn	Counter	The number of <i>References</i> that the <i>Client</i> wants the <i>Server</i> to return in the response for each <i>QueryDataSet</i> and on each subsequent continuation call response. The <i>Server</i> is allowed to further limit the response, but shall not exceed this limit. A value of 0 indicates that the <i>Client</i> is imposing no limitation. For example a result where 4 <i>Nodes</i> are being returned, but each has 100 <i>References</i> , if this limit were set to 50 then only the first 50 <i>References</i> for each <i>Node</i> would be returned on the initial call and a continuation point would be set indicating additional data.

Table 48 – QueryFirst Response parameters

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
queryDataSets []	QueryDataSet	The array of <i>QueryDataSets</i> . This array is empty if no <i>Nodes</i> or <i>References</i> met the <i>nodeTypes</i> criteria. In this case the continuationPoint parameter shall be empty. The <i>QueryDataSet</i> type is defined in 7.23.
continuationPoint	ContinuationPoint	Server-defined opaque value that identifies the continuation point. The continuation point is used only when the <i>Query</i> results are too large to be returned in a single response. "Too large" in this context means that the <i>Server</i> is not able to return a larger response or that the number of <i>QueryDataSets</i> to return exceeds the maximum number of <i>QueryDataSets</i> to return that was specified by the <i>Client</i> in the request. The continuation point is used in the <i>QueryNext Service</i> . When not used, the value of this parameter is null. If a continuation point is returned, the <i>Client</i> shall call <i>QueryNext</i> to get the next set of <i>QueryDataSets</i> or to free the resources for the continuation point in the <i>Server</i> . A continuation point shall remain active until the <i>Client</i> passes the continuation point to <i>QueryNext</i> or the session is closed. If the maximum continuation points have been reached the oldest continuation point shall be reset. The <i>ContinuationPoint</i> type is described in 7.6.
parsingResults[]	ParsingResult	List of parsing results for <i>QueryFirst</i> . The size and order of the list matches the size and order of the <i>NodeTypes</i> request parameter. This structure is defined in-line with the following indented items. This list is populated with any status codes that are related to the processing of the node types that are part of the query. The array can be empty if no errors were encountered. If any node type encountered an error all node types shall have an associated status code.
statusCode	StatusCode	Parsing result for the requested <i>NodeTypeDescription</i> .
dataStatusCodes []	StatusCode	List of results for <i>dataToReturn</i> . The size and order of the list matches the size and order of the <i>dataToReturn</i> request parameter. The array can be empty if no errors were encountered.
dataDiagnosticInfos []	DiagnosticInfo	List of diagnostic information <i>dataToReturn</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>dataToReturn</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the query request.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the requested <i>NodeTypeDescription</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the query request.
filterResult	ContentFilter Result	A structure that contains any errors associated with the filter. This structure shall be empty if no errors occurred. The <i>ContentFilterResult</i> type is defined in 7.4.2.

5.9.3.2 Service results

If the *Query* is invalid or cannot be processed, then *QueryDataSets* are not returned and only a *Service* result, *filterResult*, *parsingResults* and optional *DiagnosticInfo* is returned. Table 49 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 49 – QueryFirst Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_ContentFilterInvalid	See Table 178 for the description of this result code.
Bad_ViewIdUnknown	See Table 177 for the description of this result code.
Bad_ViewTimestampInvalid	See Table 177 for the description of this result code.
Bad_ViewParameterMismatchInvalid	See Table 177 for the description of this result code.
Bad_ViewVersionInvalid	See Table 177 for the description of this result code.
Bad_InvalidFilter	The provided filter is invalid, see the <i>filterResult</i> for specific errors
Bad_NodelistError	The <i>NodeTypes</i> provided contain an error, see the <i>parsingResults</i> for specific errors
Bad_InvalidView	The provided <i>ViewDescription</i> is not a valid <i>ViewDescription</i> .
Good_ResultsMayBeIncomplete	The <i>Server</i> should have followed a reference to a node in a remote <i>Server</i> but did not. The result set may be incomplete.

5.9.3.3 StatusCodes

Table 50 defines values for the *parsingResults statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 50 – QueryFirst Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_NotTypeDefinition	The provided <i>NodeId</i> was not a type definition <i>NodeId</i> .
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.

5.9.4 QueryNext

5.9.4.1 Descriptions

This *Service* is used to request the next set of *QueryFirst* or *QueryNext* response information that is too large to be sent in a single response. "Too large" in this context means that the *Server* is not able to return a larger response or that the number of *QueryDataSets* to return exceeds the maximum number of *QueryDataSets* to return that was specified by the *Client* in the original request. The *QueryNext* shall be submitted on the same session that was used to submit the *QueryFirst* or *QueryNext* that is being continued.

5.9.4.2 Parameters

Table 51 defines the parameters for the *Service*.

Table 51 – QueryNext Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
releaseContinuationPoint	Boolean	A <i>Boolean</i> parameter with the following values: TRUE passed <i>continuationPoint</i> shall be reset to free resources for the continuation point in the <i>Server</i>. FALSE passed <i>continuationPoint</i> shall be used to get the next set of <i>QueryDataSets</i>. A <i>Client</i> shall always use the continuation point returned by a <i>QueryFirst</i> or <i>QueryNext</i> response to free the resources for the continuation point in the <i>Server</i>. If the <i>Client</i> does not want to get the next set of <i>Query</i> information, <i>QueryNext</i> shall be called with this parameter set to TRUE. If the parameter is set to TRUE all array parameters in the response shall contain empty arrays.
continuationPoint	ContinuationPoint	Server defined opaque value that represents the continuation point. The value of the continuation point was returned to the <i>Client</i> in a previous <i>QueryFirst</i> or <i>QueryNext</i> response. This value is used to identify the previously processed <i>QueryFirst</i> or <i>QueryNext</i> request that is being continued, and the point in the result set from which the browse response is to continue. The <i>ContinuationPoint</i> type is described in 7.6.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
queryDataSets []	QueryDataSet	The array of <i>QueryDataSets</i>. The <i>QueryDataSet</i> type is defined in 7.23.
revisedContinuationPoint	ContinuationPoint	Server-defined opaque value that represents the continuation point. It is used only if the information to be returned is too large to be contained in a single response. When not used or when <i>releaseContinuationPoint</i> is set, the value of this parameter is null. The <i>ContinuationPoint</i> type is described in 7.6.

5.9.4.3 Service results

Table 52 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 52 – QueryNext Service result codes

Symbolic Id	Description
Bad_ContinuationPointInvalid	See Table 178 for the description of this result code.

5.10 Attribute Service Set

5.10.1 Overview

This *Service Set* provides *Services* to access *Attributes* that are part of *Nodes*.

5.10.2 Read

5.10.2.1 Description

This *Service* is used to read one or more *Attributes* of one or more *Nodes*. For constructed *Attribute* values whose elements are indexed, such as an array, this *Service* allows *Clients* to read the entire set of indexed values as a composite, to read individual elements or to read ranges of elements of the composite.

The *maxAge* parameter is used to direct the *Server* to access the value from the underlying data source, such as a device, if its copy of the data is older than that which the *maxAge* specifies. If the *Server* cannot meet the requested maximum age, it returns its "best effort" value rather than rejecting the request.

5.10.2.2 Parameters

Table 53 defines the parameters for the *Service*.

Table 53 – Read Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
maxAge	Duration	Maximum age of the value to be read in milliseconds. The age of the value is based on the difference between the <i>ServerTimestamp</i> and the time when the <i>Server</i> starts processing the request. For example if the <i>Client</i> specifies a <i>maxAge</i> of 500 milliseconds and it takes 100 milliseconds until the <i>Server</i> starts processing the request, the age of the returned value could be 600 milliseconds prior to the time it was requested. If the <i>Server</i> has one or more values of an <i>Attribute</i> that are within the maximum age, it can return any one of the values or it can read a new value from the data source. The number of values of an <i>Attribute</i> that a <i>Server</i> has depends on the number of <i>MonitoredItems</i> that are defined for the <i>Attribute</i>. In any case, the <i>Client</i> can make no assumption about which copy of the data will be returned. If the <i>Server</i> does not have a value that is within the maximum age, it shall attempt to read a new value from the data source. If the <i>Server</i> cannot meet the requested <i>maxAge</i>, it returns its "best effort" value rather than rejecting the request. This may occur when the time it takes the <i>Server</i> to process and return the new data value after it has been accessed is greater than the specified maximum age. If <i>maxAge</i> is set to 0, the <i>Server</i> shall attempt to read a new value from the data source. If <i>maxAge</i> is set to the max Int32 value or greater, the <i>Server</i> shall attempt to get a cached value. Negative values are invalid for <i>maxAge</i>.
timestampsToReturn	Enum TimestampsToReturn	An enumeration that specifies the <i>Timestamps</i> to be returned for each requested <i>Variable Value Attribute</i> . The <i>TimestampsToReturn</i> enumeration is defined in 7.35.
nodesToRead []	ReadValueId	List of <i>Nodes</i> and their <i>Attributes</i> to read. For each entry in this list, a <i>StatusCode</i> is returned, and if it indicates success, the <i>Attribute Value</i> is also returned. The <i>ReadValueId</i> parameter type is defined in 7.24.

Name	Type	Description
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	DataValue	List of <i>Attribute</i> values (see 7.7 for <i>DataValue</i> definition). The size and order of this list matches the size and order of the <i>nodesToRead</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToRead</i> parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of this list matches the size and order of the <i>nodesToRead</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToRead</i> parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.2.3 Service results

Table 54 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 54 – Read Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_MaxAgeInvalid	The max age parameter is invalid.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.

5.10.2.4 StatusCodes

Table 55 defines values for the operation level *statusCode* contained in the *DataValue* structure of each *results* element. Common *StatusCodes* are defined in Table 178.

Table 55 – Read Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeNoData	See Table 178 for the description of this result code.
Bad_DataEncodingInvalid	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_NotReadable	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_SecurityModelInsufficient	See Table 178 for the description of this result code.

5.10.3 HistoryRead

5.10.3.1 Description

This *Service* is used to read historical values or *Events* of one or more *Nodes*. For constructed *Attribute* values whose elements are indexed, such as an array, this *Service*

allows *Clients* to read the entire set of indexed values as a composite, to read individual elements or to read ranges of elements of the composite. *Servers* may make historical values available to *Clients* using this *Service*, although the historical values themselves are not visible in the *AddressSpace*.

The *AccessLevel Attribute* defined in IEC 62541-3 indicates a *Node's* support for historical values. Several request parameters indicate how the *Server* is to access values from the underlying history data source. The *EventNotifier Attribute* defined in IEC 62541-3 indicates a *Node's* support for historical *Events*.

The *continuationPoint* parameter in the *HistoryRead* is used to mark a point from which to continue the read if not all values could be returned in one response. The value is opaque for the *Client* and is only used to maintain the state information for the *Server* to continue from. A *Server* may use the timestamp of the last returned data item if the timestamp is unique. This can reduce the need in the *Server* to store state information for the continuation point.

In some cases it may take longer than the *Client* timeout hint to read the data for all nodes to read. Then the *Server* may return zero results with a continuation point for the affected nodes before the timeout expires. That allows the *Server* to resume the data acquisition on the next *Client* read call.

For additional details on reading historical data and historical *Events* see IEC 62541-11.

5.10.3.2 Parameters

Table 56 defines the parameters for the *Service*.

Table 56 – HistoryRead Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
historyReadDetails	Extensible Parameter HistoryReadDetails	The details define the types of history reads that can be performed. The <i>HistoryReadDetails</i> parameter type is an extensible parameter type formally defined in IEC 62541-11. The <i>ExtensibleParameter</i> type is defined in 7.12.
timestampsToReturn	Enum TimestampsToReturn	An enumeration that specifies the timestamps to be returned for each requested <i>Variable Value Attribute</i> . The <i>TimestampsToReturn</i> enumeration is defined in 7.35. Specifying a <i>TimestampsToReturn</i> of NEITHER is not valid. A <i>Server</i> shall return a <i>Bad_InvalidTimestampArgument StatusCode</i> in this case. IEC 62541-11 defines exceptions where this parameter shall be ignored.
releaseContinuation Points	Boolean	A <i>Boolean</i> parameter with the following values: TRUE passed <i>continuationPoints</i> shall be reset to free resources in the <i>Server</i> . FALSE passed <i>continuationPoints</i> shall be used to get the next set of historical information. A <i>Client</i> shall always use the continuation point returned by a <i>HistoryRead</i> response to free the resources for the continuation point in the <i>Server</i> . If the <i>Client</i> does not want to get the next set of historical information, <i>HistoryRead</i> shall be called with this parameter set to TRUE.
nodesToRead []	HistoryReadValueId	This parameter contains the list of items upon which the historical retrieval is to be performed. This structure is defined in-line with the following indented items.

Name	Type	Description
nodeId	NodeId	If the <i>HistoryReadDetails</i> is RAW, PROCESSED, MODIFIED or ATTIME: The <i>nodeId</i> of the <i>Nodes</i> whose historical values are to be read. The value returned shall always include a timestamp. If the <i>HistoryReadDetails</i> is EVENTS: The <i>NodeId</i> of the <i>Node</i> whose <i>Event</i> history is to be read. If the <i>Node</i> does not support the requested access for historical values or historical <i>Events</i> the appropriate error response for the given <i>Node</i> shall be generated.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for arrays. If a range of elements is specified, the values are returned as a composite. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is null if the value is not an array. However, if the value is an array, and this parameter is null, then all elements are to be included in the range.
dataEncoding	QualifiedName	A <i>QualifiedName</i> that specifies the data encoding to be returned for the <i>Value</i> to be read (see 7.24 for definition how to specify the data encoding). This parameter only applies if the <i>DataType</i> of the <i>Variable</i> is a subtype of <i>Structure</i>. It is an error to specify this parameter if the <i>DataType</i> of the <i>Variable</i> is not a subtype of <i>Structure</i>. The parameter is ignored when reading history of <i>Events</i>.
continuationPoint	ContinuationPoint	For each <i>NodesToRead</i> item this parameter specifies a continuation point returned from a previous <i>HistoryRead</i> call, allowing the <i>Client</i> to continue that read from the last value received. The <i>HistoryRead</i> is used to select an ordered sequence of historical values or events. A continuation point marks a point in that ordered sequence, such that the <i>Server</i> returns the subset of the sequence that follows that point. A null value indicates that this parameter is not used. See 7.6 for a general description of continuation points. This continuation point is described in more detail in IEC 62541-11.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> type).
results []	HistoryReadResult	List of read results. The size and order of the list matches the size and order of the <i>nodesToRead</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the <i>NodesToRead</i> item (see 7.34 for <i>StatusCode</i> definition).
continuationPoint	ContinuationPoint	This parameter is used only if the number of values to be returned is too large to be returned in a single response or if the timeout provided as hint by the <i>Client</i> is close to expiring and not all nodes have been processed. When this parameter is not used, its value is null. <i>Servers</i> shall support at least one continuation point per <i>Session</i>. <i>Servers</i> specify a max history continuation points per <i>Session</i> in the <i>Server capabilities Object</i> defined in IEC 62541-5. A continuation point shall remain active until the <i>Client</i> passes the continuation point to <i>HistoryRead</i> or the <i>Session</i> is closed. If the max continuation points have been reached the oldest continuation point shall be reset.
historyData	Extensible Parameter HistoryData	The history data returned for the <i>Node</i>. The <i>HistoryData</i> parameter type is an extensible parameter type formally defined in IEC 62541-11. It specifies the types of history data that can be returned. The <i>ExtensibleParameter</i> base type is defined in 7.12.
diagnosticInfos []	Diagnostic Info	List of diagnostic information. The size and order of the list matches the size and order of the <i>nodesToRead</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToRead</i> parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.3.3 Service results

Table 57 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 57 – HistoryRead Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.
Bad_HistoryOperationInvalid	See Table 178 for the description of this result code.
Bad_HistoryOperationUnsupported	See Table 178 for the description of this result code. The requested history operation is not supported by the server.

5.10.3.4 StatusCodes

Table 58 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178. History access specific *StatusCodes* are defined in IEC 62541-11.

Table 58 – HistoryRead Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_DataEncodingInvalid	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_ContinuationPointInvalid	See Table 177 for the description of this result code.
Bad_InvalidTimestampArgument	The defined timestamp to return was invalid.
Bad_HistoryOperationUnsupported	See Table 178 for the description of this result code. The requested history operation is not supported for the requested node.
Bad_NoContinuationPoints	See Table 178 for the description of this result code. See 7.6 for the rules to apply this status code.

5.10.4 Write

5.10.4.1 Description

This *Service* is used to write values to one or more *Attributes* of one or more *Nodes*. For constructed *Attribute* values whose elements are indexed, such as an array, this *Service* allows *Clients* to write the entire set of indexed values as a composite, to write individual elements or to write ranges of elements of the composite.

The values are written to the data source, such as a device, and the *Service* does not return until it writes the values or determines that the value cannot be written. In certain cases, the *Server* will successfully write to an intermediate system or *Server*, and will not know if the data source was updated properly. In these cases, the *Server* should report a success code that indicates that the write was not verified. In the cases where the *Server* is able to verify that it has successfully written to the data source, it reports an unconditional success.

The order the operations are processed in the *Server* is not defined and depends on the different data sources and the internal *Server* logic. If an *Attribute* and *Node* combination is contained in more than one operation, the order of the processing is undefined. If a *Client* requires sequential processing the *Client* needs separate *Service* calls.

It is possible that the *Server* may successfully write some *Attributes*, but not others. Rollback is the responsibility of the *Client*.

If a *Server* allows writing of *Attributes* with the *DataType LocalizedText*, the *Client* can add or overwrite the text for a locale by writing the text with the associated *LocaleId*. Writing a null *String* for the text for a locale shall delete the *String* for that locale. Writing a null *String* for the *locale* and a non-null *String* for the *text* is setting the *text* for an invariant locale. Writing a null *String* for the *text* and a null *String* for the *locale* shall delete the entries for all locales. If a *Client* attempts to write a *locale* that is either syntactically invalid or not supported, the *Server* returns *Bad_LocaleNotSupported*. The *Write* behaviour for *Value Attributes* with a *LocalizedText DataType* is *Server* specific but it is recommended to follow the same rules.

5.10.4.2 Parameters

Table 59 defines the parameters for the *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020

Table 59 – Write Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
nodesToWrite []	WriteValue	List of <i>Nodes</i> and their <i>Attributes</i> to write. This structure is defined in-line with the following indented items.
nodeId	NodeId	<i>NodeId</i> of the <i>Node</i> that contains the <i>Attributes</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>Attribute</i> id. The <i>IntegerId</i> is defined in 7.14. The <i>IntegerIds</i> for the <i>Attributes</i> are defined in IEC 62541-6.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for arrays. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is not used if the specified <i>Attribute</i> is not an array. However, if the specified <i>Attribute</i> is an array and this parameter is not used, then all elements are to be included in the range. The parameter is null if not used. A <i>Server</i> shall return a <i>Bad_WriteNotSupported</i> error if an <i>indexRange</i> is provided and writing of <i>indexRange</i> is not possible for the <i>Node</i>.
value	DataValue	The <i>Node's Attribute</i> value (see 7.7 for <i>DataValue</i> definition). If the <i>indexRange</i> parameter is specified then the <i>Value</i> shall be an array even if only one element is being written. If the <i>SourceTimestamp</i> or the <i>ServerTimestamp</i> is specified, the <i>Server</i> shall use these values. The <i>Server</i> returns a <i>Bad_WriteNotSupported</i> error if it does not support writing of timestamps. A <i>Server</i> shall return a <i>Bad_TypeMismatch</i> error if the data type of the written value is not the same type or subtype of the <i>Attribute's DataType</i>. Based on the <i>DataType</i> hierarchy, subtypes of the <i>Attribute DataType</i> shall be accepted by the <i>Server</i>. <i>Servers</i> may reject subtypes defined in newer specification versions than supported by the <i>Server</i> with <i>Bad_TypeMismatch</i>. For the <i>Value Attribute</i> the <i>DataType</i> is defined through the <i>DataType Attribute</i>. A <i>ByteString</i> is structurally the same as a one dimensional array of <i>Byte</i>. A <i>Server</i> shall accept a <i>ByteString</i> if an array of <i>Byte</i> is expected. The <i>Server</i> returns a <i>Bad_DataEncodingUnsupported</i> error if it does not support the provided data encoding. <i>Simple DataTypes</i> (see IEC 62541-3) use the same representation on the wire as their super types and therefore writing a value of a simple <i>DataType</i> cannot be distinguished from writing a value of its super type. The <i>Server</i> shall assume that by receiving the correct wire representation for a simple <i>DataType</i> the correct type was chosen. <i>Servers</i> are allowed to impose additional data validations on the value independent of the encoding (e.g. having an image in GIF format in a <i>ByteString</i>). In this case the <i>Server</i> shall return a <i>Bad_TypeMismatch</i> error if the validation fails.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of results for the <i>Nodes</i> to write (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>nodesToWrite</i> request parameter. There is one entry in this list for each <i>Node</i> contained in the <i>nodesToWrite</i> parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Nodes</i> to write (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>nodesToWrite</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.4.3 Service results

Table 60 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 60 – Write Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.10.4.4 StatusCodes

Table 61 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 61 – Write Operation Level result codes

Symbolic Id	Description
Good_CompletesAsynchronously	See Table 177 for the description of this result code. The value was successfully written to an intermediate system but the <i>Server</i> does not know if the data source was updated properly.
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeNoData	See Table 178 for the description of this result code.
Bad_WriteNotSupported	The requested write operation is not supported. If a <i>Client</i> attempts to write any value, status code, timestamp combination and the <i>Server</i> does not support the requested combination (which could be a single quantity such as just timestamp); then the <i>Server</i> shall not perform any write on this <i>Node</i> and shall return this <i>StatusCode</i> for this <i>Node</i> . It is also used if writing an <i>IndexRange</i> is not supported for a <i>Node</i> .
Bad_NotWritable	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code. The current user does not have permission to write the attribute.
Bad_OutOfRange	See Table 178 for the description of this result code. If a <i>Client</i> attempts to write a value outside the valid range like a value not contained in the enumeration data type of the <i>Node</i> , the <i>Server</i> shall return this <i>StatusCode</i> for this <i>Node</i> .
Bad_TypeMismatch	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_NoCommunication	See Table 178 for the description of this result code.
Bad_LocaleNotSupported	The locale in the requested write operation is not supported.

5.10.5 HistoryUpdate

5.10.5.1 Description

This *Service* is used to update historical values or *Events* of one or more *Nodes*. Several request parameters indicate how the *Server* is to update the historical value or *Event*. Valid actions are Insert, Replace or Delete.

5.10.5.2 Parameters

Table 62 defines the parameters for the *Service*.

Table 62 – HistoryUpdate Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
historyUpdateDetails []	Extensible Parameter HistoryUpdate Details	The details defined for this update. The <i>HistoryUpdateDetails</i> parameter type is an extensible parameter type formally defined in IEC 62541-11. It specifies the types of history updates that can be performed. The <i>ExtensibleParameter</i> type is defined in 7.12.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	HistoryUpdate Result	List of update results for the history update details. The size and order of the list matches the size and order of the details element of the <i>historyUpdateDetails</i> parameter specified in the request. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the update of the <i>Node</i> (see 7.34 for <i>StatusCode</i> definition).
operationResults []	StatusCode	List of <i>StatusCodes</i> for the operations to be performed on a <i>Node</i> . The size and order of the list matches the size and order of any list defined by the details element being reported by this result entry.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the operations to be performed on a <i>Node</i> (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of any list defined by the details element being reported by this <i>results</i> entry. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the history update details. The size and order of the list matches the size and order of the details element of the <i>historyUpdateDetails</i> parameter specified in the request. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.10.5.3 Service results

Table 63 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 63 – HistoryUpdate Service result codes

SymbolicId	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.10.5.4 StatusCodes

Table 64 defines values for the *statusCode* and *operationResults* parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178. History access specific *StatusCodes* are defined in IEC 62541-11.

Table 64 – HistoryUpdate Operation Level result codes

Symbolic Id	Description
Bad_NotWritable	See Table 178 for the description of this result code.
Bad_HistoryOperationInvalid	See Table 178 for the description of this result code.
Bad_HistoryOperationUnsupported	See Table 178 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code. The current user does not have permission to update the history.

5.11 Method Service Set

5.11.1 Overview

Methods represent the function calls of *Objects*. They are defined in IEC 62541-3. *Methods* are invoked and return only after completion (successful or unsuccessful). Execution times for *Methods* may vary, depending on the function that they perform.

The *Method Service Set* defines the means to invoke *Methods*. A *Method* shall be a component of an *Object*. Discovery is provided through the *Browse* and *Query Services*. *Clients* discover the *Methods* supported by a *Server* by browsing for the owning *Objects References* that identify their supported *Methods*.

Because *Methods* may control some aspect of plant operations, *Method* invocation may depend on environmental or other conditions. This may be especially true when attempting to re-invoke a *Method* immediately after it has completed execution. Conditions that are required to invoke the *Method* might not yet have returned to the state that permits the *Method* to start again.

5.11.2 Call

5.11.2.1 Description

This *Service* is used to call (invoke) a list of *Methods*.

This *Service* provides for passing input and output arguments to/from a *Method*. These arguments are defined by *Properties* of the *Method*.

If the *Method* is invoked in the context of a *Session* and the *Session* is terminated, the results of the *Method's* execution cannot be returned to the *Client* and are discarded. This is independent of the task actually performed at the *Server*.

The order the operations are processed in the *Server* is not defined and depends on the different tasks and the internal *Server* logic. If a *Method* is contained in more than one operation, the order of the processing is undefined. If a *Client* requires sequential processing the *Client* needs separate *Service* calls.

5.11.2.2 Parameters

Table 65 defines the parameters for the *Service*.

Table 65 – Call Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
methodsToCall []	CallMethodRequest	List of <i>Methods</i> to call. This structure is defined in-line with the following indented items.
objectId	NodeId	The <i>NodeId</i> shall be that of the <i>Object</i> or <i>ObjectType</i> on which the <i>Method</i> is invoked. In case of an <i>ObjectType</i>, the <i>ObjectType</i> or a super type of the <i>ObjectType</i> shall be the source of a <i>HasComponent Reference</i> (or subtype of <i>HasComponent Reference</i>) to the <i>Method</i> specified in <i>methodId</i>. In case of an <i>Object</i>, the <i>Object</i> or the <i>ObjectType</i> of the <i>Object</i> or a super type of that <i>ObjectType</i> shall be the source of a <i>HasComponent Reference</i> (or subtype of <i>HasComponent Reference</i>) to the <i>Method</i> specified in <i>methodId</i>. See IEC 62541-3 for a description of <i>Objects</i> and their <i>Methods</i>.
methodId	NodeId	<i>NodeId</i> of the <i>Method</i> to invoke. If the <i>objectId</i> is the <i>NodeId</i> of an <i>Object</i>, it is allowed to use the <i>NodeId</i> of a <i>Method</i> that is the target of a <i>HasComponent Reference</i> from the <i>ObjectType</i> of the <i>Object</i>.
inputArguments []	BaseDataType	List of input argument values. An empty list indicates that there are no input arguments. The size and order of this list matches the size and order of the input arguments defined by the input <i>InputArguments Property</i> of the <i>Method</i>. The name, a description and the data type of each argument are defined by the <i>Argument</i> structure in each element of the method's <i>InputArguments Property</i>.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	CallMethodResult	Result for the <i>Method</i> calls. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> of the <i>Method</i> executed in the server. This <i>StatusCode</i> is set to the <i>Bad_InvalidArgument StatusCode</i> if at least one input argument broke a constraint (e.g. wrong data type, value out of range). This <i>StatusCode</i> is set to a bad <i>StatusCode</i> if the <i>Method</i> execution failed in the server, e.g. based on an exception.
inputArgumentResults []	StatusCode	List of <i>StatusCodes</i> corresponding to the <i>inputArguments</i>. This list is empty unless the operation level result is <i>Bad_InvalidArgument</i>. If this list is populated, it has the same length as the <i>inputArguments</i> list.
inputArgumentDiagnosticInfos []	DiagnosticInfo	List of diagnostic information corresponding to the <i>inputArguments</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.
outputArguments []	BaseDataType	List of output argument values. An empty list indicates that there are no output arguments. The size and order of this list matches the size and order of the output arguments defined by the <i>OutputArguments Property</i> of the <i>Method</i>. The name, a description and the data type of each argument are defined by the <i>Argument</i> structure in each element of the methods <i>OutputArguments Property</i>.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>statusCode</i> of the <i>results</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.11.2.3 Service results

Table 66 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 66 – Call Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.11.2.4 StatusCodes

Table 67 defines values for the *statusCode* parameter and Table 68 defines values for the *inputArgumentResults* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Server vendors or OPC UA companion specifications may reuse existing *StatusCodes* for application specific error information. This is valid as long as the canonical description of the *StatusCode* does not have a different meaning than the application specific description. To eliminate any vagueness, the *Server* should include the application specific description in the *DiagnosticInfo*.

Good *StatusCodes* with sub-status shall not be used as *statusCode* since many programming language bindings would cause such codes to throw an exception.

Table 67 – Call Operation Level result codes

Symbolic Id	Description
Bad_NodeIdInvalid	See Table 178 for the description of this result code. Used to indicate that the specified <i>Object</i> is not valid.
Bad_NodeIdUnknown	See Table 178 for the description of this result code. Used to indicate that the specified <i>Object</i> is not valid.
Bad_NotExecutable	The executable <i>Attribute</i> does not allow the execution of the <i>Method</i> .
Bad_ArgumentsMissing	The <i>Client</i> did not specify all of the input arguments for the <i>Method</i> .
Bad_TooManyArguments	The <i>Client</i> specified more input arguments than defined for the <i>Method</i> .
Bad_InvalidArgument	See Table 177 for the description of this result code. Used to indicate in the operation level results that one or more of the input arguments are invalid. The <i>inputArgumentResults</i> contain the specific status code for each invalid argument.
Bad_UserAccessDenied	See Table 177 for the description of this result code.
Bad_SecurityModelInsufficient	See Table 178 for the description of this result code.
Bad_MethodInvalid	The method id does not refer to a <i>Method</i> for the specified <i>Object</i> .
Bad_NoCommunication	See Table 178 for the description of this result code.

Table 68 – Call Input Argument Result Codes

Symbolic Id	Description
Bad_OutOfRange	See Table 178 for the description of this result code. Used to indicate that an input argument is outside the acceptable range.
Bad_TypeMismatch	See Table 178 for the description of this result code. Used to indicate that an input argument does not have the correct data type. A <i>ByteString</i> is structurally the same as a one dimensional array of <i>Byte</i> . A <i>Server</i> shall accept a <i>ByteString</i> if an array of <i>Byte</i> is expected.

5.12 MonitoredItem Service Set

5.12.1 MonitoredItem model

5.12.1.1 Overview

Clients define *MonitoredItems* to subscribe to data and *Events*. Each *MonitoredItem* identifies the item to be monitored and the *Subscription* to use to send *Notifications*. The item to be monitored may be any *Node Attribute*.

Notifications are data structures that describe the occurrence of data changes and *Events*. They are packaged into *NotificationMessages* for transfer to the *Client*. The *Subscription* periodically sends *NotificationMessages* at a user-specified publishing interval, and the cycle during which these messages are sent is called a publishing cycle.

Four primary parameters are defined for *MonitoredItems* that tell the *Server* how the item is to be sampled, evaluated and reported. These parameters are the sampling interval, the monitoring mode, the filter and the queue parameter. Figure 15 illustrates these concepts.




Figure 15 – MonitoredItem model

Attributes, other than the *Value Attribute*, are only monitored for a change in value. The filter is not used for these *Attributes*. Any change in value for these *Attributes* causes a *Notification* to be generated.

The *Value Attribute* is used when monitoring *Variables*. *Variable* values are monitored for a change in value or a change in their status. The filters defined in this document (see 7.17.2) and in IEC 62541-8 are used to determine if the value change is large enough to cause a *Notification* to be generated for the *Variable*.

Objects and *views* can be used to monitor *Events*. *Events* are only available from *Nodes* where the *SubscribeToEvents* bit of the *EventNotifier Attribute* is set. The filter defined in this document (see 7.17.3) is used to determine if an *Event* received from the *Node* is sent to the *Client*. The filter also allows selecting fields of the *EventType* that will be contained in the *Event* such as *EventId*, *EventType*, *SourceNode*, *Time* and *Description*.

IEC 62541-3 describes the *Event* model and the base *EventTypes*.

The *Properties* of the base *EventTypes* and the representation of the base *EventTypes* in the *AddressSpace* are specified in IEC 62541-5.

5.12.1.2 Sampling interval

Each *MonitoredItem* created by the *Client* is assigned a sampling interval that is either inherited from the publishing interval of the *Subscription* or that is defined specifically to override that rate. A negative number indicates that the default sampling interval defined by the publishing interval of the *Subscription* is requested. The sampling interval indicates the fastest rate at which the *Server* should sample its underlying source for data changes.

A *Client* shall define a sampling interval of 0 if it subscribes for *Events*.

The assigned sampling interval defines a "best effort" cyclic rate that the *Server* uses to sample the item from its source. "Best effort" in this context means that the *Server* does its best to sample at this rate. Sampling at rates faster than this rate is acceptable, but not necessary to meet the needs of the *Client*. How the *Server* deals with the sampling rate and how often it actually polls its data source internally is a *Server* implementation detail. However, the time between values returned to the *Client* shall be greater or equal to the sampling interval.

The *Client* may also specify 0 for the sampling interval, which indicates that the *Server* should use the fastest practical rate. It is expected that *Servers* will support only a limited set of sampling intervals to optimize their operation. If the exact interval requested by the *Client* is not supported by the *Server*, then the *Server* assigns to the *MonitoredItem* the most appropriate interval as determined by the *Server*. It returns this assigned interval to the *Client*. The *Server Capabilities Object* defined in IEC 62541-5 identifies the sampling intervals supported by the *Server*.

The *Server* may support data that is collected based on a sampling model or generated based on an exception-based model. The fastest supported sampling interval may be equal to 0, which indicates that the data item is exception-based rather than being sampled at some period. An exception-based model means that the underlying system does not require sampling and reports data changes.

The *Client* may use the revised sampling interval values as a hint for setting the publishing interval as well as the keep-alive count of a *Subscription*. If, for example, the smallest revised sampling interval of the *MonitoredItems* is 5 seconds, then the time before a keep-alive is sent should be longer than 5 seconds.

Note that, in many cases, the OPC UA *Server* provides access to a decoupled system and therefore has no knowledge of the data update logic. In this case, even though the OPC UA *Server* samples at the negotiated rate, the data might be updated by the underlying system at a much slower rate. In this case, changes can only be detected at this slower rate.

If the behaviour by which the underlying system updates the item is known, it will be available via the *MinimumSamplingInterval Attribute* defined in IEC 62541-3. If the *Server* specifies a value for the *MinimumSamplingInterval Attribute* it shall always return a *revisedSamplingInterval* that is equal or higher than the *MinimumSamplingInterval* if the *Client* subscribes to the *Value Attribute*.

Clients should also be aware that the sampling by the OPC UA *Server* and the update cycle of the underlying system are usually not synchronized. This can cause additional delays in change detection, as illustrated in Figure 16.




Figure 16 – Typical delay in change detection

5.12.1.3 Monitoring mode

The monitoring mode parameter is used to enable and disable the sampling of a *MonitoredItem*, and also to provide for independently enabling and disabling the reporting of *Notifications*. This capability allows a *MonitoredItem* to be configured to sample, sample and report, or neither. Disabling sampling does not change the values of any of the other *MonitoredItem* parameter, such as its sampling interval.

When a *MonitoredItem* is enabled (i.e. when the *MonitoringMode* is changed from *DISABLED* to *SAMPLING* or *REPORTING*) or it is created in the enabled state, the *Server* shall report the first sample as soon as possible and the time of this sample becomes the starting point for the next sampling interval.

5.12.1.4 Filter

Each time a *MonitoredItem* is sampled, the *Server* evaluates the sample using the filter defined for the *MonitoredItem*. The filter parameter defines the criteria that the *Server* uses to determine if a *Notification* should be generated for the sample. The type of filter is dependent on the type of the item that is being monitored. For example, the *DataChangeFilter* and the *AggregateFilter* are used when monitoring *Variable Values* and the *EventFilter* is used when monitoring *Events*. Sampling and evaluation, including the use of filters, are described in this document. Additional filters may be defined in other parts of IEC 62541.

5.12.1.5 Queue parameters

If the sample passes the filter criteria, a *Notification* is generated and queued for transfer by the *Subscription*. The size of the queue is defined when the *MonitoredItem* is created. When the queue is full and a new *Notification* is received, the *Server* either discards the oldest *Notification* and queues the new one, or it replaces the last value added to the queue with the new one. The *MonitoredItem* is configured for one of these discard policies when the *MonitoredItem* is created. If a *Notification* is discarded for a *DataValue* and the size of the queue is larger than one, then the *Overflow* bit (flag) in the *InfoBits* portion of the *DataValue*

statusCode is set. If *discardOldest* is TRUE, the oldest value gets deleted from the queue and the next value in the queue gets the flag set. If *discardOldest* is FALSE, the last value added to the queue gets replaced with the new value. The new value gets the flag set to indicate the lost values in the next *NotificationMessage*. Figure 17 illustrates the queue overflow handling.




Figure 17 – Queue overflow handling

If the queue size is one, the queue becomes a buffer that always contains the newest *Notification*. In this case, if the sampling interval of the *MonitoredItem* is faster than the publishing interval of the *Subscription*, the *MonitoredItem* will be over sampling and the *Client* will always receive the most up-to-date value. The discard policy is ignored if the queue size is one.

On the other hand, the *Client* may want to subscribe to a continuous stream of *Notifications* without any gaps, but does not want them reported at the sampling interval. In this case, the *MonitoredItem* would be created with a queue size large enough to hold all *Notifications* generated between two consecutive publishing cycles. Then, at each publishing cycle, the *Subscription* would send all *Notifications* queued for the *MonitoredItem* to the *Client*. The *Server* shall return *Notifications* for any particular item in the same order they are in the queue.

The *Server* may be sampling at a faster rate than the sampling interval to support other *Clients*; the *Client* should only expect values at the negotiated sampling interval. The *Server* may deliver fewer values than dictated by the sampling interval, based on the filter and implementation constraints. If a *DataChangeFilter* is configured for a *MonitoredItem*, it is always applied to the newest value in the queue compared to the current sample.

If, for example, the *AbsoluteDeadband* in the *DataChangeFilter* is "10", the queue could consist of values in the following order:

- 100
- 111
- 100
- 89
- 100

Queuing of data may result in unexpected behaviour when using a *Deadband* filter and the number of encountered changes is larger than the number of values that can be maintained. The new first value in the queue may not exceed the *Deadband* limit of the previous value sent to the *Client*.

The queue size is the maximum value supported by the *Server* when monitoring *Events*. In this case, the *Server* is responsible for the *Event* buffer. If *Events* are lost, an *Event* of the type *EventQueueOverflowEventType* is placed in the queue. This *Event* is generated when the first *Event* is discarded on a *MonitoredItem* subscribing for *Events*. It is put into the Queue of the *MonitoredItem* in addition to the size of the Queue defined for this *MonitoredItem* without discarding any other *Event*. If *discardOldest* is set to TRUE it is put at the beginning of the queue and is never discarded, otherwise at the end. An aggregating *Server* shall not pass on such an *Event*. It shall be handled like other connection error scenarios.

5.12.1.6 Triggering model

The *MonitoredItems Service* allows the addition of items that are reported only when some other item (the triggering item) triggers. This is done by creating links between the triggered items and the items to report. The monitoring mode of the items to report is set to sampling-only so that it will sample and queue *Notifications* without reporting them. Figure 18 illustrates this concept.




Figure 18 – Triggering model

The triggering mechanism is a useful feature that allows *Clients* to reduce the data volume on the wire by configuring some items to sample frequently but only report when some other *Event* happens.

The following triggering behaviours are specified.

- If the monitoring mode of the triggering item is *SAMPLING*, then it is not reported when the triggering item triggers the items to report.
- If the monitoring mode of the triggering item is *REPORTING*, then it is reported when the triggering item triggers the items to report.
- If the monitoring mode of the triggering item is *DISABLED*, then the triggering item does not trigger the items to report.
- If the monitoring mode of the item to report is *SAMPLING*, then it is reported when the triggering item triggers the items to report.
- If the monitoring mode of the item to report is *REPORTING*, this effectively causes the triggering item to be ignored. All notifications of the items to report are sent after the publishing interval expires.
- If the monitoring mode of the item to report is *DISABLED*, then there will be no sampling of the item to report and therefore no notifications to report.
- The first trigger shall occur when the first notification is queued for the triggering item after the creation of the link.

Clients create and delete triggering links between a triggering item and a set of items to report. If the *MonitoredItem* that represents an item to report is deleted before its associated triggering link is deleted, the triggering link is also deleted, but the triggering item is otherwise unaffected.

Deletion of a *MonitoredItem* should not be confused with the removal of the *Attribute* that it monitors. If the *Node* that contains the *Attribute* being monitored is deleted, the *MonitoredItem* generates a *Notification* with a *StatusCode* *Bad_NodeIdUnknown* that indicates the deletion, but the *MonitoredItem* is not deleted.

5.12.2 CreateMonitoredItems

5.12.2.1 Description

This *Service* is used to create and add one or more *MonitoredItems* to a *Subscription*. A *MonitoredItem* is deleted automatically by the *Server* when the *Subscription* is deleted. Deleting a *MonitoredItem* causes its entire set of triggered item links to be deleted, but has no effect on the *MonitoredItems* referenced by the triggered items.

Calling the *CreateMonitoredItems Service* repetitively to add a small number of *MonitoredItems* each time may adversely affect the performance of the *Server*. Instead, *Clients* should add a complete set of *MonitoredItems* to a *Subscription* whenever possible.

When a *MonitoredItem* is added, the *Server* performs initialization processing for it. The initialization processing is defined by the *Notification* type of the item being monitored. *Notification* types are specified in this document and in the Access Type Specification parts of IEC 62541, such as IEC 62541-8. See IEC TR 62541-1 for a description of the Access Type Parts.

When a user adds a monitored item that the user is denied read access to, the add operation for the item shall succeed and the bad status *Bad_NotReadable* or *Bad_UserAccessDenied* shall be returned in the *Publish* response. This is the same behaviour for the case where the access rights are changed after the call to *CreateMonitoredItems*. If the access rights change to read rights, the *Server* shall start sending data for the *MonitoredItem*. The same procedure shall be applied for an *IndexRange* that does not deliver data for the current value but could deliver data in the future.

Monitored *Nodes* can be removed from the *AddressSpace* after the creation of a *MonitoredItem*. This does not affect the validity of the *MonitoredItem* but a *Bad_NodeIdUnknown* shall be returned in the *Publish* response. It is possible that the *MonitoredItem* becomes valid again if the *Node* is added again to the *AddressSpace* and the *MonitoredItem* still exists.

If a *NodeId* is known to be valid by a *Server* but the corresponding *Node Attributes* are currently not available, the *Server* may allow the creation of a *MonitoredItem* and return an appropriate Bad *StatusCode* in the *Publish* response.

The return diagnostic info setting in the request header of the *CreateMonitoredItems* or the last *ModifyMonitoredItems Service* is applied to the *Monitored Items* and is used as the diagnostic information settings when sending *Notifications* in the *Publish* response.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

It is strongly recommended by OPC UA that a *Client* reuses a *Subscription* after a short network interruption by activating the existing *Session* on a new *SecureChannel* as described in 6.7. If a *Client* called *CreateMonitoredItems* during the network interruption and the call succeeded in the *Server* but did not return to the *Client*, then the *Client* does not know if the call succeeded. The *Client* may receive data changes for these monitored items but is not

able to remove them since it does not know the *Server* handle for each monitored item. There is also no way for the *Client* to detect if the create succeeded. To delete and recreate the *Subscription* is also not an option since there may be several monitored items operating normally that should not be interrupted. To resolve this situation, the *Server Object* provides a *Method GetMonitoredItems* that returns the list of server and client handles for the monitored items in a *Subscription*. This *Method* is defined in IEC 62541-5. The *Server* shall verify that the *Method* is called within the *Session* context of the *Session* that owns the *Subscription*.

5.12.2.2 Parameters

Table 69 defines the parameters for the *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 69 – CreateMonitoredItems Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> that will report <i>Notifications</i> for this <i>MonitoredItem</i> (see 7.14 for <i>IntegerId</i> definition).
timestampsToReturn	Enum Timestamps ToReturn	An enumeration that specifies the timestamp <i>Attributes</i> to be transmitted for each <i>MonitoredItem</i> . The <i>TimestampsToReturn</i> enumeration is defined in 7.35. When monitoring <i>Events</i> , this applies only to <i>Event</i> fields that are of type <i>DataValue</i> .
itemsToCreate []	MonitoredItem CreateRequest	A list of <i>MonitoredItems</i> to be created and assigned to the specified <i>Subscription</i> . This structure is defined in-line with the following indented items.
itemToMonitor	ReadValueId	Identifies an item in the <i>AddressSpace</i> to monitor. To monitor for <i>Events</i> , the <i>attributeId</i> element of the <i>ReadValueId</i> structure is the id of the <i>EventNotifierAttribute</i> . The <i>ReadValueId</i> type is defined in 7.24.
monitoringMode	Enum MonitoringMode	The monitoring mode to be set for the <i>MonitoredItem</i> . The <i>MonitoringMode</i> enumeration is defined in 7.18.
requestedParameters	Monitoring Parameters	The requested monitoring parameters. <i>Servers</i> negotiate the values of these parameters based on the <i>Subscription</i> and the capabilities of the <i>Server</i> . The <i>MonitoringParameters</i> type is defined in 7.16.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	MonitoredItem CreateResult	List of results for the <i>MonitoredItems</i> to create. The size and order of the list matches the size and order of the <i>itemsToCreate</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	Status Code	<i>Status Code</i> for the <i>MonitoredItem</i> to create (see 7.34 for <i>Status Code</i> definition).
monitoredItemId	IntegerId	<i>Server</i> -assigned id for the <i>MonitoredItem</i> (see 7.14 for <i>IntegerId</i> definition). This id is unique within the <i>Subscription</i> , but might not be unique within the <i>Server</i> or <i>Session</i> . This parameter is present only if the <i>statusCode</i> indicates that the <i>MonitoredItem</i> was successfully created.
revisedSamplingInterval	Duration	The actual sampling interval that the <i>Server</i> will use. This value is based on a number of factors, including capabilities of the underlying system. The <i>Server</i> shall always return a <i>revisedSamplingInterval</i> that is equal or higher than the requested <i>samplingInterval</i> . If the requested <i>samplingInterval</i> is higher than the maximum sampling interval supported by the <i>Server</i> , the maximum sampling interval is returned.
revisedQueueSize	Counter	The actual queue size that the <i>Server</i> will use.
filterResult	Extensible Parameter MonitoringFilterResult	Contains any revised parameter values or error results associated with the <i>MonitoringFilter</i> specified in <i>requestedParameters</i> . This parameter may be null if no errors occurred. The <i>MonitoringFilterResult</i> parameter type is an extensible parameter type specified in 7.17.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to create (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>itemsToCreate</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.2.3 Service results

Table 70 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 70 – CreateMonitoredItems Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.12.2.4 StatusCodes

Table 71 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 71 – CreateMonitoredItems Operation Level result codes

Symbolic Id	Description
Bad_MonitoringModelInvalid	See Table 178 for the description of this result code.
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_IndexRangeNoData	See Table 178 for the description of this result code. If the <i>ArrayDimensions</i> have a fixed length that cannot change and no data exists within the range of indexes specified, <i>Bad_IndexRangeNoData</i> is returned in <i>CreateMonitoredItems</i> . Otherwise if the length of the array is dynamic, the <i>Server</i> shall return this status in a <i>Publish</i> response for the <i>MonitoredItem</i> if no data exists within the range.
Bad_DataEncodingInvalid	See Table 178 for the description of this result code.
Bad_DataEncodingUnsupported	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterInvalid	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterUnsupported	See Table 178 for the description of this result code.
Bad_FilterNotAllowed	See Table 177 for the description of this result code.
Bad_TooManyMonitoredItems	The <i>Server</i> has reached its maximum number of monitored items.

5.12.3 ModifyMonitoredItems

5.12.3.1 Description

This *Service* is used to modify *MonitoredItems* of a *Subscription*. Changes to the *MonitoredItem* settings shall be applied immediately by the *Server*. They take effect as soon as practical but not later than twice the new *revisedSamplingInterval*.

The return diagnostic info setting in the request header of the *CreateMonitoredItems* or the last *ModifyMonitoredItems Service* is applied to the *Monitored Items* and is used as the diagnostic information settings when sending Notifications in the *Publish* response.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

5.12.3.2 Parameters

Table 72 defines the parameters for the *Service*.

Table 72 – ModifyMonitoredItems Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> used to qualify the <i>monitoredItemId</i> (see 7.14 for <i>IntegerId</i> definition).
timestampsToReturn	Enum Timestamps ToReturn	An enumeration that specifies the timestamp <i>Attributes</i> to be transmitted for each <i>MonitoredItem</i> to be modified. The <i>TimestampsToReturn</i> enumeration is defined in 7.35. When monitoring <i>Events</i> , this applies only to <i>Event</i> fields that are of type <i>DataValue</i> .
itemsToModify []	MonitoredItemModifyRequest	The list of <i>MonitoredItems</i> to modify. This structure is defined in-line with the following indented items.
monitoredItemId	IntegerId	<i>Server</i> -assigned id for the <i>MonitoredItem</i> .
requestedParameters	Monitoring Parameters	The requested values for the monitoring parameters. The <i>MonitoringParameters</i> type is defined in 7.16. If the number of notifications in the queue exceeds the new queue size, the notifications exceeding the size shall be discarded following the configured discard policy.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	MonitoredItemModifyResult	List of results for the <i>MonitoredItems</i> to modify. The size and order of the list matches the size and order of the <i>itemsToModify</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for the <i>MonitoredItem</i> to be modified (see 7.34 for <i>StatusCode</i> definition).
revisedSampling Interval	Duration	The actual sampling interval that the <i>Server</i> will use. The <i>Server</i> returns the value it will actually use for the sampling interval. This value is based on a number of factors, including capabilities of the underlying system. The <i>Server</i> shall always return a <i>revisedSamplingInterval</i> that is equal or higher than the requested <i>samplingInterval</i> . If the requested <i>samplingInterval</i> is higher than the maximum sampling interval supported by the <i>Server</i> , the maximum sampling interval is returned.
revisedQueueSize	Counter	The actual queue size that the <i>Server</i> will use.
filterResult	Extensible Parameter MonitoringFilter Result	Contains any revised parameter values or error results associated with the <i>MonitoringFilter</i> specified in the request. This parameter may be null if no errors occurred. The <i>MonitoringFilterResult</i> parameter type is an extensible parameter type specified in 7.17.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to modify (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>itemsToModify</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.3.3 Service results

Table 73 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 73 – ModifyMonitoredItems Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_TimestampsToReturnInvalid	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.12.3.4 StatusCodes

Table 74 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 74 – ModifyMonitoredItems Operation Level result codes

Symbolic Id	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterInvalid	See Table 178 for the description of this result code.
Bad_MonitoredItemFilterUnsupported	See Table 178 for the description of this result code.
Bad_FilterNotAllowed	See Table 177 for the description of this result code.

5.12.4 SetMonitoringMode

5.12.4.1 Description

This *Service* is used to set the monitoring mode for one or more *MonitoredItems* of a *Subscription*. Setting the mode to DISABLED causes all queued *Notifications* to be deleted.

5.12.4.2 Parameters

Table 75 defines the parameters for the *Service*.

Table 75 – SetMonitoringMode service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> used to qualify the <i>monitoredItemIds</i> (see 7.14 for <i>IntegerId</i> definition).
monitoringMode	Enum MonitoringMode	The monitoring mode to be set for the <i>MonitoredItems</i> . The <i>MonitoringMode</i> enumeration is defined in 7.18.
monitoredItemIds []	IntegerId	List of <i>Server</i> -assigned ids for the <i>MonitoredItems</i> .
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of <i>StatusCodes</i> for the <i>MonitoredItems</i> to enable/disable (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to enable/disable (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.4.3 Service results

Table 76 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 76 – SetMonitoringMode Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_MonitoringModeInvalid	See Table 178 for the description of this result code.

5.12.4.4 StatusCodes

Table 77 defines values for the operation level *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 77 – SetMonitoringMode Operation Level result codes

Value	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.12.5 SetTriggering

5.12.5.1 Description

This *Service* is used to create and delete triggering links for a triggering item. The triggering item and the items to report shall belong to the same *Subscription*.

Each triggering link links a triggering item to an item to report. Each link is represented by the *MonitoredItem* id for the item to report. An error code is returned if this id is invalid.

See 5.12.1.6 for a description of the triggering model.

5.12.5.2 Parameters

Table 78 defines the parameters for the *Service*.

Table 78 – SetTriggering Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> that contains the triggering item and the items to report (see 7.14 for <i>IntegerId</i> definition).
triggeringItemId	IntegerId	<i>Server</i> -assigned id for the <i>MonitoredItem</i> used as the triggering item.
linksToAdd []	IntegerId	The list of <i>Server</i> -assigned ids of the items to report that are to be added as triggering links. The list of <i>linksToRemove</i> is processed before the <i>linksToAdd</i> .
linksToRemove []	IntegerId	The list of <i>Server</i> -assigned ids of the items to report for the triggering links to be deleted. The list of <i>linksToRemove</i> is processed before the <i>linksToAdd</i> .
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
addResults []	Status Code	List of <i>StatusCodes</i> for the items to add (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>linksToAdd</i> parameter specified in the request.
addDiagnosticInfos []	Diagnostic Info	List of diagnostic information for the links to add (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>linksToAdd</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.
removeResults []	Status Code	List of <i>StatusCodes</i> for the items to delete. The size and order of the list matches the size and order of the <i>linksToRemove</i> parameter specified in the request.
removeDiagnosticInfos []	Diagnostic Info	List of diagnostic information for the links to delete. The size and order of the list matches the size and order of the <i>linksToRemove</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.5.3 Service results

Table 79 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in 7.34.

Table 79 – SetTriggering Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.12.5.4 StatusCodes

Table 80 defines values for the results parameters that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 80 – SetTriggering Operation Level result codes

Symbolic Id	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.12.6 DeleteMonitoredItems

5.12.6.1 Description

This *Service* is used to remove one or more *MonitoredItems* of a *Subscription*. When a *MonitoredItem* is deleted, its triggered item links are also deleted.

Successful removal of a *MonitoredItem*, however, might not remove *Notifications* for the *MonitoredItem* that are in the process of being sent by the *Subscription*. Therefore, *Clients* may receive *Notifications* for the *MonitoredItem* after they have received a positive response that the *MonitoredItem* has been deleted.

5.12.6.2 Parameters

Table 81 defines the parameters for the *Service*.

Table 81 – DeleteMonitoredItems Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> that contains the <i>MonitoredItems</i> to be deleted (see 7.14 for <i>IntegerId</i> definition).
monitoredItemIds []	IntegerId	List of <i>Server</i> -assigned ids for the <i>MonitoredItems</i> to be deleted.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCode	List of <i>StatusCodes</i> for the <i>MonitoredItems</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>MonitoredItems</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>monitoredItemIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.12.6.3 Service results

Table 82 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 82 – DeleteMonitoredItems Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.12.6.4 StatusCodes

Table 83 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 83 – DeleteMonitoredItems Operation Level result codes

Symbolic Id	Description
Bad_MonitoredItemIdInvalid	See Table 178 for the description of this result code.

5.13 Subscription Service Set

5.13.1 Subscription model

5.13.1.1 Description

Subscriptions are used to report *Notifications* to the *Client*. Their general behaviour is summarized below. Their precise behaviour is described in 5.13.1.2.

- Subscriptions* have a set of *MonitoredItems* assigned to them by the *Client*. *MonitoredItems* generate *Notifications* that are to be reported to the *Client* by the *Subscription* (see 5.12.1 for a description of *MonitoredItems*).
- Subscriptions* have a publishing interval. The publishing interval of a *Subscription* defines the cyclic rate at which the *Subscription* executes. Each time it executes, it attempts to send a *NotificationMessage* to the *Client*. *NotificationMessages* contain *Notifications* that have not yet been reported to *Client*.
- NotificationMessages* are sent to the *Client* in response to *Publish* requests. *Publish* requests are normally queued to the *Session* as they are received, and one is de-queued and processed by a subscription related to this *Session* for each publishing cycle, if there are *Notifications* to report. When there are not, the *Publish* request is not de-queued from the *Session*, and the *Server* waits until the next cycle and checks again for *Notifications*.
- At the beginning of a cycle, if there are *Notifications* to send but there are no *Publish* requests queued, the *Server* enters a wait state for a *Publish* request to be received. When one is received, it is processed immediately without waiting for the next publishing cycle.
- NotificationMessages* are uniquely identified by sequence numbers that enable *Clients* to detect missed *Messages*. The publishing interval also defines the default sampling interval for its *MonitoredItems*.
- Subscriptions* have a keep-alive counter that counts the number of consecutive publishing cycles in which there have been no *Notifications* to report to the *Client*. When the maximum keep-alive count is reached, a *Publish* request is de-queued and used to return a keep-alive *Message*. This keep-alive *Message* informs the *Client* that the *Subscription* is still active. Each keep-alive *Message* is a response to a *Publish* request in which the *notificationMessage* parameter does not contain any *Notifications* and that contains the sequence number of the next *NotificationMessage* that is to be sent. In the clauses that follow, the term *NotificationMessage* refers to a response to a *Publish* request in which the *notificationMessage* parameter actually contains one or more *Notifications*, as opposed to a keep-alive *Message* in which this parameter contains no *Notifications*. The maximum

keep-alive count is set by the *Client* during *Subscription* creation and may be subsequently modified using the *ModifySubscription Service*. Similar to *Notification* processing described in (c) above, if there are no *Publish* requests queued, the *Server* waits for the next one to be received and sends the keep-alive immediately without waiting for the next publishing cycle.

- g) Publishing by a *Subscription* may be enabled or disabled by the *Client* when created, or subsequently using the *SetPublishingMode Service*. Disabling causes the *Subscription* to cease sending *NotificationMessages* to the *Client*. However, the *Subscription* continues to execute cyclically and continues to send keep-alive *Messages* to the *Client*.
- h) *Subscriptions* have a lifetime counter that counts the number of consecutive publishing cycles in which there have been no *Publish* requests available to send a *Publish* response for the *Subscription*. Any *Service* call that uses the *SubscriptionId* or the processing of a *Publish* response resets the lifetime counter of this *Subscription*. When this counter reaches the value calculated for the lifetime of a *Subscription* based on the *MaxKeepAliveCount* parameter in the *CreateSubscription Service* (5.13.2), the *Subscription* is closed. Closing the *Subscription* causes its *MonitoredItems* to be deleted. In addition the *Server* shall issue a *StatusChangeNotification notificationMessage* with the status code *Bad_Timeout*. The *StatusChangeNotification notificationMessage* type is defined in 7.20.4.
- i) *Sessions* maintain a retransmission queue of sent *NotificationMessages*. *NotificationMessages* are retained in this queue until they are acknowledged. The *Session* shall maintain a retransmission queue size of at least two times the number of *Publish* requests per *Session* the *Server* supports. The minimum size of the retransmission queue may be changed by a *Profile* in IEC 62541-7. The minimum number of *Publish* requests per *Session* the *Server* shall support is defined in IEC 62541-7. *Clients* are required to acknowledge *NotificationMessages* as they are received. In the case of a retransmission queue overflow, the oldest sent *NotificationMessage* gets deleted. If a *Subscription* is transferred to another *Session*, the queued *NotificationMessages* for this *Subscription* are moved from the old to the new *Session*.

The sequence number is an unsigned 32-bit integer that is incremented by one for each *NotificationMessage* sent. The value 0 is never used for the sequence number. The first *NotificationMessage* sent on a *Subscription* has a sequence number of 1. If the sequence number rolls over, it rolls over to 1.

When a *Subscription* is created, the first *Message* is sent at the end of the first publishing cycle to inform the *Client* that the *Subscription* is operational. A *NotificationMessage* is sent if there are *Notifications* ready to be reported. If there are none, a keep-alive *Message* is sent instead that contains a sequence number of 1, indicating that the first *NotificationMessage* has not yet been sent. This is the only time a keep-alive *Message* is sent without waiting for the maximum keep-alive count to be reached, as specified in (f) above.

A *Client* shall be prepared for receiving *Publish* responses for a *Subscription* more frequently than the corresponding publishing interval. One example is the situation where the number of available notifications exceeds the *Subscription* setting *maxNotificationsPerPublish*. A *Client* is always able to control the timing of the *Publish* responses by not queueing *Publish* requests. If a *Client* does not queue *Publish* requests in the *Server*, the *Server* can only send a *Publish* response if it receives a new *Publish* request. This would increase latency for delivery of notifications but allows a *Client* to throttle the number of received *Publish* responses in high load situations.

The value of the sequence number is never reset during the lifetime of a *Subscription*. Therefore, the same sequence number shall not be reused on a *Subscription* until over four billion *NotificationMessages* have been sent. At a continuous rate of one thousand *NotificationMessages* per second on a given *Subscription*, it would take roughly fifty days for the same sequence number to be reused. This allows *Clients* to safely treat sequence numbers as unique.

Sequence numbers are also used by *Clients* to acknowledge the receipt of *NotificationMessages*. *Publish* requests allow the *Client* to acknowledge all *Notifications* up to a specific sequence number and to acknowledge the sequence number of the last *NotificationMessage* received. One or more gaps may exist in between. Acknowledgements allow the *Server* to delete *NotificationMessages* from its retransmission queue.

Clients may ask for retransmission of selected *NotificationMessages* using the *Republish Service*. This *Service* returns the requested *Message*.

Subscriptions are designed to work independent of the actual communication connection between OPC UA *Client* and *Server* and independent of a *Session*. Short communication interruptions can be handled without losing data or events. To make sure that longer communication interruptions or planned disconnects can be handled without losing data or events, an OPC UA *Server* may support durable *Subscriptions*. If this feature is supported, the *Server* accepts a high *Subscription RequestedLifetimeCount* and large *MonitoredItem QueueSize* parameter settings. Subclause 6.8 describes how durable *Subscriptions* can be created and used.

5.13.1.2 State table

The state table formally describes the operation of the *Subscription*. The following model of operations is described by this state table. This description applies when publishing is enabled or disabled for the *Subscription*.

After creation of the *Subscription*, the *Server* starts the publishing timer and restarts it whenever it expires. If the timer expires the number of times defined for the *Subscription* lifetime without having received a *Subscription Service* request from the *Client*, the *Subscription* assumes that the *Client* is no longer present, and terminates.

Clients send *Publish* requests to *Servers* to receive *Notifications*. *Publish* requests are not directed to any one *Subscription* and, therefore, may be used by any *Subscription*. Each contains acknowledgements for one or more *Subscriptions*. These acknowledgements are processed when the *Publish* request is received. The *Server* then queues the request in a queue shared by all *Subscriptions*, except in the following cases.

- a) The previous *Publish* response indicated that there were still more *Notifications* ready to be transferred and there were no more *Publish* requests queued to transfer them.
- b) The publishing timer of a *Subscription* expired and there were either *Notifications* to be sent or a keep-alive *Message* to be sent.

In these cases, the newly received *Publish* request is processed immediately by the first *Subscription* to encounter either case (a) or case (b).

Each time the publishing timer expires, it is immediately reset. If there are *Notifications* or a keep-alive *Message* to be sent, it de-queues and processes a *Publish* request. When a *Subscription* processes a *Publish* request, it accesses the queues of its *MonitoredItems* and de-queues its *Notifications*, if any. It returns these *Notifications* in the response, setting the *moreNotifications* flag if it was not able to return all available *Notifications* in the response.

If there were *Notifications* or a keep-alive *Message* to be sent but there were no *Publish* requests queued, the *Subscription* assumes that the *Publish* request is late and waits for the next *Publish* request to be received, as described in case (b).

If the *Subscription* is disabled when the publishing timer expires or if there are no *Notifications* available, it enters the keep-alive state and sets the keep-alive counter to its maximum value as defined for the *Subscription*.

While in the keep-alive state, it checks for *Notifications* each time the publishing timer expires. If one or more *Notifications* have been generated, a *Publish* request is de-queued

and a *NotificationMessage* is returned in the response. However, if the publishing timer expires without a *Notification* becoming available, a *Publish* request is de-queued and a keep-alive *Message* is returned in the response. The *Subscription* then returns to the normal state of waiting for the publishing timer to expire again. If, in either of these cases, there are no *Publish* requests queued, the *Subscription* waits for the next *Publish* request to be received, as described in case (b).

The *Subscription* states are defined in Table 84.

Table 84 – Subscription states

State	Description
CLOSED	The <i>Subscription</i> has not yet been created or has terminated.
CREATING	The <i>Subscription</i> is being created.
NORMAL	The <i>Subscription</i> is cyclically checking for <i>Notifications</i> from its <i>MonitoredItems</i> . The keep-alive counter is not used in this state.
LATE	The publishing timer has expired and there are <i>Notifications</i> available or a keep-alive <i>Message</i> is ready to be sent, but there are no <i>Publish</i> requests queued. When in this state, the next <i>Publish</i> request is processed when it is received. The keep-alive counter is not used in this state.
KEEPAIVE	The <i>Subscription</i> is cyclically checking for <i>Notifications</i> from its <i>MonitoredItems</i> or for the keep-alive counter to count down to 0 from its maximum.

The state table is described in Table 85. The following rules and conventions apply.

- Events* represent the receipt of *Service* requests and the occurrence internal *Events*, such as timer expirations.
- Service* requests *Events* may be accompanied by conditions that test *Service* parameter values. Parameter names begin with a lower case letter.
- Internal *Events* may be accompanied by conditions that test state *Variable* values. State *Variables* are defined in 5.13.1.3. They begin with an upper case letter.
- Service* request and internal *Events* may be accompanied by conditions represented by functions whose return value is tested. Functions are identified by "()" after their name. They are described in 5.13.1.4.
- When an *Event* is received, the first transition for the current state is located and the transitions are searched sequentially for the first transition that meets the *Event* or conditions criteria. If none are found, the *Event* is ignored.
- Actions are described by functions and state *Variable* manipulations.
- The LifetimeTimerExpires *Event* is triggered when its corresponding counter reaches zero.

Table 85 – Subscription state table

#	Current State	Event/Conditions	Action	Next State
1	CLOSED	Receive CreateSubscription Request	CreateSubscription()	CREATING
2	CREATING	CreateSubscription fails	ReturnNegativeResponse()	CLOSED
3	CREATING	CreateSubscription succeeds	InitializeSubscription() MessageSent = FALSE ReturnResponse()	NORMAL
4	NORMAL	Receive <i>Publish</i> Request && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && MoreNotifications == FALSE))	DeleteAkedNotificationMsgs() EnqueuePublishingReq()	NORMAL
5	NORMAL	Receive <i>Publish</i> Request && PublishingEnabled == TRUE && MoreNotifications == TRUE	ResetLifetimeCounter() DeleteAkedNotificationMsgs() ReturnNotifications() MessageSent = TRUE	NORMAL
6	NORMAL	PublishingTimer Expires && PublishingReqQueued == TRUE && PublishingEnabled == TRUE && NotificationsAvailable == TRUE	ResetLifetimeCounter() StartPublishingTimer() DequeuePublishReq() ReturnNotifications() MessageSent == TRUE	NORMAL
7	NORMAL	PublishingTimer Expires && PublishingReqQueued == TRUE && MessageSent == FALSE && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	ResetLifetimeCounter() StartPublishingTimer() DequeuePublishReq() ReturnKeepAlive() MessageSent == TRUE	NORMAL
8	NORMAL	PublishingTimer Expires && PublishingReqQueued == FALSE && (MessageSent == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == TRUE))	StartPublishingTimer()	LATE

#	Current State	Event/Conditions	Action	Next State
9	NORMAL	PublishingTimer Expires && MessageSent == TRUE && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	StartPublishingTimer() ResetKeepAliveCounter()	KEEPALIVE
10	LATE	Receive <i>Publish</i> Request && PublishingEnabled == TRUE && (NotificationsAvailable == TRUE MoreNotifications == TRUE)	ResetLifetimeCounter() DeleteAkedNotificationMsgs() ReturnNotifications() MessageSent = TRUE	NORMAL
11	LATE	Receive <i>Publish</i> Request && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE && MoreNotifications == FALSE))	ResetLifetimeCounter() DeleteAkedNotificationMsgs() ReturnKeepAlive() MessageSent = TRUE	KEEPALIVE
12	LATE	PublishingTimer Expires	StartPublishingTimer()	LATE
13	KEEPALIVE	Receive <i>Publish</i> Request	DeleteAkedNotificationMsgs() EnqueuePublishingReq()	KEEPALIVE
14	KEEPALIVE	PublishingTimer Expires && PublishingEnabled == TRUE && NotificationsAvailable == TRUE && PublishingReqQueued == TRUE	ResetLifetimeCounter() StartPublishingTimer() DequeuePublishReq() ReturnNotifications() MessageSent == TRUE	NORMAL
15	KEEPALIVE	PublishingTimer Expires && PublishingReqQueued == TRUE && KeepAliveCounter == 1 && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	StartPublishingTimer() DequeuePublishReq() ReturnKeepAlive() ResetKeepAliveCounter()	KEEPALIVE

#	Current State	Event/Conditions	Action	Next State
16	KEEPALIVE	PublishingTimer Expires && KeepAliveCounter > 1 && (PublishingEnabled == FALSE (PublishingEnabled == TRUE && NotificationsAvailable == FALSE))	StartPublishingTimer() KeepAliveCounter--	KEEPALIVE
17	KEEPALIVE	PublishingTimer Expires && PublishingReqQueued == FALSE && (KeepAliveCounter == 1 (KeepAliveCounter > 1 && PublishingEnabled == TRUE && NotificationsAvailable == TRUE))	StartPublishingTimer()	LATE
18	NORMAL LATE KEEPALIVE	Receive ModifySubscription Request	ResetLifetimeCounter() UpdateSubscriptionParams() ReturnResponse()	SAME
19	NORMAL LATE KEEPALIVE	Receive SetPublishingMode Request	ResetLifetimeCounter() SetPublishingEnabled() MoreNotifications = FALSE ReturnResponse()	SAME
20	NORMAL LATE KEEPALIVE	Receive Republish Request && RequestedMessageFound == TRUE	ResetLifetimeCounter() ReturnResponse()	SAME
21	NORMAL LATE KEEPALIVE	Receive Republish Request && RequestedMessageFound == FALSE	ResetLifetimeCounter() ReturnNegativeResponse()	SAME
22	NORMAL LATE KEEPALIVE	Receive TransferSubscriptions Request && SessionChanged() == FALSE	ResetLifetimeCounter() ReturnNegativeResponse ()	SAME
23	NORMAL LATE KEEPALIVE	Receive TransferSubscriptions Request && SessionChanged() == TRUE && ClientValidated() ==TRUE	SetSession() ResetLifetimeCounter() ReturnResponse() IssueStatusChangeNotification()	SAME
24	NORMAL LATE KEEPALIVE	Receive TransferSubscriptions Request && SessionChanged() == TRUE && ClientValidated() == FALSE	ReturnNegativeResponse()	SAME
25	NORMAL LATE KEEPALIVE	Receive DeleteSubscriptions Request && SubscriptionAssignedToClient ==TRUE	DeleteMonitoredItems() DeleteClientPublReqQueue()	CLOSED

#	Current State	Event/Conditions	Action	Next State
26	NORMAL LATE KEEPALIVE	Receive DeleteSubscriptions Request && SubscriptionAssignedToClient ==FALSE	ResetLifetimeCounter() ReturnNegativeResponse()	SAME
27	NORMAL LATE KEEPALIVE	LifetimeCounter == 1 The LifetimeCounter is decremented if PublishingTimer expires and PublishingReqQueued == FALSE The LifetimeCounter is reset if PublishingReqQueued == TRUE.	DeleteMonitoredItems() IssueStatusChangeNotification()	CLOSED

5.13.1.3 State variables and parameters

The state variables are defined alphabetically in Table 86.

Table 86 – State variables and parameters

State Variable	Description
MoreNotifications	A boolean value that is set to TRUE only by the CreateNotificationMsg() when there were too many <i>Notifications</i> for a single <i>NotificationMessage</i> .
LatePublishRequest	A boolean value that is set to TRUE to reflect that, the last time the publishing timer expired, there were no <i>Publish</i> requests queued.
LifetimeCounter	A value that contains the number of consecutive publishing timer expirations without <i>Client</i> activity before the <i>Subscription</i> is terminated.
MessageSent	A boolean value that is set to TRUE to mean that either a <i>NotificationMessage</i> or a keep-alive <i>Message</i> has been sent on the <i>Subscription</i> . It is a flag that is used to ensure that either a <i>NotificationMessage</i> or a keep-alive <i>Message</i> is sent out the first time the publishing timer expires.
NotificationsAvailable	A boolean value that is set to TRUE only when there is at least one <i>MonitoredItem</i> that is in the reporting mode and that has a <i>Notification</i> queued or there is at least one item to report whose triggering item has triggered and that has a <i>Notification</i> queued. The transition of this state <i>Variable</i> from FALSE to TRUE creates the "New <i>Notification</i> Queued" <i>Event</i> in the state table.
PublishingEnabled	The parameter that requests publishing to be enabled or disabled.
PublishingReqQueued	A boolean value that is set to TRUE only when there is a <i>Publish</i> request <i>Message</i> queued to the <i>Subscription</i> .
RequestedMessageFound	A boolean value that is set to TRUE only when the <i>Message</i> requested to be retransmitted was found in the retransmission queue.
SeqNum	The value that records the value of the sequence number used in <i>NotificationMessages</i> .
SubscriptionAssignedToClient	A boolean value that is set to TRUE only when the <i>Subscription</i> requested to be deleted is assigned to the <i>Client</i> that issued the request. A <i>Subscription</i> is assigned to the <i>Client</i> that created it. That assignment can only be changed through successful completion of the <i>TransferSubscriptions Service</i> .

5.13.1.4 Functions

The action functions are defined alphabetically in Table 87.

Table 87 – Functions

Function	Description
ClientValidated()	A boolean function that returns TRUE only when the <i>Client</i> that is submitting a <i>TransferSubscriptions</i> request is operating on behalf of the same user and supports the same <i>Profiles</i> as the <i>Client</i> of the previous <i>Session</i> .
CreateNotificationMsg()	Increment the SeqNum and create a <i>NotificationMessage</i> from the <i>MonitoredItems</i> assigned to the <i>Subscription</i> . Save the newly-created <i>NotificationMessage</i> in the retransmission queue. If all available <i>Notifications</i> can be sent in the <i>Publish</i> response, the <i>MoreNotifications</i> state <i>Variable</i> is set to FALSE. Otherwise, it is set to TRUE.
CreateSubscription()	Attempt to create the <i>Subscription</i> .
DeleteAkedNotificationMsgs()	Delete the <i>NotificationMessages</i> from the retransmission queue that were acknowledged by the request.
DeleteClientPublReqQueue()	Clear the <i>Publish</i> request queue for the <i>Client</i> that is sending the <i>DeleteSubscriptions</i> request, if there are no more <i>Subscriptions</i> assigned to that <i>Client</i> .
DeleteMonitoredItems()	Delete all <i>MonitoredItems</i> assigned to the <i>Subscription</i> .
DequeuePublishReq()	De-queue a publishing request in first-in first-out order. Validate if the publish request is still valid by checking the <i>timeoutHint</i> in the <i>RequestHeader</i> . If the request timed out, send a <i>Bad_Timeout</i> service result for the request and de-queue another publish request. ResetLifetimeCounter()
EnqueuePublishingReq()	Enqueue the publishing request.
InitializeSubscription()	ResetLifetimeCounter() MoreNotifications = FALSE PublishRateChange = FALSE PublishingEnabled = value of publishingEnabled parameter in the <i>CreateSubscription</i> request PublishingReqQueued = FALSE SeqNum = 0 SetSession() StartPublishingTimer()
IssueStatusChangeNotification()	Issue a <i>StatusChangeNotification notificationMessage</i> with a status code for the status change of the <i>Subscription</i> . The <i>StatusChangeNotification notificationMessage</i> type is defined in 7.20.4. <i>Bad_Timeout</i> status code is used if the lifetime expires and <i>Good_SubscriptionTransferred</i> is used if the <i>Subscriptions</i> was transferred to another <i>Session</i> .
ResetKeepAliveCounter()	Reset the keep-alive counter to the maximum keep-alive count of the <i>Subscription</i> . The maximum keep-alive count is set by the <i>Client</i> when the <i>Subscription</i> is created and may be modified using the <i>ModifySubscription Service</i> .
ResetLifetimeCounter()	Reset the <i>LifetimeCounter Variable</i> to the value specified for the lifetime of a <i>Subscription</i> in the <i>CreateSubscription Service</i> (5.13.2).
ReturnKeepAlive()	CreateKeepAliveMsg() ReturnResponse()
ReturnNegativeResponse ()	Return a <i>Service</i> response indicating the appropriate <i>Service</i> level error. No parameters are returned other than the responseHeader that contains the <i>Service</i> level <i>StatusCode</i> .
ReturnNotifications()	CreateNotificationMsg() ReturnResponse() If (MoreNotifications == TRUE) && (PublishingReqQueued == TRUE) { DequeuePublishReq() Loop through this function again }

Function	Description
ReturnResponse()	Return the appropriate response, setting the appropriate parameter values and <i>StatusCodes</i> defined for the <i>Service</i> .
SessionChanged()	A boolean function that returns TRUE only when the <i>Session</i> used to send a <i>TransferSubscriptions</i> request is different from the <i>Client Session</i> currently associated with the <i>Subscription</i> .
SetPublishingEnabled ()	Set the <i>PublishingEnabled</i> state <i>Variable</i> to the value of the <i>publishingEnabled</i> parameter received in the request.
SetSession	Set the <i>Session</i> information for the <i>Subscription</i> to match the <i>Session</i> on which the <i>TransferSubscriptions</i> request was issued.
StartPublishingTimer()	Start or restart the publishing timer and decrement the <i>LifetimeCounter Variable</i> .
UpdateSubscriptionParams()	Negotiate and update the <i>Subscription</i> parameters. If the new keep-alive interval is less than the current value of the keep-alive counter, perform <i>ResetKeepAliveCounter()</i> and <i>ResetLifetimeCounter()</i> .

5.13.2 CreateSubscription

5.13.2.1 Description

This *Service* is used to create a *Subscription*. *Subscriptions* monitor a set of *MonitoredItems* for *Notifications* and return them to the *Client* in response to *Publish* requests.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

5.13.2.2 Parameters

Table 88 defines the parameters for the *Service*.

Table 88 – CreateSubscription Service parameters

Name	Type	Description
Request		
requestHeader	Request Header	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
requestedPublishingInterval	Duration	This interval defines the cyclic rate that the <i>Subscription</i> is being requested to return <i>Notifications</i> to the <i>Client</i> . This interval is expressed in milliseconds. This interval is represented by the publishing timer in the <i>Subscription</i> state table (see 5.13.1.2). The negotiated value for this parameter returned in the response is used as the default sampling interval for <i>MonitoredItems</i> assigned to this <i>Subscription</i> . If the requested value is 0 or negative, the <i>Server</i> shall revise with the fastest supported publishing interval.
requestedLifetimeCount	Counter	Requested lifetime count (see 7.5 for <i>Counter</i> definition). The lifetime count shall be a minimum of three times the keep keep-alive count. When the publishing timer has expired this number of times without a <i>Publish</i> request being available to send a <i>NotificationMessage</i> , then the <i>Subscription</i> shall be deleted by the <i>Server</i> .
requestedMaxKeepAliveCount	Counter	Requested maximum keep-alive count (see 7.5 for <i>Counter</i> definition). When the publishing timer has expired this number of times without requiring any <i>NotificationMessage</i> to be sent, the <i>Subscription</i> sends a keep-alive <i>Message</i> to the <i>Client</i> . The negotiated value for this parameter is returned in the response. If the requested value is 0, the <i>Server</i> shall revise with the smallest supported keep-alive count.
maxNotificationsPerPublish	Counter	The maximum number of notifications that the <i>Client</i> wishes to receive in a single <i>Publish</i> response. A value of zero indicates that there is no limit. The number of notifications per <i>Publish</i> is the sum of monitoredItems in the <i>DataChangeNotification</i> and events in the <i>EventNotificationList</i> .

Name	Type	Description
publishingEnabled	Boolean	A <i>Boolean</i> parameter with the following values: TRUE publishing is enabled for the <i>Subscription</i> . FALSE publishing is disabled for the <i>Subscription</i> . The value of this parameter does not affect the value of the monitoring mode <i>Attribute of MonitoredItems</i> .
priority	Byte	Indicates the relative priority of the <i>Subscription</i> . When more than one <i>Subscription</i> needs to send <i>Notifications</i> , the <i>Server</i> should de-queue a Publish request to the <i>Subscription</i> with the highest <i>priority</i> number. For <i>Subscriptions</i> with equal <i>priority</i> the <i>Server</i> should de-queue Publish requests in a round-robin fashion. A Client that does not require special priority settings should set this value to zero.
Response		
responseHeader	Response Header	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition). This identifier shall be unique for the entire <i>Server</i> , not just for the <i>Session</i> , in order to allow the <i>Subscription</i> to be transferred to another <i>Session</i> using the TransferSubscriptions service. After <i>Server</i> start-up the generation of <i>subscriptionIds</i> should start from a random <i>IntegerId</i> or continue from the point before the restart.
revisedPublishingInterval	Duration	The actual publishing interval that the <i>Server</i> will use, expressed in milliseconds. The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.
revisedLifetimeCount	Counter	The lifetime of the <i>Subscription</i> shall be a minimum of three times the keep-alive interval negotiated by the <i>Server</i> .
revisedMaxKeepAliveCount	Counter	The actual maximum keep-alive count (see 7.5 for <i>Counter</i> definition). The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.

5.13.2.3 Service results

Table 89 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 89 – CreateSubscription Service result codes

Symbolic Id	Description
Bad_TooManySubscriptions	The <i>Server</i> has reached its maximum number of subscriptions.

5.13.3 ModifySubscription

5.13.3.1 Description

This *Service* is used to modify a *Subscription*.

Illegal request values for parameters that can be revised do not generate errors. Instead the *Server* will choose default values and indicate them in the corresponding revised parameter.

Changes to the *Subscription* settings shall be applied immediately by the *Server*. They take effect as soon as practical but not later than twice the new *revisedPublishingInterval*.

5.13.3.2 Parameters

Table 90 defines the parameters for the *Service*.

Table 90 – ModifySubscription Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition).
requestedPublishingInterval	Duration	This interval defines the cyclic rate at which the <i>Subscription</i> is being requested to return <i>Notifications</i> to the <i>Client</i>. This interval is expressed in milliseconds. This interval is represented by the publishing timer in the <i>Subscription</i> state table (see 5.13.1.2). The negotiated value for this parameter returned in the response is used as the default sampling interval for <i>MonitoredItems</i> assigned to this <i>Subscription</i>. If the requested value is 0 or negative, the <i>Server</i> shall revise with the fastest supported publishing interval.
requestedLifetimeCount	Counter	Requested lifetime count (see 7.5 for <i>Counter</i> definition). The lifetime count shall be a minimum of three times the keep-alive count. When the publishing timer has expired this number of times without a <i>Publish</i> request being available to send a <i>NotificationMessage</i>, then the <i>Subscription</i> shall be deleted by the <i>Server</i>.
requestedMaxKeepAliveCount	Counter	Requested maximum keep-alive count (see 7.5 for <i>Counter</i> definition). When the publishing timer has expired this number of times without requiring any <i>NotificationMessage</i> to be sent, the <i>Subscription</i> sends a keep-alive <i>Message</i> to the <i>Client</i>. The negotiated value for this parameter is returned in the response. If the requested value is 0, the <i>Server</i> shall revise with the smallest supported keep-alive count.
maxNotificationsPerPublish	Counter	The maximum number of notifications that the <i>Client</i> wishes to receive in a single <i>Publish</i> response. A value of zero indicates that there is no limit.
priority	Byte	Indicates the relative priority of the <i>Subscription</i>. When more than one <i>Subscription</i> needs to send <i>Notifications</i>, the <i>Server</i> should de-queue a <i>Publish</i> request to the <i>Subscription</i> with the highest <i>priority</i> number. For <i>Subscriptions</i> with equal <i>priority</i> the <i>Server</i> should de-queue <i>Publish</i> requests in a round-robin fashion. A <i>Client</i> that does not require special priority settings should set this value to zero.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
revisedPublishingInterval	Duration	The actual publishing interval that the <i>Server</i> will use, expressed in milliseconds. The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.
revisedLifetimeCount	Counter	The lifetime of the <i>Subscription</i> shall be a minimum of three times the keep-alive interval negotiated by the <i>Server</i> .
revisedMaxKeepAliveCount	Counter	The actual maximum keep-alive count (see 7.5 for <i>Counter</i> definition). The <i>Server</i> should attempt to honour the <i>Client</i> request for this parameter, but may negotiate this value up or down to meet its own constraints.

5.13.3.3 Service results

Table 91 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 91 – ModifySubscription Service result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.13.4 SetPublishingMode

5.13.4.1 Description

This *Service* is used to enable sending of *Notifications* on one or more *Subscriptions*.

5.13.4.2 Parameters

Table 92 defines the parameters for the *Service*.

Table 92 – SetPublishingMode Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
publishingEnabled	Boolean	A <i>Boolean</i> parameter with the following values: TRUE publishing of <i>NotificationMessages</i> is enabled for the <i>Subscription</i> . FALSE publishing of <i>NotificationMessages</i> is disabled for the <i>Subscription</i> . The value of this parameter does not affect the value of the monitoring mode <i>Attribute</i> of <i>MonitoredItems</i> . Setting this value to FALSE does not discontinue the sending of keep-alive <i>Messages</i> .
subscriptionIds []	IntegerId	List of <i>Server</i> -assigned identifiers for the <i>Subscriptions</i> to enable or disable (see 7.14 for <i>IntegerId</i> definition).
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCodes	List of <i>StatusCodes</i> for the <i>Subscriptions</i> to enable/disable (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Subscriptions</i> to enable/disable (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.4.3 Service results

Table 93 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 93 – SetPublishingMode Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.13.4.4 StatusCodes

Table 94 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 94 – SetPublishingMode Operation Level result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

5.13.5 Publish

5.13.5.1 Description

This *Service* is used for two purposes. First, it is used to acknowledge the receipt of *NotificationMessages* for one or more *Subscriptions*. Second, it is used to request the *Server* to return a *NotificationMessage* or a keep-alive *Message*. Since *Publish* requests are not directed to a specific *Subscription*, they may be used by any *Subscription*. 5.13.1.2 describes the use of the *Publish Service*.

Client strategies for issuing *Publish* requests may vary depending on the networking delays between the *Client* and the *Server*. In many cases, the *Client* may wish to issue a *Publish* request immediately after creating a *Subscription*, and thereafter, immediately after receiving a *Publish* response.

In other cases, especially in high latency networks, the *Client* may wish to pipeline *Publish* requests to ensure cyclic reporting from the *Server*. Pipelining involves sending more than one *Publish* request for each *Subscription* before receiving a response. For example, if the network introduces a delay between the *Client* and the *Server* of 5 seconds and the publishing interval for a *Subscription* is one second, then the *Client* shall issue *Publish* requests every second instead of waiting for a response to be received before sending the next request.

A *Server* should limit the number of active *Publish* requests to avoid an infinite number since it is expected that the *Publish* requests are queued in the *Server*. But a *Server* shall accept more queued *Publish* requests than created *Subscriptions*. It is expected that a *Server* supports several *Publish* requests per *Subscription*. When a *Server* receives a new *Publish* request that exceeds its limit it shall de-queue the oldest *Publish* request and return a response with the result set to *Bad_TooManyPublishRequests*. If a *Client* receives this *Service* result for a *Publish* request it shall not issue another *Publish* request before one of its outstanding *Publish* requests is returned from the *Server*.

Clients can limit the size of *Publish* responses with the *maxNotificationsPerPublish* parameter passed to the *CreateSubscription Service*. However, this could still result in a message that is too large for the *Client* or *Server* to process. In this situation, the *Client* will find that either the *SecureChannel* goes into a fault state and needs to be re-established or the *Publish* response returns an error and calling the *Republish Service* also returns an error. If either situation occurs then the *Client* shall adjust its message processing limits or the parameters for the *Subscription* and/or *MonitoredItems*.

The return diagnostic info setting in the request header of the *CreateMonitoredItems* or the last *ModifyMonitoredItems Service* is applied to the *Monitored Items* and is used as the diagnostic information settings when sending *Notifications* in the *Publish* response.

5.13.5.2 Parameters

Table 95 defines the parameters for the *Service*.

Table 95 – Publish Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscription Acknowledgements []	Subscription Acknowledgement	The list of acknowledgements for one or more <i>Subscriptions</i> . This list may contain multiple acknowledgements for the same <i>Subscription</i> (multiple entries with the same <i>subscriptionId</i>). This structure is defined in-line with the following indented items.
subscriptionId	IntegerId	The <i>Server</i> assigned identifier for a <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition).
sequenceNumber	Counter	The sequence number being acknowledged (see 7.5 for <i>Counter</i> definition). The <i>Server</i> may delete the <i>Message</i> with this sequence number from its retransmission queue.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> for which <i>Notifications</i> are being returned (see 7.14 for <i>IntegerId</i> definition). The value 0 is used to indicate that there were no <i>Subscriptions</i> defined for which a response could be sent.
availableSequence Numbers []	Counter	A list of sequence number ranges that identify unacknowledged <i>NotificationMessages</i> that are available for retransmission from the <i>Subscription</i> 's retransmission queue. This list is prepared after processing the acknowledgements in the request (see 7.5 for <i>Counter</i> definition).
moreNotifications	Boolean	A <i>Boolean</i> parameter with the following values: TRUE the number of <i>Notifications</i> that were ready to be sent could not be sent in a single response. FALSE all <i>Notifications</i> that were ready are included in the response.
notificationMessage	Notification Message	The <i>NotificationMessage</i> that contains the list of <i>Notifications</i> . The <i>NotificationMessage</i> parameter type is specified in 7.21.
results []	StatusCode	List of results for the acknowledgements (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>subscriptionAcknowledgements</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the acknowledgements (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionAcknowledgements</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.5.3 Service results

Table 96 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 96 – Publish Service result codes

Symbolic Id	Description
Bad_TooManyPublishRequests	The <i>Server</i> has reached the maximum number of queued <i>Publish</i> requests.
Bad_NoSubscription	There is no <i>Subscription</i> available for this session.

5.13.5.4 StatusCodes

Table 97 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 97 – Publish Operation Level Result Codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_SequenceNumberUnknown	The sequence number is unknown to the <i>Server</i> .

5.13.6 Republish

5.13.6.1 Description

This *Service* requests the *Subscription* to republish a *NotificationMessage* from its retransmission queue. If the *Server* does not have the requested *Message* in its retransmission queue, it returns an error response.

See 5.13.1.2 for the detail description of the behaviour of this *Service*.

See 6.7 for a description of the issues and strategies regarding reconnect handling and *Republish*.

5.13.6.2 Parameters

Table 98 defines the parameters for the *Service*.

Table 98 – Republish Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionId	IntegerId	The <i>Server</i> assigned identifier for the <i>Subscription</i> to be republished (see 7.14 for <i>IntegerId</i> definition).
retransmitSequence Number	Counter	The sequence number of a specific <i>NotificationMessage</i> to be republished (see 7.5 for <i>Counter</i> definition).
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
notificationMessage	Notification Message	The requested <i>NotificationMessage</i> . The <i>NotificationMessage</i> parameter type is specified in 7.21.

5.13.6.3 Service results

Table 99 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 99 – Republish Service result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_MessageNotAvailable	The requested message is no longer available.

5.13.7 TransferSubscriptions

5.13.7.1 Description

This *Service* is used to transfer a *Subscription* and its *MonitoredItems* from one *Session* to another. For example, a *Client* may need to reopen a *Session* and then transfer its *Subscriptions* to that *Session*. It may also be used by one *Client* to take over a *Subscription* from another *Client* by transferring the *Subscription* to its *Session*.

The *authenticationToken* contained in the request header identifies the *Session* to which the *Subscription* and *MonitoredItems* shall be transferred. The *Server* shall validate that the *Client* of that *Session* is operating on behalf of the same user and that the potentially new *Client* supports the *Profiles* that are necessary for the *Subscription*. If the *Server* transfers the *Subscription*, it returns the sequence numbers of the *NotificationMessages* that are available for retransmission. The *Client* should acknowledge all *Messages* in this list for which it will not request retransmission.

If the *Server* transfers the *Subscription* to the new *Session*, the *Server* shall issue a *StatusChangeNotification notificationMessage* with the status code *Good_SubscriptionTransferred* to the old *Session*. The *StatusChangeNotification notificationMessage* type is defined in 7.20.4.

5.13.7.2 Parameters

Table 100 defines the parameters for the *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 PDF

Table 100 – TransferSubscriptions Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionIds []	IntegerId	List of identifiers for the <i>Subscriptions</i> to be transferred to the new <i>Client</i> (see 7.14 for <i>IntegerId</i> definition). These identifiers are transferred from the primary <i>Client</i> to a backup <i>Client</i> via external mechanisms.
sendInitialValues	Boolean	A <i>Boolean</i> parameter with the following values: TRUE the first Publish response(s) after the <i>TransferSubscriptions</i> call shall contain the current values of all Monitored Items in the Subscription where the Monitoring Mode is set to Reporting. If a value is queued for a data <i>MonitoredItem</i>, the next value in the queue is sent in the <i>Publish</i> response. If no value is queued for a data <i>MonitoredItem</i>, the last value sent is repeated in the <i>Publish</i> response. FALSE the first Publish response after the <i>TransferSubscriptions</i> call shall contain only the value changes since the last Publish response was sent. This parameter only applies to MonitoredItems used for monitoring Attribute changes.
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	TransferResult	List of results for the <i>Subscriptions</i> to transfer. The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	<i>StatusCode</i> for each <i>Subscription</i> to be transferred (see 7.34 for <i>StatusCode</i> definition).
availableSequence Numbers []	Counter	A list of sequence number ranges that identify <i>NotificationMessages</i> that are in the <i>Subscription</i> 's retransmission queue. This parameter is null if the transfer of the <i>Subscription</i> failed. The <i>Counter</i> type is defined in 7.5.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Subscriptions</i> to transfer (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.7.3 Service results

Table 101 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 101 – TransferSubscriptions Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.
Bad_InsufficientClientProfile	The <i>Client</i> of the current <i>Session</i> does not support one or more <i>Profiles</i> that are necessary for the <i>Subscription</i> .

5.13.7.4 StatusCodes

Table 102 defines values for the operation level *statusCode* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 102 – TransferSubscriptions Operation Level result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.
Bad_UserAccessDenied	See Table 177 for the description of this result code. The <i>Client</i> of the current <i>Session</i> is not operating on behalf of the same user as the <i>Session</i> that owns the <i>Subscription</i> .

5.13.8 DeleteSubscriptions

5.13.8.1 Description

This *Service* is invoked to delete one or more *Subscriptions* that belong to the *Client's Session*.

Successful completion of this *Service* causes all *MonitoredItems* that use the *Subscription* to be deleted. If this is the last *Subscription* for the *Session*, then all *Publish* requests still queued for that *Session* are de-queued and shall be returned with *Bad_NoSubscription*.

Subscriptions that were transferred to another *Session* shall be deleted by the *Client* that owns the *Session*.

5.13.8.2 Parameters

Table 103 defines the parameters for the *Service*.

Table 103 – DeleteSubscriptions Service parameters

Name	Type	Description
Request		
requestHeader	RequestHeader	Common request parameters (see 7.28 for <i>RequestHeader</i> definition).
subscriptionIds []	IntegerId	The <i>Server</i> -assigned identifier for the <i>Subscription</i> (see 7.14 for <i>IntegerId</i> definition).
Response		
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).
results []	StatusCodes	List of <i>StatusCodes</i> for the <i>Subscriptions</i> to delete (see 7.34 for <i>StatusCode</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information for the <i>Subscriptions</i> to delete (see 7.8 for <i>DiagnosticInfo</i> definition). The size and order of the list matches the size and order of the <i>subscriptionIds</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the request.

5.13.8.3 Service results

Table 104 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 104 – DeleteSubscriptions Service result codes

Symbolic Id	Description
Bad_NothingToDo	See Table 177 for the description of this result code.
Bad_TooManyOperations	See Table 177 for the description of this result code.

5.13.8.4 StatusCodes

Table 105 defines values for the *results* parameter that are specific to this *Service*. Common *StatusCodes* are defined in Table 178.

Table 105 – DeleteSubscriptions Operation Level result codes

Symbolic Id	Description
Bad_SubscriptionIdInvalid	See Table 177 for the description of this result code.

6 Service behaviours

6.1 Security

6.1.1 Overview

The OPC UA *Services* define a number of mechanisms to meet the security requirements outlined in IEC TR 62541-2. This clause describes a number of important security-related procedures that *OPC UA Applications* shall follow.

6.1.2 Obtaining and installing an Application Instance Certificate

All *OPC UA Applications* require an *Application Instance Certificate* which shall contain the following information:

- The network name or address of the computer where the application runs;
- The name of the organization that administers or owns the application;
- The name of the application;
- The URI of the application instance;
- The name of the *Certificate Authority* that issued the *Certificate*;
- The issue and expiry date for the *Certificate*;
- The public key issued to the application by the *Certificate Authority* (CA);
- A digital signature created by the *Certificate Authority* (CA).

In addition, each *Application Instance Certificate* has a private key which should be stored in a location that can only be accessed by the application. If this private key is compromised, the administrator shall assign a new *Application Instance Certificate* and private key to the application.

This *Certificate* may be generated automatically when the application is installed. In this situation the private key assigned to the *Certificate* shall be used to create the *Certificate* signature. *Certificates* created in this way are called self-signed *Certificates*.

If the administrator responsible for the application decides that a self-signed *Certificate* does not meet the security requirements of the organization, then the administrator should install a *Certificate* issued by a *Certification Authority*. The steps involved in requesting an *Application Instance Certificate* from a *Certificate Authority* are shown in Figure 19.




Figure 19 – Obtaining and installing an Application Instance Certificate

Figure 19 illustrates the interactions between the application, the *Administrator* and the *Certificate Authority*. The *Application* is as *OPC UA Application* installed on a single machine. The *Administrator* is the person responsible for managing the machine and the *OPC UA Application*. The *Certificate Authority* is an entity that can issue digital *Certificates* that meet the requirements of the organization deploying the *OPC UA Application*.

If the *Administrator* decides that a self-signed *Certificate* meets the security requirements for the organization, then the *Administrator* may skip Steps 3 through 5. Application vendors shall ensure that a *Certificate* is available after the installation process. Every *OPC UA Application* shall allow the *Administrators* to replace *Application Instance Certificates* with *Certificates* that meet their requirements.

When the *Administrator* requests a new *Certificate* from a *Certificate Authority*, the *Certificate Authority* may require that the *Administrator* provide proof of authorization to request *Certificates* for the organization that will own the *Certificate*. The exact mechanism used to provide this proof depends on the *Certificate Authority*.

Vendors may choose to automate the process of acquiring *Certificates* from an authority. If this is the case, the *Administrator* would still go through the steps illustrated in Figure 19, however, the installation program for the application would do them automatically and only prompt the *Administrator* to provide information about the application instance being installed.

6.1.3 Determining if a Certificate is trusted

Applications shall never communicate with another application that they do not trust. An *Application* decides if another application is trusted by checking whether the *Application Instance Certificate* for the other application is trusted. Applications shall rely on lists of *Certificates* provided by the *Administrator* to determine trust. There are two separate lists: a list of trusted *Applications* and a list of trusted *Certificate Authorities* (CAs). If an application is not directly trusted (i.e. its *Certificate* is not in the list of trusted applications) then the application shall build a chain of *Certificates* back to a trusted CA.

When building a chain each *Certificate* in the chain shall be validated. If any validation error occurs then the trust check fails. Some validation errors are non-critical which means they can be suppressed by a user of an *Application* with the appropriate privileges. Suppressed validation errors are always reported via auditing (i.e. an appropriate Audit event is raised).

Building a trust chain requires access to all *Certificates* in the chain. These *Certificates* may be stored locally or they may be provided with the application *Certificate*. Processing fails with *Bad_SecurityChecksFailed* if a CA *Certificate* cannot be found.

Table 106 specifies the steps used to validate a *Certificate* in the order that they shall be followed. These steps are repeated for each *Certificate* in the chain. Each validation step has a unique error status and audit event type that shall be reported if the check fails. The audit event is in addition to any audit event that was generated for the particular *Service* that was invoked. The *Service* audit event in its message text shall include the audit *EventId* of the *AuditCertificateEventType* (for more details, see 6.5). Processing halts if an error occurs, unless it is non-critical and it has been suppressed.

ApplicationInstanceCertificates shall not be used in a *Client* or *Server* until they have been evaluated and marked as trusted. This can happen automatically by a PKI trust chain or in an offline manner where the *Certificate* is marked as trusted by an administrator after evaluation.

Table 106 – Certificate validation steps

Step	Error/AuditEvent	Description
Certificate Structure	Bad_CertificateInvalid Bad_SecurityChecksFailed AuditCertificateInvalidEventType	The <i>Certificate</i> structure is verified. This error may not be suppressed. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .
Build Certificate Chain	Bad_CertificateChainIncomplete Bad_SecurityChecksFailed AuditCertificateInvalidEventType	The trust chain for the <i>Certificate</i> is created. An error during the chain creation may not be suppressed. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .
Signature	Bad_CertificateInvalid Bad_SecurityChecksFailed AuditCertificateInvalidEventType	A <i>Certificate</i> with an invalid signature shall always be rejected. A <i>Certificate</i> signature is invalid if the <i>Issuer Certificate</i> is unknown. A self-signed <i>Certificate</i> is its own issuer. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .
Security Policy Check	Bad_CertificatePolicyCheckFailed Bad_SecurityChecksFailed AuditCertificateInvalidEventType	A <i>Certificate</i> signature shall comply with the <i>CertificateSignatureAlgorithm</i> , <i>MinAsymmetricKeyLength</i> and <i>MaxAsymmetricKeyLength</i> requirements for the used <i>SecurityPolicy</i> defined in IEC 62541-7. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> . This error may be suppressed.
Trust List Check	Bad_CertificateUntrusted Bad_SecurityChecksFailed AuditCertificateUntrustedEventType	If the <i>Application Instance Certificate</i> is not trusted and none of the CA <i>Certificates</i> in the chain is trusted, the result of the <i>Certificate</i> validation shall be <i>Bad_CertificateUntrusted</i> . If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .

Step	Error/AuditEvent	Description
Validity Period	Bad_CertificateTimeInvalid Bad_CertificateIssuerTimeInvalid AuditCertificateExpiredEventType	The current time shall be after the start of the validity period and before the end. This error may be suppressed.
Host Name	Bad_CertificateHostNameInvalid AuditCertificateDataMismatchEventType	The HostName in the URL used to connect to the Server shall be the same as one of the HostNames specified in the <i>Certificate</i> . This check is skipped for CA <i>Certificates</i> . This check is skipped for <i>Server</i> side validation. This error may be suppressed.
URI	Bad_CertificateUriInvalid AuditCertificateDataMismatchEventType	<i>Application</i> and <i>Software Certificates</i> contain an application or product URI that shall match the URI specified in the <i>ApplicationDescription</i> provided with the <i>Certificate</i> . This check is skipped for CA <i>Certificates</i> . This error may not be suppressed. The <i>gatewayServerUri</i> is used to validate an <i>Application Certificate</i> when connecting to a <i>Gateway Server</i> (see 7.1).
Certificate Usage	Bad_CertificateUseNotAllowed Bad_CertificateIssuerUseNotAllowed AuditCertificateMismatchEventType	Each <i>Certificate</i> has a set of uses for the <i>Certificate</i> (see IEC 62541-6). These uses shall match use requested for the <i>Certificate</i> (i.e. Application, Software or CA). This error may be suppressed unless the <i>Certificate</i> indicates that the usage is mandatory.
Find Revocation List	Bad_CertificateRevocationUnknown Bad_CertificateIssuerRevocationUnknown AuditCertificateRevokedEventType	Each CA <i>Certificate</i> may have a revocation list. This check fails if this list is not available (i.e. a network interruption prevents the application from accessing the list). No error is reported if the <i>Administrator</i> disables revocation checks for a CA <i>Certificate</i> . This error may be suppressed.
Revocation Check	Bad_CertificateRevoked Bad_CertificateIssuerRevoked AuditCertificateRevokedEventType	The <i>Certificate</i> has been revoked and may not be used. This error may not be suppressed. If this check fails on the <i>Server</i> side, the error <i>Bad_SecurityChecksFailed</i> shall be reported back to the <i>Client</i> .

Certificates are usually placed in a central location called a *CertificateStore*. Figure 20 illustrates the interactions between the *Application*, the *Administrator* and the *CertificateStore*. The *CertificateStore* could be on the local machine or in some central server. The exact mechanisms used to access the *CertificateStore* depend on the application and PKI environment set up by the *Administrator*.




Figure 20 – Determining if an Application Instance Certificate is trusted

6.1.4 Creating a SecureChannel

All *OPC UA Applications* shall establish a *SecureChannel* before creating a *Session*. This *SecureChannel* requires that both applications have access to *Certificates* that can be used to encrypt and sign *Messages* exchange. The *Application Instance Certificates* installed by following the process described in 6.1.2 may be used for this purpose.

The steps involved in establishing a *SecureChannel* are shown in Figure 21.




Figure 21 – Establishing a SecureChannel

Figure 21 assumes *Client* and *Server* have online access to a *CertificateAuthority* (CA). If online access is not available and if the administrator has installed the CA public key on the local machine, then the *Client* and *Server* shall still validate the application *Certificates* using that key. The figure shows only one CA, however, there is no requirement that the *Client* and *Server* *Certificates* be issued by the same authority. A self-signed *Application Instance Certificate* does not need to be verified with a CA. Any *Certificate* shall be rejected if it is not in a trust list provided by the administrator.

Both the *Client* and *Server* shall have a list of *Certificates* that they have been configured to trust (sometimes called the *Certificate Trust List* or CTL). These trusted *Certificates* may be *Certificates* for *Certificate Authorities* or they may be *OPC UA Application Instance Certificates*. *OPC UA Applications* shall be configured to reject connections with applications that do not have a trusted *Certificate*.

Certificates can be compromised, which means they should no longer be trusted. Administrators can revoke a *Certificate* by removing it from the trust list for all applications or the CA can add the *Certificate* to the *Certificate Revocation List* (CRL) for the *Issuer Certificate*. Administrators may save a local copy of the CRL for each *Issuer Certificate* when online access is not available.

A *Client* does not need to call *GetEndpoints* each time it connects to the *Server*. This information should change rarely and the *Client* can cache it locally. If the *Server* rejects the *OpenSecureChannel* request the *Client* should call *GetEndpoints* and make sure the *Server* configuration has not changed.

There are two security risks which a *Client* shall be aware of when using the *GetEndpoints Service*. The first could come from a rogue *Discovery Server* that tries to direct the *Client* to a rogue *Server*. For this reason the *Client* shall verify that the *ServerCertificate* in the *EndpointDescription* is a trusted *Certificate* before it calls *CreateSession*.

The second security risk comes from a third party that alters the contents of the *EndpointDescriptions* as they are transferred over the network back to the *Client*. The *Client* protects itself against this by comparing the list of *EndpointDescriptions* returned from the *GetEndpoints Service* with list returned in the *CreateSession* response.

The exact mechanisms for using the *SecurityToken* to sign and encrypt *Messages* exchanged over the *SecureChannel* are described in IEC 62541-6. The process for renewing tokens is also described in detail in IEC 62541-6.

In many cases, the *Certificates* used to establish the *SecureChannel* will be the *Application Instance Certificates*. However, some *Communication Stacks* might not support *Certificates* that are specific to a single application. Instead, they expect all communication to be secured with a *Certificate* specific to a user or the entire machine. For this reason, *OPC UA Applications* will need to exchange their *Application Instance Certificates* when creating a *Session*.

6.1.5 Creating a Session

Once an OPC UA *Client* has established a *SecureChannel* with a *Server* it can create an OPC UA *Session*.

The steps involved in establishing a *Session* are shown in Figure 22.




Figure 22 – Establishing a Session

Figure 22 illustrates the interactions between a *Client*, a *Server*, a *Certificate Authority* (CA) and an identity provider. The CA is responsible for issuing the *Application Instance Certificates*. If the *Client* or *Server* does not have online access to the CA, then they shall validate the *Application Instance Certificates* using the CA public key that the administrator shall install on the local machine.

The identity provider may be a central database that can verify that user token provided by the *Client*. This identity provider may also tell the *Server* which access rights the user has. The identity provider depends on the user identity token. It could be a *Certificate Authority*, a Kerberos ticket granting service, a WS-Trust *Server* or a proprietary database of some sort.

The *Client* and *Server* shall prove possession of their *Application Instance Certificates* by signing the *Certificates* with a nonce appended. The exact mechanism used to create the proof of possession signatures is described in 5.6.2. Similarly, the *Client* shall prove possession by either providing a secret like a password in the user identity token or by creating a signature with the secret associated with a user identity token like x.509 v3.

6.1.6 Impersonating a User

Once an OPC UA *Client* has established a *Session* with a *Server* it can change the user identity associated with the *Session* by calling the *ActivateSession* service.

The steps involved in impersonating a user are shown in Figure 23.




Figure 23 – Impersonating a User

6.2 Authorization Services

6.2.1 Overview

Authorization Services provide *Access Tokens* to *Clients* on behalf of *Users* that they pass to a *Server* to be granted access to resources.

In a basic model (as shown in Figure 22) the *Server* is responsible for authorization (i.e. deciding what a user can do) while a separate identity provider (e.g. the operating system) is responsible for authentication (deciding who the user is).

In more complex models, the *Server* relies on external *Authorization Services* to provide some of its authorization requirements. These *Authorization Services* act in concert with an external identity provider which validates the user credentials before the external *Authorization Service* creates an *Access Token* that tells the *Server* what the user is allowed to do. The *Client* interactions with these services may be indirect as shown in 6.2.2 or direct as shown in 6.2.3.

Even when the *Server* requires the *Client* to use an external *Authorization Service*, the *Server* is still responsible for managing and enforcing the *Permissions* assigned to *Nodes* in its *Address Space*. The subclauses below discuss the use of an external *Authorization Service* in more detail.


6.2.2 Indirect handshake with an Identity Provider

Authorization Services (AS) provide access to identity providers which can validate the credentials provided by *Clients*. They then provide tokens which can be passed to a *Server* instead of the credentials. These tokens are passed as an *IssuedIdentityToken* defined in 7.36.6.

The protocol to request tokens depends on the *Authorization Service* (AS). Common protocols include Kerberos and OAuth2. OAuth2 supports claims based authorization as described in IEC TR 62541-2.

Servers publish the *Authorization Services* (AS) they support in the *UserTokenPolicies* list return with *GetEndpoints*. The *IssuedTokenType* field specifies the protocol used to communicate with the AS. The *IssuerEndpointUrl* field contains the information needed by the *Client* to connect to the AS using the protocol required by the AS.

The basic handshake is shown in Figure 24.



IEC

Figure 24 – Indirect handshake with an Identity Provider

6.2.3 Direct handshake with an Identity Provider

Authorization Services require that *Servers* be registered with them because the *Access Tokens* can only be used with a single *Server*. This can introduce a lot of complexity for administrators. One way to reduce this complexity is to leverage the *Server* information that is already managed by a Global Discovery Service (GDS) described in IEC 62541-12. In this model the user identities are still managed by a central *Authorization Service*. The interactions are shown in Figure 25.




Figure 25 – Direct handshake with an Identity Provider

The *UserTokenPolicy* returned from the *Server* provides the URL of the *Authorization Service* and the identity provider. If the *Application Authorization Service* is linked with the GDS, it knows of all *Servers* which have been issued *Certificates*. The *ApplicationUri* is used as the identifier for the *Server* passed to the AS. The identity provider is responsible for managing users known to the system. It validates the credentials provided by the *Client* and returns an *Identity Access Token* which identifies the user. The *Identity Access Token* is passed to the *Application Authorization Service* which validates the *Client* and *Server* applications and creates a new *Access Token* that can be used to access the *Server*.

6.3 Session-less Service invocation

6.3.1 Description

The *Session-less Service* invocation is introduced for *Services*, such as *Read*, *Write* or *Call*, that do not require any caller specific state information. It is accessible through the *SessionlessInvoke Service* which provides the context information required to call *Services* without a *Session*.

Session-less invocation is limited to *Services* of the *View Service Set* (with exception of *RegisterNodes* and *UnregisterNodes*), *Attribute Service Set*, *Method Service Set*, *NodeManagement Service Set* and *Query Service Set*. All *Services* belonging to these *Service Sets* that are supported by a *Server* via a *Session* shall also be supported via the *SessionlessInvoke Service*.

Session-less Services can be invoked via a *SecureChannel* by using the *Access Token* returned from the *Authorization Service* as the *authenticationToken* in the *requestHeader*. The *SecureChannel* shall have encryption enabled to prevent eavesdroppers from seeing the *Access Token*. The *Access Token* provides the user authentication. If application authentication through the *SecureChannel* is sufficient, *Servers* may not require the *Access Token* and assume an anonymous user. In this case the *authenticationToken* shall be null.

The *SessionlessInvoke Messages* are just an envelope for the *Service* to invoke and do not have a *RequestHeader* and *ResponseHeader* like other *Services*. Those parameters are already part of the *body* which contains the *Message* for the *Service* to invoke.

Any *Endpoint* used for normal communication could be used for *Session-less* invocation provided the *Endpoint* supports encryption. The *Server* returns *Bad_ServiceUnsupported* if it does not support *Session-less* invocation for the request specified in the *body*. If it supports

invocation but not with the combination of *Endpoint* and security settings used, it returns *Bad_SecurityModelInsufficient*.

Servers may expose *Endpoints* which are only for use with *Session-less* invocation. These *Endpoints* shall support *GetEndpoints* and *FindServers* in addition to the *SessionlessInvoke Service*. The *Server* returns *Bad_ServiceUnsupported* for the other *Services*.

A *Session* ensures that a namespace index or a server index does not change during the lifetime of a *Session*. This cannot be ensured between *Session-less Services* invocations. There are two options to ensure the namespace indices in the call match the expected namespace URIs in the *Server*. One option for the caller is to provide the list of namespace URIs used to build the namespace indices. This works best for single *Session-less Service* invocations. The second option is to provide the *UrisVersion* to ensure consistency of namespace arrays between *Client* and *Server*. The *UrisVersion* is first read from the *Server* together with the *NamespaceArray* and *ServerArray*. This reduces the overhead per call for a sequence of *Session-less Service* invocations.

6.3.2 Parameters

Table 107 defines the parameters for the *Service*.

Table 107 – SessionlessInvoke Service parameters

Name	Type	Description
Request		
urisVersion	VersionTime	The version of the <i>NamespaceArray</i> and the <i>ServerArray</i> used for the <i>Service</i> invocation. The version shall match the value of the <i>UrisVersion Property</i> that defines the version for the URI lists in the <i>NamespaceArray</i> and the <i>ServerArray Properties</i> defined in IEC 62541-5. If the <i>urisVersion</i> parameter does not match the <i>Server's UrisVersion Property</i> , the <i>Server</i> shall return <i>Bad_VersionTimeInvalid</i> . In this case the <i>Client</i> shall read the <i>UrisVersion</i> , <i>NamespaceArray</i> and the <i>ServerArray</i> from the <i>Server Object</i> to repeat the <i>Service</i> invocation with the right version. The <i>VersionTime DataType</i> is defined in 7.38. If the value is 0, the parameter is ignored and the URIs are defined by the <i>namespaceUris</i> and <i>serverUris</i> parameters in request and response. If the value is non-zero, the <i>namespaceUris</i> and <i>serverUris</i> parameters in the request are ignored by the <i>Server</i> and set to null arrays in the response.
namespaceUris []	String	A list of URIs referenced by <i>NodeIds</i> or <i>QualifiedNames</i> in the request. NamespaceIndex 0 shall not be in this list. The first entry in this list is NamespaceIndex 1. The parameter shall be ignored by the <i>Server</i> if the <i>urisVersion</i> is not 0.
serverUris []	String	A list of URIs referenced by <i>ExpandedNodeIds</i> in the request. ServerIndex 0 shall not be in this list. The first entry in this list is ServerIndex 1. The parameter shall be ignored by the <i>Server</i> if the <i>urisVersion</i> is not 0.
localeIds []	LocaleId	List of locale ids in priority order for localized strings. The first <i>LocaleId</i> in the list has the highest priority. If the <i>Server</i> returns a localized string to the <i>Client</i> , the <i>Server</i> shall return the translation with the highest priority that it can. If it does not have a translation for any of the locales identified in this list, then it shall return the string value that it has and include the locale id with the string. See IEC 62541-3 for more detail on locale ids. If <i>localeIds</i> is empty, the returned language variant is <i>Server</i> specific.
serviceId	UInt32	The numeric identifier assigned to the <i>Service</i> request <i>DataType</i> describing the body.
body	*	The body of the request. The body is an embedded structure containing the corresponding <i>Service</i> request for the <i>serviceId</i> .

Name	Type	Description
Response		
namespaceUris []	String	A list of URIs referenced by <i>NodeIds</i> or <i>QualifiedNames</i> in the response. NamespaceIndex 0 shall not be in this list. The first entry in this list is NamespaceIndex 1. An empty array shall be returned if the <i>urisVersion</i> is not 0.
serverUris []	String	A list of URIs referenced by <i>ExpandedNodeIds</i> in the response. ServerIndex 0 shall not be in this list. The first entry in this list is ServerIndex 1. An empty array shall be returned if the <i>urisVersion</i> is not 0.
serviceld	UInt32	The numeric identifier assigned to the <i>Service</i> response <i>DataType</i> describing the body.
body	*	The body of the response. The body is an embedded structure containing the corresponding <i>Service</i> response for the <i>serviceld</i> .

6.3.3 Service results

Table 108 defines the *Service* results specific to this *Service*. Common *StatusCodes* are defined in Table 177.

Table 108 – SessionlessInvoke Service result codes

Symbolic Id	Description
Bad_VersionTimeInvalid	The provided version time is no longer valid.

6.4 Software Certificates

NOTE Details on *SoftwareCertificates* will be defined in a future edition of IEC 62541-4.

6.5 Auditing

6.5.1 Overview

Auditing is a requirement in many systems. It provides a means of tracking activities that occur as part of normal operation of the system. It also provides a means of tracking abnormal behaviour. It is also a requirement from a security standpoint. For more information on the security aspects of auditing, see IEC TR 62541-2. Subclause 6.5 describes what is expected of an OPC UA *Server* and *Client* with respect to auditing and it details the audit requirements for each service set. Auditing can be accomplished using one or both of the following methods:

- The *OPC UA Application* that generates the audit event can log the audit entry in a log file or other storage location;
- The *OPC UA Server* that generates the audit event can publish the audit event using the OPC UA event mechanism. This allows an external OPC UA *Client* to subscribe to and log the audit entries to a log file or other storage location.

6.5.2 General audit logs

Each OPC UA *Service* request contains a string parameter that is used to carry an audit record id. A *Client* or any *Server* operating as a *Client*, such as an aggregating *Server*, can create a local audit log entry for a request that it submits. This parameter allows this *Client* to pass the identifier for this entry with the request. If this *Server* also maintains an audit log, it should include this id in its audit log entry that it writes. When this log is examined and that

entry is found, the examiner will be able to relate it directly to the audit log entry created by the *Client*. This capability allows for traceability across audit logs within a system.

6.5.3 General audit Events

A *Server* that maintains an audit log shall provide the audit log entries via *Event Messages*. The *AuditEventType* and its sub-types are defined in IEC 62541-3. An audit *Event Message* also includes the audit record Id. The details of the *AuditEventType* and its subtypes are defined in IEC 62541-5. A *Server* that is an aggregating *Server* that supports auditing shall also subscribe for audit events for all of the *Servers* that it is aggregating (assuming they provide auditing). The combined stream should be available from the aggregating *Server*.

6.5.4 Auditing for Discovery Service Set

This *Service Set* can be separated into two groups: *Services* that are called by OPC UA *Clients* and *Services* that are invoked by OPC UA *Servers*. The *FindServers* and *GetEndpoints* *Services* that are called by OPC UA *Clients* may generate audit entries for failed *Service* invocations. The *RegisterServer* *Service* that is invoked by OPC UA *Servers* shall generate audit entries for all new registrations and for failed *Service* invocations. These audit entries shall include the *Server* URI, *Server* names, *Discovery* URIs and isOnline status. Audit entries should not be generated for *RegisterServer* invocation that does not cause changes to the registered *Servers*.

6.5.5 Auditing for SecureChannel Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for failed service invocations and for successful invocation of the *OpenSecureChannel* and *CloseSecureChannel* *Services*. The *Client* generated audit entries should be setup prior to the actual call, allowing the correct audit record Id to be provided. The *OpenSecureChannel* *Service* shall generate an audit *Event* of type *AuditOpenSecureChannelEventType* or a subtype of it for the *requestType* ISSUE_0. Audit *Events* for the *requestType* RENEW_1 are only created if the renew fails. The *CloseSecureChannel* service shall generate an audit *Event* of type *AuditChannelEventType* or a subtype of it. Both of these *Event* types are subtypes of the *AuditChannelEventType*. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure cases the *Message* for *Events* of this type should include a description of why the service failed. This description should be more detailed than what was returned to the client. From a security point of view a *Client* only needs to know that it failed, but from an *Auditing* point of view the exact details of the failure need to be known. In the case of *Certificate* validation errors the description should include the audit *EventId* of the specific *AuditCertificateEventType* that was generated to report the *Certificate* error. The *AuditCertificateEventType* shall also contain the detailed *Certificate* validation error. The additional parameters should include the details of the request. It is understood that these events may be generated by the underlying *Communication Stacks* in many cases, but they shall be made available to the *Server* and the *Server* shall report them.

6.5.6 Auditing for Session Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed *Service* invocations. These *Services* shall generate an audit *Event* of type *AuditSessionEventType* or a subtype of it. In particular, they shall generate the base *EventType* or the appropriate subtype, depending on the service that was invoked. The *CreateSession* service shall generate *AuditCreateSessionEventType* events or sub-types of it. The *ActivateSession* service shall generate *AuditActivateSessionEventType* events or subtypes of it. When the *ActivateSession* *Service* is called to change the user identity then the *Server* shall generate *AuditActivateSessionEventType* events or subtypes of it. The *CloseSession* service shall generate *AuditSessionEventType* events or subtypes of it. It shall always be generated if a *Session* is terminated like *Session* timeout expiration or *Server* shutdown. The *SourceName* for *Events* of this type shall be "Session/Timeout" for a *Session* timeout, "Session/CloseSession" for a *CloseSession* *Service* call and "Session/Terminated" for all

other cases. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case the *Message* for *Events* of this type should include a description of why the *Service* failed. The additional parameters should include the details of the request.

This *Service Set* shall also generate additional audit events in the cases when *Certificate* validation errors occur. These audit *Events* are generated in addition to the *AuditSessionEventType Events*. See IEC 62541-3 for the definition of *AuditCertificateEventType* and its subtypes.

For *Clients*, that support auditing, accessing the services in the *Session Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. These audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.7 Auditing for NodeManagement Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed *Service* invocations. These *Services* shall generate an audit *Event* of type *AuditNodeManagementEventType* or subtypes of it. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case, the *Message* for *Events* of this type should include a description of why the service failed. The additional parameters should include the details of the request.

For *Clients* that support auditing, accessing the *Services* in the *NodeManagement Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.8 Auditing for Attribute Service Set

The *Write* or *HistoryUpdate* *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed *Service* invocations. These *Services* shall generate an audit *Event* of type *AuditUpdateEventType* or subtypes of it. In particular, the *Write Service* shall generate an audit event of type *AuditWriteUpdateEventType* or a subtype of it. The *HistoryUpdate Service* shall generate an audit *Event* of type *AuditHistoryUpdateEventType* or a subtype of it. Three subtypes of *AuditHistoryUpdateEventType* are defined as *AuditHistoryEventUpdateEventType*, *AuditHistoryValueUpdateEventType* and *AuditHistoryDeleteEventType*. The subtype depends on the type of operation being performed, historical event update, historical data value update or a historical delete. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case the *Message* for *Events* of this type should include a description of why the *Service* failed. The additional parameters should include the details of the request.

The *Read* and *HistoryRead Services* may generate audit entries and audit *Events* for failed *Service* invocations. These *Services* should generate an audit *Event* of type *AuditEventType* or a subtype of it. See IEC 62541-5 for the detailed assignment of the *SourceNode*, *SourceName* and additional parameters. The *Message* for *Events* of this type should include a description of why the *Service* failed.

For *Clients* that support auditing, accessing the *Write* or *HistoryUpdate* services in the *Attribute Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. Invocations of the other *Services* in this *Service Set* may generate audit entries. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.9 Auditing for Method Service Set

All *Services* in this *Service Set* for *Servers* that support auditing may generate audit entries and shall generate audit *Events* for both successful and failed service invocations if the invocation modifies the *AddressSpace*, writes a value or modifies the state of the system (alarm acknowledge, batch sequencing or other system changes). These method calls shall generate an audit *Event* of type *AuditUpdateMethodEventType* or subtypes of it. Methods that do not modify the *AddressSpace*, write values or modify the state of the system may generate events. See IEC 62541-5 for the detailed assignment of the *SourceNode*, *SourceName* and additional parameters.

For *Clients* that support auditing, accessing the *Method Service Set* shall generate audit entries for both successful and failed invocations of the *Service*, if the invocation modifies the *AddressSpace*, writes a value or modifies the state of the system (alarm acknowledge, batch sequencing or other system changes). Invocations of the other *Methods* may generate audit entries. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.5.10 Auditing for View, Query, MonitoredItem and Subscription Service Set

All of the *Services* in these four *Service Sets* only provide the *Client* with information, with the exception of the *TransferSubscriptions Service* in the *Subscription Service Set*. In general, these services will not generate audit entries or audit *Event Messages*. The *TransferSubscriptions Service* shall generate an audit *Event* of type *AuditSessionEventType* or subtypes of it for both successful and failed *Service* invocations. See IEC 62541-5 for the detailed assignment of the *SourceNode*, the *SourceName* and additional parameters. For the failure case, the *Message* for *Events* of this type should include a description of why the service failed.

For *Clients* that support auditing, accessing the *TransferSubscriptions Service* in the *Subscription Service Set* shall generate audit entries for both successful and failed invocations of the *Service*. Invocations of the other *Services* in this *Service Set* do not require audit entries. All audit entries should be setup prior to the actual *Service* invocation, allowing the invocation to contain the correct audit record id.

6.6 Redundancy

6.6.1 Redundancy overview

OPC UA enables *Servers*, *Clients* and networks to be redundant. OPC UA provides the data structures and *Services* by which *Redundancy* may be achieved in a standardized manner.

Server Redundancy allows *Clients* to have multiple sources from which to obtain the same data. *Server Redundancy* can be achieved in multiple manners, some of which require *Client* interaction, others that require no interaction from a *Client*. Redundant *Servers* could exist in systems without redundant networks or *Clients*. Redundant *Servers* could also coexist in systems with network and *Client Redundancy*. *Server Redundancy* is formally defined in 6.6.2.

Client Redundancy allows identically configured *Clients* to behave as if they were single *Clients*, but not all *Clients* are obtaining data at a given time. Ideally there should be no loss of information when a *Client Failover* occurs. Redundant *Clients* could exist in systems without redundant networks or *Servers*. Redundant *Clients* could also coexist in systems with network and *Server Redundancy*. *Client Redundancy* is formally defined in 6.6.3.

Network Redundancy allows a *Client* and *Server* to have multiple communication paths to obtain the same data. Redundant networks could exist in systems without redundant *Servers* or *Clients*. Redundant networks could also coexist in systems with *Client* and *Server Redundancy*. *Network Redundancy* is formally defined in 6.6.4.

6.6.2 Server Redundancy

6.6.2.1 General

There are two general modes of *Server Redundancy*, transparent and non-transparent.

In transparent *Redundancy* the *Failover* of *Server* responsibilities from one *Server* to another is transparent to the *Client*. The *Client* is unaware that a *Failover* has occurred and the *Client* has no control over the *Failover* behaviour. Furthermore, the *Client* does not need to perform any actions to continue to send or receive data.

In non-transparent *Redundancy* the *Failover* from one *Server* to another and actions to continue to send or receive data are performed by the *Client*. The *Client* shall be aware of the *Redundant Server Set* and shall perform the required actions to benefit from the *Server Redundancy*.

The *ServerRedundancy Object* defined in IEC 62541-5 indicates the mode supported by the *Server*. The *ServerRedundancyType ObjectType* and its subtypes *TransparentRedundancyType* and *NonTransparentRedundancyType* defined in IEC 62541-5 specify information for the supported *Redundancy* mode.

6.6.2.2 Redundant Server Set requirements

OPC UA *Servers* that are part of a *Redundant Server Set* have certain *AddressSpace* requirements. These requirements allow a *Client* to consistently access information from *Servers* in a *Redundant Server Set* and to make intelligent choices related to the health and availability of *Servers* in the *Redundant Server Set*.

Servers in the *Redundant Server Set* shall have an identical *AddressSpace* including:

- identical *NodeIds*
- identical browse paths and structure of the *AddressSpace*
- identical logic for setting the *ServiceLevel*

The only *Nodes* that can differ between *Servers* in a *Redundant Server Set* are the *Nodes* that are in the local *Server* namespace like the *Server* diagnostic *Nodes*. A *Client* that fails over shall not be required to translate browse paths or otherwise resolve *NodeIds*. *Servers* are allowed to add and delete *Nodes* as long as all *Servers* in the *Redundant Server Set* will be updated with the same *Node* changes.

All *Servers* in a *Redundant Server Set* shall be synchronized with respect to time. This may mean installing a NTP service or a PTP service.

There are other important considerations for a redundant system regarding synchronization:

- **EventIds:** Each UA *Server* in a *Transparent and HotAndMirrored Redundant Server Set* shall synchronize *EventIds* to prevent a *Client* from mistakenly processing the same event multiple times simply because the *EventIds* are different. This is very important for Alarms & Conditions. For Cold, Warm, and Hot *Redundant Server Sets*, *Clients* shall be able to handle *EventIds* that are not synchronized. Following any *Failover* the *Client* shall call *ConditionRefresh* defined in IEC 62541-9.
- **Timestamp (Source/Server):** If a *Server* is exposing data from a downstream device (PLC, DCS, etc.), then the *SourceTimestamp* and *ServerTimestamp* reported by all redundant *Servers* should match as closely as possible. *Clients* should favour the use of the *SourceTimestamp*.
- **ContinuationPoints:** Behaviour of continuation points does not change, in that *Clients* shall be prepared for lost continuation points. *Servers* in *Transparent and HotAndMirrored Redundancy* sets shall synchronize continuation points and they may do so in other modes.

6.6.2.3 Transparent Redundancy

6.6.2.3.1.1 Client behaviour

To a *Client* the transparent *Redundant Server Set* appears as if it is just a single *Server* and the *Client* has no *Failover* actions to perform. All *Servers* in the *Redundant Server Set* have an identical *ServerUri* and an identical *EndpointUrl*.

Figure 26 shows a typical transparent *Redundancy* setup.




Figure 26 – Transparent Redundancy setup example

For transparent *Redundancy*, OPC UA provides data structures to allow *Clients* to identify which *Servers* are in the *Redundant Server Set*, the *ServiceLevel* of each *Server*, and which *Server* is currently responsible for the *Client Session*. This information is specified in *TransparentRedundancyType ObjectType* defined in IEC 62541-5. Since the *ServerUri* is identical for all *Servers* in the *Redundant Server Set*, the *Servers* are identified with a *ServerId* contained in the information provided in the *TransparentRedundancyType Object*.

In transparent *Redundancy*, a *Client* is not able to control which physical *Server* it actually connects to. *Failover* is controlled by the *Redundant Server Set* and a *Client* is also not able to actively *Failover* to another *Server* in the *Redundant Server Set*.

6.6.2.3.1.2 Server requirements

All OPC UA interactions within a given *Session* shall be supported by one *Server* and the *Client* is able to identify which *Server* that is, allowing a complete audit trail for the data. It is the responsibility of the *Servers* to ensure that information is synchronised between the *Servers*. A functional *Server* will take over the *Session* and *Subscriptions* from the *Failed Server*. *Failover* may require a reconnection of the *Client's SecureChannel* but the *EndpointUrl* of the *Server* and the *ServerUri* shall not change. The *Client* shall be able to continue communication with the *Sessions* and *Subscriptions* created on the previously used *Server*.

Figure 26 provides an abstract view of a transparent *Redundant Server Set*. The two or more *Servers* in the *Redundant Server Set* share a virtual network address and therefore all *Servers* have the identical *EndpointUrl*. How this virtual network address is created and managed is vendor specific. There may be special hardware that mediates the network address displayed to the rest of the network. There may be custom hardware, where all components are redundant and *Failover* at a hardware level automatically. There may even be software based systems where all the transparency is governed completely by software.

6.6.2.4 Non-transparent Redundancy

6.6.2.4.1 Overview

For non-transparent *Redundancy*, OPC UA provides the data structures to allow the *Client* to identify what *Servers* are available in the *Redundant Server Set* and also *Server* information which tells the *Client* what modes of *Failover* the *Server* supports. This information allows the *Client* to determine what actions it may need to take in order to accomplish *Failover*. This information is specified in *NonTransparentRedundancyType ObjectType* defined in IEC 62541-5.

Figure 27 shows a typical non-transparent *Redundancy* setup.




Figure 27 – Non-Transparent Redundancy setup

For non-transparent *Redundancy*, the *Servers* will have unique IP addresses. The *Server* also has additional *Failover* modes of *Cold*, *Warm*, *Hot* and *HotAndMirrored*. The *Client* shall be aware of the *Redundant Server Set* and shall be required to perform some actions depending on the *Failover* mode. These actions are described in Table 111 and additional examples and

explanations are provided in 6.6.2.4.5.2 for *Cold*, 6.6.2.4.5.3 for *Warm*, 6.6.2.4.5.4 for *Hot* and 6.6.2.4.5.5 for *HotAndMirrored*.

A *Client* needs to be able to expect that the *SourceTimestamp* associated with a value is approximately the same from all *Servers* in the *Redundant Server Set* for the same value.

6.6.2.4.2 ServiceLevel

The *ServiceLevel* provides information to a *Client* regarding the health of a *Server* and its ability to provide data. See IEC 62541-5 for a formal definition for *ServiceLevel*. The *ServiceLevel* is a byte with a range of 0 to 255, where the values fall into the sub-ranges defined in Table 109.

The algorithm used by a *Server* to determine its *ServiceLevel* within each sub-range is *Server* specific. However, all *Servers* in a *Redundant Server Set* shall use the same algorithm to determine the *ServiceLevel*. All *Servers*, regardless of *Redundant Server Set* membership, shall adhere to the sub-ranges defined in Table 109.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 DRLV

Table 109 – ServiceLevel ranges

Sub-range	Name	Description
0-0	Maintenance	The <i>Failed Server</i> is in maintenance sub-range. Therefore, new <i>Clients</i> shall not connect and currently connected <i>Clients</i> shall disconnect. The <i>Server</i> should expose a target time at which the <i>Clients</i> are able to reconnect. See <i>EstimatedReturnTime</i> defined in IEC 62541-5 for additional information. A <i>Server</i> that has been set to <i>Maintenance</i> is typically undergoing some maintenance or updates. The main goal for the <i>Maintenance ServiceLevel</i> is to ensure that <i>Clients</i> do not generate load on the <i>Server</i> and allow time for the <i>Server</i> to complete any actions that are required. This load includes even simple connections attempts or monitoring of the <i>ServiceLevel</i>. The <i>EstimatedReturnTime</i> indicates when the <i>Client</i> should check to see if the <i>Server</i> is available. If updates or patches are taking longer than expected the <i>Client</i> may discover that the <i>EstimatedReturnTime</i> has been extended further into the future. If the <i>Server</i> does not provide the <i>EstimatedReturnTime</i>, or if the time has lapsed, the <i>Client</i> should use a much longer interval between reconnects to a <i>Server</i> in the <i>Maintenance</i> sub-range than its normal reconnect interval.
1-1	NoData	The <i>Failed Server</i> is not operational. Therefore, a <i>Client</i> is not able to exchange any information with it. The <i>Server</i> most likely has no data other than <i>ServiceLevel</i>, <i>ServerStatus</i> and diagnostic information available. A <i>Failed Server</i> in this sub-range has no data available. <i>Clients</i> may connect to it to obtain <i>ServiceLevel</i>, <i>ServerStatus</i> and other diagnostic information. If the underlying system has failed, typically the <i>ServerStatus</i> would indicate COMMUNICATION_FAULT_6. The <i>Client</i> may monitor this <i>Server</i> for a <i>ServerStatus</i> and <i>ServiceLevel</i> change, which would indicate that normal communication could be resumed.
2-199	Degraded	The <i>Server</i> is partially operational, but is experiencing problems such that portions of the <i>AddressSpace</i> are out of service or unavailable. To understand <i>Client</i> options, see Degraded <i>Servers</i> discussion in this section. An example usage of this <i>ServiceLevel</i> sub-range would be if 3 of 10 devices connected to a <i>Server</i> are unavailable. <i>Servers</i> that report a <i>ServiceLevel</i> in the <i>Degraded</i> sub-range are partially able to service <i>Client</i> requests. The degradation could be caused by loss of connection to underlying systems. Alternatively, it could be that the <i>Server</i> is overloaded to the point that it is unable to reliably deliver data to <i>Clients</i> in a timely manner. If <i>Clients</i> are experiencing difficulties obtaining required data, they shall switch to another <i>Server</i> if any <i>Servers</i> in the <i>Healthy</i> range are available. If no <i>Servers</i> are available in the <i>Healthy</i> range, then <i>Clients</i> may switch to a <i>Server</i> with a higher <i>ServiceLevel</i> or one that provides the required data. Some <i>Clients</i> may also be configured for higher priority data and may check all <i>Degraded Servers</i>, to see if any of the <i>Servers</i> are able to report as good quality the high priority data, but this functionality would be <i>Client</i> specific. In some cases a <i>Client</i> may connect to multiple <i>Degraded Servers</i> to maximize the available information.
200-255	Healthy	The <i>Server</i> is fully operational. Therefore, a <i>Client</i> can obtain all information from this <i>Server</i>. The sub-range allows a <i>Server</i> to provide information that can be used by <i>Clients</i> to load balance. An example usage of this <i>ServiceLevel</i> sub-range would be to reflect the <i>Server's</i> CPU load where data is delivered as expected. <i>Servers</i> in the <i>Healthy ServiceLevel</i> sub-range are able to deliver information in a timely manner. This <i>ServiceLevel</i> may change for an internal <i>Server</i> reason or it may be used for load balancing described in 6.6.2.4.3. <i>Clients</i> shall connect to the <i>Server</i> with the highest <i>ServiceLevel</i>. Once connected, the <i>ServiceLevel</i> may change, but a <i>Client</i> shall not <i>Failover</i> to a different <i>Server</i> as long as the <i>ServiceLevel</i> of the <i>Server</i> is accessible and in the <i>Healthy</i> sub-range.

6.6.2.4.3 Load balancing

In systems where multiple Hot *Servers* (see 6.6.2.4.5.4) are available, the *Servers* in the *Redundant Server Set* can share the load generated by *Clients* by setting the *ServiceLevel* in the *Healthy* sub-range based on the current load. *Clients* are expected to connect to the *Server* with the highest *ServiceLevel*. *Clients* shall not *Failover* to a different *Server* in the *Redundant Server Set* of *Servers* as long as the *Server* is in the *Healthy* sub-range. This is the normal behaviour for all *Clients*, when communicating with redundant *Servers*. *Servers* can adjust their *ServiceLevel* based on the number of *Clients* that are connected, CPU loading, memory utilization, or any other *Server* specific criteria.

For example in a system with 3 *Servers*, all *Servers* are initially at *ServiceLevel* 255, but when a *Client* connects, the *Server* with the *Client* connection sets its level to 254. The next *Client* would connect to a different *Server* since both of the other *Servers* are still at 255.

It is up to the *Server* vendor to define the logic for spreading the load and the number of expected *Clients*, CPU load or other criteria on each *Server* before the *ServiceLevel* is decremented. It is envisioned that some *Servers* would be able to accomplish this without any communication between the *Servers*.

6.6.2.4.4 Server Failover modes

The *Failover* mode of a *Server* is provided in the *ServerRedundancy Object* defined in IEC 62541-5. The different Failover modes for non-transparent Redundancy are described in Table 110.

Table 110 – Server Failover modes

Name	Description
Cold	<i>Cold Failover</i> mode is where only one <i>Server</i> can be active at a time. This may mean that redundant <i>Servers</i> are unavailable (not powered up) or are available but not running (PC is running, but application is not started)
Warm	<i>Warm Failover</i> mode is where the backup <i>Server(s)</i> can be active, but cannot connect to actual data points (typically, a system where the underlying devices are limited to a single connection). Underlying devices, such as PLCs, may have limited resources that permit a single <i>Server</i> connection. Therefore, only a single <i>Server</i> will be able to consume data. The <i>ServiceLevel Variable</i> defined in IEC 62541-5 indicates the ability of the <i>Server</i> to provide its data to the <i>Client</i> .
Hot	<i>Hot Failover</i> mode is where all <i>Servers</i> are powered-on, and are up and running. In scenarios where <i>Servers</i> acquire data from a downstream device, such as a PLC, then one or more <i>Servers</i> are actively connected to the downstream device(s) in parallel. These <i>Servers</i> have minimal knowledge of the other <i>Servers</i> in their group and are independently functioning. When a <i>Server</i> fails or encounters a serious problem then its <i>ServiceLevel</i> drops. On recovery, the <i>Server</i> returns to the <i>Redundant Server Set</i> with an appropriate <i>ServiceLevel</i> to indicate that it is available.
HotAndMirrored	<i>HotAndMirrored Failover</i> mode is where <i>Failovers</i> are for <i>Servers</i> that are mirroring their internal states to all <i>Servers</i> in the <i>Redundant Server Set</i> and more than one <i>Server</i> can be active and fully operational. Mirroring state minimally includes <i>Sessions</i> , <i>Subscriptions</i> , registered <i>Nodes</i> , <i>ContinuationPoints</i> , sequence numbers, and sent <i>Notifications</i> . The <i>ServiceLevel Variable</i> defined in IEC 62541-5 should be used by the <i>Client</i> to find the <i>Servers</i> with the highest <i>ServiceLevel</i> to achieve load balancing.

6.6.2.4.5 Client Failover behaviour

6.6.2.4.5.1 General

Each *Server* maintains a list of *ServerUris* for all redundant *Servers* in the *Redundant Server Set*. The list is provided together with the *Failover* mode in the *ServerRedundancy Object* defined in IEC 62541-5. To enable *Clients* to connect to all *Servers* in the list, each *Server* in the list shall provide the *ApplicationDescription* for all *Servers* in the *Redundant Server Set* through the *FindServers Service*. This information is needed by the *Client* to translate the *ServerUri* into information needed to connect to the other *Servers* in the *Redundant Server Set*. Therefore a *Client* needs to connect to only one of the redundant *Servers* to find the other *Servers* based on the provided information. A *Client* should persist information about other *Servers* in the *Redundant Server Set*.

Table 111 defines a list of *Client* actions for initial connections and *Failovers*.

Table 111 – Redundancy Failover actions

Failover mode and <i>Client</i> options	Cold	Warm	Hot (a)	Hot (b)	HotAndMirrored
On initial connection in addition to actions on Active Server:					
Connect to more than one OPC UA Server.		X	X	X	Optional for status check
Create <i>Subscriptions</i> and add monitored items.		X	X	X	
Activate sampling on the <i>Subscriptions</i> .			X	X	
Activate publishing.				X	
At Failover:					
OpenSecureChannel to backup OPC UA Server	X				X
CreateSession on backup OPC UA Server	X				
ActivateSession on backup OPC UA Server	X				X
Create <i>Subscriptions</i> and add monitored items.	X				
Activate sampling on the <i>Subscriptions</i> .	X	X			
Activate publishing.	X	X	X		

Clients communicating with a non-transparent *Redundant Server Set* of *Servers* require some additional logic to be able to handle *Server* failures and to *Failover* to another *Server* in the *Redundant Server Set*. Figure 28 provides an overview of the steps a *Client* typically performs when it is first connecting to a *Redundant Server Set*. The figure does not cover all possible error scenarios.




Figure 28 – Client Start-up steps

The initial *Server* may be obtained via standard discovery or from a persisted list of *Servers* in the *Redundant Server Set*. But in any case the *Client* needs to check which *Server* in the *Server* set it should connect to. Individual actions will depend on the *Server Failover* mode the *Server* provides and the *Failover* mode the *Client* will use.

Clients once connected to a redundant *Server* shall be aware of the modes of *Failover* supported by a *Server* since this support affects the available options related to *Client* behaviour. A *Client* may always treat a *Server* using a lesser *Failover* mode, i.e. for a *Server* that provides *Hot Redundancy*, a *Client* might connect and choose to treat it as if the *Server* was running in *Warm Redundancy* or *Cold Redundancy*. This choice is up to the client. In the case of *Failover* mode *HotAndMirrored*, the *Client* shall not use *Failover* mode *Hot* or *Warm* as it would generate unnecessary load on the *Servers*.

6.6.2.4.5.2 Cold

A *Cold Failover* mode is where the *Client* can only connect to one *Server* at a time. When the *Client* loses connectivity with the *Active Server* it will attempt a connection to the redundant *Server(s)* which may or may not be available. In this situation the *Client* may need to wait for the redundant *Server* to become available and then create *Subscriptions* and *MonitoredItems* and activate publishing. The *Client* shall cache any information that is required related to the list of available *Servers* in the *Redundant Server Set*. Figure 29 illustrates the action a *Client* would take if it is talking to a *Server* using *Cold Failover* mode.




Figure 29 – Cold Failover

NOTE There may be a loss of data from the time the connection to the *Active Server* is interrupted until the time the *Client* gets *Publish Responses* from the backup *Server*.

6.6.2.4.5.3 Warm

A *Warm Failover* mode is where the *Client* should connect to one or more *Servers* in the *Redundant Server Set* primarily to monitor the *ServiceLevel*. A *Client* can connect and create *Subscriptions* and *MonitoredItems* on more than one *Server*, but sampling and publishing can only be active on one *Server*. However, the active *Server* will return actual data, whereas the other *Servers* in the *Redundant Server Set* will return an appropriate error for the *MonitoredItems* in the *Publish* response such as *Bad_NoCommunication*. The one *Active Server* can be found by reading the *ServiceLevel Variable* from all *Servers*. The *Server* with the highest *ServiceLevel* is the *Active Server*. For *Failover* the *Client* activates sampling and publishing on the *Server* with the highest *ServiceLevel*. Figure 30 illustrates the steps a *Client* would perform when communicating with a *Server* using *Warm Failover* mode.




Figure 30 – Warm Failover

NOTE There may be a temporary loss of data from the time the connection to the *Active Server* is interrupted until the time the *Client* gets *Publish Responses* from the backup *Server*.

6.6.2.4.5.4 Hot

A *Hot Failover* mode is where the *Client* should connect to two or more *Servers* in the *Redundant Server Set* and to subscribe to the *ServiceLevel* variable defined in IEC 62541-5 to find the highest *ServiceLevel* to achieve load balancing; this means that *Clients* should issue *Service* requests such as *Browse*, *Read*, *Write* to the *Server* with the highest *ServiceLevel*. *Subscription* related activities will need to be invoked for each connected *Server*. *Clients* have the following choices for implementing subscription behaviour in a *Hot Failover* mode:

- The *Client* connects to multiple *Servers* and establishes subscription(s) in each where only one is *Reporting*; the others are *Sampling* only. The *Client* should setup the queue size for the *MonitoredItems* such that it can buffer all changes during the *Failover* time. The *Failover* time is the time between the connection interruption and the time the *Client* gets *Publish Responses* from the backup *Server*. On a *Failover* the *Client* shall enable *Reporting* on the *Server* with the next highest availability.
- The *Client* connects to multiple *Servers* and establishes subscription(s) in each where all subscriptions are *Reporting*. The *Client* is responsible for handling/processing multiple subscription streams concurrently.

Figure 31 illustrate the functionality a *Client* would perform when communicating with a *Server* using *Hot Failover* mode (Figure 31 includes both (a) and (b) options).




Figure 31 – Hot Failover

Clients are not expected to automatically switch over to a *Server* that has recovered from a failure, but the *Client* should establish a connection to it.

6.6.2.4.5.5 HotAndMirrored

A *HotAndMirrored Failover* mode is where a *Client* only connects to one *Server* in the *Redundant Server Set* because the *Server* will share this session/state information with the other *Servers*. In order to validate the capability to connect to other redundant *Servers* it is allowed to create *Sessions* with other *Servers* and maintain the open connections by periodically reading the *ServiceLevel*. A *Client* shall not create *Subscriptions* on the backup *Servers* for status monitoring (to prevent excessive load on the *Servers*). This mode allows *Clients* to fail over without creating a new context for communication. On a *Failover* the *Client* will simply create a new *SecureChannel* on an alternate *Server* and then call *ActivateSession*; all *Client* activities (browsing, subscriptions, history reads, etc.) will then resume. Figure 32

illustrates the behaviour a *Client* would perform when communicating to a *Server* in *HotAndMirrored Failover* mode.




Figure 32 – HotAndMirrored Failover

This *Failover* mode is similar to the transparent *Redundancy*. The advantage is that the *Client* has full control over selecting the *Server*. The disadvantage is that the *Client* needs to be able to handle *Failovers*.

6.6.2.5 Hiding Failover with a Server Proxy

A vendor can use the non-transparent *Redundancy* features to create a *Server* proxy running on the *Client* machine to provide transparent *Redundancy* to the client. This reduces the amount of functionality that needs to be designed into the *Client* and to enable simpler *Clients* to take advantage of non-transparent *Redundancy*. The *Server* proxy simply duplicates *Subscriptions* and modifications to *Subscriptions*, by passing the calls on to both *Servers*, but only enabling publishing and sampling on one *Server*. When the proxy detects a failure, it enables publishing and/or sampling on the backup *Server*, just as the *Client* would if it were a *Redundancy* aware *Client*.

Figure 33 shows the *Server* proxy used to provide transparent *Redundancy*.




Figure 33 – Server proxy for Redundancy

6.6.3 Client Redundancy

Client Redundancy is supported in OPC UA by the *TransferSubscriptions Service* and by exposing *Client* information in the *Server* diagnostic information. Since *Subscription* lifetime is not tied to the *Session* in which it was created, backup *Clients* may use standard diagnostic information available to monitor the active *Client's Session* with the *Server*. Upon detection of an active *Client* failure, a backup *Client* would then instruct the *Server* to transfer the *Subscriptions* to its own session. If the *Subscription* is crafted carefully, with sufficient resources to buffer data during the change-over, data loss from a *Client Failover* can be prevented.

OPC UA does not provide a standardized mechanism for conveying the *SessionId* and *SubscriptionIds* from the active *Client* to the backup *Clients*, but as long as the backup *Clients* know the *Client* name of the active *Client*, this information is readily available using the *SessionDiagnostics* and *SubscriptionDiagnostics* portions of the *ServerDiagnostics* data. This information is available for authorized users and for the user active on the *Session*. *TransferSubscriptions* requires the same user on all redundant *Clients* to succeed.

6.6.4 Network Redundancy

6.6.4.1 Overview

Redundant networks can be used with OPC UA in either transparent or non-transparent *Redundancy*.

Network Redundancy can be combined with *Server* and *Client Redundancy*.

6.6.4.2 Transparent

In the transparent network use-case a single *Server Endpoint* can be reached through different network paths. This case is completely handled by the network infrastructure. The selected network path and *Failover* are transparent to the *Client* and the *Server*. Transparent network *Redundancy* is illustrated in Figure 34.





Figure 34 – Transparent network Redundancy

Examples:

- A physical appliance/device such as a router or gateway which automatically changes the network routing to maintain communications.
- A virtual adapter which automatically changes the network adapter to maintain communications.

6.6.4.3 Non-transparent

In the non-transparent network use-case the *Server* provides different *Endpoints* for the different network paths. This requires both the *Server* and the *Client* to support multiple network connections. In this case the *Client* is responsible for selecting the *Endpoint* and for *Failover*. For *Failover* the normal reconnect scenario described in 6.7 can be used. Only the *SecureChannel* is created with another *Endpoint*. *Sessions* and *Subscriptions* can be reused. Non-transparent network *Redundancy* is illustrated in Figure 35.



IEC

Figure 35 – Non-transparent network Redundancy

The information about the different network paths is specified in *NonTransparentRedundancyType* *ObjectType* defined in IEC 62541-5.

6.6.5 Manually forcing Failover

In redundant systems, it is common to require that a particular *Server* in the *Redundant Server Set* be taken out of the *Redundant Server Set* for a period of time. Some items that could cause this may include:

- certificate update;
- security reconfiguration;
- rebooting or restarting of the machine for
 - software updates and patches,
 - installation of new software;
- reconfiguration of the *AddressSpace*.

The removal from the *Redundant Server Set* can be done through a complete shutdown or by setting the *ServiceLevel* of the *Server* to *Maintenance* sub-range. This can be done through a *Server* specific configuration tool or through the *Method RequestServerStateChange* on the *ServerRedundancyType*. The Method is formally defined in IEC 62541-5.

This *Method* requires that the *Client* provide credentials with administrative rights on the *Server*.

6.7 Re-establishing connections

After a *Client* establishes a connection to a *Server* and creates a *Subscription*, the *Client* monitors the connection status. Figure 36 shows the steps to connect a *Client* to a *Server* and the general logic for reconnect handling. Not all possible error scenarios are covered.

The preferred mechanism for a *Client* to monitor the connection status is through the keep-alive of the *Subscription*. A *Client* should subscribe for the *State Variable* in the *ServerStatus*

to detect shutdown or other failure states. If no *Subscription* is created or the *Server* does not support *Subscriptions*, the connection can be monitored by periodically reading the *State Variable*.




Figure 36 – Reconnect sequence

When a *Client* loses the connection to the *Server*, the goal is to reconnect without losing information. To do this the *Client* shall re-establish the connection by creating a new *SecureChannel* and activating the *Session* with the *Service ActivateSession*. This assigns the new *SecureChannel* to the existing *Session* and allows the *Client* to reuse the *Session* and *Subscriptions* in the *Server*. To re-establish the *SecureChannel* and activate the *Session*, the *Client* shall use the same security policy, application instance certificate and the same user credential used to create the original *SecureChannel*. This will result in the *Client* receiving data and event *Notifications* without losing information provided the queues in the *MonitoredItems* do not overflow.

The *Client* shall only create a new *Session* if *ActivateSession* fails. *TransferSubscriptions* is used to transfer the *Subscription* to the new *Session*. If *TransferSubscriptions* fails, the *Client* needs to create a new *Subscription*.

When the connection is lost, *Publish* responses may have been sent but not received by the *Client*.

After re-establishing the connection the *Client* shall call *Republish* in a loop, starting with the next expected sequence number and incrementing the sequence number until the *Server*

returns the status *Bad_MessageNotAvailable*. After receiving this status, the *Client* shall start sending *Publish* requests with the normal *Publish* handling. This sequence ensures that the lost *NotificationMessages* queued in the *Server* are not overwritten by new *Publish* responses.

If the *Client* detects missing sequence numbers in the *Publish* and is not able to get the lost *NotificationMessages* through *Republish*, the *Client* should use the *Method ResendData* or should read the values of all data *MonitoredItems* to make sure the *Client* has the latest values for all *MonitoredItems*.

The *Server Object* provides a *Method ResendData* that initiates resending of all data monitored items in a *Subscription*. This *Method* is defined in IEC 62541-5. If this *Method* is called, subsequent *Publish* responses shall contain the current values of all data *MonitoredItems* in the *Subscription* where the *MonitoringMode* is set to *Reporting*. If a value is queued for a data *MonitoredItem*, the next value in the queue is sent in the *Publish* response. If no value is queued for a data *MonitoredItem*, the last value sent is repeated in the *Publish* response. The *Server* shall verify that the *Method* is called within the *Session* context of the *Session* that owns the *Subscription*.

Independent of the detailed recovery strategy, the *Client* should make sure that it does not overwrite newer data in the *Client* with older values provided through *Republish*.

If the *Republish* returns *Bad_SubscriptionIdInvalid*, then the *Client* needs to create a new *Subscription*.

Re-establishing the connection by creating a new *SecureChannel* may be rejected, because of a new *Server Application Instance Certificate* or other security errors. In case of security failures, the *Client* shall use the *GetEndpoints Service* to fetch the most up to date security information from the *Server*.

IEC 62541-6 defines a reverse connect mechanism where the *Server* initiates the logical connection. All subsequent steps like creating a *SecureChannel* are initiated by the *Client*. In this scenario the *Client* is only able to initiate a reconnect if the *Server* initiates a new logical connection after a connection interruption. The *Client* side reconnect handling described in Figure 36 applies also to the reverse connect case. A *Server* is not able to actively check the connection status; therefore the *Server* shall initiate a new connection in a configurable interval, even if a connection to the *Client* is established. This ensures that an initiated connection is available for the reconnect handling in addition to other scenarios where the *Client* needs more than one connection.

6.8 Durable Subscriptions

MonitoredItems are used to monitor *Variable Values* for data changes and event notifier *Objects* for new *Events*. Subscriptions are used to combine data changes and events of the assigned *MonitoredItems* to an optimized stream of network messages. A reliable delivery is ensured as long as the lifetime of the *Subscription* and the queues in the *MonitoredItems* are long enough for a network interruption between OPC UA *Client* and *Server*. All queues that ensure reliable delivery are normally kept in memory and a *Server* restart would delete them.

There are use cases where OPC UA *Clients* have no permanent network connection to the OPC UA *Server* or where reliable delivery of data changes and events is necessary even if the OPC UA *Server* is restarted or the network connection is interrupted for a longer time.

To ensure this reliable delivery, the OPC UA *Server* shall store collected data and events in non-volatile memory until the OPC UA *Client* has confirmed reception. It is possible that there will be data lost if the *Server* is not shut down gracefully or in case of power failure. But the OPC UA *Server* should store the queues frequently even if the *Server* is not shut down.

The *Method SetSubscriptionDurable* defined in IEC 62541-5 is used to set a *Subscription* into this durable mode and to allow much longer lifetimes and queue sizes than for normal

Subscriptions. The *Method* shall be called before the *MonitoredItems* are created in the durable *Subscription*. The *Server* shall verify that the *Method* is called within the *Session* context of the *Session* that owns the *Subscription*.

A value of 0 for the parameter *lifetimeInHours* requests the highest lifetime supported by the *Server*.

An OPC UA *Server* providing durable *Subscriptions* shall:

- support the *SetSubscriptionDurable Method* defined in IEC 62541-5;
- support *Service TransferSubscriptions*;
- support long *Subscription* lifetimes, minimum requirement is defined in IEC 62541-7;
- support large *MonitoredItem* queues, minimum requirement is defined in IEC 62541-7;
- store *Subscriptions* settings and sent notification messages with sequence numbers;
- store *MonitoredItem* settings and queues.

An OPC UA *Client* using durable *Subscriptions* shall:

- use the *SetSubscriptionDurable Method* defined in IEC 62541-5 to create a durable *Subscription*;
- close *Sessions* for planned communication interruptions;
- use the *Service TransferSubscriptions* to assign the durable *Subscription* to a new *Session* for data transfer;
- store *SubscriptionId*, *MonitoredItem* client and server handles and the last confirmed sequence number.

7 Common parameter type definitions

7.1 ApplicationDescription

The components of this parameter are defined in Table 112.

Table 112 – ApplicationDescription

Name	Type	Description
ApplicationDescription	structure	Specifies an application that is available.
applicationUri	String	The globally unique identifier for the application instance. This URI is used as <i>ServerUri</i> in <i>Services</i> if the application is a <i>Server</i> .
productUri	String	The globally unique identifier for the product.
applicationName	LocalizedText	A localized descriptive name for the application.
applicationType	Enum ApplicationType	The type of application. This value is an enumeration with one of the following values: SERVER_0 The application is a <i>Server</i>. CLIENT_1 The application is a <i>Client</i>. CLIENTANDSERVER_2 The application is a <i>Client</i> and a <i>Server</i>. DISCOVERYSERVER_3 The application is a <i>DiscoveryServer</i>.
gatewayServerUri	String	A URI that identifies the <i>Gateway Server</i> associated with the <i>discoveryUrls</i> . This value is not specified if the <i>Server</i> can be accessed directly. This field is not used if the <i>applicationType</i> is CLIENT_1.
discoveryProfileUri	String	A URI that identifies the discovery profile supported by the URLs provided. This field is not used if the <i>applicationType</i> is CLIENT_1. If this value is not specified then the Endpoints shall support the Discovery Services defined in 5.4. Alternate discovery profiles are defined in IEC 62541-7.
discoveryUrls []	String	A list of URLs for the <i>DiscoveryEndpoints</i> provided by the application. If the <i>applicationType</i> is CLIENT_1, this field shall contain an empty list.

7.2 ApplicationInstanceCertificate

An *ApplicationInstanceCertificate* is a *ByteString* containing an encoded *Certificate*. The encoding of an *ApplicationInstanceCertificate* depends on the security technology mapping and is defined completely in IEC 62541-6. Table 113 specifies the information that shall be contained in an *ApplicationInstanceCertificate*.

Table 113 – ApplicationInstanceCertificate

Name	Type	Description
ApplicationInstanceCertificate	structure	<i>ApplicationInstanceCertificate</i> with signature created by a <i>Certificate Authority</i> .
version	String	An identifier for the version of the <i>Certificate</i> encoding.
serialNumber	ByteString	A unique identifier for the <i>Certificate</i> assigned by the Issuer.
signatureAlgorithm	String	The algorithm used to sign the <i>Certificate</i> . The syntax of this field depends on the <i>Certificate</i> encoding.
signature	ByteString	The signature created by the Issuer.
issuer	Structure	A name that identifies the <i>Issuer Certificate</i> used to create the signature.
validFrom	UtcTime	When the <i>Certificate</i> becomes valid.
validTo	UtcTime	When the <i>Certificate</i> expires.
subject	Structure	A name that identifies the application instance that the <i>Certificate</i> describes. This field shall contain the <i>productName</i> and the name of the organization responsible for the application instance.
applicationUri	String	The <i>applicationUri</i> specified in the <i>ApplicationDescription</i> . The <i>ApplicationDescription</i> is described in 7.1.
hostnames []	String	The name of the machine where the application instance runs. A machine may have multiple names if it is accessible via multiple networks. The hostname may be a numeric network address or a descriptive name. <i>Server Certificates</i> shall have at least one hostname defined.
publicKey	ByteString	The public key associated with the <i>Certificate</i> .
keyUsage []	String	Specifies how the <i>Certificate</i> key may be used. <i>ApplicationInstanceCertificates</i> shall support Digital Signature, Non-Repudiation Key Encryption, Data Encryption and Client/Server Authorization. The contents of this field depend on the <i>Certificate</i> encoding.

7.3 BrowseResult

The components of this parameter are defined in Table 114.

Table 114 – BrowseResult

Name	Type	Description
BrowseResult	structure	The results of a Browse operation.
statusCode	StatusCode	The status for the <i>BrowseDescription</i> . This value is set to <i>Good</i> if there are still references to return for the <i>BrowseDescription</i> .
continuationPoint	ContinuationPoint	A <i>Server</i> defined opaque value that identifies the continuation point. The <i>ContinuationPoint</i> type is defined in 7.6.
References []	ReferenceDescription	The set of references that meet the criteria specified in the <i>BrowseDescription</i> . Empty, if no <i>References</i> met the criteria. The Reference Description type is defined in 7.25.

7.4 ContentFilter

7.4.1 ContentFilter structure

The ContentFilter structure defines a collection of elements that define filtering criteria. Each element in the collection describes an operator and an array of operands to be used by the operator. The operators that can be used in a ContentFilter are described in Table 119. The filter is evaluated by evaluating the first entry in the element array starting with the first operand in the operand array. The operands of an element may contain References to sub-elements resulting in the evaluation continuing to the referenced elements in the element array. The evaluation shall not introduce loops. For example evaluation starting from element "A" shall never be able to return to element "A". However there may be more than one path leading to another element "B". If an element cannot be traced back to the starting element it is ignored. Extra operands for any operator shall result in an error. Annex B provides examples using the ContentFilter structure.

Table 115 defines the *ContentFilter* structure.

Table 115 – ContentFilter structure

Name	Type	Description
ContentFilter	structure	
elements []	ContentFilterElement	List of operators and their operands that compose the filter criteria. The filter is evaluated by starting with the first entry in this array. This structure is defined in-line with the following indented items.
filterOperator	Enum FilterOperator	Filter operator to be evaluated. The <i>FilterOperator</i> enumeration is defined in Table 119.
filterOperands []	Extensible Parameter FilterOperand	Operands used by the selected operator. The number and use depend on the operators defined in Table 119. This array needs at least one entry. This extensible parameter type is the <i>FilterOperand</i> parameter type specified in 7.4.4. It specifies the list of valid <i>FilterOperand</i> values.

7.4.2 ContentFilterResult

The components of this data type are defined in Table 116.

Table 116 – ContentFilterResult structure

Name	Type	Description
ContentFilterResult	structure	A structure that contains any errors associated with the filter.
elementResults []	ContentFilter ElementResult	A list of results for individual elements in the filter. The size and order of the list matches the size and order of the elements in the <i>ContentFilter</i> parameter. This structure is defined in-line with the following indented items.
statusCode	StatusCode	The status code for a single element.
operandStatusCodes []	StatusCode	A list of status codes for the operands in an element. The size and order of the list matches the size and order of the operands in the <i>ContentFilterElement</i> . This list is empty if no operand errors occurred.
operandDiagnosticInfos []	DiagnosticInfo	A list of diagnostic information for the operands in an element. The size and order of the list matches the size and order of the operands in the <i>ContentFilterElement</i> . This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the operands.
elementDiagnosticInfos []	DiagnosticInfo	A list of diagnostic information for individual elements in the filter. The size and order of the list matches the size and order of the elements in the <i>filter</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the elements.

Table 117 defines values for the *StatusCode* parameter that are specific to this structure. Common *StatusCodes* are defined in Table 178.

Table 117 – ContentFilterResult result codes

Symbolic Id	Description
Bad_FilterOperandCountMismatch	The number of operands provided for the filter operator was less than expected for the operand provided.
Bad_FilterOperatorInvalid	An unrecognized operator was provided in a filter.
Bad_FilterOperatorUnsupported	A valid operator was provided, but the <i>Server</i> does not provide support for this filter operator.

Table 118 defines values for the *operandStatusCodes* parameter that are specific to this structure. Common *StatusCodes* are defined in Table 178.

Table 118 – ContentFilterResult Operand result codes

Symbolic Id	Description
Bad_FilterOperandInvalid	See Table 178 for the description of this result code.
Bad_FilterElementInvalid	The referenced element is not a valid element in the content filter.
Bad_FilterLiteralInvalid	The referenced literal is not a valid <i>BaseDataType</i> .
Bad_AttributeIdInvalid	The attribute id is not a valid attribute id in the system.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_NodeIdInvalid	See Table 178 for the description of this result code.
Bad_NodeIdUnknown	See Table 178 for the description of this result code.
Bad_NotTypeDefinition	The provided <i>NodeId</i> was not a type definition <i>NodeId</i> .

7.4.3 FilterOperator

Table 119 defines the basic operators that can be used in a *ContentFilter*. See Table 120 for a description of advanced operators. See 7.4.4 for a definition of operands.

Table 119 – Basic FilterOperator definition

Operator	Number of Operands	Description
Equals_0	2	TRUE if operand[0] is equal to operand[1]. If the operands are of different types, the system shall perform any implicit conversion to a common type. This operator resolves to FALSE if no implicit conversion is available and the operands are of different types. This operator returns FALSE if the implicit conversion fails. See the discussion on data type precedence in Table 123 for more information how to convert operands of different types.
IsNull_1	1	TRUE if operand[0] is a null value.
GreaterThan_2	2	TRUE if operand[0] is greater than operand[1]. The following restrictions apply to the operands: [0]: Any operand that resolves to an ordered value. [1]: Any operand that resolves to an ordered value. The same conversion rules as defined for <i>Equals</i> apply.

Operator	Number of Operands	Description
LessThan_3	2	TRUE if operand[0] is less than operand[1]. The same conversion rules and restrictions as defined for <i>GreaterThan</i> apply.
GreaterThanOrEqual_4	2	TRUE if operand[0] is greater than or equal to operand[1]. The same conversion rules and restrictions as defined for <i>GreaterThan</i> apply.
LessThanOrEqual_5	2	TRUE if operand[0] is less than or equal to operand[1]. The same conversion rules and restrictions as defined for <i>GreaterThan</i> apply.
Like_6	2	TRUE if operand[0] matches a pattern defined by operand[1]. See Table 121 for the definition of the pattern syntax. The following restrictions apply to the operands: [0]: Any operand that resolves to a String. [1]: Any operand that resolves to a String. This operator resolves to FALSE if no operand can be resolved to a string.
Not_7	1	TRUE if operand[0] is FALSE. The following restrictions apply to the operands: [0]: Any operand that resolves to a Boolean. If the operand cannot be resolved to a Boolean, the result is a NULL. See below for a discussion on the handling of NULL.
Between_8	3	TRUE if operand[0] is greater or equal to operand[1] and less than or equal to operand[2]. The following restrictions apply to the operands: [0]: Any operand that resolves to an ordered value. [1]: Any operand that resolves to an ordered value. [2]: Any operand that resolves to an ordered value. If the operands are of different types, the system shall perform any implicit conversion to match all operands to a common type. If no implicit conversion is available and the operands are of different types, the particular result is FALSE. See the discussion on data type precedence in Table 123 for more information how to convert operands of different types.
InList_9	2..n	TRUE if operand[0] is equal to one or more of the remaining operands. The Equals Operator is evaluated for operand[0] and each remaining operand in the list. If any Equals evaluation is TRUE, InList returns TRUE.
And_10	2	TRUE if operand[0] and operand[1] are TRUE. The following restrictions apply to the operands: [0]: Any operand that resolves to a Boolean. [1]: Any operand that resolves to a Boolean. If any operand cannot be resolved to a Boolean it is considered a NULL. See below for a discussion on the handling of NULL.
Or_11	2	TRUE if operand[0] or operand[1] are TRUE. The following restrictions apply to the operands: [0]: Any operand that resolves to a Boolean. [1]: Any operand that resolves to a Boolean. If any operand cannot be resolved to a Boolean it is considered a NULL. See below for a discussion on the handling of NULL.

Operator	Number of Operands	Description
Cast_12	2	Converts operand[0] to a value with a data type with a NodeId identified by operand[1]. The following restrictions apply to the operands: [0]: Any operand. [1]: Any operand that resolves to a NodeId or ExpandedNodeId where the <i>Node</i> is of the <i>NodeClass DataType</i>. If there is any error in conversion or in any of the parameters then the Cast Operation evaluates to a NULL. See below for a discussion on the handling of NULL.
BitwiseAnd_16	2	The result is an integer which matches the size of the largest operand and contains a bitwise And operation of the two operands where both have been converted to the same size (largest of the two operands). The following restrictions apply to the operands: [0]: Any operand that resolves to an integer. [1]: Any operand that resolves to an integer. If any operand cannot be resolved to an integer it is considered a NULL. See below for a discussion on the handling of NULL.
BitwiseOr_17	2	The result is an integer which matches the size of the largest operand and contains a bitwise Or operation of the two operands where both have been converted to the same size (largest of the two operands). The following restrictions apply to the operands: [0]: Any operand that resolves to an integer. [1]: Any operand that resolves to an integer. If any operand cannot be resolved to an integer it is considered a NULL. See below for a discussion on the handling of NULL.

Many operands have restrictions on their type. This requires the operand to be evaluated to determine what the type is. In some cases the type is specified in the operand (i.e. a *LiteralOperand*). In other cases the type requires that the value of an attribute be read. An *ElementOperand* evaluates to a Boolean value unless the operator is a Cast or a nested *RelatedTo* operator.

Table 120 defines complex operators that require a target node (i.e. row) to evaluate. These operators shall be re-evaluated for each possible target node in the result set.

Table 120 – Complex FilterOperator definition

Operator	Number of Operands	Description
InView_13	1	TRUE if the target <i>Node</i> is contained in the <i>View</i> defined by operand[0]. The following restrictions apply to the operands: [0]: Any operand that resolves to a <i>NodeId</i> that identifies a <i>View Node</i>. If operand[0] does not resolve to a <i>NodeId</i> that identifies a <i>View Node</i>, this operation shall always be False.
OfType_14	1	TRUE if the target <i>Node</i> is of type operand[0] or of a subtype of operand[0]. The following restrictions apply to the operands: [0]: Any operand that resolves to a <i>NodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i>. If operand[0] does not resolve to a <i>NodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i>, this operation shall always be False.
RelatedTo_15	6	TRUE if the target <i>Node</i> is of type operand[0] and is related to a <i>NodeId</i> of the type defined in operand[1] by the <i>Reference</i> type defined in operand[2]. operand[0] or operand[1] can also point to an element <i>Reference</i> where the referred to element is another <i>RelatedTo</i> operator. This allows chaining of relationships (e.g. A is related to B is related to C), where the relationship is defined by the <i>ReferenceType</i> defined in operand[2]. In this case, the referred to element returns a list of <i>NodeIds</i> instead of TRUE or FALSE. In this case if any errors occur or any of the operands cannot be resolved to an appropriate value, the result of the chained relationship is an empty list of nodes. Operand[3] defines the number of hops for which the relationship should be followed. If operand[3] is 1, then objects shall be directly related. If a hop is greater than 1, then a <i>NodeId</i> of the type described in operand[1] is checked for at the depth specified by the hop. In this case, the type of the intermediate <i>Node</i> is undefined, and only the <i>Reference</i> type used to reach the end <i>Node</i> is defined. If the requested number of hops cannot be followed, then the result is FALSE, i.e., an empty <i>Node</i> list. If operand[3] is 0, the relationship is followed to its logical end in a forward direction and each <i>Node</i> is checked to be of the type specified in operand[1]. If any <i>Node</i> satisfies this criterion, then the result is TRUE, i.e., the <i>NodeId</i> is included in the sub-list. Operand [4] defines if operands [0] and [1] should include support for subtypes of the types defined by these operands. A TRUE indicates support for subtypes operand [5] defines if operand [2] should include support for subtypes of the reference type. A TRUE indicates support for subtypes. The following restrictions apply to the operands: [0]: Any operand that resolves to a <i>NodeId</i> or <i>ExpandedNodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i> or a reference to another element which is a <i>RelatedTo</i> operator. [1]: Any operand that resolves to a <i>NodeId</i> or <i>ExpandedNodeId</i> that identifies an <i>ObjectType</i> or <i>VariableType Node</i> or a reference to another element which is a <i>RelatedTo</i> operator. [2]: Any operand that resolves to a <i>NodeId</i> that identifies a <i>ReferenceType Node</i>. [3]: Any operand that resolves to a value implicitly convertible to UInt32. [4]: Any operand that resolves to a value implicitly convertible to a Boolean; if this operand does not resolve to a Boolean, then a value of FALSE is used. [5]: Any operand that resolves to a value implicitly convertible to a Boolean; if this operand does not resolve to a Boolean, then a value of FALSE is used. If none of the operands [0],[1],[2],[3] resolves to an appropriate value then the result of this operation shall always be False (or an Empty set in the case of a nested <i>RelatedTo</i> operand). See examples for <i>RelatedTo</i> in B.2.

The *RelatedTo* operator can be used to identify if a given type, set as operand[1], is a subtype of another type set as operand[0] by setting operand[2] to the *HasSubtype ReferenceType* and operand[3] to 0.

The *Like* operator can be used to perform wildcard comparisons. Several special characters can be included in the second operand of the *Like* operator. The valid characters are defined in Table 121. The wildcard characters can be combined in a single string (i.e. 'Th[ia][ts]%' would match 'That is fine', 'This is fine', 'That as one', 'This it is', 'Then at any', etc.). The *Like* operator is case sensitive.

Table 121 – Wildcard characters

Special Character	Description
%	Match any string of zero or more characters (i.e. 'main%' would match any string that starts with 'main', '%en%' would match any string that contains the letters 'en' such as 'entail', 'green' and 'content'.) If a '%' sign is intend in a string the list operand can be used (i.e. 5[%] would match '5%').
_	Match any single character (i.e. '_ould' would match 'would', 'could'). If the '_' is intended in a string then the list operand can be used (i.e. 5[_] would match '5_').
\	Escape character allows literal interpretation (i.e. \\ is \, \% is %, _ is _)
[]	Match any single character in a list (i.e. 'abc[13-68]' would match 'abc1', 'abc3', 'abc4', 'abc5', 'abc6', and 'abc8'. 'xyz[c-f]' would match 'xyzc', 'xyzd', 'xyze', 'xyzf').
[^]	Not Matching any single character in a list. The ^ shall be the first character inside on the []. (i.e. 'ABC[^13-5]' would NOT match 'ABC1', 'ABC3', 'ABC4', and 'ABC5'. xyz[^dgh] would NOT match 'xyzd', 'xyzg', 'xyzh'.)

Table 122 defines the conversion rules for the operand values. The types are automatically converted if an implicit conversion exists (I). If an explicit conversion exists (E) then type can be converted with the cast operator. If no conversion exists (X) the then types cannot be converted, however, some servers may support application specific explicit conversions. The types used in the table are defined in IEC 62541-3. A data type that is not in the table does not have any defined conversions.

Table 122 – Conversion rules

Source type (from)	Target type (to)																				
	Boolean	Byte	ByteString	DateTime	Double	ExpandedNodeId	Float	Guid	Int16	Int32	Int64	NodeId	SByte	StatusCode	String	LocalizedText	QualifiedName	UInt16	UInt32	UInt64	XmlElement
Boolean	-	I	X	X	I	X	I	X	I	I	I	X	I	X	E	X	X	I	I	I	X
Byte	E	-	X	X	I	X	I	X	I	I	I	X	I	X	E	X	X	I	I	I	X
ByteString	X	X	-	X	X	X	X	E	X	X	X	X	X	X	X	X	X	X	X	X	X
DateTime	X	X	X	-	X	X	X	X	X	X	X	X	X	X	E	X	X	X	X	X	X
Double	E	E	X	X	-	X	E	X	E	E	E	X	E	X	E	X	X	E	E	E	X
ExpandedNodeId	X	X	X	X	X	-	X	X	X	X	X	E	X	X	I	X	X	X	X	X	X
Float	E	E	X	X	I	X	-	X	E	E	E	X	E	X	E	X	X	E	E	E	X
Guid	X	X	E	X	X	X	X	-	X	X	X	X	X	X	E	X	X	X	X	X	X
Int16	E	E	X	X	I	X	I	X	-	I	I	X	E	X	E	X	X	E	I	I	X
Int32	E	E	X	X	I	X	I	X	E	-	I	X	E	E	E	X	X	E	E	I	X
Int64	E	E	X	X	I	X	I	X	E	E	X	X	E	E	E	X	X	E	E	E	X
NodeId	X	X	X	X	X	I	X	X	X	X	X	-	X	X	I	X	X	X	X	X	X
SByte	E	E	X	X	I	X	I	X	I	I	I	X	-	X	E	X	X	I	I	I	X
StatusCode	X	X	X	X	X	X	X	X	X	I	I	X	X	-	X	X	X	E	I	I	X
String	I	I	X	E	I	E	I	I	I	I	I	E	I	X	-	E	E	I	I	I	X
LocalizedText	X	X	X	X	X	X	X	X	X	X	X	X	X	X	I	-	X	X	X	X	X
QualifiedName	X	X	X	X	X	X	X	X	X	X	X	X	X	X	I	I	-	X	X	X	X
UInt16	E	E	X	X	I	X	I	X	I	I	I	X	E	I	E	X	X	-	I	I	X
UInt32	E	E	X	X	I	X	I	X	E	I	I	X	E	E	E	X	X	E	-	I	X
UInt64	E	E	X	X	I	X	I	X	E	E	I	X	E	E	E	X	X	E	E	-	X
XmlElement	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	-

Arrays of a source type can be converted to arrays of the target type by converting each element. A conversion error for any element causes the entire conversion to fail.

Arrays of length 1 can be implicitly converted to a scalar value of the same type.

Guid, *NodeId* and *ExpandedNodeId* are converted to and from *String* using the syntax defined in IEC 62541-6.

Floating point values are rounded by adding 0,5 and truncating when they are converted to integer values.

Converting a negative value to an unsigned type causes a conversion error. If the conversion fails the result is a null value.

Converting a value that is outside the range of the target type causes a conversion error. If the conversion fails the result is a null value.

ByteString is converted to *String* by formatting the bytes as a sequence of hexadecimal digits.

LocalizedText values are converted to *Strings* by dropping the *Locale*. *Strings* are converted to *LocalizedText* values by setting the *Locale* to "".

QualifiedName values are converted to *Strings* by dropping the *NamespaceIndex*. *Strings* are converted to *QualifiedName* values by setting the *NamespaceIndex* to 0.

A *StatusCode* can be converted to and from a *UInt32* and *Int32* by copying the bits. Only the top 16-bits of the *StatusCode* are copied when it is converted to and from a *UInt16* or *Int16* value.

Boolean values are converted to '1' when true and '0' when false. Non zero numeric values are converted to true *Boolean* values. Numeric values of 0 are converted to false *Boolean* values. *String* values containing "true", "false", "1" or "0" can be converted to *Boolean* values. Other string values cause a conversion error. In this case *Strings* are case-insensitive.

It is sometimes possible to use implicit casts when operands with different data types are used in an operation. In this situation the precedence rules defined in Table 123 are used to determine which implicit conversion to use. The first data type in the list (top down) has the most precedence. If a data type is not in this table then it cannot be converted implicitly while evaluating an operation.

For example, assume that A = 1,1 (*Float*) and B = 1 (*Int32*) and that these values are used with an *Equals* operator. This operation would be evaluated by casting the *Int32* value to a *Float* since the *Float* data type has more precedence.

Table 123 – Data Precedence rules

Rank	Data Type
1	Double
2	Float
3	Int64
4	UInt64
5	Int32
6	UInt32
7	StatusCode
8	Int16
9	UInt16
10	SByte
11	Byte
12	Boolean
13	Guid
14	String
15	ExpandedNodeId
16	NodeId
17	LocalizedText
18	QualifiedName

Operands may contain null values (i.e. values which do not exist). When this happens, the element always evaluates to NULL (unless the *IsNull_1* operator has been specified). Table 124 defines how to combine elements that evaluate to NULL with other elements in a logical AND operation.

Table 124 – Logical AND Truth table

	TRUE	FALSE	NULL
TRUE	TRUE	FALSE	NULL
FALSE	FALSE	FALSE	FALSE
NULL	NULL	FALSE	NULL

Table 125 defines how to combine elements that evaluate to NULL with other elements in a logical OR operation.

Table 125 – Logical OR Truth table

	TRUE	FALSE	NULL
TRUE	TRUE	TRUE	TRUE
FALSE	TRUE	FALSE	NULL
NULL	TRUE	NULL	NULL

The NOT operator always evaluates to NULL if applied to a NULL operand.

A *ContentFilter* which evaluates to NULL after all elements are evaluated is evaluated as false.

7.4.4 FilterOperand parameters

7.4.4.1 Overview

The *ContentFilter* structure specified in 7.4 defines a collection of elements that makes up filter criteria and contains different types of *FilterOperands*. The *FilterOperand* parameter is an extensible parameter. This parameter is defined in Table 126. The *ExtensibleParameter* type is defined in 7.12.

Table 126 – FilterOperand parameter Typelds

Symbolic Id	Description
Element	Specifies an index into the array of elements. This type is used to build a logic tree of sub-elements by linking the operand of one element to a sub-element.
Literal	Specifies a literal value.
Attribute	Specifies any <i>Attribute</i> of an <i>Object</i> or <i>Variable Node</i> using a <i>Node</i> in the type system and relative path constructed from <i>ReferenceTypes</i> and <i>BrowseNames</i> .
SimpleAttribute	Specifies any <i>Attribute</i> of an <i>Object</i> or <i>Variable Node</i> using a <i>TypeDefinition</i> and a relative path constructed from <i>BrowseNames</i> .

7.4.4.2 ElementOperand

The *ElementOperand* provides the linking to sub-elements within a *ContentFilter*. The link is in the form of an integer that is used to index into the array of elements contained in the *ContentFilter*. An index is considered valid if its value is greater than the element index it is part of and it does not *Reference* a non-existent element. *Clients* shall construct filters in this way to avoid circular and invalid *References*. *Servers* should protect against invalid indexes by verifying the index prior to using it.

Table 127 defines the *ElementOperand* type.

Table 127 – ElementOperand

Name	Type	Description
ElementOperand	structure	ElementOperand value.
index	UInt32	Index into the element array.

7.4.4.3 LiteralOperand

Table 128 defines the *LiteralOperand* type.

Table 128 – LiteralOperand

Name	Type	Description
LiteralOperand	structure	LiteralOperand value.
value	BaseDataType	A literal value.

7.4.4.4 AttributeOperand

Table 129 defines the *AttributeOperand* type.

Table 129 – AttributeOperand

Name	Type	Description
AttributeOperand	structure	Attribute of a <i>Node</i> in the <i>AddressSpace</i> .
nodeId	NodeId	<i>NodeId</i> of a <i>Node</i> from the type system.
alias	String	An optional parameter used to identify or refer to an alias. An alias is a symbolic name that can be used to alias this operand and use it in other locations in the filter structure.
browsePath	RelativePath	Browse path relative to the <i>Node</i> identified by the <i>nodeId</i> parameter. See 7.26 for the definition of <i>RelativePath</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>AttributeId</i> . The <i>IntegerId</i> is defined in 7.14. The <i>IntegerIds</i> for the <i>Attributes</i> are defined in IEC 62541-6.
indexRange	NumericRange	This parameter is used to identify a single element of an array or a single range of indexes for an array. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is not used if the specified <i>Attribute</i> is not an array. However, if the specified <i>Attribute</i> is an array and this parameter is not used, then all elements are to be included in the range. The parameter is null if not used.

7.4.4.5 SimpleAttributeOperand

The *SimpleAttributeOperand* is a simplified form of the *AttributeOperand* and all of the rules that apply to the *AttributeOperand* also apply to the *SimpleAttributeOperand*. The examples provided in B.1 only use *AttributeOperand*, however, the *AttributeOperand* can be replaced by a *SimpleAttributeOperand* whenever all *ReferenceTypes* in the *RelativePath* are subtypes of *HierarchicalReferences* and the targets are *Object* or *Variable Nodes* and an *Alias* is not required.

Table 130 defines the *SimpleAttributeOperand* type.

Table 130 – SimpleAttributeOperand

Name	Type	Description
SimpleAttributeOperand	structure	Attribute of a <i>Node</i> in the <i>AddressSpace</i> .
typedId	NodeId	<i>NodeId</i> of a <i>TypeDefinitionNode</i> . This parameter restricts the operand to instances of the <i>TypeDefinitionNode</i> or one of its subtypes.
browsePath []	QualifiedName	A relative path to a <i>Node</i> . This parameter specifies a relative path using a list of <i>BrowseNames</i> instead of the <i>RelativePath</i> structure used in the <i>AttributeOperand</i> . The list of <i>BrowseNames</i> is equivalent to a <i>RelativePath</i> that specifies forward references which are subtypes of the <i>HierarchicalReferences ReferenceType</i> . All <i>Nodes</i> followed by the <i>browsePath</i> shall be of the <i>NodeClass Object</i> or <i>Variable</i> . If this list is empty the <i>Node</i> is the instance of the <i>TypeDefinition</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . The <i>IntegerId</i> is defined in 7.14. The <i>Value Attribute</i> shall be supported by all <i>Servers</i> . The support of other <i>Attributes</i> depends on requirements set in Profiles or other parts of this specification.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for an array. The first element is identified by index 0 (zero). This parameter is ignored if the selected <i>Node</i> is not a <i>Variable</i> or the <i>Value</i> of a <i>Variable</i> is not an array. The parameter is null if not specified. All values in the array are used if this parameter is not specified. The <i>NumericRange</i> type is defined in 7.22.

7.5 Counter

This primitive data type is a UInt32 that represents the value of a counter. The initial value of a counter is specified by its use. Modulus arithmetic is used for all calculations, where the modulus is max value + 1. Therefore,

$$x + y = (x + y) \bmod (\text{max value} + 1)$$

For example:

$$\text{max value} + 1 = 0$$

$$\text{max value} + 2 = 1$$

7.6 ContinuationPoint

A *ContinuationPoint* is used to pause a *Browse*, *QueryFirst* or *HistoryRead* operation and allow it to be restarted later by calling *BrowseNext*, *QueryNext* or *HistoryRead*. Operations are paused when the number of results found exceeds the limits set by either the *Client* or the *Server*.

The *Client* specifies the maximum number of results per operation in the request message. A *Server* shall not return more than this number of results but it may return fewer results. The *Server* allocates a *ContinuationPoint* if there are more results to return.

Servers shall support at least one *ContinuationPoint* per *Session*. *Servers* specify a maximum number of *ContinuationPoints* per *Session* in the *ServerCapabilities Object* defined in IEC 62541-5. *ContinuationPoints* remain active until the *Client* retrieves the remaining results, the *Client* releases the *ContinuationPoint* or the *Session* is closed. A *Server* shall automatically free *ContinuationPoints* from prior requests from a *Session* if they are needed to

process a new request from this *Session*. The *Server* returns a *Bad_ContinuationPointInvalid* error if a *Client* tries to use a *ContinuationPoint* that has been released. A *Client* can avoid this situation by completing paused operations before starting new operations.

Requests will often specify multiple operations that may or may not require a *ContinuationPoint*. A *Server* shall process the operations until it uses the maximum number of continuation points in this response. Once that happens the *Server* shall return a *Bad_NoContinuationPoints* error for any remaining operations. A *Client* can avoid this situation by sending requests with a number of operations that do not exceed the maximum number of *ContinuationPoints* per *Session* defined for the *Service* in the *ServerCapabilities Object* defined in IEC 62541-5.

A *Client* restarts an operation by passing the *ContinuationPoint* back to the *Server*. *Server* should always be able to reuse the *ContinuationPoint* provided so *Servers* shall never return *Bad_NoContinuationPoints* error when continuing a previously halted operation.

A *ContinuationPoint* is a subtype of the *ByteString* data type.

7.7 DataValue

7.7.1 General

The components of this parameter are defined in Table 131.

Table 131 – DataValue

Name	Type	Description
DataValue	structure	The value and associated information.
value	BaseDataType	The data value. If the <i>StatusCode</i> indicates an error then the value is to be ignored and the <i>Server</i> shall set it to null.
statusCode	StatusCode	The <i>StatusCode</i> that defines with the <i>Server's</i> ability to access/provide the value. The <i>StatusCode</i> type is defined in 7.34
sourceTimestamp	UtcTime	The source timestamp for the value.
sourcePicoSeconds	UInteger	Specifies the number of 10 picoseconds (1,0 e-11 seconds) intervals which shall be added to the sourceTimestamp.
serverTimestamp	UtcTime	The <i>Server</i> timestamp for the value.
serverPicoSeconds	UInteger	Specifies the number of 10 picoseconds (1,0 e-11 seconds) intervals which shall be added to the serverTimestamp.

7.7.2 PicoSeconds

Some applications require high resolution timestamps. The *PicoSeconds* fields allow applications to specify timestamps with a resolution of 10 picoseconds. The actual size of the *PicoSeconds* field depends on the resolution of the *UtcTime DataType*. For example, if the *UtcTime DataType* has a resolution of 100 nanoseconds then the *PicoSeconds* field would need to store values up to 10 000 in order to provide the resolution of 10 picoseconds. The resolution of the *UtcTime DataType* depends on the *Mappings* defined in IEC 62541-6.

7.7.3 SourceTimestamp

The *sourceTimestamp* is used to reflect the timestamp that was applied to a *Variable* value by the data source. Once a value has been assigned a source timestamp, the source timestamp for that value instance never changes. In this context, "value instance" refers to the value received, independent of its actual value.

The *sourceTimestamp* shall be UTC time and should indicate the time of the last change of the *value* or *statusCode*.

The *sourceTimestamp* should be generated as close as possible to the source of the value but the timestamp needs to be set always by the same physical clock. In the case of redundant sources, the clocks of the sources should be synchronised.

If the OPC UA Server receives the *Variable* value from another OPC UA Server, then the OPC UA Server shall always pass the source timestamp without changes. If the source that applies the timestamp is not available, the source timestamp is set to null. For example, if a value could not be read because of some error during processing like invalid arguments passed in the request then the *sourceTimestamp* shall be null.

In the case of a bad or uncertain status *sourceTimestamp* is used to reflect the time that the source recognized the non-good status or the time the Server last tried to recover from the bad or uncertain status.

The *sourceTimestamp* is only returned with a *Value Attribute*. For all other *Attributes* the returned *sourceTimestamp* is set to null.

7.7.4 ServerTimestamp

The *serverTimestamp* is used to reflect the time that the Server received a *Variable* value or knew it to be accurate.

In the case of a bad or uncertain status, *serverTimestamp* is used to reflect the time that the Server received the status or that the Server last tried to recover from the bad or uncertain status.

In the case where the OPC UA Server subscribes to a value from another OPC UA Server, each Server applies its own *serverTimestamp*. This is in contrast to the *sourceTimestamp* in which only the originator of the data is allowed to apply the *sourceTimestamp*.

If the Server subscribes to the value from another Server every ten seconds and the value changes, then the *serverTimestamp* is updated each time a new value is received. If the value does not change, then new values will not be received on the *Subscription*. However, in the absence of errors, the receiving Server applies a new *serverTimestamp* every ten seconds because not receiving a value means that the value has not changed. Thus, the *serverTimestamp* reflects the time at which the Server knew the value to be accurate.

This concept also applies to OPC UA Servers that receive values from exception-based data sources. For example, suppose that a Server is receiving values from an exception-based device, and that

- a) the device is checking values every 0,5 seconds,
- b) the connection to the device is good,
- c) the device sent an update 3 minutes ago with a value of 1,234.

In this case, the Server value would be 1,234 and the *serverTimestamp* would be updated every 0,5 seconds after the receipt of the value.

7.7.5 StatusCode assigned to a value

The *StatusCode* is used to indicate the conditions under which a *Variable* value was generated, and thereby can be used as an indicator of the usability of the value. The *StatusCode* is defined in 7.34.

Overall condition (severity)

- A *StatusCode* with severity Good means that the value is of good quality.

- A *StatusCode* with severity Uncertain means that the quality of the value is uncertain for reasons indicated by the *SubCode*.
- A *StatusCode* with severity Bad means that the value is not usable for reasons indicated by the *SubCode*.

Rules

- The *StatusCode* indicates the usability of the value. Therefore, It is required that *Clients* minimally check the *StatusCode Severity* of all results, even if they do not check the other fields, before accessing and using the value.
- A *Server*, which does not support status information, shall return a severity code of Good. It is also acceptable for a *Server* to simply return a severity and a non-specific (0) *SubCode*.
- If the *Server* has no known value - in particular when *Severity* is BAD, it shall return a NULL value.

7.8 DiagnosticInfo

The components of this parameter are defined in Table 132.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Table 132 – DiagnosticInfo

Name	Type	Description
DiagnosticInfo	structure	Vendor-specific diagnostic information.
namespaceUri	Int32	The <i>symbolicId</i> is defined within the context of a namespace. This namespace is represented as a string and is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>namespaceIndex</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified. The <i>namespaceUri</i> shall not be the standard OPC UA namespace. There are no <i>symbolicIds</i> provided for standard <i>StatusCodes</i>.
symbolicId	Int32	The <i>symbolicId</i> shall be used to identify a vendor-specific error or condition; typically the result of some <i>Server</i> internal operation. The maximum length of this string is 32 characters. <i>Servers</i> wishing to return a numeric return code should convert the return code into a string and use this string as <i>symbolicId</i> (e.g., "0xC0040007" or "-4"). This symbolic identifier string is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>symbolicId</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified. The <i>symbolicId</i> shall not contain <i>StatusCodes</i>. If the <i>localizedText</i> contains a translation for the description of a <i>StatusCode</i>, the <i>symbolicId</i> is -1.
locale	Int32	The locale part of the vendor-specific localized text describing the symbolic id. This localized text string is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>locale</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified.
localizedText	Int32	A vendor-specific localized text string describes the symbolic id. The maximum length of this text string is 256 characters. This localized text string is conveyed to the <i>Client</i> in the <i>stringTable</i> parameter of the <i>ResponseHeader</i> parameter defined in 7.29. The <i>localizedText</i> parameter contains the index into the <i>stringTable</i> for this string. -1 indicates that no string is specified. The <i>localizedText</i> refers to the <i>symbolicId</i> if present or the string that describes the standard <i>StatusCode</i> if the <i>Server</i> provides translations. If the index is -1, the <i>Server</i> has no translation to return and the <i>Client</i> should use the invariant <i>StatusCode</i> description from the specification.
additionalInfo	String	Vendor-specific diagnostic information.
innerStatusCode	StatusCode	The <i>StatusCode</i> from the inner operation. Many applications will make calls into underlying systems during OPC UA request processing. An OPC UA <i>Server</i> has the option of reporting the status from the underlying system in the diagnostic info.
innerDiagnosticInfo	DiagnosticInfo	The diagnostic info associated with the inner <i>StatusCode</i> .

7.9 DiscoveryConfiguration parameters

7.9.1 Overview

The *DiscoveryConfiguration* structure used in the *RegisterServer2* Service allows *Servers* to provide additional configuration parameters to *Discovery Servers* for registration. Table 133 defines the current set of discovery configuration options. The *ExtensibleParameter* type is defined in 7.12.

Table 133 – DiscoveryConfiguration parameterTypeIds

Symbolic Id	Description
MdnsDiscoveryConfiguration	Configuration parameters for mDNS discovery.

7.9.2 MdnsDiscoveryConfiguration

Table 134 defines the *MdnsDiscoveryConfiguration* parameter.

Table 134 – MdnsDiscoveryConfiguration

Name	Type	Description
MdnsDiscoveryConfiguration	structure	mDNS discovery configuration.
mdnsServerName	String	The name of the <i>Server</i> when it is announced via mDNS. See IEC 62541-12 for the details about mDNS. This string shall be less than 64 bytes. If not specified the first element of the <i>serverNames</i> array is used (truncated to 63 bytes if necessary).
serverCapabilities []	String	The set of <i>Server</i> capabilities supported by the <i>Server</i> . A <i>Server</i> capability is a short identifier for a feature. The set of allowed <i>Server</i> capabilities are defined in IEC 62541-12.

7.10 EndpointDescription

The components of this parameter are defined in Table 135.

Table 135 – EndpointDescription

Name	Type	Description
EndpointDescription	structure	Describes an <i>Endpoint</i> for a <i>Server</i> .
endpointUrl	String	The URL for the <i>Endpoint</i> described.
server	ApplicationDescription	The description for the <i>Server</i> that the <i>Endpoint</i> belongs to. The <i>ApplicationDescription</i> type is defined in 7.1.
serverCertificate	ApplicationInstance Certificate	The <i>Application Instance Certificate</i> issued to the <i>Server</i> . The <i>ApplicationInstanceCertificate</i> type is defined in 7.2.
securityMode	Enum MessageSecurityMode	The type of security to apply to the messages. The type <i>MessageSecurityMode</i> type is defined in 7.15. A <i>SecureChannel</i> may need to be created even if the <i>securityMode</i> is NONE. The exact behaviour depends on the mapping used and is described in the IEC 62541-6.
securityPolicyUri	String	The URI for <i>SecurityPolicy</i> to use when securing messages. The set of known URIs and the <i>SecurityPolicies</i> associated with them are defined in IEC 62541-7.
userIdentityTokens []	UserTokenPolicy	The user identity tokens that the <i>Server</i> will accept. The <i>Client</i> shall pass one of the <i>UserIdentityTokens</i> in the <i>ActivateSession</i> request. The <i>UserTokenPolicy</i> type is described in 7.37.
transportProfileUri	String	The URI of the <i>Transport Profile</i> supported by the <i>Endpoint</i> . IEC 62541-7 defines URIs for the <i>Transport Profiles</i> .
securityLevel	Byte	A numeric value that indicates how secure the <i>EndpointDescription</i> is compared to other <i>EndpointDescriptions</i> for the same <i>Server</i> . A value of 0 indicates that the <i>EndpointDescription</i> is not recommended and is only supported for backward compatibility. A higher value indicates better security.

7.11 ExpandedNodeId

The components of this parameter are defined in Table 136. *ExpandedNodeId* allows the namespace to be specified explicitly as a string or with an index in the *Server's* namespace table.

Table 136 – ExpandedNodeId

Name	Type	Description
ExpandedNodeId	structure	The <i>NodeId</i> with the namespace expanded to its string representation.
serverIndex	Index	Index that identifies the <i>Server</i> that contains the <i>TargetNode</i> . This <i>Server</i> may be the local <i>Server</i> or a remote <i>Server</i> . This index is the index of that <i>Server</i> in the local <i>Server's</i> <i>Server</i> table. The index of the local <i>Server</i> in the <i>Server</i> table is always 0. All remote <i>Servers</i> have indexes greater than 0. The <i>Server</i> table is contained in the <i>Server Object</i> in the <i>AddressSpace</i> (see IEC 62541-3 and IEC 62541-5). The <i>Client</i> may read the <i>Server</i> table <i>Variable</i> to access the description of the target <i>Server</i> .
namespaceUri	String	The URI of the namespace. If this parameter is specified then the namespace index is ignored. 5.4 and IEC 62541-12 describe discovery mechanism that can be used to resolve URIs into URLs.
namespaceIndex	Index	The index in the <i>Server's</i> namespace table. This parameter shall be 0 and is ignored in the <i>Server</i> if the namespace URI is specified.
identifierType	IdType	Type of the identifier element of the <i>NodeId</i> .
identifier	*	The identifier for a <i>Node</i> in the <i>AddressSpace</i> of an OPC UA <i>Server</i> (see <i>NodeId</i> definition in IEC 62541-3).

7.12 ExtensibleParameter

The extensible parameter types can only be extended by additional parts of IEC 62541.

The *ExtensibleParameter* defines a data structure with two elements. The *parameterTypeId* specifies the data type encoding of the second element. Therefore the second element is specified as "--". The *ExtensibleParameter* base type is defined in Table 137.

Concrete extensible parameters that are common to OPC UA are defined in Clause 7. Additional parts of IEC 62541 can define additional extensible parameter types.

Table 137 – ExtensibleParameter base type

Name	Type	Description
ExtensibleParameter	structure	Specifies the details of an extensible parameter type.
parameterTypeId	NodeId	Identifies the data type of the parameter that follows.
parameterData	--	The details for the extensible parameter type.

7.13 Index

This primitive data type is a UInt32 that identifies an element of an array.

7.14 IntegerId

This primitive data type is a UInt32 that is used as an identifier, such as a handle. All values, except for 0, are valid.

7.15 MessageSecurityMode

The *MessageSecurityMode* is an enumeration that specifies what security should be applied to messages exchanges during a Session. The possible values are described in Table 138.

Table 138 – MessageSecurityMode values

Value	Description
INVALID_0	The MessageSecurityMode is invalid. This value is the default value to avoid an accidental choice of no security is applied. This choice will always be rejected.
NONE_1	No security is applied.
SIGN_2	All messages are signed but not encrypted.
SIGNANDENCRYPT_3	All messages are signed and encrypted.

7.16 MonitoringParameters

The components of this parameter are defined in Table 139.

Table 139 – MonitoringParameters

Name	Type	Description
MonitoringParameters	structure	Parameters that define the monitoring characteristics of a <i>MonitoredItem</i> .
clientHandle	IntegerId	<i>Client</i> -supplied id of the <i>MonitoredItem</i> . This id is used in <i>Notifications</i> generated for the list <i>Node</i> . The <i>IntegerId</i> type is defined in 7.14.
samplingInterval	Duration	The interval that defines the fastest rate at which the <i>MonitoredItem</i> (s) should be accessed and evaluated. This interval is defined in milliseconds. The value 0 indicates that the <i>Server</i> should use the fastest practical rate. The value -1 indicates that the default sampling interval defined by the publishing interval of the <i>Subscription</i> is requested. A different sampling interval is used if the publishing interval is not a supported sampling interval. Any negative number is interpreted as -1. The sampling interval is not changed if the publishing interval is changed by a subsequent call to the <i>ModifySubscription</i> Service. The <i>Server</i> uses this parameter to assign the <i>MonitoredItems</i> to a sampling interval that it supports. The assigned interval is provided in the <i>revisedSamplingInterval</i> parameter. The <i>Server</i> shall always return a <i>revisedSamplingInterval</i> that is equal or higher than the requested <i>samplingInterval</i> . If the requested <i>samplingInterval</i> is higher than the maximum sampling interval supported by the <i>Server</i> , the maximum sampling interval is returned.
filter	Extensible Parameter MonitoringFilter	A filter used by the <i>Server</i> to determine if the <i>MonitoredItem</i> should generate a <i>Notification</i> . If not used, this parameter is null. The <i>MonitoringFilter</i> parameter type is an extensible parameter type specified in 7.17. It specifies the types of filters that can be used.

Name	Type	Description												
queueSize	Counter	The requested size of the <i>MonitoredItem</i> queue. The following values have special meaning for data monitored items: <table><tr><th>Value</th><th>Meaning</th></tr><tr><td>0 or 1</td><td>the <i>Server</i> returns the default queue size which shall be 1 as <i>revisedQueueSize</i> for data monitored items. The queue has a single entry, effectively disabling queuing.</td></tr></table> For values larger than one a first-in-first-out queue is to be used. The <i>Server</i> may limit the size in <i>revisedQueueSize</i>. In the case of a queue overflow, the <i>Overflow</i> bit (flag) in the <i>InfoBits</i> portion of the <i>DataValue statusCode</i> is set in the new value. The following values have special meaning for event monitored items: <table><tr><th>Value</th><th>Meaning</th></tr><tr><td>0</td><td>the <i>Server</i> returns the default queue size for <i>Event Notifications</i> as <i>revisedQueueSize</i> for event monitored items.</td></tr><tr><td>1</td><td>the <i>Server</i> returns the minimum queue size the <i>Server</i> requires for <i>Event Notifications</i> as <i>revisedQueueSize</i>.</td></tr><tr><td>MaxUInt32</td><td>the <i>Server</i> returns the maximum queue size that the <i>Server</i> can support for <i>Event Notifications</i> as <i>revisedQueueSize</i>.</td></tr></table> If a <i>Client</i> chooses a value between the minimum and maximum settings of the <i>Server</i> the value shall be returned in the <i>revisedQueueSize</i>. If the requested <i>queueSize</i> is outside the minimum or maximum, the <i>Server</i> shall return the corresponding bounding value. In the case of a queue overflow, an <i>Event</i> of the type <i>EventQueueOverflowEventType</i> is generated.	Value	Meaning	0 or 1	the <i>Server</i> returns the default queue size which shall be 1 as <i>revisedQueueSize</i> for data monitored items. The queue has a single entry, effectively disabling queuing.	Value	Meaning	0	the <i>Server</i> returns the default queue size for <i>Event Notifications</i> as <i>revisedQueueSize</i> for event monitored items.	1	the <i>Server</i> returns the minimum queue size the <i>Server</i> requires for <i>Event Notifications</i> as <i>revisedQueueSize</i> .	MaxUInt32	the <i>Server</i> returns the maximum queue size that the <i>Server</i> can support for <i>Event Notifications</i> as <i>revisedQueueSize</i> .
Value	Meaning													
0 or 1	the <i>Server</i> returns the default queue size which shall be 1 as <i>revisedQueueSize</i> for data monitored items. The queue has a single entry, effectively disabling queuing.													
Value	Meaning													
0	the <i>Server</i> returns the default queue size for <i>Event Notifications</i> as <i>revisedQueueSize</i> for event monitored items.													
1	the <i>Server</i> returns the minimum queue size the <i>Server</i> requires for <i>Event Notifications</i> as <i>revisedQueueSize</i> .													
MaxUInt32	the <i>Server</i> returns the maximum queue size that the <i>Server</i> can support for <i>Event Notifications</i> as <i>revisedQueueSize</i> .													
discardOldest	Boolean	Specifies the discard policy when the queue is full and a new <i>Notification</i> is to be queued. It has the following values: <table><tr><td>TRUE</td><td>the oldest (first) <i>Notification</i> in the queue is discarded. The new <i>Notification</i> is added to the end of the queue.</td></tr><tr><td>FALSE</td><td>the last <i>Notification</i> added to the queue gets replaced with the new <i>Notification</i>.</td></tr></table>	TRUE	the oldest (first) <i>Notification</i> in the queue is discarded. The new <i>Notification</i> is added to the end of the queue.	FALSE	the last <i>Notification</i> added to the queue gets replaced with the new <i>Notification</i> .								
TRUE	the oldest (first) <i>Notification</i> in the queue is discarded. The new <i>Notification</i> is added to the end of the queue.													
FALSE	the last <i>Notification</i> added to the queue gets replaced with the new <i>Notification</i> .													

7.17 MonitoringFilter parameters

7.17.1 Overview

The *CreateMonitoredItem Service* allows specifying a filter for each *MonitoredItem*. The *MonitoringFilter* is an extensible parameter whose structure depends on the type of item being monitored. The *parameterTypeIds* are defined in Table 140. Other types can be defined by additional parts of this multi-part specification or other specifications based on OPC UA. The *ExtensibleParameter* type is defined in 7.12.

Each *MonitoringFilter* may have an associated *MonitoringFilterResult* structure which returns revised parameters and/or error information to clients in the response. The result structures, when they exist, are described in the section that defines the *MonitoringFilter*.

Table 140 – MonitoringFilter parameterTypeIds

Symbolic Id	Description
DataChangeFilter	The change in a data value that shall cause a <i>Notification</i> to be generated.
EventFilter	If a <i>Notification</i> conforms to the <i>EventFilter</i> , the <i>Notification</i> is sent to the <i>Client</i> .
AggregateFilter	The <i>Aggregate</i> and its intervals when it will be calculated and a <i>Notification</i> is generated.

7.17.2 DataChangeFilter

The *DataChangeFilter* defines the conditions under which a *DataChange Notification* should be reported and, optionally, a range or band for value changes where no *DataChange Notification* is generated. This range is called *Deadband*. The *DataChangeFilter* is defined in Table 141.

Table 141 – DataChangeFilter

Name	Type	Description								
DataChangeFilter	structure									
trigger	Enum DataChangeTrigger	Specifies the conditions under which a data change notification should be reported. It has the following values: STATUS_0 Report a notification ONLY if the <i>StatusCode</i> associated with the value changes. See Table 178 for <i>StatusCodes</i> defined in this document. IEC 62541-8 specifies additional <i>StatusCodes</i> that are valid in particular for device data. STATUS_VALUE_1 Report a notification if either the <i>StatusCode</i> or the value change. The <i>Deadband</i> filter can be used in addition for filtering value changes. For floating point values a <i>Server</i> shall check for NaN and only report a single notification with NaN when the value enters the NaN state. This is the default setting if no filter is set. STATUS_VALUE_TIMESTAMP_2 Report a notification if either <i>StatusCode</i>, value or the <i>SourceTimestamp</i> change. If a <i>Deadband</i> filter is specified, this trigger has the same behaviour as STATUS_VALUE_1. If the <i>DataChangeFilter</i> is not applied to the monitored item, STATUS_VALUE_1 is the default reporting behaviour.								
deadbandType	UInt32	A value that defines the <i>Deadband</i> type and behaviour. <table><thead><tr><th>Value</th><th>deadbandType</th></tr></thead><tbody><tr><td>None_0</td><td>No <i>Deadband</i> calculation should be applied.</td></tr><tr><td>Absolute_1</td><td>AbsoluteDeadband (see below)</td></tr><tr><td>Percent_2</td><td>PercentDeadband (This type is specified in IEC 62541-8).</td></tr></tbody></table>	Value	deadbandType	None_0	No <i>Deadband</i> calculation should be applied.	Absolute_1	AbsoluteDeadband (see below)	Percent_2	PercentDeadband (This type is specified in IEC 62541-8).
Value	deadbandType									
None_0	No <i>Deadband</i> calculation should be applied.									
Absolute_1	AbsoluteDeadband (see below)									
Percent_2	PercentDeadband (This type is specified in IEC 62541-8).									
deadbandValue	Double	The <i>Deadband</i> is applied only if * the <i>trigger</i> includes value changes and * the <i>deadbandType</i> is set appropriately. Deadband is ignored if the status of the data item changes. DeadbandType = AbsoluteDeadband: For this type the <i>deadbandValue</i> contains the absolute change in a data value that shall cause a <i>Notification</i> to be generated. This parameter applies only to <i>Variables</i> with any <i>Number</i> data type. An exception that causes a <i>DataChange Notification</i> based on an AbsoluteDeadband is determined as follows: Generate a Notification if (absolute value of (last cached value - current value) > AbsoluteDeadband) The last cached value is defined as the last value pushed to the queue. If the item is an array of values, the entire array is returned if any array element exceeds the AbsoluteDeadband, or the size or dimension of the array changes. DeadbandType = PercentDeadband: This type is specified in IEC 62541-8								

The *DataChangeFilter* does not have an associated result structure.

7.17.3 EventFilter

The *EventFilter* provides for the filtering and content selection of *Event Subscriptions*.

If an *Event Notification* conforms to the filter defined by the *where* parameter of the *EventFilter*, then the *Notification* is sent to the *Client*.

Each *Event Notification* shall include the fields defined by the *selectClauses* parameter of the *EventFilter*. The defined *EventTypes* are specified in IEC 62541-5.

The *selectClauses* and *whereClause* parameters are specified with the *SimpleAttributeOperand* structure (see 7.4.4.5). This structure requires the *NodeId* of an *EventType* supported by the *Server* and a path to an *InstanceDeclaration*. An *InstanceDeclaration* is a *Node* which can be found by following forward hierarchical references from the fully inherited *EventType* where the *Node* is also the source of a *HasModellingRule* reference. *EventTypes*, *InstanceDeclarations* and *Modelling Rules* are described completely in IEC 62541-3.

In some cases the same *BrowsePath* will apply to multiple *EventTypes*. If the *Client* specifies the *BaseEventType* in the *SimpleAttributeOperand* then the *Server* shall evaluate the *BrowsePath* without considering the *Type*.

Each *InstanceDeclaration* in the path shall be *Object* or *Variable Node*. The final *Node* in the path may be an *Object Node*; however, *Object Nodes* are only available for *Events* which are visible in the *Server's AddressSpace*.

The *SimpleAttributeOperand* structure allows the *Client* to specify any *Attribute*; however, the *Server* is only required to support the *Value Attribute* for *Variable Nodes* and the *NodeId Attribute* for *Object Nodes*. That said, profiles defined in IEC 62541-7 may make support for additional *Attributes* mandatory.

The *SimpleAttributeOperand* structure is used in the *selectClauses* to select the value to return if an *Event* meets the criteria specified by the *whereClause*. A null value is returned in the corresponding event field in the Publish response if the selected *field* is not part of the *Event* or an error was returned in the *selectClauseResults* of the *EventFilterResult*. If the selected *field* is supported but not available at the time of the event notification, the event field shall contain a *StatusCode* that indicates the reason for the unavailability. For example, the *Server* shall set the event field to *Bad_UserAccessDenied* if the value is not accessible to the user associated with the *Session*. If a *Value Attribute* has an uncertain or bad *StatusCode* associated with it then the *Server* shall provide the *StatusCode* instead of the *Value Attribute*. The *Server* shall set the event field to *Bad_EncodingLimitsExceeded* if a value exceeds the *maxResponseMessageSize*. The *EventId*, *EventType* and *ReceiveTime* cannot contain a *StatusCode* or a null value.

The *Server* shall validate the *selectClauses* when a *Client* creates or updates the *EventFilter*. Any errors which are true for all possible *Events* are returned in the *selectClauseResults* parameter described in Table 143. Some *Servers*, like aggregating *Servers*, may not know all possible *EventTypes* at the time the *EventFilter* is set. These *Servers* do not return errors for unknown *EventTypes* or *BrowsePaths*. The *Server* shall not report errors that might occur depending on the state or the *Server* or type of *Event*. For example, a *selectClauses* that requests a single element in an array would always produce an error if the *DataType* of the *Attribute* is a scalar. However, even if the *DataType* is an array an error could occur if the requested index does not exist for a particular *Event*, the *Server* would not report an error in the *selectClauseResults* parameter if the latter situation existed.

The *SimpleAttributeOperand* is used in the *whereClause* to select a value which forms part of a logical expression. These logical expressions are then used to determine whether a particular *Event* should be reported to the *Client*. The *Server* shall use a null value if any error occurs when a *whereClause* is evaluated for a particular *Event*. If a *Value Attribute* has an

uncertain or bad *StatusCode* associated with it, then the *Server* shall use a null value instead of the *Value*.

Any basic *FilterOperator* in Table 119 may be used in the *whereClause*, however, only the *OfType_14 FilterOperator* from Table 120 is permitted.

The *Server* shall validate the *whereClause* when a *Client* creates or updates the *EventFilter*. Any structural errors in the construction of the filter and any errors which are true for all possible *Events* are returned in the *whereClauseResult* parameter described in Table 143. Errors that could occur depending on the state of the *Server* or the *Event* are not reported. Some *Servers*, like aggregating *Servers*, may not know all possible *EventTypes* at the time the *EventFilter* is set. These *Servers* do not return errors for unknown *EventTypes* or *BrowsePaths*.

EventQueueOverflowEventType Events are special *Events* which are used to provide control information to the *Client*. These *Events* are only published to the *MonitoredItems* in the *Subscription* that produced the *EventQueueOverflowEventType Event*. These *Events* bypass the *whereClause*.

Table 142 defines the *EventFilter* structure.

Table 142 – EventFilter structure

Name	Type	Description
EventFilter	structure	
selectClauses []	SimpleAttribute Operand	List of the values to return with each <i>Event</i> in a <i>Notification</i> . At least one valid clause shall be specified. See 7.4.4.5 for the definition of <i>SimpleAttributeOperand</i> .
whereClause	ContentFilter	Limit the <i>Notifications</i> to those <i>Events</i> that match the criteria defined by this <i>ContentFilter</i> . The <i>ContentFilter</i> structure is described in 7.4. The <i>AttributeOperand</i> structure may not be used in an <i>EventFilter</i> .

Table 143 defines the *EventFilterResult* structure. This is the *MonitoringFilterResult* associated with the *EventFilter MonitoringFilter*.

Table 143 – EventFilterResult structure

Name	Type	Description
EventFilterResult	structure	
selectClauseResults []	StatusCode	List of status codes for the elements in the select clause. The size and order of the list matches the size and order of the elements in the <i>selectClauses</i> request parameter. The <i>Server</i> returns null for unavailable or rejected <i>Event</i> fields.
selectClauseDiagnosticInfos []	DiagnosticInfo	A list of diagnostic information for individual elements in the select clause. The size and order of the list matches the size and order of the elements in the <i>selectClauses</i> request parameter. This list is empty if diagnostics information was not requested in the request header or if no diagnostic information was encountered in processing of the select clauses.
whereClauseResult	ContentFilter Result	Any results associated with the <i>whereClause</i> request parameter. The <i>ContentFilterResult</i> type is defined in 7.4.2.

Table 144 defines values for the *selectClauseResults* parameter. Common *StatusCodes* are defined in Table 178.

Table 144 – EventFilterResult result codes

Symbolic Id	Description
Bad_TypeDefinitionInvalid	See Table 178 for the description of this result code. The typeId is not the NodeId for BaseEventType or a subtype of it.
Bad_NodeIdUnknown	See Table 178 for the description of this result code. The browsePath is specified but it will never exist in any Event.
Bad_BrowseNameInvalid	See Table 178 for the description of this result code. The browsePath is specified and contains a null element.
Bad_AttributeIdInvalid	See Table 178 for the description of this result code. The node specified by the browse path will never allow the given <i>AttributeId</i> to be returned.
Bad_IndexRangeInvalid	See Table 178 for the description of this result code.
Bad_TypeMismatch	See Table 178 for the description of this result code. The indexRange is valid but the value of the Attribute is never an array.

7.17.4 AggregateFilter

The *AggregateFilter* defines the *Aggregate* function that should be used to calculate the values to be returned. See IEC 62541-13 for details on possible *Aggregate* functions. It specifies a *startTime* of the first *Aggregate* to be calculated. The *samplingInterval* of the *MonitoringParameters* (see 7.16) defines how the *Server* should internally sample the underlying data source. The *processingInterval* specifies the size of a time-period where the *Aggregate* is calculated. The *queueSize* from the *MonitoringAttributes* specifies the number of processed values that should be kept.

The intention of the *AggregateFilter* is not to read historical data, the *HistoryRead* service should be used for this purpose. However, it is allowed that the *startTime* is set to a time that is in the past when received from the server. The number of *Aggregates* to be calculated in the past should not exceed the *queueSize* defined in the *MonitoringAttributes* since the values exceeding the *queueSize* would directly be discharged and never returned to the client.

The *startTime* and the *processingInterval* can be revised by the server, but the *startTime* should remain in the same boundary ($\text{startTime} + \text{revisedProcessingInterval} * n = \text{revisedStartTime}$). That behaviour simplifies accessing historical values of the *Aggregates* using the same boundaries by calling the *HistoryRead* service. The extensible Parameter *AggregateFilterResult* is used to return the revised values for the *AggregateFilter*.

Some underlying systems may poll data and produce multiple samples with the same value. Other systems may only report changes to the values. The definition for each *Aggregate* type explains how to handle the two different scenarios.

The *MonitoredItem* only reports values for intervals that have completed when the publish timer expires. Unused data is carried over and used to calculate a value returned in the next publish.

The *ServerTimestamp* for each interval shall be the time of the end of the processing interval.

The *AggregateFilter* is defined in Table 145.

Table 145 – AggregateFilter structure

Name	Type	Description
AggregateFilter	structure	
startTime	UtcTime	Beginning of period to calculate the <i>Aggregate</i> the first time. The size of each period used to calculate the <i>Aggregate</i> is defined by the <i>samplingInterval</i> of the <i>MonitoringParameters</i> (see 7.16).
aggregateType	NodeId	The <i>NodeId</i> of the <i>AggregateFunctionType Object</i> that indicates the <i>Aggregate</i> to be used when retrieving processed data. See IEC 62541-13 for details.
processingInterval	Duration	The period be used to compute the <i>Aggregate</i> .
aggregateConfiguration	Aggregate Configuration	This parameter allows <i>Clients</i> to override the <i>Aggregate</i> configuration settings supplied by the <i>AggregateConfiguration Object</i> on a per monitored item basis. See IEC 62541-13 for more information on <i>Aggregate</i> configurations. If the <i>Server</i> does not support the ability to override the <i>Aggregate</i> configuration settings it shall return a <i>Status Code</i> of <i>Bad_AggregateListMismatch</i> . This structure is defined in-line with the following indented items.
useServerCapabilities Defaults	Boolean	If value = TRUE use <i>Aggregate</i> configuration settings as outlined by the <i>AggregateConfiguration</i> object. If value=FALSE use configuration settings as outlined in the following <i>aggregateConfiguration</i> parameters. Default is TRUE.
treatUncertainAsBad	Boolean	As described in IEC 62541-13.
percentDataBad	Byte	As described in IEC 62541-13.
percentDataGood	Byte	As described in IEC 62541-13.
useSloped Extrapolation	Boolean	As described in IEC 62541-13.

The *AggregateFilterResult* defines the revised *AggregateFilter* the *Server* can return when an *AggregateFilter* is defined for a *MonitoredItem* in the *CreateMonitoredItems* or *ModifyMonitoredItems* Services. The *AggregateFilterResult* is defined in Table 146. This is the *MonitoringFilterResult* associated with the *AggregateFilter MonitoringFilter*.

Table 146 – AggregateFilterResult structure

Name	Type	Description
AggregateFilterResult	structure	
revisedStartTime	UtcTime	The actual <i>StartTime</i> interval that the <i>Server</i> shall use. This value is based on a number of factors, including capabilities of the <i>Server</i> to access historical data. The <i>revisedStartTime</i> should remain in the same boundary as the <i>startTime</i> ($\text{startTime} + \text{samplingInterval} * n = \text{revisedStartTime}$).
revisedProcessingInterval	Duration	The actual <i>processingInterval</i> that the <i>Server</i> shall use. The <i>revisedProcessingInterval</i> shall be at least twice the <i>revisedSamplingInterval</i> for the <i>MonitoredItem</i> .

7.18 MonitoringMode

The *MonitoringMode* is an enumeration that specifies whether sampling and reporting are enabled or disabled for a *MonitoredItem*. The value of the publishing enabled parameter for a *Subscription* does not affect the value of the monitoring mode for a *MonitoredItem* of the *Subscription*. The values of this parameter are defined in Table 147.

Table 147 – MonitoringMode values

Value	Description
DISABLED_0	The item being monitored is not sampled or evaluated, and <i>Notifications</i> are not generated or queued. <i>Notification</i> reporting is disabled.
SAMPLING_1	The item being monitored is sampled and evaluated, and <i>Notifications</i> are generated and queued. <i>Notification</i> reporting is disabled.
REPORTING_2	The item being monitored is sampled and evaluated, and <i>Notifications</i> are generated and queued. <i>Notification</i> reporting is enabled.

7.19 NodeAttributes parameters

7.19.1 Overview

The *AddNodes Service* allows specifying the *Attributes* for the *Nodes* to add. The *NodeAttributes* is an extensible parameter whose structure depends on the type of the *NodeClass* being added. It identifies the *NodeClass* that defines the structure of the *Attributes* that follow. The *parameterTypeIds* are defined in Table 148. The *ExtensibleParameter* type is defined in 7.12.

Table 148 – NodeAttributes parameterTypeIds

Symbolic Id	Description
ObjectAttributes	Defines the <i>Attributes</i> for the <i>Object NodeClass</i> .
VariableAttributes	Defines the <i>Attributes</i> for the <i>Variable NodeClass</i> .
MethodAttributes	Defines the <i>Attributes</i> for the <i>Method NodeClass</i> .
ObjectTypeAttributes	Defines the <i>Attributes</i> for the <i>ObjectType NodeClass</i> .
VariableTypeAttributes	Defines the <i>Attributes</i> for the <i>VariableType NodeClass</i> .
ReferenceTypeAttributes	Defines the <i>Attributes</i> for the <i>ReferenceType NodeClass</i> .
DataTypeAttributes	Defines the <i>Attributes</i> for the <i>DataType NodeClass</i> .
ViewAttributes	Defines the <i>Attributes</i> for the <i>View NodeClass</i> .
GenericAttributes	Defines an id and value list for passing in any number of <i>Attribute</i> values. It should be used instead of the <i>NodeClass</i> specific structures since it allows the handling of additional <i>Attributes</i> defined in future specification versions.

Table 149 defines the bit mask used in the *NodeAttributes* parameters to specify which *Attributes* are set by the *Client*.

Table 149 – Bit mask for specified Attributes

Field	Bit	Description
AccessLevel	0	Indicates if the AccessLevel Attribute is set.
ArrayDimensions	1	Indicates if the ArrayDimensions Attribute is set.
Reserved	2	Reserved to be consistent with WriteMask defined in IEC 62541-3.
ContainsNoLoops	3	Indicates if the ContainsNoLoops Attribute is set.
DataType	4	Indicates if the DataType Attribute is set.
Description	5	Indicates if the Description Attribute is set.
DisplayName	6	Indicates if the DisplayName Attribute is set.
EventNotifier	7	Indicates if the EventNotifier Attribute is set.
Executable	8	Indicates if the Executable Attribute is set.
Historizing	9	Indicates if the Historizing Attribute is set.
InverseName	10	Indicates if the InverseName Attribute is set.
IsAbstract	11	Indicates if the IsAbstract Attribute is set.
MinimumSamplingInterval	12	Indicates if the MinimumSamplingInterval Attribute is set.
Reserved	13	Reserved to be consistent with WriteMask defined in IEC 62541-3.
Reserved	14	Reserved to be consistent with WriteMask defined in IEC 62541-3.
Symmetric	15	Indicates if the Symmetric Attribute is set.
UserAccessLevel	16	Indicates if the UserAccessLevel Attribute is set.
UserExecutable	17	Indicates if the UserExecutable Attribute is set.
UserWriteMask	18	Indicates if the UserWriteMask Attribute is set.
ValueRank	19	Indicates if the ValueRank Attribute is set.
WriteMask	20	Indicates if the WriteMask Attribute is set.
Value	21	Indicates if the Value Attribute is set.
Reserved	22:32	Reserved for future use. Shall always be zero.

7.19.2 ObjectAttributes parameter

Table 150 defines the *ObjectAttributes* parameter.

Table 150 – ObjectAttributes

Name	Type	Description
ObjectAttributes	structure	Defines the <i>Attributes</i> for the <i>Object NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
eventNotifier	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.3 VariableAttributes parameter

Table 151 defines the *VariableAttributes* parameter.

Table 151 – VariableAttributes

Name	Type	Description
VariableAttributes	structure	Defines the <i>Attributes</i> for the <i>Variable NodeClass</i>
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
value	Defined by the <i>DataType Attribute</i>	See IEC 62541-3 for the description of this <i>Attribute</i> .
dataType	NodeId	See IEC 62541-3 for the description of this <i>Attribute</i> .
valueRank	Int32	See IEC 62541-3 for the description of this <i>Attribute</i> .
arrayDimensions	UInt32 []	See IEC 62541-3 for the description of this <i>Attribute</i> .
accessLevel	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
userAccessLevel	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
minimumSamplingInterval	Duration	See IEC 62541-3 for the description of this <i>Attribute</i> .
historizing	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.4 MethodAttributes parameter

Table 152 defines the *MethodAttributes* parameter.

Table 152 – MethodAttributes

Name	Type	Description
BaseAttributes	structure	Defines the <i>Attributes</i> for the <i>Method NodeClass</i>
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
executable	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
userExecutable	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.5 ObjectTypeAttributes parameter

Table 153 defines the *ObjectTypeAttributes* parameter.

Table 153 – ObjectTypeAttributes

Name	Type	Description
ObjectTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>ObjectType NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.6 VariableTypeAttributes parameter

Table 154 defines the *VariableTypeAttributes* parameter.

Table 154 – VariableTypeAttributes

Name	Type	Description
VariableTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>VariableType NodeClass</i>
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
value	Defined by the <i>Data Type Attribute</i>	See IEC 62541-3 for the description of this <i>Attribute</i> .
dataType	NodeId	See IEC 62541-3 for the description of this <i>Attribute</i> .
valueRank	Int32	See IEC 62541-3 for the description of this <i>Attribute</i> .
arrayDimensions	UInt32 []	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.7 ReferenceTypeAttributes parameter

Table 155 defines the *ReferenceTypeAttributes* parameter.

Table 155 – ReferenceTypeAttributes

Name	Type	Description
ReferenceTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>ReferenceType NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
symmetric	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
inverseName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.8 DataTypeAttributes parameter

Table 156 defines the *DataTypeAttributes* parameter.

Table 156 – DataTypeAttributes

Name	Type	Description
DataTypeAttributes	structure	Defines the <i>Attributes</i> for the <i>DataType NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
isAbstract	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.9 ViewAttributes parameter

Table 157 defines the *ViewAttributes* parameter.

Table 157 – ViewAttributes

Name	Type	Description
ViewAttributes	structure	Defines the <i>Attributes</i> for the <i>View NodeClass</i> .
specifiedAttributes	UInt32	A bit mask that indicates which fields contain valid values. A field shall be ignored if the corresponding bit is set to 0. The bit values are defined in Table 149.
displayName	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
description	LocalizedText	See IEC 62541-3 for the description of this <i>Attribute</i> .
containsNoLoops	Boolean	See IEC 62541-3 for the description of this <i>Attribute</i> .
eventNotifier	Byte	See IEC 62541-3 for the description of this <i>Attribute</i> .
writeMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .
userWriteMask	UInt32	See IEC 62541-3 for the description of this <i>Attribute</i> .

7.19.10 GenericAttributes parameter

This structure should be used instead of the *NodeClass* specific structures defined in the other subclauses of 7.19 since it allows the handling of additional *Attributes* defined in future specification versions.

Table 158 defines the *GenericAttributes* parameter.

Table 158 – GenericAttributes

Name	Type	Description
GenericAttributes	structure	Defines a generic structure for passing in any number of <i>Attributes</i> .
attributeValues	GenericAttributeValue []	Defines one <i>attributeId</i> and <i>value</i> combination.
attributeId	IntegerId	Id of the <i>Attribute</i> specified.
value	BaseDataType	Value of the <i>Attribute</i> specified.

7.20 NotificationData parameters

7.20.1 Overview

The *NotificationMessage* structure used in the *Subscription Service* set allows specifying different types of *NotificationData*. The *NotificationData* parameter is an extensible parameter whose structure depends on the type of *Notification* being sent. This parameter is defined in Table 159. Other types can be defined by additional parts of IEC 62541 or other specifications based on OPC UA. The *ExtensibleParameter* type is defined in 7.12.

There may be multiple notifications for a single *MonitoredItem* in a single *NotificationData* structure. When that happens the *Server* shall ensure the notifications appear in the same order that they were queued in the *MonitoredItem*. These notifications do not need to appear as a contiguous block.

Table 159 – NotificationData parameterTypeIds

Symbolic Id	Description
DataChange	<i>Notification</i> data parameter used for data change <i>Notifications</i> .
Event	<i>Notification</i> data parameter used for <i>Event Notifications</i> .
StatusChange	<i>Notification</i> data parameter used for Subscription status change <i>Notifications</i> .

7.20.2 DataChangeNotification parameter

Table 160 defines the *NotificationData* parameter used for data change notifications. This structure contains the monitored data items that are to be reported. Monitored data items are reported under two conditions:

- if the *MonitoringMode* is set to REPORTING and a change in value or its status (represented by its *StatusCode*) is detected;
- if the *MonitoringMode* is set to SAMPLING, the *MonitoredItem* is linked to a triggering item and the triggering item triggers.

See 5.12 for a description of the *MonitoredItem Service* set, and in particular the *MonitoredItem model* and the *Triggering model*.

After creating a *MonitoredItem*, the current value or status of the monitored Attribute shall be queued without applying the filter. If the current value is not available after the first sampling interval the first *Notification* shall be queued after getting the initial value or status from the data source.

Table 160 – DataChangeNotification

Name	Type	Description
DataChangeNotification	structure	Data change <i>Notification</i> data.
monitoredItems []	MonitoredItem Notification	The list of <i>MonitoredItems</i> for which a change has been detected. This structure is defined in-line with the following indented items.
clientHandle	IntegerId	<i>Client</i> -supplied handle for the <i>MonitoredItem</i> . The <i>IntegerId</i> type is defined in 7.14
Value	DataValue	The <i>StatusCode</i> , value and timestamp(s) of the monitored <i>Attribute</i> depending on the sampling and queuing configuration. If the <i>StatusCode</i> indicates an error then the value is to be ignored. If not every detected change has been returned since the <i>Server's</i> queue buffer for the <i>MonitoredItem</i> reached its limit and had to purge out data and the size of the queue is larger than one, the <i>Overflow</i> bit in the <i>DataValue InfoBits</i> of the <i>statusCode</i> is set. <i>DataValue</i> is a common type defined in 7.7.
diagnosticInfos []	DiagnosticInfo	List of diagnostic information. The size and order of this list matches the size and order of the <i>monitoredItems</i> parameter. There is one entry in this list for each <i>Node</i> contained in the <i>monitoredItems</i> parameter. This list is empty if diagnostics information was not requested or is not available for any of the <i>MonitoredItems</i> . <i>DiagnosticInfo</i> is a common type defined in 7.8.

7.20.3 EventNotificationList parameter

Table 161 defines the *NotificationData* parameter used for *Event* notifications.

The EventNotificationList defines a table structure that is used to return *Event* fields to a *Client Subscription*. The structure is in the form of a table consisting of one or more *Events*, each containing an array of one or more fields. The selection and order of the fields returned for each *Event* is identical to the selected parameter of the *EventFilter*.

Table 161 – EventNotificationList

Name	Type	Description
EventNotificationList	structure	Event <i>Notification</i> data.
events []	EventFieldList	The list of <i>Events</i> being delivered. This structure is defined in-line with the following indented items.
clientHandle	IntegerId	<i>Client</i> -supplied handle for the <i>MonitoredItem</i> . The <i>IntegerId</i> type is defined in 7.14.
eventFields []	BaseDataType	List of selected <i>Event</i> fields. This shall be a one to one match with the fields selected in the <i>EventFilter</i> . 7.17.3 specifies how the <i>Server</i> shall deal with error conditions.

7.20.4 StatusChangeNotification parameter

Table 162 defines the *NotificationData* parameter used for a *StatusChangeNotification*.

The *StatusChangeNotification* informs the *Client* about a change in the status of a *Subscription*.

Table 162 – StatusChangeNotification

Name	Type	Description
StatusChangeNotification	structure	Event <i>Notification</i> data
status	StatusCode	The <i>StatusCode</i> that indicates the status change.
diagnosticInfo	DiagnosticInfo	Diagnostic information for the status change

7.21 NotificationMessage

The components of this parameter are defined in Table 163.

Table 163 – NotificationMessage

Name	Type	Description
NotificationMessage	structure	The <i>Message</i> that contains one or more <i>Notifications</i> .
sequenceNumber	Counter	The sequence number of the <i>NotificationMessage</i> .
publishTime	UtcTime	The time that this <i>Message</i> was sent to the <i>Client</i> . If this <i>Message</i> is retransmitted to the <i>Client</i> , this parameter contains the time it was first transmitted to the <i>Client</i> .
notificationData []	Extensible Parameter NotificationData	The list of <i>NotificationData</i> structures. The <i>NotificationData</i> parameter type is an extensible parameter type specified in 7.20. It specifies the types of <i>Notifications</i> that can be sent. The <i>ExtensibleParameter</i> type is specified in 7.12. Notifications of the same type should be grouped into one <i>NotificationData</i> element. If a <i>Subscription</i> contains <i>MonitoredItems</i> for events and data, this array should have not more than 2 elements. If the <i>Subscription</i> contains <i>MonitoredItems</i> only for data or only for events, the array size should always be one for this <i>Subscription</i> .

7.22 NumericRange

This parameter is defined in Table 164. A formal BNF definition of the numeric range can be found in Clause A.3.

The syntax for the string contains one of the following two constructs. The first construct is the string representation of an individual integer. For example, "6" is valid, but "6,0" and "3,2" are not. The minimum and maximum values that can be expressed are defined by the use of this parameter and not by this parameter type definition. The second construct is a range represented by two integers separated by the colon (":") character. The first integer shall always have a lower value than the second. For example, "5:7" is valid, while "7:5" and "5:5" are not. The minimum and maximum values that can be expressed by these integers are defined by the use of this parameter, and not by this parameter type definition. No other characters, including white-space characters, are permitted.

Multi-dimensional arrays can be indexed by specifying a range for each dimension separated by a ','. For example, a 2x2 block in a 4x4 matrix could be selected with the range "1:2,0:1". A single element in a multi-dimensional array can be selected by specifying a single number instead of a range. For example, "1,1" selects the [1,1] element in a two dimensional array.

Dimensions are specified in the order that they appear in the *ArrayDimensions* Attribute. All dimensions shall be specified for a *NumericRange* to be valid.

All indexes start with 0. The maximum value for any index is one less than the length of the dimension.

When reading a value and any of the lower bounds of the indexes is out of range, the *Server* shall return a *Bad_IndexRangeNoData*. If any of the upper bounds of the indexes is out of range, the *Server* shall return partial results.

Bad_IndexRangeInvalid is only used for invalid syntax of the *NumericRange*. All other invalid requests with a valid syntax shall result in *Bad_IndexRangeNoData*.

When writing a value, the size of the array shall match the size specified by the *NumericRange*. The *Server* shall return an error if it cannot write all elements specified by the *Client*.

The *NumericRange* can also be used to specify substrings for *ByteString* and *String* values. Arrays of *ByteString* and *String* values are treated as two dimensional arrays where the final index specifies the substring range within the *ByteString* or *String* value. The entire *ByteString* or *String* value is selected if the final index is omitted.

Table 164 – NumericRange

Name	Type	Description
NumericRange	String	A number or a numeric range. A null string indicates that this parameter is not used.

7.23 QueryDataSet

The components of this parameter are defined in Table 165.

Table 165 – QueryDataSet

Name	Type	Description
QueryDataSet	structure	Data related to a <i>Node</i> returned in a Query response.
nodeId	ExpandedNodeId	The <i>NodeId</i> for this <i>Node</i> description.
typeDefinitionNode	ExpandedNodeId	The <i>NodeId</i> for the type definition for this <i>Node</i> description.
values []	BaseDataType	Values for the selected <i>Attributes</i>. The order of returned items matches the order of the requested items. There is an entry for each requested item for the given <i>TypeDefinitionNode</i> that matches the selected instance, this includes any related nodes that were specified using a relative path from the selected instance's <i>TypeDefinitionNode</i>. If no values were found for a given requested item a null value is returned for that item. If a value has a bad status, the <i>StatusCode</i> is returned instead of the value. If multiple values exist for a requested item then an array of values is returned. If the requested item is a reference then a <i>ReferenceDescription</i> or array of <i>ReferenceDescription</i> is returned for that item. If the <i>QueryDataSet</i> is returned in a <i>QueryNext</i> to continue a list of <i>ReferenceDescription</i>, the <i>values</i> array will have the same size but the other values already returned are null.

7.24 ReadValueId

The components of this parameter are defined in Table 166.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020

Table 166 – ReadValueId

Name	Type	Description
ReadValueId	structure	Identifier for an item to read or to monitor.
nodeId	NodeId	<i>NodeId</i> of a <i>Node</i> .
attributeId	IntegerId	Id of the <i>Attribute</i> . This shall be a valid <i>Attribute</i> id. The <i>IntegerId</i> is defined in 7.14. The <i>IntegerIds</i> for the <i>Attributes</i> are defined in IEC 62541-6.
indexRange	NumericRange	This parameter is used to identify a single element of an array, or a single range of indexes for arrays. If a range of elements is specified, the values are returned as a composite. The first element is identified by index 0 (zero). The <i>NumericRange</i> type is defined in 7.22. This parameter is null if the specified <i>Attribute</i> is not an array. However, if the specified <i>Attribute</i> is an array, and this parameter is null, then all elements are to be included in the range.
dataEncoding	QualifiedName	This parameter specifies the <i>BrowseName</i> of the <i>DataTypeEncoding</i> that the <i>Server</i> should use when returning the <i>Value Attribute</i> of a <i>Variable</i>. It is an error to specify this parameter for other <i>Attributes</i>. This parameter only applies if the <i>DataType</i> of the <i>Variable</i> is a subtype of <i>Structure</i>. It is an error to specify this parameter if the <i>DataType</i> of the <i>Variable</i> is not a subtype of <i>Structure</i>. A <i>Client</i> can discover what <i>DataTypeEncodings</i> are available by following the <i>HasEncoding Reference</i> from the <i>DataType Node</i> for a <i>Variable</i>. OPC UA defines <i>BrowseNames</i> which <i>Servers</i> shall recognize even if the <i>DataType Nodes</i> are not visible in the <i>Server AddressSpace</i>. These <i>BrowseNames</i> are: Default Binary The default or native binary (or non-XML) encoding. Default XML The default XML encoding. Default JSON The default JSON encoding Each <i>DataType</i> shall support at least one of these encodings. <i>DataTypes</i> that do not have a true binary encoding (e.g. they only have a non-XML text encoding) should use the Default Binary name to identify the encoding that is considered to be the default non-XML encoding. <i>DataTypes</i> that support at least one XML-based encoding shall identify one of the encodings as the Default XML encoding. Other standards bodies may define other well-known data encodings that could be supported. If this parameter is not specified then the <i>Server</i> shall choose the default according to what <i>Message</i> encoding (see IEC 62541-6) is used for the <i>Session</i>. If the <i>Server</i> does not support the encoding that matches the <i>Message</i> encoding then the <i>Server</i> shall choose the default encoding that it does support.

7.25 ReferenceDescription

The components of this parameter are defined in Table 167.

Table 167 – ReferenceDescription

Name	Type	Description
ReferenceDescription	structure	Reference parameters returned for the <i>Browse Service</i> .
referenceTypeid	Nodeid	<i>Nodeid</i> of the <i>ReferenceType</i> that defines the <i>Reference</i> .
isForward	Boolean	If the value is TRUE, the <i>Server</i> followed a forward <i>Reference</i> . If the value is FALSE, the <i>Server</i> followed an inverse <i>Reference</i> .
nodeid	Expanded Nodeid	<i>Nodeid</i> of the <i>TargetNode</i> as assigned by the <i>Server</i> identified by the <i>Server</i> index. The <i>ExpandedNodeid</i> type is defined in 7.11. If the <i>serverIndex</i> indicates that the <i>TargetNode</i> is a remote <i>Node</i> , then the <i>nodeid</i> shall contain the absolute namespace URI. If the <i>TargetNode</i> is a local <i>Node</i> the <i>nodeid</i> shall contain the namespace index.
browseName ¹⁾	QualifiedName	The <i>BrowseName</i> of the <i>TargetNode</i> .
displayName	LocalizedText	The <i>DisplayName</i> of the <i>TargetNode</i> .
nodeClass ¹⁾	NodeClass	<i>NodeClass</i> of the <i>TargetNode</i> .
typeDefinition ¹⁾	Expanded Nodeid	Type definition <i>Nodeid</i> of the <i>TargetNode</i> . Type definitions are only available for the <i>NodeClasses Object</i> and <i>Variable</i> . For all other <i>NodeClasses</i> a null <i>Nodeid</i> shall be returned.
¹⁾ If the <i>Server</i> index indicates that the <i>TargetNode</i> is a remote <i>Node</i> , then the <i>browseName</i> , <i>nodeClass</i> and <i>typeDefinition</i> may be null or empty. If they are not, they might not be up to date because the local <i>Server</i> might not continuously monitor the remote <i>Server</i> for changes. The <i>displayName</i> shall be provided for remote <i>Nodes</i> .		

7.26 RelativePath

The components of this parameter are defined in Table 168.

Table 168 – RelativePath

Name	Type	Description
RelativePath	structure	Defines a sequence of <i>References</i> and <i>BrowseNames</i> to follow.
elements []	RelativePath Element	A sequence of <i>References</i> and <i>BrowseNames</i> to follow. This structure is defined in-line with the following indented items. Each element in the sequence is processed by finding the targets and then using those targets as the starting nodes for the next element. The targets of the final element are the target of the <i>RelativePath</i> .
referenceTypeid	Nodeid	The type of reference to follow from the current node. The current path cannot be followed any further if the <i>referenceTypeid</i> is not available on the <i>Node</i> instance. If not specified then all <i>References</i> are included and the parameter <i>includeSubtypes</i> is ignored.
isInverse	Boolean	Only inverse references shall be followed if this value is TRUE. Only forward references shall be followed if this value is FALSE.
includeSubtypes	Boolean	Indicates whether subtypes of the <i>ReferenceType</i> should be followed. Subtypes are included if this value is TRUE.
targetName	QualifiedName	The <i>BrowseName</i> of the target node. The final element may have an empty <i>targetName</i> . In this situation all targets of the references identified by the <i>referenceTypeid</i> are the targets of the <i>RelativePath</i> . The <i>targetName</i> shall be specified for all other elements. The current path cannot be followed any further if no targets with the specified <i>BrowseName</i> exist.

A *RelativePath* can be applied to any starting *Node*. The targets of the *RelativePath* are the set of *Nodes* that are found by sequentially following the elements in *RelativePath*.

A text format for the *RelativePath* can be found in Clause A.2. This format is used in examples that explain the *Services* that make use of the *RelativePath* structure.

7.27 RegisteredServer

The components of this parameter are defined in Table 169.

Table 169 – RegisteredServer

Name	Type	Description
RegisteredServer	structure	The <i>Server</i> to register.
serverUri	String	The globally unique identifier for the <i>Server</i> instance. The <i>serverUri</i> matches the <i>applicationUri</i> from the <i>ApplicationDescription</i> defined in 7.1.
productUri	String	The globally unique identifier for the <i>Server</i> product.
serverNames []	LocalizedText	A list of localized descriptive names for the <i>Server</i> . The list shall have at least one valid entry.
serverType	Enum ApplicationType	The type of application. The enumeration values are defined in Table 112. The value "CLIENT_1" (the application is a <i>Client</i>) is not allowed. The <i>Service</i> result shall be <i>Bad_InvalidArgument</i> in this case.
gatewayServerUri	String	The URI of the <i>Gateway Server</i> associated with the <i>discoveryUrls</i> . This value is only specified by <i>Gateway Servers</i> that wish to register the <i>Servers</i> that they provide access to. For <i>Servers</i> that do not act as a <i>Gateway Server</i> this parameter shall be null.
discoveryUrls []	String	A list of <i>DiscoveryEndpoints</i> for the <i>Server</i> . The list shall have at least one valid entry.
semaphoreFilePath	String	The path to the semaphore file used to identify an automatically-launched <i>Server</i> instance; manually-launched servers will not use this parameter. If a Semaphore file is provided, the <i>isOnline</i> flag is ignored. If a Semaphore file is provided and exists, the <i>LocalDiscoveryServer</i> shall save the registration information in a persistent data store that it reads whenever the <i>LocalDiscoveryServer</i> starts. If a Semaphore file is specified but does not exist, the <i>Discovery Server</i> shall remove the registration from any persistent data store. If the <i>Server</i> has registered with a <i>semaphoreFilePath</i> , the <i>Discovery Server</i> shall check that this file exists before returning the <i>ApplicationDescription</i> to the client. If the <i>Server</i> did not register with a <i>semaphoreFilePath</i> (it is null or empty), then the <i>Discovery Server</i> does not attempt to verify the existence of the file before returning the <i>ApplicationDescription</i> to the client.
isOnline	Boolean	True if the <i>Server</i> is currently able to accept connections from <i>Clients</i> . The <i>Discovery Server</i> shall return <i>ApplicationDescriptions</i> to the <i>Client</i> . The <i>Server</i> is expected to periodically re-register with the <i>Discovery Server</i> . False if the <i>Server</i> is currently unable to accept connections from <i>Clients</i> . The <i>Discovery Server</i> shall NOT return <i>ApplicationDescriptions</i> to the <i>Client</i> . This parameter is ignored if a <i>semaphoreFilePath</i> is provided.

7.28 RequestHeader

The components of this parameter are defined in Table 170.

Table 170 – RequestHeader

Name	Type	Description																																				
RequestHeader	structure	Common parameters for all requests submitted on a <i>Session</i> .																																				
authenticationToken	Session AuthenticationToken	The secret <i>Session</i> identifier used to verify that the request is associated with the <i>Session</i> . The <i>SessionAuthenticationToken</i> type is defined in 7.31.																																				
timestamp	UtcTime	The time the <i>Client</i> sent the request. The parameter is only used for diagnostic and logging purposes in the server.																																				
requestHandle	IntegerId	A <i>requestHandle</i> associated with the request. This <i>Client</i> defined handle can be used to cancel the request. It is also returned in the response.																																				
returnDiagnostics	UInt32	A bit mask that identifies the types of vendor-specific diagnostics to be returned in <i>diagnosticInfo</i> response parameters. The value of this parameter may consist of zero, one or more of the following values. No value indicates that diagnostics are not to be returned. <table><thead><tr><th>Bit Value</th><th>Diagnostics to return</th></tr></thead><tbody><tr><td>0x0000 0001</td><td>ServiceLevel / SymbolicId</td></tr><tr><td>0x0000 0002</td><td>ServiceLevel / LocalizedText</td></tr><tr><td>0x0000 0004</td><td>ServiceLevel / AdditionalInfo</td></tr><tr><td>0x0000 0008</td><td>ServiceLevel / Inner <i>StatusCode</i></td></tr><tr><td>0x0000 0010</td><td>ServiceLevel / Inner Diagnostics</td></tr><tr><td>0x0000 0020</td><td>OperationLevel / SymbolicId</td></tr><tr><td>0x0000 0040</td><td>OperationLevel / LocalizedText</td></tr><tr><td>0x0000 0080</td><td>OperationLevel / AdditionalInfo</td></tr><tr><td>0x0000 0100</td><td>OperationLevel / Inner <i>StatusCode</i></td></tr><tr><td>0x0000 0200</td><td>OperationLevel / Inner Diagnostics</td></tr></tbody></table> Each of these values is composed of two components, <i>level</i> and <i>type</i>, as described below. If none are requested, as indicated by a 0 value, or if no diagnostic information was encountered in processing of the request, then diagnostics information is not returned. Level: <table><tbody><tr><td>ServiceLevel</td><td>return diagnostics in the <i>diagnosticInfo</i> of the <i>Service</i>.</td></tr><tr><td>OperationLevel</td><td>return diagnostics in the <i>diagnosticInfo</i> defined for individual operations requested in the <i>Service</i>.</td></tr></tbody></table> Type: <table><tbody><tr><td>SymbolicId</td><td>return a namespace-qualified, symbolic identifier for an error or condition. The maximum length of this identifier is 32 characters.</td></tr><tr><td>LocalizedText</td><td>return up to 256 bytes of localized text that describes the symbolic id.</td></tr><tr><td>AdditionalInfo</td><td>return a byte string that contains additional diagnostic information, such as a memory image. The format of this byte string is vendor-specific, and may depend on the type of error or condition encountered.</td></tr><tr><td>InnerStatusCode</td><td>return the inner <i>StatusCode</i> associated with the operation or <i>Service</i>.</td></tr><tr><td>InnerDiagnostics</td><td>return the inner diagnostic info associated with the operation or <i>Service</i>. The contents of the inner diagnostic info structure are determined by other bits in the mask. Note that setting this bit could cause multiple levels of nested diagnostic info structures to be returned.</td></tr></tbody></table>	Bit Value	Diagnostics to return	0x0000 0001	ServiceLevel / SymbolicId	0x0000 0002	ServiceLevel / LocalizedText	0x0000 0004	ServiceLevel / AdditionalInfo	0x0000 0008	ServiceLevel / Inner <i>StatusCode</i>	0x0000 0010	ServiceLevel / Inner Diagnostics	0x0000 0020	OperationLevel / SymbolicId	0x0000 0040	OperationLevel / LocalizedText	0x0000 0080	OperationLevel / AdditionalInfo	0x0000 0100	OperationLevel / Inner <i>StatusCode</i>	0x0000 0200	OperationLevel / Inner Diagnostics	ServiceLevel	return diagnostics in the <i>diagnosticInfo</i> of the <i>Service</i> .	OperationLevel	return diagnostics in the <i>diagnosticInfo</i> defined for individual operations requested in the <i>Service</i> .	SymbolicId	return a namespace-qualified, symbolic identifier for an error or condition. The maximum length of this identifier is 32 characters.	LocalizedText	return up to 256 bytes of localized text that describes the symbolic id.	AdditionalInfo	return a byte string that contains additional diagnostic information, such as a memory image. The format of this byte string is vendor-specific, and may depend on the type of error or condition encountered.	InnerStatusCode	return the inner <i>StatusCode</i> associated with the operation or <i>Service</i> .	InnerDiagnostics	return the inner diagnostic info associated with the operation or <i>Service</i> . The contents of the inner diagnostic info structure are determined by other bits in the mask. Note that setting this bit could cause multiple levels of nested diagnostic info structures to be returned.
Bit Value	Diagnostics to return																																					
0x0000 0001	ServiceLevel / SymbolicId																																					
0x0000 0002	ServiceLevel / LocalizedText																																					
0x0000 0004	ServiceLevel / AdditionalInfo																																					
0x0000 0008	ServiceLevel / Inner <i>StatusCode</i>																																					
0x0000 0010	ServiceLevel / Inner Diagnostics																																					
0x0000 0020	OperationLevel / SymbolicId																																					
0x0000 0040	OperationLevel / LocalizedText																																					
0x0000 0080	OperationLevel / AdditionalInfo																																					
0x0000 0100	OperationLevel / Inner <i>StatusCode</i>																																					
0x0000 0200	OperationLevel / Inner Diagnostics																																					
ServiceLevel	return diagnostics in the <i>diagnosticInfo</i> of the <i>Service</i> .																																					
OperationLevel	return diagnostics in the <i>diagnosticInfo</i> defined for individual operations requested in the <i>Service</i> .																																					
SymbolicId	return a namespace-qualified, symbolic identifier for an error or condition. The maximum length of this identifier is 32 characters.																																					
LocalizedText	return up to 256 bytes of localized text that describes the symbolic id.																																					
AdditionalInfo	return a byte string that contains additional diagnostic information, such as a memory image. The format of this byte string is vendor-specific, and may depend on the type of error or condition encountered.																																					
InnerStatusCode	return the inner <i>StatusCode</i> associated with the operation or <i>Service</i> .																																					
InnerDiagnostics	return the inner diagnostic info associated with the operation or <i>Service</i> . The contents of the inner diagnostic info structure are determined by other bits in the mask. Note that setting this bit could cause multiple levels of nested diagnostic info structures to be returned.																																					

Name	Type	Description
auditEntryId	String	An identifier that identifies the <i>Client</i> 's security audit log entry associated with this request. An empty string value means that this parameter is not used. The <i>auditEntryId</i> typically contains who initiated the action and from where it was initiated. The <i>auditEntryId</i> is included in the <i>AuditEvent</i> to allow human readers to correlate an <i>Event</i> with the initiating action. More details of the <i>Audit</i> mechanisms are defined in 6.5 and in IEC 62541-3.
timeoutHint	UInt32	This timeout in milliseconds is used in the <i>Client</i> side <i>Communication Stack</i> to set the timeout on a per-call base. For a <i>Server</i> this timeout is only a hint and can be used to cancel long running operations to free resources. If the <i>Server</i> detects a timeout, he can cancel the operation by sending the <i>Service</i> result <i>Bad_Timeout</i> . The <i>Server</i> should wait at minimum the timeout after he received the request before cancelling the operation. The <i>Server</i> shall check the <i>timeoutHint</i> parameter of a <i>Publish</i> request before processing a <i>Publish</i> response. If the request timed out, a <i>Bad_Timeout Service</i> result is sent and another <i>Publish</i> request is used. The value of 0 indicates no timeout.
additionalHeader	Extensible Parameter AdditionalHeader	Reserved for future use. Applications that do not understand the header should ignore it.

7.29 ResponseHeader

The components of this parameter are defined in Table 171.

Table 171 – ResponseHeader

Name	Type	Description
ResponseHeader	structure	Common parameters for all responses.
timestamp	UtcTime	The time the <i>Server</i> sent the response.
requestHandle	IntegerId	The requestHandle given by the <i>Client</i> to the request.
serviceResult	StatusCode	OPC UA-defined result of the <i>Service</i> invocation. The <i>StatusCode</i> type is defined in 7.34.
serviceDiagnostics	DiagnosticInfo	Diagnostic information for the <i>Service</i> invocation. This parameter is empty if diagnostics information was not requested in the request header. The <i>DiagnosticInfo</i> type is defined in 7.8.
stringTable []	String	There is one string in this list for each unique namespace, symbolic identifier, and localized text string contained in all of the diagnostics information parameters contained in the response (see 7.8). Each is identified within this table by its zero-based index.
additionalHeader	Extensible Parameter AdditionalHeader	Reserved for future use. Applications that do not understand the header should ignore it.

7.30 ServiceFault

The components of this parameter are defined in Table 172.

The *ServiceFault* parameter is returned instead of the *Service* response message when a service level error occurs. The *requestHandle* in the *ResponseHeader* should be set to what was provided in the *RequestHeader* even if these values were not valid. The level of diagnostics returned in the *ResponseHeader* is specified by the *returnDiagnostics* parameter in the *RequestHeader*.

The exact use of this parameter depends on the mappings defined in IEC 62541-6.

Table 172 – ServiceFault

Name	Type	Description
ServiceFault	structure	An error response sent when a service level error occurs.
responseHeader	ResponseHeader	Common response parameters (see 7.29 for <i>ResponseHeader</i> definition).

7.31 SessionAuthenticationToken

The *SessionAuthenticationToken* type is an opaque identifier that is used to identify requests associated with a particular *Session*. This identifier is used in conjunction with the *SecureChannelId* or *Client Certificate* to authenticate incoming messages. It is the secret form of the *sessionId* for internal use in the *Client* and *Server Applications*.

A *Server* returns a *SessionAuthenticationToken* in the *CreateSession* response. The *Client* then sends this value with every request which allows the *Server* to verify that the sender of the request is the same as the sender of the original *CreateSession* request.

For the purposes of this discussion, a *Server* consists of application (code) and a *Communication Stack* as shown in Figure 37. The security provided by the *SessionAuthenticationToken* depends on a trust relationship between the *Server* application and the *Communication Stack*. The *Communication Stack* shall be able to verify the sender of the message and it uses the *SecureChannelId* or the *Client Certificate* to identify the sender to the *Server*. In these cases, the *SessionAuthenticationToken* is a UInt32 identifier that allows the *Server* to distinguish between different *Sessions* created by the same sender.



Figure 37 – Logical layers of a Server

In some cases, the application and the *Communication Stack* cannot exchange information at runtime which means the application will not have access to the *SecureChannelId* or the *Certificate* used to create the *SecureChannel*. In these cases the application shall create a random *ByteString* value that is at least 32 bytes long. This value shall be kept secret and shall always be exchanged over a *SecureChannel* with encryption enabled. The Administrator is responsible for ensuring that encryption is enabled. The *Profiles* in IEC 62541-7 may define additional requirements for a *ByteString SessionAuthenticationToken*.

Client and *Server* applications should be written to be independent of the *SecureChannel* implementation. Therefore, they should always treat the *SessionAuthenticationToken* as secret information even if it is not required when using some *SecureChannel* implementations.

Figure 38 illustrates the information exchanged between the *Client*, the *Server* and the *Server Communication Stack* when the *Client* obtains a *SessionAuthenticationToken*. In this figure the *GetSecureChannelInfo* step represents an API that depends on the *Communication Stack* implementation.




Figure 38 – Obtaining a SessionAuthenticationToken

The *SessionAuthenticationToken* is a subtype of the *NodeId* data type; however, it is never used to identify a *Node* in the *AddressSpace*. Servers may assign a value to the *NamespaceIndex*; however, its meaning is *Server* specific.

7.32 SignatureData

The components of this parameter are defined in Table 173.

Table 173 – SignatureData

Name	Type	Description
SignatureData	structure	Contains a digital signature created with a <i>Certificate</i> .
algorithm	String	A string containing the URI of the <i>algorithm</i> . The URI string values are defined as part of the security profiles specified in IEC 62541-7.
signature	ByteString	This is a signature generated with the private key associated with a <i>Certificate</i> .

7.33 SignedSoftwareCertificate

NOTE Details on SoftwareCertificates will be defined in a future edition of IEC 62541-4.

Table 174 specifies *SignedSoftwareCertificate Structure*.

Table 174 – SignedSoftwareCertificate

Name	Type	Description
SignedSoftwareCertificate	structure	
certificateData	ByteString	The certificate data serialized as a ByteString.
signature	ByteString	The signature for the certificateData.

7.34 StatusCode

7.34.1 General

A *StatusCode* in OPC UA is numerical value that is used to report the outcome of an operation performed by an OPC UA *Server*. This code may have associated diagnostic information that describes the status in more detail; however, the code by itself is intended to provide *Client* applications with enough information to make decisions on how to process the results of an OPC UA *Service*.

The *StatusCode* is a 32-bit unsigned integer. The top 16 bits represent the numeric value of the code that shall be used for detecting specific errors or conditions. The bottom 16 bits are bit flags that contain additional information but do not affect the meaning of the *StatusCode*.

All OPC UA *Clients* shall always check the *StatusCode* associated with a result before using it. Results that have an uncertain/warning status associated with them shall be used with care since these results might not be valid in all situations. Results with a bad/failed status shall never be used.

OPC UA *Servers* should return good/success *StatusCodes* if the operation completed normally and the result is always valid. Different *StatusCode* values can provide additional information to the *Client*.

OPC UA *Servers* should use uncertain/warning *StatusCodes* if they could not complete the operation in the manner requested by the *Client*, however, the operation did not fail entirely.

The list of *StatusCodes* is managed by OPC UA. The complete list of *StatusCodes* is defined in IEC 62541-6. *Servers* shall not define their own *StatusCodes*. OPC UA companion working groups may request additional *StatusCodes* from the OPC Foundation to be added to the list in IEC 62541-6.

The exact bit assignments are shown in Table 175.

Table 175 – StatusCode bit assignments

Field	Bit Range	Description
Severity	30:31	Indicates whether the <i>StatusCode</i> represents a good, bad or uncertain condition. These bits have the following meanings: Good 00 Indicates that the operation was successful and the associated results may be used. Uncertain 01 Indicates that the operation was partially successful and that associated results might not be suitable for some purposes. Warning Bad 10 Indicates that the operation failed and any associated results cannot be used. Failure Reserved 11 Reserved for future use. All <i>Clients</i> should treat a <i>StatusCode</i> with this severity as "Bad".
Reserved	29:29	Reserved for use in OPC UA application specific APIs. This bit shall always be zero on the wire but may be used by OPC UA application specific APIs for API specific status codes.
Reserved	28:28	Reserved for future use. Shall always be zero.
SubCode	16:27	The code is a numeric value assigned to represent different conditions. Each code has a symbolic name and a numeric value. All descriptions in the OPC UA specification refer to the symbolic name. IEC 62541-6 maps the symbolic names onto a numeric value.
StructureChanged	15:15	Indicates that the structure of the associated data value has changed since the last <i>Notification</i> . <i>Clients</i> should not process the data value unless they re-read the metadata. <i>Servers</i> shall set this bit if the <i>DataTypeEncoding</i> used for a <i>Variable</i> changes. 7.24 describes how the <i>DataTypeEncoding</i> is specified for a <i>Variable</i>. <i>Servers</i> shall also set this bit if the <i>EnumStrings Property</i> of the <i>DataType</i> of the <i>Variable</i> changes. This bit is provided to warn <i>Clients</i> that parse complex data values that their parsing routines could fail because the serialized form of the data value has changed. This bit has meaning only for <i>StatusCodes</i> returned as part of a data change <i>Notification</i> or the <i>HistoryRead</i>. <i>StatusCodes</i> used in other contexts shall always set this bit to zero.
SemanticsChanged	14:14	Indicates that the semantics of the associated data value have changed. <i>Clients</i> should not process the data value until they re-read the metadata associated with the <i>Variable</i> . <i>Servers</i> should set this bit if the metadata has changed in way that could cause application errors if the <i>Client</i> does not re-read the metadata. For example, a change to the engineering units could create problems if the <i>Client</i> uses the value to perform calculations. IEC 62541-8 defines the conditions where a <i>Server</i> shall set this bit for a DA <i>Variable</i>. Other specifications may define additional conditions. A <i>Server</i> may define other conditions that cause this bit to be set. This bit has meaning only for <i>StatusCodes</i> returned as part of a data change <i>Notification</i> or the <i>HistoryRead</i>. <i>StatusCodes</i> used in other contexts shall always set this bit to zero.
Reserved	12:13	Reserved for future use. Shall always be zero.
InfoType	10:11	The type of information contained in the info bits. These bits have the following meanings: NotUsed 00 The info bits are not used and shall be set to zero. DataValue 01 The <i>StatusCode</i> and its info bits are associated with a data value returned from the <i>Server</i>. The info bits are defined in Table 176. Reserved 1X Reserved for future use. The info bits shall be ignored.
InfoBits	0:9	Additional information bits that qualify the <i>StatusCode</i> . The structure of these bits depends on the Info Type field.

Table 176 describes the structure of the *InfoBits* when the Info Type is set to *DataValue* (01).

Table 176 – DataValue InfoBits

Info Type	Bit Range	Description																					
LimitBits	8:9	The limit bits associated with the data value. The limits bits have the following meanings: <table> <tr> <th>Limit</th><th>Bits</th><th>Description</th></tr> <tr> <td>None</td><td>00</td><td>The value is free to change.</td></tr> <tr> <td>Low</td><td>01</td><td>The value is at the lower limit for the data source.</td></tr> <tr> <td>High</td><td>10</td><td>The value is at the higher limit for the data source.</td></tr> <tr> <td>Constant</td><td>11</td><td>The value is constant and cannot change.</td></tr> </table>	Limit	Bits	Description	None	00	The value is free to change.	Low	01	The value is at the lower limit for the data source.	High	10	The value is at the higher limit for the data source.	Constant	11	The value is constant and cannot change.						
Limit	Bits	Description																					
None	00	The value is free to change.																					
Low	01	The value is at the lower limit for the data source.																					
High	10	The value is at the higher limit for the data source.																					
Constant	11	The value is constant and cannot change.																					
Overflow	7	This bit shall only be set if the <i>MonitoredItem</i> queue size is greater than 1. If this bit is set, not every detected change has been returned since the <i>Server's</i> queue buffer for the <i>MonitoredItem</i> reached its limit and had to purge out data.																					
Reserved	5:6	Reserved for future use. Shall always be zero.																					
HistorianBits	0:4	These bits are set only when reading historical data. They indicate where the data value came from and provide information that affects how the <i>Client</i> uses the data value. The historian bits have the following meaning: <table> <tr> <td>Raw</td><td>XXX00</td><td>A raw data value.</td></tr> <tr> <td>Calculated</td><td>XXX01</td><td>A data value which was calculated.</td></tr> <tr> <td>Interpolated</td><td>XXX10</td><td>A data value which was interpolated.</td></tr> <tr> <td>Reserved</td><td>XXX11</td><td>Undefined.</td></tr> <tr> <td>Partial</td><td>XX1XX</td><td>A data value which was calculated with an incomplete interval.</td></tr> <tr> <td>Extra Data</td><td>X1XXX</td><td>A raw data value that hides other data at the same timestamp.</td></tr> <tr> <td>Multi Value</td><td>1XXXX</td><td>Multiple values match the <i>Aggregate</i> criteria (i.e. multiple minimum values at different timestamps within the same interval).</td></tr> </table> IEC 62541-11 describes how these bits are used in more detail.	Raw	XXX00	A raw data value.	Calculated	XXX01	A data value which was calculated.	Interpolated	XXX10	A data value which was interpolated.	Reserved	XXX11	Undefined.	Partial	XX1XX	A data value which was calculated with an incomplete interval.	Extra Data	X1XXX	A raw data value that hides other data at the same timestamp.	Multi Value	1XXXX	Multiple values match the <i>Aggregate</i> criteria (i.e. multiple minimum values at different timestamps within the same interval).
Raw	XXX00	A raw data value.																					
Calculated	XXX01	A data value which was calculated.																					
Interpolated	XXX10	A data value which was interpolated.																					
Reserved	XXX11	Undefined.																					
Partial	XX1XX	A data value which was calculated with an incomplete interval.																					
Extra Data	X1XXX	A raw data value that hides other data at the same timestamp.																					
Multi Value	1XXXX	Multiple values match the <i>Aggregate</i> criteria (i.e. multiple minimum values at different timestamps within the same interval).																					

7.34.2 Common StatusCodes

Table 177 defines the common *StatusCodes* for all *Service* results used in more than one service. It does not provide a complete list. These *StatusCodes* may also be used as operation level result code. IEC 62541-6 maps the symbolic names to a numeric value and provides a complete list of *StatusCodes* including codes defines in other parts.

Table 177 – Common Service result codes

Symbolic Id	Description
Good	The operation was successful.
Good_CompletesAsynchronously	The processing will complete asynchronously.
Good_SubscriptionTransferred	The subscription was transferred to another session.
Bad_CertificateHostNameInvalid	The <i>HostName</i> used to connect to a <i>Server</i> does not match a <i>HostName</i> in the <i>Certificate</i> .
Bad_CertificateChainIncomplete	The <i>Certificate</i> chain is incomplete.
Bad_CertificateIssuerRevocationUnknown	It was not possible to determine if the <i>Issuer Certificate</i> has been revoked.
Bad_CertificateIssuerUseNotAllowed	The <i>Issuer Certificate</i> may not be used for the requested operation.
Bad_CertificateIssuerTimeInvalid	An <i>Issuer Certificate</i> has expired or is not yet valid.
Bad_CertificateIssuerRevoked	The <i>Issuer Certificate</i> has been revoked.
Bad_CertificateInvalid	The <i>Certificate</i> provided as a parameter is not valid.
Bad_CertificateRevocationUnknown	It was not possible to determine if the <i>Certificate</i> has been revoked.
Bad_CertificateRevoked	The <i>Certificate</i> has been revoked.
Bad_CertificateTimeInvalid	The <i>Certificate</i> has expired or is not yet valid.
Bad_CertificateUriInvalid	The URI specified in the <i>ApplicationDescription</i> does not match the URI in the <i>Certificate</i> .
Bad_CertificateUntrusted	The <i>Certificate</i> is not trusted.
Bad_CertificateUseNotAllowed	The <i>Certificate</i> may not be used for the requested operation.
Bad_CommunicationError	A low level communication error occurred.
Bad_DataTypeIdUnknown	The <i>ExtensionObject</i> cannot be (de)serialized because the data type id is not recognized.
Bad_DecodingError	Decoding halted because of invalid data in the stream.
Bad_EncodingError	Encoding halted because of invalid data in the objects being serialized.
Bad_EncodingLimitsExceeded	The message encoding/decoding limits imposed by the <i>Communication Stack</i> have been exceeded.
Bad_IdentityTokenInvalid	The user identity token is not valid.
Bad_IdentityTokenRejected	The user identity token is valid but the <i>Server</i> has rejected it.
Bad_InternalError	An internal error occurred as a result of a programming or configuration error.
Bad_InvalidArgument	One or more arguments are invalid. Each service defines parameter-specific <i>StatusCodes</i> and these <i>StatusCodes</i> shall be used instead of this general error code. This error code shall be used only by the <i>Communication Stack</i> and in services where it is defined in the list of valid <i>StatusCodes</i> for the service.
Bad_InvalidState	The operation cannot be completed because the object is closed, uninitialized or in some other invalid state.
Bad_InvalidTimestamp	The timestamp is outside the range allowed by the <i>Server</i> .
Bad_LicenseExpired	The UA Server requires a license to operate in general or to perform a service or operation, but the existing license is expired.
Bad_LicenseLimitsExceeded	The UA Server has limits on number of allowed operations / objects, based on installed licenses, and these limits were exceeded.
Bad_LicenseNotAvailable	The UA Server does not have a license which is required to operate in general or to perform a service or operation.
Bad_NothingToDo	There was nothing to do because the <i>Client</i> passed a list of operations with no elements.
Bad_OutOfMemory	Not enough memory to complete the operation.

Symbolic Id	Description
Bad_RequestCancelledByClient	The request was cancelled by the client.
Bad_RequestTooLarge	The request message size exceeds limits set by the <i>Server</i> .
Bad_ResponseTooLarge	The response message size exceeds limits set by the client.
Bad_RequestHeaderInvalid	The header for the request is missing or invalid.
Bad_ResourceUnavailable	An operating system resource is not available.
Bad_SecureChannelIdInvalid	The specified secure channel is no longer valid.
Bad_SecurityChecksFailed	An error occurred while verifying security.
Bad_ServerHalted	The <i>Server</i> has stopped and cannot process any requests.
Bad_ServerNotConnected	The operation could not complete because the <i>Client</i> is not connected to the <i>Server</i> .
Bad_ServerUriInvalid	The <i>Server</i> URI is not valid.
Bad_ServiceUnsupported	The <i>Server</i> does not support the requested service.
Bad_SessionIdInvalid	The <i>Session</i> id is not valid.
Bad_SessionClosed	The <i>Session</i> was closed by the client.
Bad_SessionNotActivated	The <i>Session</i> cannot be used because <i>ActivateSession</i> has not been called.
Bad_Shutdown	The operation was cancelled because the application is shutting down.
Bad_SubscriptionIdInvalid	The subscription id is not valid.
Bad_Timeout	The operation timed out.
Bad_TimestampsToReturnInvalid	The timestamps to return parameter is invalid.
Bad_TooManyOperations	The request could not be processed because it specified too many operations.
Bad_UnexpectedError	An unexpected error occurred.
Bad_UnknownResponse	An unrecognized response was received from the <i>Server</i> .
Bad_UserAccessDenied	User does not have permission to perform the requested operation.
Bad_ViewIdUnknown	The view id does not refer to a valid view Node.
Bad_ViewTimestampInvalid	The view timestamp is not available or not supported.
Bad_ViewParameterMismatchInvalid	The view parameters are not consistent with each other.
Bad_ViewVersionInvalid	The view version is not available or not supported.

Table 178 defines the common *StatusCodes* for all operation level results used in more than one service. It does not provide a complete list. IEC 62541-6 maps the symbolic names to a numeric value and provides a complete list of *StatusCodes* including codes defines in other parts. The common *Service* result codes can be also contained in the operation level.

Table 178 – Common Operation Level result codes

Symbolic Id	Description
Good_Clamped	The value written was accepted but was clamped.
Good_Overload	Sampling has slowed down due to resource limitations.
Uncertain	The value is uncertain but no specific reason is known.
Bad	The value is bad but no specific reason is known.
Bad_AttributeIdInvalid	The attribute is not supported for the specified node.
Bad_BrowseDirectionInvalid	The browse direction is not valid.
Bad_BrowseNameInvalid	The browse name is invalid.
Bad_ContentFilterInvalid	The content filter is not valid.
Bad_ContinuationPointInvalid	The continuation point provided is no longer valid. This status is returned if the continuation point was deleted or the address space was changed between the browse calls.
Bad_DataEncodingInvalid	The data encoding is invalid. This result is used if no <i>dataEncoding</i> can be applied because an <i>Attribute</i> other than <i>Value</i> was requested or the <i>DataType</i> of the <i>Value Attribute</i> is not a subtype of the <i>Structure DataType</i> .
Bad_DataEncodingUnsupported	The <i>Server</i> does not support the requested data encoding for the node. This result is used if a <i>dataEncoding</i> can be applied but the passed data encoding is not known to the <i>Server</i> .
Bad_EventFilterInvalid	The event filter is not valid.
Bad_FilterNotAllowed	A monitoring filter cannot be used in combination with the attribute specified.
Bad_FilterOperandInvalid	The operand used in a content filter is not valid.
Bad_HistoryOperationInvalid	The history details parameter is not valid.
Bad_HistoryOperationUnsupported	The <i>Server</i> does not support the requested operation.
Bad_IndexRangeInvalid	The syntax of the index range parameter is invalid.
Bad_IndexRangeNoData	No data exists within the range of indexes specified.
Bad_MonitoredItemFilterInvalid	The monitored item filter parameter is not valid.
Bad_MonitoredItemFilterUnsupported	The <i>Server</i> does not support the requested monitored item filter.
Bad_MonitoredItemIdInvalid	The monitoring item id does not refer to a valid monitored item.
Bad_MonitoringModeInvalid	The monitoring mode is invalid.
Bad_NoCommunication	Communication with the data source is defined, but not established, and there is no last known value available. This status/sub-status is used for cached values before the first value is received or for Write and Call if the communication is not established.
Bad_NoContinuationPoints	The operation could not be processed because all continuation points have been allocated.
Bad_NodeClassInvalid	The node class is not valid.
Bad_NodeIdInvalid	The syntax of the node id is not valid.
Bad_NodeIdUnknown	The node id refers to a node that does not exist in the <i>Server</i> address space.
Bad_NoDeleteRights	The <i>Server</i> will not allow the node to be deleted.
Bad_NodeNotInView	The <i>nodesToBrowse</i> is not part of the view.
Bad_NotFound	A requested item was not found or a search operation ended without success.
Bad_NotImplemented	Requested operation is not implemented.
Bad_NotReadable	The access level does not allow reading or subscribing to the <i>Node</i> .

Symbolic Id	Description
Bad_NotSupported	The requested operation is not supported.
Bad_NotWritable	The access level does not allow writing to the <i>Node</i> .
Bad_ObjectDeleted	The <i>Object</i> cannot be used because it has been deleted.
Bad_OutOfRange	The value was out of range.
Bad_ReferenceTypeInvalid	The reference type id does not refer to a valid reference type node.
Bad_SecurityModeInsufficient	The <i>SecurityPolicy</i> and/or <i>MessageSecurityMode</i> do not match the <i>Server</i> requirements to complete the operation. For example, a user may have the right to receive the data but the data can only be transferred through an encrypted channel with an appropriate <i>SecurityPolicy</i> .
Bad_SourceNodeIdInvalid	The source node id does not refer to a valid node.
Bad_StructureMissing	A mandatory structured parameter was missing or null.
Bad_TargetNodeIdInvalid	The target node id does not refer to a valid node.
Bad_TypeDefinitionInvalid	The type definition node id does not reference an appropriate type node.
Bad_TypeMismatch	The value supplied for the attribute is not of the same type as the attribute's value.
Bad_WaitingForInitialData	Waiting for the <i>Server</i> to obtain values from the underlying data source. After creating a <i>MonitoredItem</i> or after setting the <i>MonitoringMode</i> from DISABLED to REPORTING or SAMPLING, it may take some time for the <i>Server</i> to actually obtain values for these items. In such cases the <i>Server</i> can send a <i>Notification</i> with this status prior to the <i>Notification</i> with the first value or status from the data source.

7.35 TimestampsToReturn

The *TimestampsToReturn* is an enumeration that specifies the *Timestamp Attributes* to be transmitted for *MonitoredItems* or *Nodes* in *Read* and *HistoryRead*. The values of this parameter are defined in Table 179.

Table 179 – TimestampsToReturn values

Value	Description
SOURCE_0	Return the source timestamp.
SERVER_1	Return the <i>Server</i> timestamp.
BOTH_2	Return both the source and <i>Server</i> timestamps.
NEITHER_3	Return neither timestamp. This is the default value for <i>MonitoredItems</i> if a <i>Variable</i> value is not being accessed. For <i>HistoryRead</i> this is not a valid setting.

7.36 UserIdentityToken parameters

7.36.1 Overview

The *UserIdentityToken* structure used in the *Server Service Set* allows *Clients* to specify the identity of the user they are acting on behalf of. The exact mechanism used to identify users depends on the system configuration. The different types of identity tokens are based on the most common mechanisms that are used in systems today. Table 180 defines the current set of user identity tokens. The *ExtensibleParameter* type is defined in 7.12.

Table 180 – UserIdentityToken parameterTypeIds

Symbolic Id	Description
AnonymousIdentityToken	No user information is available.
UserNameIdentityToken	A user identified by user name and password.
X509IdentityToken	A user identified by an X.509 v3 <i>Certificate</i> .
IssuedIdentityToken	A user identified by a token issued by an external <i>Authorization Service</i> .

7.36.2 Token Encryption and Proof of Possession

7.36.2.1 Overview

The *Client* shall always prove possession of a *UserIdentityToken* when it passes it to the *Server*. Some tokens include a secret such as a password which the *Server* will accept as proof. In order to protect these secrets the *Token* may be encrypted before it is passed to the *Server*. Other types of tokens allow the *Client* to create a signature with the secret associated with the *Token*. In these cases, the *Client* proves possession of a *UserIdentityToken* by creating a signature with the secret and passing it to the *Server*.

Each *UserIdentityToken* allowed by an *Endpoint* shall have a *UserTokenPolicy* specified in the *EndpointDescription*. The *UserTokenPolicy* specifies what *SecurityPolicy* to use when encrypting or signing. If this *SecurityPolicy* is omitted then the *Client* uses the *SecurityPolicy* in the *EndpointDescription*. If the matching *SecurityPolicy* is set to *None* then no encryption or signature is required. The possible *SecurityPolicies* are defined in IEC 62541-7.

It is recommended that applications never set the *SecurityPolicy* to *None* for *UserIdentityTokens* that include a secret because these secrets could be used by an attacker to gain access to the system.

Clients shall validate the *Server Certificate* and ensure it is trusted before sending a *UserIdentityToken* encrypted with the *Certificate*.

The encrypted secret and *Signature* are embedded in a *ByteString* which is part of the *UserIdentityToken*. The format of this *ByteString* depends on the type of *UserIdentityToken* and the *SecurityPolicy*.

The legacy token secret format defined in 7.36.2.2 is not extensible and provides only encryption but the encrypted data is not signed. It is used together with the *USERNAME_1 UserIdentityToken*. The password secret exchanged with this format shall not exceed 64 bytes.

The *EncryptedSecret* format defined in 7.36.2.3 provides an extensible secret format together with the definition how the secret is signed and encrypted. It allows for the layout to be updated as new token types are defined or new *SecurityPolicies* are added.

The *UserIdentityToken* types and the token formats supported by the *Endpoint* are identified by the *UserTokenPolicy* defined in 7.37.

To prevent the leakage of information useful to attackers, *Servers* shall ensure that the process of validating *UserIdentityTokens* completes in a fixed interval independent of whether an error occurs or not. The process of validation includes decrypting, check for padding and checking for a valid nonce. If any errors occur the return code is *Bad_IdentityTokenInvalid*.

Servers shall log details of any failure to validate a *UserIdentityToken* and shall lock out *Client* applications after five failures.

7.36.2.2 Legacy Encrypted Token Secret Format

When encrypting a *UserIdentityToken*, the Client appends the last *ServerNonce* to the secret. The data is then encrypted with the public key from the *Server's Certificate*.

A *Client* should not add any padding after the secret. If a *Client* adds padding then all bytes shall be zero. A *Server* shall check for padding added by *Clients* and ensure that all padding bytes are zeros. *Servers* shall reject *UserIdentityTokens* with invalid padding. Administrators shall be able to configure *Servers* to accept *UserIdentityTokens* with invalid padding.

If no encryption is applied, the structure is not used and only the secret without any *Nonce* is passed to the *Server*.

Table 181 describes how to serialize *UserIdentityTokens* before applying encryption.

Table 181 – Legacy UserIdentityToken Encrypted Token Secret Format

Name	Type	Description
Length	Byte [4]	The length of the data to be encrypted including the <i>ServerNonce</i> but excluding the <i>length</i> field. This field is a 4-byte unsigned integer encoded with the least significant bytes appearing first.
tokenData	Byte [*]	The token data.
serverNonce	Byte [*]	The last <i>ServerNonce</i> returned by the <i>Server</i> in the <i>CreateSession</i> or <i>ActivateSession</i> response.

7.36.2.3 EncryptedSecret Format

The *EncryptedSecret* uses an extensible format which has the *TypeId* of a *DataType Node* as a prefix as defined for the *ExtensionObject* encoding in IEC 62541-6. The general layout of the *EncryptedSecret* is shown in Figure 39.




Figure 39 – EncryptedSecret layout

The *TypeId* specifies how the *EncryptedSecret* is serialized and secured. For example, *RsaEncryptedSecrets* require that the policy header be encrypted with the public key associated with the *EncryptingCertificate* before it is serialized. Other *TypeIds* may or may not require the policy header to be encrypted.

The *SecurityPolicyUri* is used to determine what algorithms were used to encrypt and sign the data.

The payload is always encrypted using the symmetric encryption algorithm specified by the *SecurityPolicyUri*. The *EncryptingKey* and *InitializationVector* are used to initialize this algorithm. The mechanisms used to initialize the symmetric encryption algorithm depend on the *TypeId*. The lengths of the fields are specified by the *SecurityPolicyUri*.

The *Signature* is used for the symmetric signature algorithm specified by the *SecurityPolicyUri*. The *SymmetricKey* and *InitializationVector* are used to initialize this algorithm. The mechanisms used to initialize the symmetric signature algorithm depend on the *TypeId*. The length of the *SymmetricKey* is specified by the *SecurityPolicyUri*.

The *EncryptedSecret* is secured and serialized as follows:

- serialize the common header;
- serialize the policy header;
- encrypt the policy header and append the result to the common header;
- update the *PolicyHeaderLength* with the length of the encrypted header;
- append the *Nonce* and the *Secret* to the encrypted policy header;
- calculate padding required on the payload and append after the *Secret*;
- encrypt the payload;
- Calculate a *Signature* on {common header | encrypted policy header | encrypted payload};
- append the *Signature*.

Individual fields are serialized using the UA Binary encoding (see IEC 62541-6) for the *DataType* specified in Table 182. The *Padding* is used to ensure there is enough data to fill an integer multiple of encryption blocks. The size of the encryption block depends on the encryption algorithm. Two separate padding operations are needed because two different encryption algorithms may be used. The total length of the *Padding*, not including the *PaddingSize*, is encoded as a *UInt16*. The individual bytes of the *Padding* are set to the least significant byte of the *PaddingSize*.

The *EncryptedSecret* is deserialized and validated as follows:

- deserialize the common header;
- decrypt the policy header using the private key associated with the *EncryptingCertificate*;
- verify the padding in the policy header;
- verify the *Signature* with the *SigningKey*;
- decrypt the payload;
- verify the padding on the payload;
- extract the *Secret*;

If the *TypeId* does not require the policy header to be encrypted then the padding on the policy header is omitted and the *PolicyHeaderLength* specifies the length of the unencrypted data.

The fields in the *EncryptedSecret* are described in Table 182. The first three fields *TypeId*, *EncodingMask* and *Length* belong to the *ExtensionObject* encoding defined in IEC 62541-6.

Table 182 – EncryptedSecret layout

Name	Type	Description
TypeId	NodeId	The <i>NodeId</i> of the <i>DataType Node</i> .
EncodingMask	Byte	This value is always 1.
Length	Int32	The length of the data that follows including the <i>Signature</i> .
SecurityPolicyUri	String	The URI for the <i>SecurityPolicy</i> used to apply security.
EncryptingCertificate	ByteString	The SHA1 thumbprint of the DER form of the encrypting <i>Certificate</i> .
SigningTime	DateTime	When the <i>Signature</i> was created.
PolicyHeaderLength	UInt16	The length of the policy header that follows. If the policy header is encrypted this is the length of the encrypted data. Otherwise, it is the length of the unencrypted data.
SigningKey	ByteString	The key data used to create the <i>Signature</i> . The <i>TypeId</i> specifies how this data is used.
EncryptingKey	ByteString	The key data used to encrypt the payload. The <i>TypeId</i> specifies how this data is used.
InitializationVector	ByteString	The data used to initialize the algorithm used to encrypt the payload. The <i>TypeId</i> specifies how this data is used.
Nonce	ByteString	This is the last <i>serverNonce</i> returned in the <i>CreateSession</i> or <i>ActivateSession Response</i> when a <i>UserIdentityToken</i> is passed with the <i>ActivateSession Request</i> . If used outside of an <i>ActivateSession</i> call, the source of the <i>Nonce</i> is defined by the context in which this <i>EncryptedSecret</i> is used. The length of the <i>Nonce</i> shall equal the <i>SecureChannelNonceLength</i> specified by the <i>SecurityPolicy</i> .
Secret	ByteString	The <i>tokenData</i> that depends on the <i>IssuedIdentityToken</i> . If the <i>tokenData</i> is a <i>String</i> it is encoded using UTF-8 first.
PayloadPadding	Byte[*]	Additional padding added to ensure the size of the encrypted payload is an integer multiple of the input block size for the symmetric encryption algorithm specified by the <i>SecurityPolicyUri</i> . The value of each byte is the least significant byte of the <i>PayloadPaddingSize</i> .
PayloadPaddingSize	UInt16	The size of the padding added to the payload.
Signature	Byte[*]	The signature calculated using the symmetric signing algorithm specified by the <i>SecurityPolicyUri</i> . The length of the signature is specified by the <i>SecurityPolicyUri</i> .

The currently available *EncryptedSecret DataTypes* are defined in Table 183.

Table 183 – EncryptedSecret DataTypes

Type Name	When to Use
RsaEncryptedSecret	Used when the <i>SecurityPolicy</i> requires the use of RSA based asymmetric encryption. It is described in 7.36.2.4.

7.36.2.4 RsaEncryptedSecret DataType

The *RsaEncryptedSecret* uses RSA based asymmetric encryption to encrypt the policy header.

Additional semantics for the fields in the *EncryptedSecret* layout for the *RsaEncryptedSecret Structure* are described in Table 184.

Table 184 – RsaEncryptedSecret structure

Name	Type	Description
TypeId	NodeId	The <i>NodeId</i> of the <i>RsaEncryptedSecret DataType Node</i> .
EncodingMask	Byte	See Table 182.
Length	UInt32	See Table 182.
SecurityPolicyUri	String	See Table 182.
EncryptingCertificate	ByteString	See Table 182.
SigningTime	DateTime	See Table 182.
PolicyHeaderLength	UInt16	See Table 182.
SigningKey	ByteString	The key used to compute the <i>Signature</i> . See Table 182 for additional details.
EncryptingKey	ByteString	The key used to encrypt payload. See Table 182 for additional details.
InitializationVector	ByteString	The initialization vector used with the <i>EncryptingKey</i> . See Table 182 for additional details.
Nonce	ByteString	See Table 182.
Secret	ByteString	See Table 182.
PayloadPadding	Byte[*]	See Table 182.
PayloadPaddingSize	UInt16	See Table 182.
Signature	Byte[*]	See Table 182.

7.36.3 AnonymousIdentityToken

The *AnonymousIdentityToken* is used to indicate that the Client has no user credentials.

Table 185 defines the *AnonymousIdentityToken* parameter.

Table 185 – AnonymousIdentityToken

Name	Type	Description
AnonymousIdentityToken	Structure	An anonymous user identity.
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.

7.36.4 UserNameIdentityToken

The *UserNameIdentityToken* is used to pass simple username/password credentials to the *Server*.

This token shall be encrypted by the *Client* if required by the *SecurityPolicy* of the *UserTokenPolicy*. The *Server* should specify a *SecurityPolicy* for the *UserTokenPolicy* if the *SecureChannel* has a *SecurityPolicy* of *None* and no transport layer encryption is available. If *None* is specified for the *UserTokenPolicy* and *SecurityPolicy* is *None* then the password only contains the UTF-8 encoded password. The *SecurityPolicy* of the *SecureChannel* is used if no *SecurityPolicy* is specified in the *UserTokenPolicy*.

If the token is to be encrypted the password shall be converted to a UTF-8 *ByteString*, encrypted and then serialized as shown in Table 181.

The *Server* shall decrypt the password and verify the *ServerNonce*.

If the *SecurityPolicy* is *None* then the password only contains the UTF-8 encoded password. This configuration should not be used unless the network is encrypted in some other manner such as a VPN. The use of this configuration without network encryption would result in a serious security fault, in that it would cause the appearance of a secure user access, but it would make the password visible in clear text.

Table 186 defines the *UserNameIdentityToken* parameter.

Table 186 – UserNameIdentityToken

Name	Type	Description
UserNameIdentityToken	Structure	UserName value.
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.
userName	String	A string that identifies the user.
password	ByteString	The password for the user. The password can be an empty string. This parameter shall be encrypted with the Server's public key using the algorithm specified by the <i>SecurityPolicy</i> . The format used for the encrypted data is described in 7.36.2.2.
encryptionAlgorithm	String	A string containing the URI of the <i>AsymmetricEncryptionAlgorithm</i> . The URI string values are defined names that may be used as part of the security profiles specified in IEC 62541-7. This parameter is null if the password is not encrypted.

Table 187 describes the dependencies for selecting the *AsymmetricEncryptionAlgorithm* for the *UserNameIdentityToken*. The *SecureChannel SecurityPolicy* URI is specified in the *EndpointDescription* and used in subsequent *OpenSecureChannel* requests. The *UserTokenPolicy SecurityPolicy* URI is specified in the *EndpointDescription*. The *encryptionAlgorithm* is specified in the *UserNameIdentityToken* or *IssuedIdentityToken* provided by the *Client* in the *ActivateSession* call. The *SecurityPolicy* Other in Table 187 refers to any *SecurityPolicy* other than *None*. The selection of the *EncryptionAlgorithm* is based on the *UserTokenPolicy*. The *SecureChannel SecurityPolicy* is used if the *UserTokenPolicy* is null or empty.

Table 187 – EncryptionAlgorithm selection

SecureChannel SecurityPolicy	UserTokenPolicy SecurityPolicy	UserIdentityToken EncryptionAlgorithm
Security Policy - None	Null or empty	No encryption
Security Policy - None	Security Policy - None	No encryption
Security Policy - None	Security Policy - Other	Asymmetric algorithm for "Other"
Security Policy - Other	Null or empty	Asymmetric algorithm for "Other"
Security Policy - Other	Security Policy - Yet another	Asymmetric algorithm for "Yet another"
Security Policy - Other	Security Policy - Other	Asymmetric algorithm for "Other"
Security Policy - Other	Security Policy - None	No encryption

7.36.5 X509IdentityTokens

The *X509IdentityToken* is used to pass an X.509 v3 *Certificate* which is issued by the user.

This token shall always be accompanied by a *Signature* in the *userTokenSignature* parameter of *ActivateSession* if required by the *SecurityPolicy*. The *Server* should specify a *SecurityPolicy* for the *UserTokenPolicy* if the *SecureChannel* has a *SecurityPolicy* of *None*.

Table 188 defines the *X509IdentityToken* parameter.

Table 188 – X.509 v3 Identity Token

Name	Type	Description
X509IdentityToken	structure	X.509 v3 value.
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.
certificateData	ByteString	The X.509 v3 <i>Certificate</i> in DER format.

7.36.6 IssuedIdentityToken

The *IssuedIdentityToken* is used to pass *SecurityTokens* issued by an external *Authorization Service* to the *Server*. These tokens may be text or binary.

OAuth2 defines a standard for *Authorization Services* that produce JSON Web Tokens (JWT). These JWTs are passed as an *Issued Token* to an OPC UA Server which uses the signature contained in the JWT to validate the token. IEC 62541-6 describes OAuth2 and JWTs in more detail. If the token is encrypted, it shall use the *EncryptedSecret* format defined in 7.36.2.3.

This token shall be encrypted by the *Client* if required by the *SecurityPolicy* of the *UserTokenPolicy*. The *Server* should specify a *SecurityPolicy* for the *UserTokenPolicy* if the *SecureChannel* has a *SecurityPolicy* of *None* and no transport layer encryption is available. The *SecurityPolicy* of the *SecureChannel* is used if no *SecurityPolicy* is specified in the *UserTokenPolicy*.

If the *SecurityPolicy* is not *None*, the *tokenData* shall be encoded in UTF-8 (if it is not already binary), signed and encrypted according to the rules specified for the *tokenType* of the associated *UserTokenPolicy* (see 7.37).

If the *SecurityPolicy* is *None* then the *tokenData* only contains the UTF-8 encoded *tokenData*. This configuration should not be used unless the network is encrypted in some other manner such as a VPN. The use of this configuration without network encryption would result in a serious security fault, in that it would cause the appearance of a secure user access, but it would make the token visible in clear text.

Table 189 defines the *IssuedIdentityToken* parameter.

Table 189 – IssuedIdentityToken

Name	Type	Description
IssuedIdentityToken	structure	The token provided by an <i>Authorization Service</i> .
policyId	String	An identifier for the <i>UserTokenPolicy</i> that the token conforms to. The <i>UserTokenPolicy</i> structure is defined in 7.37.
tokenData	ByteString	The text or binary representation of the token. The format of the data depends on the associated <i>UserTokenPolicy</i> .
encryptionAlgorithm	String	The URI of the <i>AsymmetricEncryptionAlgorithm</i> . The list of OPC UA-defined names that may be used is specified in IEC 62541-7. See Table 187 for details on picking the correct URI. This parameter is null if the <i>tokenData</i> is not encrypted or if the <i>EncryptedSecret</i> format is used.

7.37 UserTokenPolicy

The components of this parameter are defined in Table 190.

Table 190 – UserTokenPolicy

Name	Type	Description
UserTokenPolicy	structure	Specifies a <i>UserIdentityToken</i> that a <i>Server</i> will accept.
policyId	String	An identifier for the <i>UserTokenPolicy</i> assigned by the <i>Server</i> . The <i>Client</i> specifies this value when it constructs a <i>UserIdentityToken</i> that conforms to the policy. This value is only unique within the context of a single <i>Server</i> .
tokenType	Enum <i>UserIdentityTokenType</i>	The type of user identity token required. This value is an enumeration with one of the following values: ANONYMOUS_0 No token is required. USERNAME_1 A username/password token. CERTIFICATE_2 An X.509 v3 <i>Certificate</i> token. ISSUEDTOKEN_3 Any token issued by an <i>Authorization Service</i> . A <i>tokenType</i> of ANONYMOUS indicates that the <i>Server</i> does not require any user identification. In this case, the <i>Client Application Instance Certificate</i> is used as the user identification.
issuedTokenType	String	A URI for the type of token. IEC 62541-6 defines URIs for common issued token types. Vendors may specify their own token types. This field may only be specified if <i>TokenType</i> is ISSUEDTOKEN_3.
issuerEndpointUrl	String	An optional string which depends on the <i>Authorization Service</i> . The meaning of this value depends on the <i>issuedTokenType</i> . Further details for the different token types are defined in IEC 62541-6. For Kerberos this string is the name of the Service Principal Name (SPN). For JWTs this is a JSON object with fields defined in IEC 62541-6.
securityPolicyUri	String	The security policy to use when encrypting or signing the <i>UserIdentityToken</i> when it is passed to the <i>Server</i> in the <i>ActivateSession</i> request. Clause 7.36 describes how this parameter is used. The security policy for the <i>SecureChannel</i> is used if this value is null or empty.

7.38 VersionTime

This primitive data type is a UInt32 that represents the time in seconds since the year 2000. The epoch date is midnight UTC (00:00) on January 1, 2000.

It is used as version number based on the last change time. If the version is updated, the new value shall be greater than the previous value.

If a *Variable* is initialized with a *VersionTime* value, the value shall be either loaded from persisted configuration or time synchronization shall be available to ensure a unique version is applied.

The value 0 is used to indicate that no version information is available.

7.39 ViewDescription

The components of this parameter are defined in Table 191.

Table 191 – ViewDescription

Name	Type	Description
ViewDescription	structure	Specifies a <i>View</i> .
viewId	NodeId	<i>NodeId</i> of the <i>View</i> to <i>Query</i> . A null value indicates the entire <i>AddressSpace</i> .
timestamp	UtcTime	The time date desired. The corresponding version is the one with the closest previous creation timestamp. Either the <i>Timestamp</i> or the <i>viewVersion</i> parameter may be set by a <i>Client</i> , but not both. If <i>ViewVersion</i> is set this parameter shall be null.
viewVersion	UInt32	The version number for the <i>View</i> desired. When <i>Nodes</i> are added to or removed from a <i>View</i> , the value of a <i>View</i> 's <i>ViewVersion Property</i> is updated. Either the <i>Timestamp</i> or the <i>viewVersion</i> parameter may be set by a <i>Client</i> , but not both. The <i>ViewVersion Property</i> is defined in IEC 62541-3. If <i>timestamp</i> is set this parameter shall be 0. The current view is used if <i>timestamp</i> is null and <i>viewVersion</i> is 0.

Annex A (informative)

BNF definitions

A.1 Overview over BNF

The BNF (Backus-Naur form) used in this annex uses '<' and '>' to mark symbols, '[' and ']' to identify optional paths and '|' to identify alternatives. If the '(' and ')' symbols are used, it indicates sets.

A.2 BNF of RelativePath

A *RelativePath* is a structure that describes a sequence of *References* and *Nodes* to follow. This annex describes a text format for a *RelativePath* that can be used in documentation or in files used to store configuration information.

The components of a *RelativePath* text format are specified in Table A.1.

Table A.1 – RelativePath

Symbol	Meaning										
/	The forward slash character indicates that the <i>Server</i> is to follow any subtype of <i>HierarchicalReferences</i> .										
.	The period (dot) character indicates that the <i>Server</i> is to follow any subtype of a <i>Aggregates ReferenceType</i> .										
<[#!ns:]ReferenceType>	A string delimited by the '<' and '>' symbols specifies the <i>BrowseName</i> of a <i>ReferenceType</i> to follow. By default, any <i>References</i> of the subtypes the <i>ReferenceType</i> are followed as well. A '#' placed in front of the <i>BrowseName</i> indicates that subtypes should not be followed. A '!' in front of the <i>BrowseName</i> is used to indicate that the inverse <i>Reference</i> should be followed. The <i>BrowseName</i> may be qualified with a namespace index (indicated by a numeric prefix followed by a colon). This namespace index is used specify the namespace component of the <i>BrowseName</i> for the <i>ReferenceType</i>. If the namespace prefix is omitted then namespace index 0 is used.										
[ns:]BrowseName	A string that follows a '/', '.', or '>' symbol specifies the <i>BrowseName</i> of a target <i>Node</i> to return or follow. This <i>BrowseName</i> may be prefixed by its namespace index. If the namespace prefix is omitted then namespace index 0 is used. Omitting the final <i>BrowseName</i> from a path is equivalent to a wildcard operation that matches all <i>Nodes</i> which are the target of the <i>Reference</i> specified by the path.										
&	The & sign character is the escape character. It is used to specify reserved characters that appear within a <i>BrowseName</i>. A reserved character is escaped by inserting the '&' in front of it. Examples of <i>BrowseNames</i> with escaped characters are: <table> <tr> <th>Received browse path name</th><th>Resolves to</th></tr> <tr> <td>"&/Name_1"</td><td>"/Name_1"</td></tr> <tr> <td>"&.Name_2"</td><td>".Name_2"</td></tr> <tr> <td>"&:Name_3"</td><td>":Name_3"</td></tr> <tr> <td>"&&Name_4"</td><td>"&Name_4"</td></tr> </table>	Received browse path name	Resolves to	"&/Name_1"	"/Name_1"	"&.Name_2"	".Name_2"	"&:Name_3"	":Name_3"	"&&Name_4"	"&Name_4"
Received browse path name	Resolves to										
"&/Name_1"	"/Name_1"										
"&.Name_2"	".Name_2"										
"&:Name_3"	":Name_3"										
"&&Name_4"	"&Name_4"										

Table A.2 provides *RelativePaths* examples in text format.

Table A.2 – RelativePath Examples

Browse Path	Description
"/2:Block&.Output"	Follows any forward hierarchical <i>Reference</i> with target <i>BrowseName</i> = "2:Block.Output".
"/3:Truck.0:NodeVersion"	Follows any forward hierarchical <i>Reference</i> with target <i>BrowseName</i> = "3:Truck" and from there a forward <i>Aggregates Reference</i> to a target with <i>BrowseName</i> "0:NodeVersion".
"<1:ConnectedTo>1:Boiler/1:HeatSensor"	Follows any forward <i>Reference</i> with a <i>BrowseName</i> = '1:ConnectedTo' and finds targets with <i>BrowseName</i> = '1:Boiler'. From there follows any hierarchical <i>Reference</i> and find targets with <i>BrowseName</i> = '1:HeatSensor'.
"<1:ConnectedTo>1:Boiler/"	Follows any forward <i>Reference</i> with a <i>BrowseName</i> = '1:ConnectedTo' and finds targets with <i>BrowseName</i> = '1:Boiler'. From there it finds all targets of hierarchical <i>References</i> .
"<0:HasChild>2:Wheel"	Follows any forward <i>Reference</i> with a <i>BrowseName</i> = 'HasChild' and qualified with the default OPC UA namespace. Then find targets with <i>BrowseName</i> = 'Wheel' qualified with namespace index '2'.
"<!HasChild>Truck"	Follows any inverse <i>Reference</i> with a <i>BrowseName</i> = 'HasChild'. Then find targets with <i>BrowseName</i> = 'Truck'. In both cases, the namespace component of the <i>BrowseName</i> is assumed to be 0.
"<0:HasChild>"	Finds all targets of forward <i>References</i> with a <i>BrowseName</i> = 'HasChild' and qualified with the default OPC UA namespace.

The following BNF describes the syntax of the *RelativePath* text format.

```

<relative-path> ::= <reference-type> <browse-name> [relative-path]
<reference-type> ::= '/' | '.' | '<' | '#' | '!' | <browse-name> '>'
<browse-name> ::= [<namespace-index> ':' ] <name>
<namespace-index> ::= <digit> [<digit>]
<digit> ::= '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9'
<name> ::= (<name-char> | '&' <reserved-char>) [<name>]
<reserved-char> ::= '/' | '.' | '<' | '>' | ':' | '#' | '!' | '&'
<name-char> ::= All valid characters for a String (see IEC 62541-3) excluding reserved-chars.

```

A.3 BNF of NumericRange

The following BNF describes the syntax of the *NumericRange* parameter type.

```

<numeric-range> ::= <dimension> [',' <dimension>]
<dimension> ::= <index> [':' <index>]
<index> ::= <digit> [<digit>]
<digit> ::= '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9'

```

Annex B (informative)

ContentFilter and Query examples

B.1 Simple ContentFilter examples

B.1.1 Overview

These examples provide fairly simple *ContentFilters*. Filter similar to these examples may be used in processing events.

The following conventions apply to these examples with regard to how *Attribute* operands are used (for a definition of this operand see 7.4.4):

- **AttributeOperand:** Refers to a *Node*, an *Attribute* of a *Node* or the *Value Attribute* of a *Property* associated with a *Node*. In the examples, the character names of NodeIds are used instead of an actual nodeId, this also applies to Attribute Ids.
- The string representation of relative paths is used instead of the actual structure.
- The NamespaceIndex used in all examples is 12 (it could just as easily have been 4 or 23 or any value). For more information about NamespaceIndex, see IEC 62541-3. The use of the NamespaceIndex illustrates that the information model being used in the examples is not a model defined by this document, but one created for the examples.

B.1.2 Example 1

For example the logic described by '(((AType.A = 5) or InList(BType.B, 3,5,7)) and BaseObjectType.displayName LIKE "Main%")' would result in a logic tree as shown in Figure B.1 and a *ContentFilter* as shown in Table B.1. For this example to return anything AType and BType both shall be subtypes of BaseObjectType, or the resulting "And" operation would always be false.




Figure B.1 – Filter logic tree example

Table B.1 describes the elements, operators and operands used in the example.

Table B.1 – ContentFilter example

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	And	ElementOperand = 1	Element Operand = 4		
1	Or	ElementOperand = 2	Element Operand = 3		
2	Equals	AttributeOperand = NodeId: AType, BrowsePath: ".12:A", Attribute:value	LiteralOperand = '5'		
3	InList	AttributeOperand = NodeId: BType, BrowsePath: ".12:B", Attribute:value	LiteralOperand = '3'	LiteralOperand = '5'	LiteralOperand = '7'
4	Like	AttributeOperand = NodeId: BaseObjectType, BrowsePath: ".", Attribute: displayName	LiteralOperand = "Main%"		

B.1.3 Example 2

As another example a filter to select all *SystemEvents* (including derived types) that are contained in the *Area1 View* or the *Area2 View* would result in a logic tree as shown in Figure B.2 and a ContentFilter as shown in Table B.2.


**Figure B.2 – Filter logic tree example**

Table B.2 describes the elements, operators and operands used in the example.

Table B.2 – ContentFilter example

Element[]	Operator	Operand[0]	Operand[1]
0	And	ElementOperand = 1	ElementOperand = 4
1	Or	ElementOperand = 2	ElementOperand = 3
2	InView	AttributeOperand = NodeId: Area1, BrowsePath: ".", Attribute: NodeId	
3	InView	AttributeOperand = NodeId: Area2, BrowsePath: ".", Attribute: NodeId	
4	OfType	AttributeOperand = NodeId: SystemEventType, BrowsePath: ".", Attribute: NodeId	

B.2 Complex examples of Query filters

B.2.1 Overview

These Query examples illustrate complex filters. The following conventions apply to these examples with regard to Attribute operands (for a definition of these operands, see 7.4.4).

- **AttributeOperand:** Refers to a *Node*, an *Attribute* of a *Node* or the *Value Attribute* of a *Property* associated with a *Node*. In the examples character names of *ExpandedNodeId* are used instead of an actual *ExpandedNodeId*, this also applies to *Attribute* Ids.
- The string representation of relative paths is used instead of the actual structure.
- The *NamespaceIndex* used in all examples is 12 (it could just as easily have been 4 or 23 or any value). For more information about *NamespaceIndex*, see IEC 62541-3. The use of the *NamespaceIndex* illustrates that the information model being used in the examples is not a model defined by this document, but one created for the examples.

B.2.2 Used type model

The following examples use the type model described below. All *Property* values are assumed to be string unless otherwise noted.

New Reference types:

- "HasChild" derived from HierarchicalReference.
- "HasAnimal" derived from HierarchicalReference.
- "HasPet" derived from HasAnimal.
- "HasFarmAnimal" derived from HasAnimal.
- "HasSchedule" derived from HierarchicalReference.

PersonType derived from BaseObjectType adds:

- HasProperty "LastName".
- HasProperty "FirstName".
- HasProperty "StreetAddress".
- HasProperty "City".
- HasProperty "ZipCode".
- May have HasChild reference to a node of type PersonType.
- May have HasAnimal reference to a node of type AnimalType (or a subtype of this Reference type).

AnimalType derived from BaseObjectType adds:

- May have HasSchedule reference to a node of type FeedingScheduleType.
- HasProperty "Name".

DogType derived from AnimalType adds:

- HasProperty "NickName".
- HasProperty "DogBreed".
- HasProperty "License".

CatType derived from AnimalType adds:

HasProperty "NickName".

HasProperty "CatBreed".

PigType derived from AnimalType adds:

HasProperty "PigBreed".

ScheduleType derived from BaseObjectType adds:

HasProperty "Period".

FeedingScheduleType derived from ScheduleType adds:

HasProperty "Food".

HasProperty "Amount" (Stored as an *Int32*).

AreaType derived from *BaseObjectType* is just a simple *Folder* and contains no *Properties*.

This example type system is shown in Figure B.3. In Figure B.3, the OPC UA notation is used for all *References* to *ObjectTypes*, *Variables*, *Properties* and subtypes. Additionally, supported *References* are contained in an inner box. The actual references only exist in the instances, thus, no connections to other *Objects* are shown in Figure B.3 and they are subtypes of the listed *Reference*.





Figure B.3 – Example Type Nodes

A corresponding example set of instances is shown in Figure B.4. These instances include a type *Reference* for *Objects*. Properties also have type *References*, but the *References* are omitted for simplicity. The name of the *Object* is provided in the box and a numeric instance *NodeId* in brackets. Standard *ReferenceTypes* use the OPC UA notation, custom *ReferenceTypes* are listed as a named *Reference*. For *Properties*, the *BrowseName*, *NodeId*, and *Value* are shown. The *Nodes* that are included in a *View* (View1) are enclosed in the coloured box. Two *Area* nodes are included for grouping of the existing person nodes. All custom nodes are defined in namespace 12 which is not included in Figure B.4.



IEC

Figure B.4 – Example Instance Nodes

B.2.3 Example Notes

For all of the examples in 7.4.4, the type definition *Node* is listed in its symbolic form, in the actual call it would be the *ExpandedNodeId* assigned to the *Node*. The *Attribute* is the symbolic name of the *Attribute*, in the actual call they would be translated to the *IntegerId* of the *Attribute*. Also in all of the examples the *BrowseName* is included in the result table for clarity; normally this would not be returned.

All of the examples include the following items:

- an English description of the object of the query;
- an SQL like description of the query;
- a table that has a *NodeTypeDescription* of the items that are to be returned;
- a figure illustrating the query filter;
- a table describing the content filter;
- a table describing the resulting dataset.

The examples assume namespace 12 is the namespace for all of the custom definitions described for the examples.

B.2.4 Example 1

This example requests a simple layered filter, a person has a pet and the pet has a schedule.

Example 1: Get PersonType.LastName, AnimalType.Name, ScheduleType.Period where the Person Has a Pet and that Pet Has a Schedule.

The *NodeTypeDescription* parameters used in the example are described in Table B.3.

Table B.3 – Example 1 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	value	N/A
		"<12:HasPet>12:AnimalType. 12:Name"	value	N/A
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:Schedule. 12:Period"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.5.




Figure B.5 – Example 1 Filter

Table B.4 describes the *ContentFilter* elements, operators and operands used in the example.

Table B.4 – Example 1 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
1	RelatedTo_15	AttributeOperand = NodeId: PersonType, BrowsePath ".", Attribute: NodeId	ElementOperand = 2	AttributeOperand = NodeId: HasPet, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	RelatedTo_15	AttributeOperand = NodeId: AnimalType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: ScheduleType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

Table B.5 describes the *QueryDataSet* that results from this query if it were executed against the instances described in Figure B.4

Table B.5 – Example 1 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (JFamily1)	PersonType	".12:LastName"	Jones
		"<12:HasPet>12:AnimalType. 12:Name"	Rosemary
			Basil
		"<12:HasPet>12:AnimalType<12:HasSchedule> 12:Schedule.12:Period"	Hourly
			Daily
12:42(HFamily1)	PersonType	".12:LastName"	Hervey
		"<12:HasPet>12:AnimalType. 12:Name"	Oliver
		"<12:HasPet>12:AnimalType<12:HasSchedule> 12:Schedule.12:Period"	Daily

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the QueryDataSet and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

The Value column is returned as an array for each *Node* description, where the order of the items in the array would correspond to the order of the items that were requested for the given Node Type. In Addition, if a single *Attribute* has multiple values then it would be returned as an array within the larger array, for example in this table Rosemary and Basil would be returned in a array for the .<HasPet>.AnimalType.Name item. They are show as separate rows for ease of viewing. The actual value array for JFamily1 would be ("Jones", {"RoseMary", "Basil"}, {"Hourly", "Daily"})

B.2.5 Example 2

The second example illustrates receiving a list of disjoint *Nodes* and also illustrates that an array of results can be received.

Example 2: Get PersonType.LastName, AnimalType.Name where a person has a child or (a pet is of type cat and has a feeding schedule).

The NodeTypeDescription parameters used in the example are described in Table B.6.

Table B.6 – Example 2 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	Value	N/A
AnimalType	TRUE	".12:Name"	Value	N/A

The corresponding ContentFilter is illustrated in Figure B.6.




Figure B.6 – Example 2 Filter logic tree

Table B.7 describes the elements, operators and operands used in the example. It is worth noting that a CatType is a subtype of AnimalType.

Table B.7 – Example 2 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	Or	ElementOperand=1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodId: PersonType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: PersonType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: HasChild, BrowsePath ".", Attribute: NodId	LiteralOperand = '1'
2	RelatedTo	AttributeOperand = NodId: CatType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: FeedingScheduleType, BrowsePath ".", Attribute: NodId	AttributeOperand = NodId: HasSchedule, BrowsePath ".", Attribute: NodId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.8.

Table B.8 – Example 2 QueryDataSets

NodId	TypeDefinition NodId	RelativePath	Value
12:30 (Jfamily1)	PersonType	. 12:LastName	Jones
12:42 (HFamily1)	PersonType	. 12:LastName	Hervey
12:48 (HFamily2)	PersonType	. 12:LastName	Hervey
12:70 (Cat1)	CatType	. 12:Name	Rosemary
12:74 (Cat2)	CatType	. 12:Name	Basil

NOTE The relative path column and browse name (in parentheses in the *NodId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodId* would be a *NodId* not the symbolic name that is included in the table.

B.2.6 Example 3

The third example provides a more complex *Query* in which the results are filtered on multiple criteria.

Example 3: Get PersonType.LastName, AnimalType.Name, ScheduleType.Period where a person has a pet and the animal has a feeding schedule and the person has a Zipcode = '02138' and (the Schedule.Period is Daily or Hourly) and Amount to feed is > 10.

Table B.9 describes the NodeTypeIdDescription parameters used in the example.

Table B.9 – Example 3 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		RelativePath	Attribute	Index Range
PersonType	FALSE	"12:LastName"	Value	N/A
		"<12:HasPet>12:AnimalType. 12:Name"	Value	N/A
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:FeedingSchedule.Period"	Value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.7.


**Figure B.7 – Example 3 Filter logic tree**

Table B.10 describes the elements, operators and operands used in the example.

Table B.10 – Example 3 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	And	Element Operand = 1	ElementOperand = 2		
1	And	ElementOperand = 4	ElementOperand = 6		
2	And	ElementOperand = 3	ElementOperand = 9		
3	Or	ElementOperand = 7	ElementOperand = 8		
4	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	ElementOperand = 5	AttributeOperand = NodeId: 12:HasPet, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
5	RelatedTo	AttributeOperand = Node: 12:AnimalType, BrowsePath ".", Attribute: NodeId Alias: AT	AttributeOperand = NodeId: 12:FeedingScheduleType, BrowsePath ".", Attribute: NodeId Alias: FST	AttributeOperand = NodeId: 12:HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
6	Equals	AttributeOperand = NodeId: 12:PersonType BrowsePath 12:Zipcode ".", Attribute: Value	LiteralOperand = '02138'		
7	Equals	AttributeOperand = NodeId: 12:PersonType BrowsePath "12:HasPet>12:AnimalType<12: HasSchedule>12: FeedingSchedule/12:Period", Attribute: Value Alias: FST	LiteralOperand = 'Daily'		
8	Equals	AttributeOperand = NodeId: 12:PersonType BrowsePath "12:HasPet>12:AnimalType<12: HasSchedule>12: FeedingSchedule/12:Period", Attribute: Value Alias: FST	LiteralOperand = 'Hourly'		
9	Greater Than	AttributeOperand = NodeId: 12:PersonType BrowsePath "12:HasPet>12:AnimalType<12: HasSchedule>12: FeedingSchedule/12:Amount", Attribute: Value Alias: FST	ElementOperand = 10		
10	Cast	LiteralOperand = 10	AttributeOperand = NodeId: Int32, BrowsePath ".", Attribute: NodeId		

The results from this query would contain the *QueryDataSets* shown in Table B.11.

Table B.11 – Example 3 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (JFamily1)	PersonType	".12:LastName"	Jones
		"<12:HasPet>12:PersonType. 12:Name"	Rosemary
			Basil
		"<12:HasPet>12:AnimalType<12:HasSchedule>12:FeedingSchedule. 12:Period"	Hourly
			Daily

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

B.2.7 Example 4

The fourth example provides an illustration of the Hop parameter that is part of the RelatedTo Operator.

Example 4: Get PersonType.LastName where a person has a child who has a child who has a pet.

Table B.12 describes the NodeTypeDescription parameters used in the example.

Table B.12 – Example 4 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.8.


**Figure B.8 – Example 4 Filter logic tree**

Table B.13 describes the elements, operators and operands used in the example.

Table B.13 – Example 4 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	Element Operand = 1	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '2'
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:AnimalType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasPet, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.14. It is worth noting that the pig "Pig1" is referenced as a pet by Sara, but is referenced as a farm animal by Sara's parent Paul.

Table B.14 – Example 4 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	".12:LastName"	Hervey

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

B.2.8 Example 5

The fifth example provides an illustration of the use of alias.

Example 5: Get the last names of children that have the same first name as a parent of theirs.

Table B.15 describes the *NodeTypeDescription* parameters used in the example.

Table B.15 – Example 5 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	"<12:HasChild>12:PersonType. 12:LastName"	Value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.9.




Figure B.9 – Example 5 Filter logic tree

In this example, one *Reference* to *PersonType* is aliased to "Parent" and another *Reference* to *PersonType* is aliased to "Child". The value of *Parent.firstName* and *Child.firstName* are then compared. Table B.16 describes the elements, operators and operands used in the example.

Table B.16 – Example 5 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	And	ElementOperand = 1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId, Alias: "Parent"	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId, Alias: "Child"	AttributeOperand = NodeId: 12:HasChild, Attribute: NodeId	LiteralOperand = "1"
2	Equals	AttributeOperand = NodeId: 12:PersonType, BrowsePath ""/12:FirstName", Attribute: Value, Alias: "Parent"	AttributeOperand = NodeId: 12:PersonType, BrowsePath ""/12:FirstName", Attribute: Value, Alias: "Child"		

The results from this query would contain the *QueryDataSets* shown in Table B.17.

Table B.17 – Example 5 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	"<12:HasChild>12:PersonType.12:LastName"	Hervy

NOTE The *RelativePath* column and browse name (in parentheses in the *NodeId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

B.2.9 Example 6

The sixth example provides an illustration a different type of request, one in which the *Client* is interested in displaying part of the *AddressSpace* of the *Server*. This request includes listing a *Reference* as something that is to be returned.

Example 6: Get *PersonType.NodeId*, *AnimalType.NodeId*, *PersonType.HasChild* Reference, *PersonType.HasAnimal* Reference where a person has a child who has a *Animal*.

Table B.18 describes the *NodeTypeDescription* parameters used in the example.

Table B.18 – Example 6 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	". 12:NodeId"	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>12:AnimalType.NodeId	value	N/A
		<12:HasChild>	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.10.




Figure B.10 – Example 6 Filter logic tree

Table B.19 describes the elements, operators and operands used in the example.

Table B.19 – Example 6 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	ElementOperand = 1	AttributeOperand = Node: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperan d = '1'
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:AnimalType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasAnimal, BrowsePath ".", Attribute: NodeId	LiteralOperan d = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.20.

Table B.20 – Example 6 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	".NodeId"	12:42 (HFamily1)
		<12:HasChild>12:PersonType<12:HasAnimal> 12:AnimalType.NodeId	12:91 (Pig1)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasChild>12:PersonType<12:HasAnimal>	HasFarmAnimal <i>ReferenceDescription</i>
12:48 (HFamily2)	PersonType	".NodeId"	12:48 (HFamily2)
		<12:HasChild>12:PersonType<12:HasAnimal> 12:AnimalType.NodeId	12:91 (Pig1)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasChild>12:PersonType<12:HasAnimal>	HasPet <i>ReferenceDescription</i>

NOTE The *RelativePath* and browse name (in parentheses) is not in the *QueryDataSet* and is only shown here for clarity and the *TypeDefinition NodeId* would be an integer, not the symbolic name that is included in the table. The value field would in this case be the *NodeId* where it was requested, but for the example the browse name is

provided in parentheses and in the case of *Reference* types on the browse name is provided. For the *References* listed in Table B.20, the value would be a *ReferenceDescription* which are described in 7.25.

Table B.21 provides an example of the same *QueryDataSet* as shown in Table B.20 without any additional fields and minimal symbolic Ids. There is an entry for each requested Attribute, in the cases where an Attribute would return multiple entries the entries are separated by comas. If a structure is being returned then the structure is enclosed in square brackets. In the case of a *ReferenceDescription* the structure contains a structure and *DisplayName* and *BrowseName* are assumed to be the same and defined in Figure B.4.

Table B.21 – Example 6 QueryDataSets without additional information

NodeId	TypeDefinition NodeId	Value
12:42	PersonType	12:42
		12:91
		[HasChild,TRUE,[48,HFamily2,HFamily2,PersonType]],
		[HasFarmAnimal,TRUE[91,Pig1,Pig1,PigType]]
12:48	PersonType	12:54
		12:91
		[HasChild,TRUE,[54,HFamily3,HFamily3,PersonType]]
		[HasPet, TRUE,[91,Pig1,Pig1,PigType]]

The PersonType, HasChild, PigType, HasPet, HasFarmAnimal identifiers used in the above table would be translated to actual *ExpandedNodeId*.

B.2.10 Example 7

The seventh example provides an illustration a request in which a *Client* wants to display part of the *AddressSpace* based on a starting point that was obtained via browsing. This request includes listing *References* as something that is to be returned. In this case the Person Browsed to Area2 and wanted to *Query* for information below this starting point.

Example 7: Get PersonType.NodeId, AnimalType.NodeId, PersonType.HasChild Reference, PersonType.HasAnimal Reference where the person is in Area2 (Cleveland nodes) and the person has a child.

Table B.22 describes the *NodeTypeDescription* parameters used in the example.

Table B.22 – Example 7 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".NodeId"	Value	N/A
		<12:HasChild>	Value	N/A
		<12:HasAnimal>NodeId	Value	N/A
		<12:HasAnimal>	Value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.11. Note that the *Browse* call would typically return a *NodeId*, thus the first filter is for the *BaseObjectType* with a *NodeId* of 95 where 95 is the *NodeId* associated with the Area2 node, all *Nodes* descend from *BaseObjectType*, and *NodeId* is a base *Property* so this filter will work for all *Queries* of this nature.




Figure B.11 – Example 7 Filter logic tree

Table B.23 describes the elements, operators and operands used in the example.

Table B.23 – Example 7 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	RelatedTo	ElementOperand = 2	ElementOperand = 1	AttributeOperand = Node: HierarchicalReference, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	Equals	AttributeOperand = NodeId: BaseObjectType, BrowsePath ".", Attribute: NodeId,	LiteralOperand = '95'		

The results from this Query would contain the QueryDataSets shown in Table B.24.

Table B.24 – Example 7 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:42 (HFamily1)	PersonType	".NodeId"	12:42 (HFamily1)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasAnimal>12:AnimalType.NodeId	NULL
		<12:HasAnimal>	HasFarmAnimal <i>ReferenceDescription</i>
12:48 (HFamily2)	PersonType	".NodeId"	12:48 (HFamily2)
		<12:HasChild>	HasChild <i>ReferenceDescription</i>
		<12:HasAnimal>12:AnimalType.NodeId	12:91 (Pig1)
		<12:HasAnimal>	HasFarmAnimal <i>ReferenceDescription</i>

NOTE The *RelativePath* and browse name (in parentheses) is not in the *QueryDataSet* and is only shown here for clarity and the *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table. The value field would in this case be the *NodeId* where it was requested, but for the example the browse name is provided in parentheses and in the case of *Reference* types on the browse name is provided. For the *References* listed in Table B.24, the value would be a *ReferenceDescription* which are described in 7.25.

B.2.11 Example 8

The eighth example provides an illustration of a request in which the *AddressSpace* is restricted by a *Server* defined *View*. This request is the same as in the second example which illustrates receiving a list of disjoint *Nodes* and also illustrates that an array of results can be received. It is important to note that all of the parameters and the *ContentFilter* are the same, only the *View* description would be specified as "View1".

Example 8: Get *PersonType.LastName*, *AnimalType.Name* where a person has a child or (a pet is of type cat and has a feeding schedule) limited by the *AddressSpace* in View1.

The *NodeTypeDescription* parameters used in the example are described in Table B.25

Table B.25 – Example 8 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".12:LastName"	value	N/A
AnimalType	TRUE	"12.Name"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.12.




Figure B.12 – Example 8 Filter logic tree

Table B.26 describes the elements, operators and operands used in the example. It is worth noting that a *CatType* is a subtype of *AnimalType*.

Table B.26 – Example 8 ContentFilter

Element[<i>i</i>]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	Or	ElementOperand=1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	RelatedTo	AttributeOperand = NodeId: 12:CatType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:FeedingScheduleType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this query would contain the *QueryDataSets* shown in Table B.27. If this is compared to the result set from example 2, the only difference is the omission of the *Cat Nodes*. These *Nodes* are not in the *View* and thus are not included in the result set.

Table B.27 – Example 8 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (Jfamily1)	Persontype	.12:LastName	Jones

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the *QueryDataSet* and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

B.2.12 Example 9

The ninth example provides a further illustration for a request in which the *AddressSpace* is restricted by a *Server* defined *View*. This request is similar to the second example except that some of the requested nodes are expressed in terms of a relative path. It is important to note that the *ContentFilter* is the same, only the *View* description would be specified as "View1".

Example 9: Get *PersonType.LastName*, *AnimalType.Name* where a person has a child or (a pet is of type cat and has a feeding schedule) limited by the *AddressSpace* in *View1*.

Table B.28 describes the *NodeTypeDescription* parameters used in the example.

Table B.28 – Example 9 NodeTypeDescription

Type Definition Node	Include Subtypes	QueryDataDescription		
		Relative Path	Attribute	Index Range
PersonType	FALSE	".NodeId"	value	N/A
		<12:HasChild>12:PersonType<12:HasAnimal>12:AnimalType.NodeId	value	N/A
		<12:HasChild>	value	N/A
		<12:HasChild>12:PersonType <12:HasAnimal>	value	N/A
PersonType	FALSE	".12:LastName"	value	N/A
		<12:HasAnimal>12:AnimalType. 12:Name	value	N/A
AnimalType	TRUE	".12:name"	value	N/A

The corresponding *ContentFilter* is illustrated in Figure B.13.


**Figure B.13 – Example 9 Filter logic tree**

Table B.29 describes the elements, operators and operands used in the example.

Table B.29 – Example 9 ContentFilter

Element[]	Operator	Operand[0]	Operand[1]	Operand[2]	Operand[3]
0	Or	ElementOperand=1	ElementOperand = 2		
1	RelatedTo	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:PersonType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasChild, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'
2	RelatedTo	AttributeOperand = NodeId: 12:CatType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:FeedingScheduleType, BrowsePath ".", Attribute: NodeId	AttributeOperand = NodeId: 12:HasSchedule, BrowsePath ".", Attribute: NodeId	LiteralOperand = '1'

The results from this *Query* would contain the *QueryDataSets* shown in Table B.30. If this is compared to the result set from example 2, the Pet *Nodes* are included in the list, even though they are outside of the *View*. This is possible since the name referenced via the relative path and the root *Node* is in the *View*.

Table B.30 – Example 9 QueryDataSets

NodeId	TypeDefinition NodeId	RelativePath	Value
12:30 (Jfamily1)	PersonType	. 12:LastName	Jones
		<12:HasAnimal>12:AnimalType. 12:Name	Rosemary
		<12:HasAnimal>12:AnimalType. 12:Name	Basil

NOTE The RelativePath column and browse name (in parentheses in the *NodeId* column) are not in the QueryDataSet and are only shown here for clarity. The *TypeDefinition NodeId* would be an integer not the symbolic name that is included in the table.

SOMMAIRE

AVANT-PROPOS	241
1 Domaine d'application	244
2 Références normatives	244
3 Termes, définitions, termes abrégés et conventions	245
3.1 Termes et définitions	245
3.2 Termes abrégés	246
3.3 Conventions pour la définition des services	246
4 Vue d'ensemble	248
4.1 Modèle de jeu de services	248
4.2 Procédures des services Demande/Réponse	251
5 Jeux de services	252
5.1 Généralités	252
5.2 En-tête de demande et de réponse de services	252
5.3 Résultats de service	252
5.4 Jeu de services Discovery	253
5.4.1 Vue d'ensemble	253
5.4.2 FindServers	255
5.4.3 FindServersOnNetwork	256
5.4.4 GetEndpoints	258
5.4.5 RegisterServer	261
5.4.6 RegisterServer2	264
5.5 Jeu de services SecureChannel	266
5.5.1 Vue d'ensemble	266
5.5.2 OpenSecureChannel	267
5.5.3 CloseSecureChannel	271
5.6 Jeu de services Session	272
5.6.1 Vue d'ensemble	272
5.6.2 CreateSession	272
5.6.3 ActivateSession	277
5.6.4 CloseSession	281
5.6.5 Cancel	282
5.7 Jeu de services NodeManagement	282
5.7.1 Vue d'ensemble	282
5.7.2 AddNodes	282
5.7.3 AddReferences	285
5.7.4 DeleteNodes	286
5.7.5 DeleteReferences	288
5.8 Jeu de services View	290
5.8.1 Vue d'ensemble	290
5.8.2 Browse	290
5.8.3 BrowseNext	293
5.8.4 TranslateBrowsePathsToNodeIds	295
5.8.5 RegisterNodes	297
5.8.6 UnregisterNodes	298
5.9 Jeu de services Query	299
5.9.1 Vue d'ensemble	299

5.9.2	Vues de requêtes	300
5.9.3	QueryFirst	300
5.9.4	QueryNext	304
5.10	Jeu de services Attribute.....	305
5.10.1	Vue d'ensemble	305
5.10.2	Read	306
5.10.3	HistoryRead.....	308
5.10.4	Write.....	311
5.10.5	HistoryUpdate.....	314
5.11	Jeu de services Method	316
5.11.1	Vue d'ensemble	316
5.11.2	Call.....	316
5.12	Jeu de services MonitoredItem.....	319
5.12.1	Modèle MonitoredItem	319
5.12.2	CreateMonitoredItems	325
5.12.3	ModifyMonitoredItems.....	328
5.12.4	SetMonitoringMode.....	330
5.12.5	SetTriggering.....	332
5.12.6	DeleteMonitoredItems	333
5.13	Jeu de services Subscription.....	334
5.13.1	Modèle Subscription	334
5.13.2	CreateSubscription	344
5.13.3	ModifySubscription	346
5.13.4	SetPublishingMode.....	348
5.13.5	Publish	349
5.13.6	Republish	351
5.13.7	TransferSubscriptions	352
5.13.8	DeleteSubscriptions.....	354
6	Comportements des services.....	355
6.1	Sécurité	355
6.1.1	Vue d'ensemble	355
6.1.2	Obtenir et installer un certificat d'instance d'application	355
6.1.3	Déterminer si un Certificat est fiable	356
6.1.4	Créer un SecureChannel	359
6.1.5	Créer une Session	361
6.1.6	Se substituer à un utilisateur	362
6.2	Services d'autorisation	362
6.2.1	Vue d'ensemble	362
6.2.2	Poignée de mains indirecte avec un fournisseur d'identité	363
6.2.3	Poignée de mains directe avec un fournisseur d'identité	363
6.3	Invocation de service sans session	364
6.3.1	Description	364
6.3.2	Paramètres.....	365
6.3.3	Résultats de service	366
6.4	Certificats logiciels.....	366
6.5	Audits	366
6.5.1	Vue d'ensemble	366
6.5.2	Journaux d'audit généraux.....	367
6.5.3	Evénements d'audit généraux.....	367

6.5.4	Audits pour le Jeu de services Discovery	367
6.5.5	Audits pour le Jeu de services SecureChannel	367
6.5.6	Audits pour le Jeu de services Session.....	368
6.5.7	Audits pour le Jeu de services NodeManagement	368
6.5.8	Audits pour le Jeu de services Attribut.....	368
6.5.9	Audits pour le Jeu de services Method.....	369
6.5.10	Audits pour les Jeux de services View, Query, MonitoredItem et Subscription	369
6.6	Redondance	370
6.6.1	Vue d'ensemble de la redondance	370
6.6.2	Redondance des serveurs	370
6.6.3	Redondance des clients.....	383
6.6.4	Redondance du réseau	383
6.6.5	Basculement forcé manuellement	385
6.7	Rétablissement des connexions	385
6.8	Abonnements durables	387
7	Définitions de types de paramètres communs	388
7.1	ApplicationDescription	388
7.2	ApplicationInstanceCertificate	389
7.3	BrowseResult.....	390
7.4	ContentFilter	391
7.4.1	Structure de ContentFilter.....	391
7.4.2	ContentFilterResult.....	391
7.4.3	FilterOperator	393
7.4.4	Paramètres FilterOperand	400
7.5	Counter.....	402
7.6	ContinuationPoint	403
7.7	DataValue	403
7.7.1	Généralités	403
7.7.2	PicoSeconds.....	404
7.7.3	SourceTimestamp.....	404
7.7.4	ServerTimestamp	404
7.7.5	StatusCode affecté à une valeur	405
7.8	DiagnosticInfo	405
7.9	Paramètres DiscoveryConfiguration	406
7.9.1	Vue d'ensemble	406
7.9.2	MdnsDiscoveryConfiguration	407
7.10	EndpointDescription.....	407
7.11	ExpandedNodeId	408
7.12	ExtensibleParameter.....	409
7.13	Index	409
7.14	IntegerId	409
7.15	MessageSecurityMode	410
7.16	MonitoringParameters.....	410
7.17	Paramètres MonitoringFilter	411
7.17.1	Vue d'ensemble	411
7.17.2	DataChangeFilter	412
7.17.3	EventFilter	413
7.17.4	AggregateFilter.....	416

7.18	MonitoringMode	417
7.19	Paramètres NodeAttributes	418
7.19.1	Vue d'ensemble	418
7.19.2	Paramètre ObjectAttributes	419
7.19.3	Paramètre VariableAttributes	419
7.19.4	Paramètre MethodAttributes	420
7.19.5	Paramètre ObjectTypeAttributes	420
7.19.6	Paramètre VariableTypeAttributes	421
7.19.7	Paramètre ReferenceTypeAttributes	421
7.19.8	Paramètre DataTypeAttributes	422
7.19.9	Paramètre ViewAttributes	422
7.19.10	Paramètre GenericAttributes	423
7.20	Paramètres NotificationData	423
7.20.1	Vue d'ensemble	423
7.20.2	Paramètre DataChangeNotification	424
7.20.3	Paramètre EventNotificationList	425
7.20.4	Paramètre StatusChangeNotification	425
7.21	NotificationMessage	425
7.22	NumericRange	426
7.23	QueryDataSet	427
7.24	ReadValued	427
7.25	ReferenceDescription	428
7.26	RelativePath	429
7.27	RegisteredServer	430
7.28	RequestHeader	430
7.29	ResponseHeader	432
7.30	ServiceFault	432
7.31	SessionAuthenticationToken	433
7.32	SignatureData	434
7.33	SignedSoftwareCertificate	434
7.34	StatusCode	435
7.34.1	Généralités	435
7.34.2	StatusCodes communs	437
7.35	TimestampsToReturn	441
7.36	Paramètres UserIdentityToken	441
7.36.1	Vue d'ensemble	441
7.36.2	Chiffrement de jeton et preuve de possession	442
7.36.3	AnonymousIdentityToken	447
7.36.4	UserNameIdentityToken	447
7.36.5	X509IdentityTokens	449
7.36.6	IssuedIdentityToken	449
7.37	UserTokenPolicy	450
7.38	VersionTime	451
7.39	ViewDescription	451
Annexe A (informative)	Définitions BNF	453
A.1	Vue d'ensemble de la notation BNF	453
A.2	Notation BNF de RelativePath	453
A.3	Notation BNF de NumericRange	454
Annexe B (informative)	Exemples d'éléments ContentFilter et Query	455

B.1	Exemples de filtres ContentFilter simples.....	455
B.1.1	Vue d'ensemble	455
B.1.2	Exemple 1	455
B.1.3	Exemple 2	456
B.2	Exemples de filtres Query complexes.....	457
B.2.1	Vue d'ensemble	457
B.2.2	Modèle de type utilisé	457
B.2.3	Remarques relatives aux exemples	460
B.2.4	Exemple 1	461
B.2.5	Exemple 2	462
B.2.6	Exemple 3	463
B.2.7	Exemple 4	466
B.2.8	Exemple 5	467
B.2.9	Exemple 6	468
B.2.10	Exemple 7	471
B.2.11	Exemple 8	472
B.2.12	Exemple 9	474
Figure 1	– Jeu de services Discovery	248
Figure 2	– Jeu de services SecureChannel	248
Figure 3	– Jeu de services Session.....	249
Figure 4	– Jeu de services NodeManagement.....	249
Figure 5	– Jeu de services View	249
Figure 6	– Jeu de services Attribute	250
Figure 7	– Jeu de services Method	250
Figure 8	– Jeux de services MonitoredItem et Subscription	251
Figure 9	– Processus de découverte	254
Figure 10	– Utilisation d'un Serveur passerelle	260
Figure 11	– Processus d'enregistrement – Serveurs lancés manuellement.....	262
Figure 12	– Processus d'enregistrement – Serveurs lancés automatiquement.....	263
Figure 13	– Services SecureChannel et Session	267
Figure 14	– Multiplexage d'utilisateurs sur une Session	274
Figure 15	– Modèle MonitoredItem.....	320
Figure 16	– Retard type pour détecter des modifications.....	322
Figure 17	– Gestion du dépassement de capacité de la file d'attente	323
Figure 18	– Modèle de déclenchement.....	324
Figure 19	– Obtenir et installer un certificat d'instance d'application	356
Figure 20	– Déterminer si un Certificat d'instance d'application est fiable.....	359
Figure 21	– Etablir un SecureChannel.....	360
Figure 22	– Etablir une Session	361
Figure 23	– Se substituer à un utilisateur	362
Figure 24	– Poignée de mains indirecte avec un fournisseur d'identité.....	363
Figure 25	– Poignée de mains directe avec un fournisseur d'identité	364
Figure 26	– Exemple de configuration de redondance transparente	372
Figure 27	– Configuration de redondance non transparente	373

Figure 28 – Etapes de démarrage du Client	378
Figure 29 – Basculement Cold	379
Figure 30 – Basculement Warm	380
Figure 31 – Basculement Hot	381
Figure 32 – Basculement HotAndMirrored	382
Figure 33 – Serveur proxy pour la redondance	383
Figure 34 – Redondance d'un réseau transparent	384
Figure 35 – Redondance d'un réseau non transparent	385
Figure 36 – Séquence de reconnexion	386
Figure 37 – Couches logiques d'un <i>Serveur</i>	433
Figure 38 – Obtention d'un SessionAuthenticationToken	434
Figure 39 – Présentation EncryptedSecret	444
Figure B.1 – Exemple d'arbre logique de filtre	455
Figure B.2 – Exemple d'arbre logique de filtre	456
Figure B.3 – Exemples de nœuds de type	459
Figure B.4 – Exemple de nœuds d'instance	460
Figure B.5 – Exemple 1 Filtre	461
Figure B.6 – Exemple 2 Arbre logique d'un filtre	463
Figure B.7 – Exemple 3 Arbre logique d'un filtre	464
Figure B.8 – Exemple 4 Arbre logique d'un filtre	466
Figure B.9 – Exemple 5 Arbre logique d'un filtre	468
Figure B.10 – Exemple 6 Arbre logique d'un filtre	469
Figure B.11 – Exemple 7 Arbre logique d'un filtre	471
Figure B.12 – Exemple 8 Arbre logique d'un filtre	473
Figure B.13 – Exemple 9 Arbre logique d'un filtre	474
Tableau 1 – Tableau de définition des services	247
Tableau 2 – Types de paramètres définis dans l'IEC 62541-3	247
Tableau 3 – Paramètres du service FindServers	256
Tableau 4 – Paramètres du service FindServersOnNetwork	258
Tableau 5 – Paramètres du service GetEndpoints	261
Tableau 6 – Paramètres du service RegisterServer	264
Tableau 7 – Codes de résultats du service RegisterServer	264
Tableau 8 – RegisterServer2	265
Tableau 9 – Codes de résultats du service RegisterServer2	265
Tableau 10 – Codes de résultats de niveau opération RegisterServer2	265
Tableau 11 – Paramètres du service OpenSecureChannel	269
Tableau 12 – Codes de résultats du service OpenSecureChannel	271
Tableau 13 – Paramètres du service CloseSecureChannel	272
Tableau 14 – Codes de résultats du service CloseSecureChannel	272
Tableau 15 – Paramètres du service CreateSession	275
Tableau 16 – Codes de résultats du service CreateSession	277
Tableau 17 – Paramètres du service ActivateSession	280

Tableau 18 – Codes de résultats du service ActivateSession	281
Tableau 19 – Paramètres du service CloseSession	281
Tableau 20 – Codes de résultats du service CloseSession	282
Tableau 21 – Paramètres du service Cancel	282
Tableau 22 – Paramètres du service AddNodes	283
Tableau 23 – Codes de résultats du service AddNodes	284
Tableau 24 – Codes de résultats de niveau opération AddNodes	284
Tableau 25 – Paramètres du service AddReferences	285
Tableau 26 – Codes de résultats du service AddReferences	286
Tableau 27 – Codes de résultats de niveau opération AddReferences	286
Tableau 28 – Paramètres du service DeleteNodes	287
Tableau 29 – Codes de résultats du service DeleteNodes	287
Tableau 30 – Codes de résultat de niveau opération DeleteNodes	288
Tableau 31 – Paramètres du service DeleteReferences	289
Tableau 32 – Codes de résultats du service DeleteReferences	289
Tableau 33 – Codes de résultats de niveau opération DeleteReferences	290
Tableau 34 – Paramètres du service Browse	291
Tableau 35 – Codes de résultats du service Browse	292
Tableau 36 – Codes de résultats de niveau opération Browse	293
Tableau 37 – Paramètres du service BrowseNext	294
Tableau 38 – Codes de résultats du service BrowseNext	294
Tableau 39 – Codes de résultats de niveau opération BrowseNext	295
Tableau 40 – Paramètres du service TranslateBrowsePathsToNodeIds	296
Tableau 41 – Codes de résultats du service TranslateBrowsePathsToNodeIds	297
Tableau 42 – Codes de résultats de niveau opération TranslateBrowsePathsToNodeIds	297
Tableau 43 – Paramètres du service RegisterNodes	298
Tableau 44 – Codes de résultats du service RegisterNodes	298
Tableau 45 – Paramètres du service UnregisterNodes	299
Tableau 46 – Codes de résultats du service UnregisterNodes	299
Tableau 47 – Paramètres de demande QueryFirst	301
Tableau 48 – Paramètres de réponse QueryFirst	302
Tableau 49 – Codes de résultats du service QueryFirst	303
Tableau 50 – Codes de résultats de niveau opération QueryFirst	304
Tableau 51 – Paramètres du service QueryNext	305
Tableau 52 – Codes de résultats du service QueryNext	305
Tableau 53 – Paramètres du service Read	306
Tableau 54 – Codes de résultats du service Read	307
Tableau 55 – Codes de résultats de niveau opération Read	307
Tableau 56 – Paramètres du service HistoryRead	308
Tableau 57 – Codes de résultats du service HistoryRead	310
Tableau 58 – Codes de résultats de niveau opération HistoryRead	311
Tableau 59 – Paramètres du service Write	312

Tableau 60 – Codes de résultats du service Write	313
Tableau 61 – Codes de résultats de niveau opération Write	314
Tableau 62 – Paramètres du service HistoryUpdate	315
Tableau 63 – Codes de résultats du service HistoryUpdate	315
Tableau 64 – Codes de résultat de niveau d'opération HistoryUpdate	316
Tableau 65 – Paramètres du service Call	317
Tableau 66 – Codes de résultats du service Call	318
Tableau 67 – Codes de résultats de niveau opération Call	319
Tableau 68 – Codes de résultats des arguments d'entrée "Call Input"	319
Tableau 69 – Paramètres du service CreateMonitoredItems	326
Tableau 70 – Codes de résultats du service CreateMonitoredItems	327
Tableau 71 – Codes de résultats de niveau opération CreateMonitoredItems	328
Tableau 72 – Paramètres du service ModifyMonitoredItems	329
Tableau 73 – Codes de résultats du service ModifyMonitoredItems	330
Tableau 74 – Codes de résultats de niveau opération ModifyMonitoredItems	330
Tableau 75 – Paramètres du service SetMonitoringMode	331
Tableau 76 – Codes de résultats du service SetMonitoringMode	331
Tableau 77 – Codes de résultats de niveau opération SetMonitoringMode	331
Tableau 78 – Paramètres du service SetTriggering	332
Tableau 79 – Codes de résultats du service SetTriggering	333
Tableau 80 – Codes de résultats de niveau opération SetTriggering	333
Tableau 81 – Paramètres du service DeleteMonitoredItems	334
Tableau 82 – Codes de résultats du service DeleteMonitoredItems	334
Tableau 83 – Codes de résultats de niveau opération DeleteMonitoredItems	334
Tableau 84 – Etats d'un abonnement	338
Tableau 85 – Table d'états de Subscription	339
Tableau 86 – Variables d'état et paramètres	342
Tableau 87 – Fonctions	343
Tableau 88 – Paramètres du service CreateSubscription	344
Tableau 89 – Codes de résultats du service CreateSubscription	346
Tableau 90 – Paramètres du service ModifySubscription	347
Tableau 91 – Codes de résultats du service ModifySubscription	348
Tableau 92 – Paramètres du service SetPublishingMode	348
Tableau 93 – Codes de résultats du service SetPublishingMode	349
Tableau 94 – Codes de résultats de niveau opération SetPublishingMode	349
Tableau 95 – Paramètres du service Publish	350
Tableau 96 – Codes de résultats du service Publish	351
Tableau 97 – Codes de résultats de niveau opération Publish	351
Tableau 98 – Paramètres du service Republish	352
Tableau 99 – Codes de résultats du service Republish	352
Tableau 100 – Paramètres du service TransferSubscriptions	353
Tableau 101 – Codes de résultats du service TransferSubscriptions	353
Tableau 102 – Codes de résultats de niveau opération TransferSubscriptions	354

Tableau 103 – Paramètres du service DeleteSubscriptions	354
Tableau 104 – Codes de résultats du service DeleteSubscriptions	355
Tableau 105 – Codes de résultats de niveau opération DeleteSubscriptions"	355
Tableau 106 – Etapes de validation d'un certificat	357
Tableau 107 – Paramètres du service SessionlessInvoke	365
Tableau 108 – Codes de résultats du service SessionlessInvoke	366
Tableau 109 – Plages de ServiceLevel	375
Tableau 110 – Modes de Basculement des serveurs	376
Tableau 111 – Actions de Basculement de redondance	377
Tableau 112 – ApplicationDescription	389
Tableau 113 – ApplicationInstanceCertificate	390
Tableau 114 – BrowseResult	390
Tableau 115 – Structure de ContentFilter	391
Tableau 116 – Structure de ContentFilterResult	392
Tableau 117 – Codes de résultats "ContentFilterResult"	392
Tableau 118 – Codes de résultats d'opérande "ContentFilterResult"	392
Tableau 119 – Définition de FilterOperator de base	393
Tableau 120 – Définition de FilterOperator complexe	396
Tableau 121 – Caractères génériques	397
Tableau 122 – Règles de conversion	398
Tableau 123 – Règles de priorité des données	399
Tableau 124 – Table de vérité des opérations AND logiques	400
Tableau 125 – Table de vérité des opérations OR logiques	400
Tableau 126 – Typelds du paramètre FilterOperand	400
Tableau 127 – ElementOperand	401
Tableau 128 – LiteralOperand	401
Tableau 129 – AttributeOperand	401
Tableau 130 – SimpleAttributeOperand	402
Tableau 131 – DataValue	403
Tableau 132 – DiagnosticInfo	406
Tableau 133 – Typelds du paramètre DiscoveryConfiguration	407
Tableau 134 – MdnsDiscoveryConfiguration	407
Tableau 135 – EndpointDescription	408
Tableau 136 – ExpandedNodeId	409
Tableau 137 – Type de base ExtensibleParameter	409
Tableau 138 – Valeurs de MessageSecurityMode	410
Tableau 139 – MonitoringParameters	410
Tableau 140 – Typelds du paramètre MonitoringFilter	412
Tableau 141 – DataChangeFilter	412
Tableau 142 – Structure d'EventFilter	415
Tableau 143 – Structure d'EventFilterResult	415
Tableau 144 – Codes de résultats "EventFilterResult"	415
Tableau 145 – Structure d'AggregateFilter	417

Tableau 146 – Structure d'AggregateFilterResult	417
Tableau 147 – Valeurs de MonitoringMode	418
Tableau 148 – Typelds du paramètre NodeAttributes	418
Tableau 149 – Masque de bits pour les attributs spécifiés	419
Tableau 150 – ObjectAttributes	419
Tableau 151 – VariableAttributes	420
Tableau 152 – MethodAttributes	420
Tableau 153 – ObjectTypeAttributes	421
Tableau 154 – VariableTypeAttributes	421
Tableau 155 – ReferenceTypeAttributes	422
Tableau 156 – DataTypeAttributes	422
Tableau 157 – ViewAttributes	423
Tableau 158 – GenericAttributes	423
Tableau 159 – Typelds du paramètre NotificationData	424
Tableau 160 – DataChangeNotification	424
Tableau 161 – EventNotificationList	425
Tableau 162 – StatusChangeNotification	425
Tableau 163 – NotificationMessage	426
Tableau 164 – NumericRange	427
Tableau 165 – QueryDataSet	427
Tableau 166 – ReadValueId	428
Tableau 167 – ReferenceDescription	429
Tableau 168 – RelativePath	429
Tableau 169 – RegisteredServer	430
Tableau 170 – RequestHeader	431
Tableau 171 – ResponseHeader	432
Tableau 172 – ServiceFault	433
Tableau 173 – SignatureData	434
Tableau 174 – SignedSoftwareCertificate	435
Tableau 175 – Affectations de bits de StatusCode	436
Tableau 176 – InfoBits de DataValue	437
Tableau 177 – Codes de résultats de services communs	438
Tableau 178 – Codes de résultats de niveau opération communs	440
Tableau 179 – Valeurs de TimestampsToReturn	441
Tableau 180 – Typelds du paramètre UserIdentityToken	442
Tableau 181 – Ancien format secret de jeton chiffré UserIdentityToken	443
Tableau 182 – Présentation EncryptedSecret	446
Tableau 183 – DataTypes EncryptedSecret	446
Tableau 184 – Structure de RsaEncryptedSecret	447
Tableau 185 – AnonymousIdentityToken	447
Tableau 186 – UserNameIdentityToken	448
Tableau 187 – Sélection d'EncryptionAlgorithm	449
Tableau 188 – Jeton d'identité X.509 v3	449

Tableau 189 – IssuedIdentityToken.....	450
Tableau 190 – UserTokenPolicy	451
Tableau 191 – ViewDescription.....	452
Tableau A.1 – RelativePath	453
Tableau A.2 – Exemples de <i>RelativePaths</i>	454
Tableau B.1 – Exemple de ContentFilter	456
Tableau B.2 – Exemple de ContentFilter	456
Tableau B.3 – Exemple 1 NodeTypeDescription.....	461
Tableau B.4 – Exemple 1 ContentFilter	461
Tableau B.5 – Exemple 1 QueryDataSets	462
Tableau B.6 – Exemple 2 NodeTypeDescription.....	462
Tableau B.7 – Exemple 2 ContentFilter	463
Tableau B.8 – Exemple 2 QueryDataSets	463
Tableau B.9 – Exemple 3 NodeTypeDescription.....	464
Tableau B.10 – Exemple 3 ContentFilter	465
Tableau B.11 – Exemple 3 QueryDataSets	466
Tableau B.12 – Exemple 4 NodeTypeDescription.....	466
Tableau B.13 – Exemple 4 ContentFilter	467
Tableau B.14 – Exemple 4 QueryDataSets	467
Tableau B.15 – Exemple 5 NodeTypeDescription.....	467
Tableau B.16 – Exemple 5 ContentFilter	468
Tableau B.17 – Exemple 5 QueryDataSets	468
Tableau B.18 – Exemple 6 NodeTypeDescription.....	469
Tableau B.19 – Exemple 6 ContentFilter	469
Tableau B.20 – Exemple 6 QueryDataSets	470
Tableau B.21 – Exemple 6 QueryDataSets sans informations supplémentaires.....	470
Tableau B.22 – Exemple 7 NodeTypeDescription.....	471
Tableau B.23 – Exemple 7 ContentFilter	472
Tableau B.24 – Exemple 7 QueryDataSets	472
Tableau B.25 – Exemple 8 NodeTypeDescription.....	473
Tableau B.26 – Exemple 8 ContentFilter	473
Tableau B.27 – Exemple 8 QueryDataSets	473
Tableau B.28 – Exemple 9 NodeTypeDescription.....	474
Tableau B.29 – Exemple 9 ContentFilter	475
Tableau B.30 – Exemple 9 QueryDataSets	475

COMMISSION ÉLECTROTECHNIQUE INTERNATIONALE

ARCHITECTURE UNIFIÉE OPC –

Partie 4: Services

AVANT-PROPOS

- 1) La Commission Electrotechnique Internationale (IEC) est une organisation mondiale de normalisation composée de l'ensemble des comités électrotechniques nationaux (Comités nationaux de l'IEC). L'IEC a pour objet de favoriser la coopération internationale pour toutes les questions de normalisation dans les domaines de l'électricité et de l'électronique. A cet effet, l'IEC – entre autres activités – publie des Normes internationales, des Spécifications techniques, des Rapports techniques, des Spécifications accessibles au public (PAS) et des Guides (ci-après dénommés "Publication(s) de l'IEC"). Leur élaboration est confiée à des comités d'études, aux travaux desquels tout Comité national intéressé par le sujet traité peut participer. Les organisations internationales, gouvernementales et non gouvernementales, en liaison avec l'IEC, participent également aux travaux. L'IEC collabore étroitement avec l'Organisation Internationale de Normalisation (ISO), selon des conditions fixées par accord entre les deux organisations.
- 2) Les décisions ou accords officiels de l'IEC concernant les questions techniques représentent, dans la mesure du possible, un accord international sur les sujets étudiés, étant donné que les Comités nationaux de l'IEC intéressés sont représentés dans chaque comité d'études.
- 3) Les Publications de l'IEC se présentent sous la forme de recommandations internationales et sont agréées comme telles par les Comités nationaux de l'IEC. Tous les efforts raisonnables sont entrepris afin que l'IEC s'assure de l'exactitude du contenu technique de ses publications; l'IEC ne peut pas être tenue responsable de l'éventuelle mauvaise utilisation ou interprétation qui en est faite par un quelconque utilisateur final.
- 4) Dans le but d'encourager l'uniformité internationale, les Comités nationaux de l'IEC s'engagent, dans toute la mesure possible, à appliquer de façon transparente les Publications de l'IEC dans leurs publications nationales et régionales. Toutes divergences entre toutes Publications de l'IEC et toutes publications nationales ou régionales correspondantes doivent être indiquées en termes clairs dans ces dernières.
- 5) L'IEC elle-même ne fournit aucune attestation de conformité. Des organismes de certification indépendants fournissent des services d'évaluation de conformité et, dans certains secteurs, accèdent aux marques de conformité de l'IEC. L'IEC n'est responsable d'aucun des services effectués par les organismes de certification indépendants.
- 6) Tous les utilisateurs doivent s'assurer qu'ils sont en possession de la dernière édition de cette publication.
- 7) Aucune responsabilité ne doit être imputée à l'IEC, à ses administrateurs, employés, auxiliaires ou mandataires, y compris ses experts particuliers et les membres de ses comités d'études et des Comités nationaux de l'IEC, pour tout préjudice causé en cas de dommages corporels et matériels, ou de tout autre dommage de quelque nature que ce soit, directe ou indirecte, ou pour supporter les coûts (y compris les frais de justice) et les dépenses découlant de la publication ou de l'utilisation de cette Publication de l'IEC ou de toute autre Publication de l'IEC, ou au crédit qui lui est accordé.
- 8) L'attention est attirée sur les références normatives citées dans cette publication. L'utilisation de publications référencées est obligatoire pour une application correcte de la présente publication.
- 9) L'attention est attirée sur le fait que certains des éléments de la présente Publication de l'IEC peuvent faire l'objet de droits de brevet. L'IEC ne saurait être tenue pour responsable de ne pas avoir identifié de tels droits de brevets et de ne pas avoir signalé leur existence.

La Norme internationale IEC 62541-4 a été établie par le sous-comité 65E: Les dispositifs et leur intégration dans les systèmes de l'entreprise, du comité d'études 65 de l'IEC: Mesure, commande et automation dans les processus industriels.

Cette troisième édition annule et remplace la deuxième édition parue en 2015. Cette édition constitue une révision technique.

Cette édition inclut les modifications techniques majeures suivantes par rapport à l'édition précédente:

- a) ajout de la capacité à renvoyer toutes les données des éléments surveillés dans un abonnement à l'aide de la méthode ResendData;
- b) ajout de la prise en charge pour les abonnements durables (durée de vie en heures ou en jours);

- c) ajout des services Register2 et FindServersOnNetwork pour prendre en charge le mécanisme de découverte sur l'ensemble du réseau à l'aide de filtres de capacités;
- d) suppression de la définition des certificats de logiciel (seront définis dans une édition ultérieure);
- e) enrichissement et révision partielle de la définition de la redondance; ajout de définitions des sous-plages pour ServiceLevel et ajout de termes supplémentaires pour la redondance;
- f) ajout d'un paragraphe expliquant comment utiliser les Services d'autorisation pour demander des jetons d'accès utilisateur;
- g) ajout des jetons JSON Web (JWT) comme nouveau jeton d'utilisateur;
- h) ajout du concept d'invocation de service sans session;
- i) ajout d'une structure générique permettant de transmettre n'importe quel nombre d'attributs dans le service AddNotes;
- j) ajout d'une exigence relative à la protection contre les attaques de jeton d'identité utilisateur;
- k) ajout du nouveau format EncryptedSecret pour les jetons d'identité utilisateur.

Le texte de cette norme est issu des documents suivants:

FDIS	Rapport de vote
65E/716/FDIS	65E/732/RVD

Le rapport de vote indiqué dans le tableau ci-dessus donne toute information sur le vote ayant abouti à l'approbation de cette Norme internationale.

Ce document a été rédigé selon les Directives ISO/IEC, Partie 2.

Dans l'ensemble du présent document et dans les autres parties de la série IEC 62541, certaines conventions de document sont utilisées:

Le format *italique* est utilisé pour mettre en évidence un terme défini ou une définition qui apparaît à l'Article 3 dans l'une des parties de la série.

Le format *italique* est également utilisé, à quelques exceptions près, pour mettre en évidence le nom d'un paramètre d'entrée ou de sortie de service, ou le nom d'une structure ou d'un élément de structure habituellement défini dans les tableaux.

Par ailleurs, les *termes* et les *noms en italique* sont souvent écrits en camel-case (pratique qui consiste à joindre, sans espace, les éléments des mots ou expressions composés, la première lettre de chaque élément étant en majuscule). Par exemple, le terme défini est *AddressSpace* et non Espace d'adressage. Cela permet de mieux comprendre qu'il existe une définition unique pour *AddressSpace*, et non deux définitions distinctes pour Espace et pour Adressage.

Une liste de toutes les parties de la série IEC 62541, publiées sous le titre général *Architecture unifiée OPC*, peut être consultée sur le site web de l'IEC.

Le comité a décidé que le contenu de ce document ne sera pas modifié avant la date de stabilité indiquée sur le site web de l'IEC sous "<http://webstore.iec.ch>" dans les données relatives au document recherché. A cette date, le document sera

- reconduit,
- supprimé,
- remplacé par une édition révisée, ou
- amendé.

IMPORTANT – Le logo "*colour inside*" qui se trouve sur la page de couverture de cette publication indique qu'elle contient des couleurs qui sont considérées comme utiles à une bonne compréhension de son contenu. Les utilisateurs devraient, par conséquent, imprimer cette publication en utilisant une imprimante couleur.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RV

ARCHITECTURE UNIFIÉE OPC –

Partie 4: Services

1 Domaine d'application

La présente partie de l'IEC 62541 définit le modèle de communication *Services* de l'architecture unifiée OPC (OPC UA). Les *Services* définis sont le recueil d'appels de procédures abstraites distantes (RPC, *Remote Procedure Call*) qui sont mises en œuvre par les *Serveurs* OPC UA et qui sont appelées par les *Clients* OPC UA. Toutes les interactions entre *Clients* et *Serveurs* OPC UA ont lieu via ces *Services*. Les *Services* définis sont dits abstraits, car aucun mécanisme RPC particulier n'est spécifié dans le présent document pour leur mise en œuvre. L'IEC 62541-6 spécifie un ou plusieurs mécanismes concrets de mapping pour la mise en œuvre. Par exemple, dans l'IEC 62541-6, l'un des mécanismes de mapping repose sur l'utilisation des Services Web XML. Dans ce cas, les *Services* décrits dans le présent document apparaissent comme les méthodes de services Web dans le contrat WSDL.

Il n'est pas nécessaire que l'ensemble des *Serveurs* OPC UA mettent en œuvre toutes les correspondances avec les messages et les protocoles de transport. L'IEC 62541-7 définit les *Profils* qui dictent les *Services* qu'il est nécessaire de mettre en œuvre afin d'être conforme à un *Profil* particulier.

2 Références normatives

Les documents suivants sont cités dans le texte de sorte qu'ils constituent, pour tout ou partie de leur contenu, des exigences du présent document. Pour les références datées, seule l'édition citée s'applique. Pour les références non datées, la dernière édition du document de référence s'applique (y compris les éventuels amendements).

IEC TR 62541-1, *OPC Unified Architecture – Part 1: Overview and Concepts* (disponible en anglais seulement)

IEC TR 62541-2, *OPC Unified Architecture – Part 2: Security Model* (disponible en anglais seulement)

IEC 62541-3, *Architecture unifiée OPC – Partie 3: Modèle d'espace d'adressage*

IEC 62541-5, *Architecture unifiée OPC – Partie 5: Modèle d'information*

IEC 62541-6, *Architecture unifiée OPC – Partie 6: Mappings*

IEC 62541-7, *Architecture unifiée OPC – Partie 7: Profils*

IEC 62541-8, *Architecture unifiée OPC – Partie 8: Accès aux données*

IEC 62541-11, *Architecture unifiée OPC – Partie 11: Accès à l'historique*

IEC 62541-12¹, *Architecture unifiée OPC – Partie 12: Services globaux et de découverte*

IEC 62541-13, *Architecture unifiée OPC – Partie 13: Agrégats*

¹ En cours d'élaboration. Stade au moment de la publication: IEC CDV 62541-12:2018.

3 Termes, définitions, termes abrégés et conventions

3.1 Termes et définitions

Pour les besoins du présent document, les termes et définitions de l'IEC TR 62541-1, l'IEC TR 62541-2 et l'IEC 62541-3 ainsi que les suivants, s'appliquent.

L'ISO et l'IEC tiennent à jour des bases de données terminologiques destinées à être utilisées en normalisation, consultables aux adresses suivantes:

- IEC Electropedia: disponible à l'adresse <http://www.electropedia.org/>
- ISO Online browsing platform: disponible à l'adresse <http://www.iso.org/obp>

3.1.1

Serveur Actif

Serveur qui fournit actuellement les informations

Note 1 à l'article: Dans les systèmes redondants OPC UA, un *Serveur Actif* est le *Serveur* qu'un *Client* utilise comme source de données.

3.1.2

Deadband

plage admise pour les variations de valeurs sans déclenchement d'une *Notification* de modification de données

Note 1 à l'article: *Deadband* peut s'appliquer comme un filtre lors de l'abonnement à des *Variables*, et sert à empêcher que des signaux de bruit ne mettent inutilement un *Client* à jour. Le présent document définit *AbsoluteDeadband* comme un filtre commun. L'IEC 62541-8 définit un filtre *Deadband* supplémentaire.

3.1.3

DiscoveryEndpoint

Endpoint qui permet aux *Clients* d'accéder aux *Services Découverte* sans sécurité

Note 1 à l'article: Un *DiscoveryEndpoint* permet d'accéder aux *Services Découverte* sans *Session* et sans sécurité de message.

3.1.4

Endpoint

adresse physique disponible sur un réseau qui permet à des *Clients* d'accéder à un ou plusieurs *Services* fournis par un *Serveur*

Note 1 à l'article: Chaque *Serveur* peut avoir plusieurs *Endpoints*. Chaque *Endpoint* inclut un *HostName*.

3.1.5

serveur défaillant

Serveur qui n'est pas opérationnel

Note 1 à l'article: Dans les systèmes redondants OPC UA, un *Serveur défaillant* est un *Serveur* qui n'est pas disponible ou qui ne peut pas fournir des données.

3.1.6

Basculement

action de modifier la source ou la cible des informations

Note 1 à l'article: Dans les systèmes redondants OPC UA, un *Basculement* est l'action entreprise par un *Client* pour passer d'un *Serveur* défaillant ou dégradé à un autre *Serveur* dans un jeu redondant (*Basculement Serveur*). Dans certains cas, un *Client* peut ignorer qu'un *Basculement* a lieu (redondance transparente). Un *Basculement Client* est l'action d'un autre *Client* en remplacement d'une connexion *Client* défaillante ou dégradée à un *Serveur*.

3.1.7

serveur passerelle

Serveur qui agit comme un intermédiaire pour un ou plusieurs *Serveurs*

Note 1 à l'article: Les *Serveurs passerelles* peuvent être déployés pour limiter l'accès externe, assurer la conversion de protocoles ou fournir des caractéristiques que les *Serveurs* sous-jacents ne prennent pas en charge.

3.1.8

HostName

identificateur unique d'une machine sur un réseau

Note 1 à l'article: Cet identificateur est unique au sein d'un réseau local, mais il peut toutefois être également unique au niveau global. L'identificateur peut être une adresse IP.

3.1.9

redondance

présence de composants dupliqués qui assurent la continuité des opérations après la défaillance d'un composant d'OPC UA

Note 1 à l'article: Ceci peut s'appliquer aux *Serveurs*, *Clients* ou Réseaux.

3.1.10

jeu de serveurs redondants

deux ou plusieurs *Serveurs* qui sont redondants les uns avec les autres

Note 1 à l'article: Un *Jeu de serveurs redondants* est un groupe de *Serveurs* qui sont configurés pour procurer la *Redondance*. Ces *Serveurs* ont des exigences d'espace d'adressage et assurent des *Basculements*.

3.1.11

SecurityToken

identificateur d'un jeu de clés cryptographiques

Note 1 à l'article: Tous les *SecurityTokens* s'inscrivent dans un contexte de sécurité. Pour l'OPC UA, le contexte de sécurité est le *SecureChannel*.

3.1.12

SoftwareCertificate

certificat numérique pour un produit logiciel qui peut être installé sur plusieurs hôtes afin de décrire les fonctions du produit logiciel

Note 1 à l'article: Des installations différentes du même produit logiciel peuvent avoir le même certificat de logiciel. Les certificats de logiciel ne sont pas pertinents pour la sécurité. Ils sont utilisés pour identifier un produit logiciel et ses caractéristiques prises en charge. Les *SoftwareCertificates* sont décrits en 6.4.

3.2 Termes abrégés

API	Application Programming Interface (Interface de programmation d'application)
BNF	Backus-Naur Form (Forme de Backus-Naur)
CA	Certificate Authority (Autorité de certification)
CRL	Certificate Revocation List (Liste de révocation de certificat)
CTL	Certificate Trust List (Liste de certificats de confiance)
DA	Data Access (Accès aux données)
NAT	Network Address Translation (Traduction d'adresse réseau)
UA	Unified Architecture (Architecture unifiée)
URI	Uniform Resource Identifier (Identifiant uniforme de ressource)
URL	Uniform Resource Locator (Localisateur uniforme de ressource)

3.3 Conventions pour la définition des services

Les *Services* OPC UA ont des paramètres qui sont acheminés entre le *Client* et le *Serveur*. Les spécifications des *services* OPC UA utilisent des tableaux pour décrire les paramètres des *services* (voir Tableau 1). Dans ce tableau, les paramètres sont organisés en deux groupes (paramètres de demande et paramètres de réponse).

Tableau 1 – Tableau de définition des services

Nom	Type	Description
Demande		Définit les paramètres de demande du <i>service</i>
Nom de paramètre simple		Description du paramètre
Nom du paramètre construit		Description du paramètre construit
Nom du paramètre constitutif		Description du paramètre constitutif
Réponse		Définit les paramètres de réponse du <i>service</i>

Les colonnes Nom, Type et Description spécifient le nom, le type de données et la description de chaque paramètre. Tous les paramètres sont obligatoires, même si certains peuvent être inutilisés dans certains cas. La colonne Description spécifie la valeur à fournir lorsqu'un paramètre est inutilisé.

Deux types de paramètres sont définis dans ces tableaux: simple et construit. Les paramètres simples ont un type de données simple (*Boolean* ou *String*).

Les paramètres construits sont des paramètres composés de deux paramètres constitutifs ou plus, qui peuvent être simples ou construits. Les noms des paramètres constitutifs sont indentés sous le nom du paramètre construit.

Les types de données utilisés dans ces tableaux peuvent être des types de base, des types communs à plusieurs *services* ou des types spécifiques à chaque *service*. Les types de données de base sont définis dans l'IEC 62541-3. Les types de base utilisés dans les *services* sont répertoriés dans le Tableau 2. Les types de données qui sont communs à plusieurs *services* sont définis à l'Article 7. Les types de données qui sont spécifiques à un *service* sont définis dans le tableau de paramètres du *service*.

Tableau 2 – Types de paramètres définis dans l'IEC 62541-3

Type de paramètre
BaseDataType
Boolean
ByteString
Double
Duration
Guid
Int32
LocaleId
NodeId
QualifiedName
String
UInt16
UInt32
UInteger
UtcTime
XmlElement

Les paramètres des primitives des services Demande et Indication sont représentés dans le Tableau 1 comme des paramètres Demande. De même, les paramètres des primitives des services Réponse et Confirmation sont représentés dans le Tableau 1 comme des paramètres Réponse. Tous les paramètres de demande et de réponse sont inchangés et acheminés entre l'expéditeur et le récepteur. Par conséquent, il n'est pas nécessaire d'avoir des colonnes distinctes pour les valeurs des paramètres de demande, d'indication, de réponse et de confirmation. Elles n'ont donc pas été prises en compte afin d'améliorer la lisibilité.

4 Vue d'ensemble

4.1 Modèle de jeu de services

Le présent article spécifie les *Services OPC UA*. Les définitions de *services OPC UA* sont des descriptions abstraites, et ne constituent donc pas une spécification de mise en œuvre. Le mapping entre les descriptions abstraites et la *Pile de communication* issue de ces *Services* est défini dans l'IEC 62541-6. Dans le cas d'une mise en œuvre en tant que *services Web*, les *Services OPC UA* correspondent au service Web et un *Services OPC UA* correspond à une opération du service Web.

Ces *Services* sont organisés en *Jeux de services*. Chaque *Jeu de services* définit un ensemble de *services* connexes. L'organisation en *Jeu de services* est un regroupement logique utilisé dans le présent document, mais qui n'est pas utilisé pour la mise en œuvre.

Le *Jeu de services Discovery*, décrit à la Figure 1, définit les *services* qui permettent à un *Client* de découvrir les *Endpoints* implémentés par un *Serveur* et de connaître la configuration de sécurité de chacun de ces *Endpoints*.




Figure 1 – Jeu de services Discovery

Le *Jeu de services SecureChannel*, décrit à la Figure 2, définit les *Services* qui permettent à un *Client* d'établir un canal de communication pour assurer la *confidentialité* et l'*intégrité* des *Messages* échangés avec le *Serveur*.




Figure 2 – Jeu de services SecureChannel

Le *Jeu de services Session*, décrit à la Figure 3, définit les *Services* qui permettent au *Client* d'authentifier l'utilisateur au nom duquel il agit et de gérer des *Sessions*.




Figure 3 – Jeu de services Session

Le *Jeu de services NodeManagement*, décrit à la Figure 4, définit les *Services* qui permettent au *Client* d'ajouter, de modifier et de supprimer des *Nœuds* dans l'*AddressSpace*.




Figure 4 – Jeu de services NodeManagement

Le *Jeu de services View*, décrit à la Figure 5, définit les *Services* qui permettent aux *Clients* de parcourir l'*AddressSpace* ou des sous-ensembles de l'*AddressSpace* appelés *Vues*. Le *Jeu de services Query* permet aux *Clients* de récupérer un sous-ensemble de données de l'*AddressSpace* ou de la *Vue*.




Figure 5 – Jeu de services View

Le *Jeu de services Attribute* est représenté à la Figure 6. Il définit les *Services* qui permettent aux *Clients* de lire et d'écrire les *Attributs de nœuds*, y compris les valeurs de l'historique. La valeur d'une *Variable* étant modélisée comme *Attribut*, ces *Services* permettent aux *Clients* de lire et d'écrire les valeurs des *Variables*.




Figure 6 – Jeu de services Attributs

Le *Jeu de services Method* est représenté à la Figure 7. Il définit les *Services* qui permettent aux *Clients* d'appeler des méthodes. Les méthodes s'exécutent totalement lorsqu'elles sont appelées. Elles peuvent être appelées avec des paramètres d'entrée spécifiques à la méthode et peuvent retourner des paramètres de sortie spécifiques à la méthode.




Figure 7 – Jeu de services Method

Le *Jeu de services MonitoredItem* et le *Jeu de services Subscription*, représentés à la Figure 8, sont utilisés en commun pour s'abonner à des *nœuds* de l'*AddressSpace* d'OPC UA.

Le *Jeu de services MonitoredItem* définit les *Services* qui permettent aux *Clients* de créer, modifier et supprimer des *MonitoredItems* utilisés pour surveiller les *Attributs* pour détecter des modifications de valeur et les *Objets* pour détecter des *Événements*.

Ces *Notifications* sont mises en file d'attente en vue de leur transfert vers le *Client* via les *Abonnements*.

Le *Jeu de services Subscription* définit les *services* qui permettent aux *Clients* de créer, modifier et supprimer des *abonnements*. Les *abonnements* envoient des *Notifications* générées par des *MonitoredItems* au *Client*. Les *services Subscription* prévoient aussi la reprise du *Client* à la suite de *messages* manqués et de défaillances de communication.




Figure 8 – Jeux de services MonitoredItem et Subscription

4.2 Procédures des services Demande/Réponse

Les procédures des *services Demande/Réponse* décrivent le traitement des demandes reçues par le *Serveur* et des réponses qui lui sont renvoyées. Les procédures commencent par la présentation au *Serveur* d'un *Message* de demande de *Service* de la part du *Client* demandeur.

A réception de la demande, le *Serveur* traite le *Message* en deux étapes. A la première étape, il tente de décoder et de localiser le *Service* à exécuter. La gestion des erreurs pour cette étape est spécifique à la technologie de communication utilisée et est décrite dans l'IEC 62541-6.

S'il réussit, il tente d'accéder à chaque opération identifiée dans la demande, puis l'exécute. Pour chaque opération indiquée dans la demande, il génère un code de succès/échec distinct qu'il insère dans un *Message* de réponse positive, accompagné des éventuelles données à retourner.

Pour exécuter ces opérations, le *Client* et le *Serveur* peuvent se servir de l'API de la *Pile de communication* pour construire et interpréter les *Messages* et accéder à l'opération demandée.

La mise en œuvre de la gestion de chaque demande de service ou de chaque réponse doit vérifier que chaque paramètre de service s'inscrit bien dans la plage spécifiée pour chaque paramètre.

5 Jeux de services

5.1 Généralités

Le présent article définit les *Jeux de services* OPC UA et leurs *Services*. L'Article 7 contient les définitions des paramètres communs utilisés par ces *Services*. Les exigences relatives à l'audit sont décrites en 6.5 pour tous les services.

Le Profil définit si un *Serveur* prend ou non en charge un *Jeu de services*, ou un *Service* inclus dans un *Jeu de services*. Les *Profils* sont décrits dans l'IEC 62541-7.

5.2 En-tête de demande et de réponse de services

Chaque demande de *service* a un *RequestHeader* et chaque réponse relative à un *service* a un *ResponseHeader*.

La structure du *RequestHeader* est définie en 7.28 et contient les paramètres de demande communs, comme *authenticationToken*, *timestamp* et *requestHandle*.

La structure du *ResponseHeader* est définie en 7.29 et contient les paramètres de réponse communs, comme *serviceResult* et *diagnosticInfo*.

5.3 Résultats de service

Les résultats de *Service* sont retournés à deux niveaux dans les réponses OPC UA, l'un indiquant le statut de l'appel de *Service* et l'autre le statut de chaque opération exigée par le *Service*.

Les résultats de *Service* sont définis via le *StatusCode* (voir 7.34).

Le statut de l'appel *Service* est représenté par le *serviceResult* contenu dans le *ResponseHeader* (voir 7.29). Le mécanisme pour retourner ce paramètre est spécifique à la technologie de communication utilisée pour acheminer la réponse de *Service*, et est défini dans l'IEC 62541-6.

Le statut des opérations individuelles d'une demande est représenté par des *StatusCodes* individuels.

Les cas ci-après définissent l'utilisation de ces paramètres.

- a) Un code *Bad* est retourné dans le *serviceResult* si le *Service* lui-même a failli. Dans ce cas, un *ServiceFault* est renvoyé. Le *ServiceFault* est défini en 7.30.
- b) Le code *Good* est retourné dans le *serviceResult* si le *Service* a réussi, partiellement ou totalement. Dans ce cas, d'autres paramètres de réponse sont retournés. Le *Client* doit toujours vérifier les paramètres de réponse, notamment tous les *StatusCodes* associés à chaque opération. Ces *StatusCodes* peuvent indiquer des résultats incorrects ou incertains pour une ou plusieurs opérations demandées lors de l'appel de *Service*.

Tous les *Services* dont la demande inclut des matrices d'opération doivent retourner un code *Bad* dans le *serviceResult* lorsque la matrice est vide.

Les *Services* définissent divers *StatusCodes* spécifiques, qui doivent être utilisés par le *Serveur* selon la description du *Service*. Il convient qu'un *Client* soit capable de gérer les *StatusCodes* spécifiques à ce *Service*. En outre, un *Client* doit prendre en compte d'autres *StatusCodes* communs définis dans le Tableau 177 et dans le Tableau 178. Des informations complémentaires pour la gestion de *StatusCodes* spécifiques par le *Client* peuvent être définies dans l'IEC 62541-7.

Si le *Serveur* découvre par des mécanismes hors bande que l'intégrité des authentifiants de l'application ou de l'utilisateur utilisés pour créer une *Session* ou un *SecureChannel* a été compromise, il convient que le *Serveur* interrompe immédiatement toutes les sessions et les canaux qui utilisent ces authentifiants. Dans ce cas, il convient que le code de résultat du *Service* soit *Bad_IdentityTokenRejected* ou *Bad_CertificateUntrusted*.

L'analyse de messages peut échouer en raison d'erreurs de syntaxe ou si les données contenues dans le message sortent de la plage prise en charge par le récepteur. Lorsque cela se produit, les messages doivent être rejetés par le récepteur. Si le récepteur est un *Serveur*, il doit retourner un *ServiceFault* avec le code de résultat *Bad_DecodingError* ou *Bad_EncodingLimitsExceeded*. Si le récepteur est le *Client*, il convient que la *Pile de communication* signale ces erreurs à l'application *Client*.

De nombreuses applications imposent des limites à la taille des messages et/ou des éléments de données contenus dans ces messages. Par exemple, un *Serveur* peut rejeter des demandes contenant des valeurs de chaîne supérieures à une longueur donnée. Ces limites sont généralement établies par des administrateurs et s'appliquent à l'ensemble des connexions entre un *Client* et un *Serveur*.

Il sera vraisemblablement nécessaire que les *Clients* qui reçoivent des défauts *Bad_EncodingLimitsExceeded* de la part du *Serveur* reformulent leurs demandes. Il peut être nécessaire que l'administrateur augmente les limites pour le *Client* s'il reçoit une réponse du *Serveur* comportant ce défaut.

Dans certains cas, les erreurs d'analyse sont fatales et il n'est pas possible de retourner un défaut. Par exemple, le message entrant peut dépasser la capacité du tampon du récepteur. Dans ces cas, ces erreurs peuvent être traitées comme un défaut de communication qui exige le rétablissement du *SecureChannel* (voir 5.5).

Le *Client* et le *Serveur* réduisent les risques d'erreur fatale en échangeant leurs limites de taille de message dans le service *CreateSession*. Cela permet à chaque partie prenante d'éviter d'envoyer un message qui génère un défaut de communication. Il convient que le *Serveur* retourne une erreur *Bad_ResponseTooLarge* si un message de réponse sérialisé dépasse la taille de message spécifiée par le *Client*. De même, il convient que la *Pile de communication du Client* rapporte une erreur *Bad_RequestTooLarge* à l'application avant d'envoyer un message qui dépasse la limite du *Serveur*.

Noter que les limites de taille de message s'appliquent seulement au corps de message. Elles n'incluent pas les en-têtes ou les effets de l'application d'une quelconque politique de sécurité. Cela signifie qu'un corps de message dont la longueur est inférieure au maximum spécifié peut tout de même générer une erreur fatale.

5.4 Jeu de services Discovery

5.4.1 Vue d'ensemble

Ce *Jeu de services* définit les *Services* utilisés pour découvrir les *Endpoints* implémentés par un *Serveur* et pour lire la configuration de sécurité pour ces *Endpoints*. Les *Services Découverte* sont mis en œuvre par des *Serveurs* individuels et par des *Serveurs Découverte* spécialisés. L'IEC 62541-12 décrit comment utiliser les *Services Découverte* avec des *Serveurs Découverte* dédiés.

Chaque *Serveur* doit avoir un *DiscoveryEndpoint* auquel les *Clients* peuvent accéder sans établir une *Session*. Cet *Endpoint* peut ne pas être le même *Endpoint de session* que les *Clients* utilisent pour établir un *SecureChannel*. Les *Clients* accèdent aux informations de sécurité nécessaires pour établir un *SecureChannel* en appelant le *Service GetEndpoints* sur le *DiscoveryEndpoint*.

En outre, les *Serveurs* peuvent s'enregistrer auprès d'un *serveur Discovery* connu en utilisant le service *RegisterServer*. Les *Clients* peuvent accéder ultérieurement à la liste de tous les *Serveurs* enregistrés en appelant le service *FindServers* sur le *serveur Discovery*.

Le processus de découverte à l'aide de *FindServers* est décrit à la Figure 9. Pour des questions de clarté, la mise en place d'un *SecureChannel* (avec *MessageSecurityMode NONE*) pour *FindServers* et *GetEndpoints* n'est pas mentionnée dans la figure.




Figure 9 – Processus de découverte

L'URL pour un *DiscoveryEndpoint* doit fournir toutes les informations nécessaires au Client pour se connecter au *DiscoveryEndpoint*.

Lorsqu'un *Client* récupère les *Endpoints*, le *Client* peut sauvegarder cette information et l'utiliser pour se reconnecter directement au *Serveur* sans passer par le processus de découverte. Si le *Client* s'aperçoit qu'il ne peut pas se connecter, la configuration du *Serveur* peut avoir été modifiée. Il est alors nécessaire que le *Client* repasse par le processus de découverte.

Les *DiscoveryEndpoints* ne doivent nécessiter aucune sécurité pour les messages, mais peuvent exiger une sécurité de la couche transport. Pour des raisons de sécurité, les Administrateurs peuvent désactiver la découverte dans les systèmes de production. Les *Clients* doivent alors s'appuyer sur les *EndpointDescriptions* placées en cache. Pour prendre en charge des systèmes dont les *services Découverte* ont été désactivés, les *Clients* doivent permettre aux *Administrateurs* de mettre à jour manuellement les *EndpointDescriptions* utilisées pour se connecter à un *Serveur*. Les *Serveurs* doivent permettre aux *Administrateurs* de désactiver le *DiscoveryEndpoint*.

Le *Client* doit être prudent lorsqu'il utilise les informations retournées par un *DiscoveryEndpoint*, car elles présentent une sécurité nulle. Pour ce faire, le *Client* compare les informations retournées par le *DiscoveryEndpoint* aux informations retournées dans la réponse *CreateSession*. Le *Client* doit vérifier que:

- L'*ApplicationUri* spécifié dans le *Certificat de Serveur* est le même que l'*ApplicationUri* fourni dans l'*EndpointDescription*.
- Le *Certificat* du *Serveur* retourné dans la réponse *CreateSession* est le même que le *Certificat* utilisé pour créer le *SecureChannel*.
- Les *EndpointDescriptions* retournées par le *DiscoveryEndpoint* sont les mêmes que les *EndpointDescriptions* retournées dans la réponse *CreateSession*.

Si le *Client* détecte que l'une des exigences ci-dessus n'est pas remplie, le *Client* doit alors fermer le *SecureChannel* et émettre un rapport d'erreur.

Le *Client* doit vérifier que le *HostName* spécifié dans le *Certificat du serveur* est le même que le *HostName* contenu dans l'*endpointUrl* fourni dans l'*EndpointDescription* retournée par *CreateSession*. Si ce n'est pas le cas, le *Client* doit rapporter la différence et peut fermer le *SecureChannel*. Les *Serveurs* doivent ajouter tous les *HostNames* possibles, comme *MyHost* et *MyHost.local*, dans le *Certificat du Serveur*. Ceci comprend les adresses IP de l'hôte ou l'*HostName* présenté par un routeur NAT utilisé pour établir la connexion avec le *Serveur*.

5.4.2 FindServers

5.4.2.1 Description

Ce *service* retourne les *Serveurs* connus auprès d'un *Serveur* ou d'un *serveur Discovery*. Le comportement des *serveurs Discovery* est décrit en détail dans l'IEC 62541-12.

Le *Client* peut réduire le nombre de résultats retournés en spécifiant des critères de filtrage. Un *serveur Discovery* retourne une liste vide si aucun *Serveur* ne correspond aux critères spécifiés par le *Client*. Les critères de filtrage pris en charge par *Service* sont décrits en 5.4.2.2.

Chaque *Serveur* doit fournir un *DiscoveryEndpoint* qui prend en charge ce *Service*. Le *Serveur* doit toujours retourner un enregistrement qui le décrit, mais plusieurs enregistrements peuvent être retournés dans certains cas. Les *Serveurs passerelles* doivent retourner un enregistrement pour chaque *Serveur* auquel ils accordent l'accès, en plus d'un éventuel enregistrement permettant au *Serveur passerelle* d'être accessible comme un *Serveur OPC UA* ordinaire. Les *Serveurs* redondants non transparents doivent fournir un enregistrement pour chaque *Serveur* du *Jeu de serveurs redondants*.

Chaque *Serveur* doit avoir un identificateur unique au niveau global, appelé *ServerUri*. Il convient que cet identificateur soit un nom de domaine qualifié complet; cependant, cela peut être un GUID ou une syntaxe similaire, qui garantit son caractère unique au niveau global. Le *ServerUri* retourné par ce *Service* doit avoir la valeur qui correspond à l'indice 0 de la propriété *ServerArray* (voir IEC 62541-5). Le *ServerUri* est retourné dans le champ *applicationUri* dans l'*ApplicationDescription* (voir 7.1).

Chaque *Serveur* doit aussi avoir un identificateur interprétable par l'utilisateur appelé *ServerName*, et qui n'est pas nécessairement unique au niveau global. Cet identificateur peut être disponible dans plusieurs paramètres de lieu.

Un *Serveur* peut avoir plusieurs *HostNames*. Pour cette raison, le *Client* doit transmettre à ce *Service* l'URL qu'il a utilisée pour se connecter à l'*Endpoint*. L'implémentation de ce *Service* doit utiliser cette information pour retourner des réponses qui sont accessibles au *Client* via l'URL fournie.

Ce *Service* ne doit nécessiter aucune sécurité pour les messages, mais peut exiger une sécurité de la couche transport.

Certains *Serveurs* peuvent être accessibles via un *Serveur passerelle* et doivent avoir une valeur spécifiée pour *gatewayServerUri* dans leur *ApplicationDescription* (voir 7.1). Les *discoveryUrls* fournis dans l'*ApplicationDescription* doivent appartenir au *Serveur passerelle*. Certains *serveurs Discovery* peuvent retourner plusieurs enregistrements pour le même *Serveur*, si le *Serveur* considéré est accessible par plusieurs chemins.

Ce *Service* peut être utilisé sans sécurité. Il est donc vulnérable aux attaques par déni de service (DOS). Il convient qu'un *Serveur* réduise le plus possible la durée de traitement exigée pour envoyer la réponse à ce *Service*. Cela peut être fait en préparant le résultat à

l'avance. Il convient que le *Serveur* ajoute également un faible retard avant de commencer le traitement d'une demande si les demandes sont nombreuses.

5.4.2.2 Paramètres

Le Tableau 3 définit les paramètres du *Service*.

Tableau 3 – Paramètres du service FindServers

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. S'il est fourni, l' <i>authenticationToken</i> doit être ignoré. Le type <i>RequestHeader</i> est défini en 7.28.
endpointUrl	String	Adresse réseau que le <i>Client</i> a utilisée pour accéder au <i>DiscoveryEndpoint</i> . Le <i>Serveur</i> utilise cette information pour des diagnostics et pour déterminer les URL à retourner dans la réponse. Il convient que le <i>Serveur</i> retourne une URL adéquate par défaut s'il ne reconnaît pas le <i>HostName</i> dans l'URL.
localeIds []	LocaleId	Liste des paramètres de lieu à utiliser. Il convient que le <i>Serveur</i> retourne l' <i>ApplicationName</i> en utilisant l' <i>ApplicationDefinition</i> définie en 7.1, en utilisant un des paramètres de lieu spécifiés. Si le <i>Serveur</i> prend en charge plus d'un des paramètres de lieu demandés, alors le <i>Serveur</i> doit utiliser celui qui apparaît le premier dans la liste. Si le <i>Serveur</i> ne prend en charge aucun des paramètres de lieu demandés, il choisit un paramètre de lieu adéquat par défaut. Le <i>Serveur</i> choisit un paramètre de lieu adéquat par défaut si cette liste est vide.
serverUris []	String	Liste des <i>Serveurs</i> à retourner. Tous les <i>Serveurs</i> connus sont retournés si la liste est vide. Un <i>serverUri</i> correspond à l' <i>applicationUri</i> dans l' <i>ApplicationDescription</i> définie en 7.1.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs. Le type <i>ResponseHeader</i> est défini en 7.29.
servers []	ApplicationDescription	Liste des <i>Serveurs</i> qui répondent aux critères spécifiés dans la demande. Cette liste est vide si aucun <i>Serveur</i> ne répond aux critères. Le type <i>ApplicationDescription</i> est défini en 7.1.

5.4.2.3 Résultats de service

Les *StatusCodes* communs sont définis dans le Tableau 177.

5.4.3 FindServersOnNetwork

5.4.3.1 Description

Ce *Service* retourne les *Serveurs* connus auprès d'un *serveur Discovery*. Contrairement au *FindServers*, ce *Service* n'est mis en œuvre que par les *serveurs Discovery*.

Le *Client* peut réduire le nombre de résultats retournés en spécifiant des critères de filtrage. Une liste vide est retournée si aucun *Serveur* ne correspond aux critères spécifiés par le *Client*.

Ce *Service* ne doit nécessiter aucune sécurité pour les messages, mais peut exiger une sécurité de la couche transport.

Chaque fois que le *Serveur Découverte* crée ou met à jour un enregistrement dans son cache, il doit affecter à l'enregistrement un identifiant qui augmente de façon monotone. Ceci permet aux *Clients* de demander des enregistrements en lots en spécifiant l'identifiant pour le dernier enregistrement reçu dans le dernier appel à *FindServersOnNetwork*. Pour le prendre en charge, le *serveur Discovery* doit retourner les enregistrements par ordre numérique en commençant par le plus faible identifiant d'enregistrement. Le *Serveur Découverte* doit également retourner l'heure à laquelle le compteur a été réinitialisé pour la dernière fois, par exemple en raison d'un redémarrage du *Serveur Découverte*. Si un *Client* détecte que cette heure est plus récente que la dernière fois où le *Client* a appelé le *Service*, il doit rappeler le *Service* avec un *startingRecordId* de 0.

Ce *Service* peut être utilisé sans sécurité. Il est donc vulnérable aux attaques par déni de service (DOS). Il convient qu'un *Serveur* réduise le plus possible la durée de traitement exigée pour envoyer la réponse à ce *Service*. Cela peut être fait en préparant le résultat à l'avance.

5.4.3.2 Paramètres

Le Tableau 4 définit les paramètres du *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020

Tableau 4 – Paramètres du service FindServersOnNetwork

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. S'il est fourni, l' <i>authenticationToken</i> doit être ignoré. Le type <i>RequestHeader</i> est défini en 7.28.
startingRecordId	Counter	Seuls les enregistrements avec un identifiant supérieur à ce nombre sont retournés. Spécifier 0 pour démarrer avec le premier enregistrement en cache.
maxRecordsToReturn	UInt32	Nombre maximal d'enregistrements à retourner en réponse. 0 indique qu'il n'y a pas de limite.
serverCapabilityFilter []	String	Liste des filtres de capacités de <i>Serveur</i> . L'ensemble des capacités de <i>Serveur</i> autorisées est défini dans l'IEC 62541-12. Seuls les enregistrements affichant l'ensemble des capacités de <i>Serveur</i> spécifiées sont retournés. La comparaison est sensible à la casse. Si cette liste est vide, alors aucun filtrage n'est effectué.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs. Le type <i>ResponseHeader</i> est défini en 7.29.
lastCounterResetTime	UtcTime	Heure où les compteurs ont été réinitialisés pour la dernière fois.
servers []	ServerOnNetwork	Liste des enregistrements du service DNS qui répondent aux critères spécifiés dans la demande. Cette liste est vide si aucun <i>Serveur</i> ne répond aux critères.
recordId	UInt32	Identificateur unique pour l'enregistrement. Il peut être utilisé pour rechercher le prochain lot de <i>Serveurs</i> dans un appel consécutif au <i>FindServersOnNetwork</i> .
serverName	String	Nom du <i>Serveur</i> spécifié dans l'annonce mDNS (voir IEC 62541-12). Celui-ci peut être le même que l' <i>ApplicationName</i> du <i>Serveur</i> .
discoveryUrl	String	URL du <i>DiscoveryEndpoint</i> .
serverCapabilities	String []	Ensemble des capacités de <i>Serveur</i> prises en charge par le <i>Serveur</i> . L'ensemble des capacités de <i>Serveur</i> autorisées est défini dans l'IEC 62541-12.

5.4.3.3 Résultats de service

Les *StatusCodes* communs sont définis dans le Tableau 177.

5.4.4 GetEndpoints

5.4.4.1 Description

Ce *Service* retourne les *Endpoints* pris en charge par un *Serveur* et la totalité des informations de configuration nécessaires pour établir un *SecureChannel* et une *Session*.

Ce *Service* ne doit nécessiter aucune sécurité pour les messages, mais peut exiger une sécurité de la couche transport.

Un *Client* peut réduire le nombre de résultats retournés en spécifiant des critères de filtrage sur la base des valeurs *LocaleIds* et des URI des *Profils de transport*. Le *Serveur* retourne une liste vide si aucun *Endpoint* ne correspond aux critères spécifiés par le *Client*. Les critères de filtrage pris en charge par *Service* sont décrits en 5.4.4.2.

Un *Serveur* peut prendre en charge plusieurs configurations de sécurité pour le même *Endpoint*. Dans cette situation, le *Serveur* doit retourner des enregistrements d'*EndpointDescription* distincts pour chaque configuration disponible. Il convient que les *Clients* traitent chacune de ces configurations comme des *Endpoints* distincts même si l'URL physique se trouve être la même.

La configuration de sécurité pour un *Endpoint* comporte quatre composants:

- le Certificat d'instance d'application du *Serveur*;
- le mode de sécurité des messages;
- la politique de sécurité;
- les jetons d'identité utilisateur pris en charge.

L'*ApplicationInstanceCertificate* est utilisé pour sécuriser la demande d'*OpenSecureChannel* (voir 5.5.2). Le *MessageSecurityMode* et la *SecurityPolicy* indiquent au *Client* comment sécuriser les messages envoyés via le *SecureChannel*. Les *UserIdentityTokens* indiquent au *Client* le type d'identifiant d'utilisateur qui doit être transmis au *Serveur* dans la demande *ActivateSession* (voir 5.6.3).

Si le *securityPolicyUri* est NONE et qu'aucune *UserTokenPolicy* n'exige de chiffrement, le *Client* doit ignorer l'*ApplicationInstanceCertificate*.

Chaque *EndpointDescription* spécifie aussi un URI pour le *Profil de Transport* que l'*Endpoint* prend en charge. Les *Profils de transport* spécifient les informations comme le format d'encodage des messages et la version des protocoles, ils sont définis dans l'IEC 62541-7.

Les messages sont sécurisés par l'application d'algorithmes de chiffrement normalisés avant envoi sur le réseau. Le jeu exact d'algorithmes utilisés dépend de la *SecurityPolicy* pour l'*Endpoint*. L'IEC 62541-7 définit des *Profils* pour des *SecurityPolicies* communes et leur affecte un URI unique. Les applications sont supposées connaître les *SecurityPolicies* qu'elles prennent en charge. Par conséquent, seul l'URI de profil pour la *SecurityPolicy* est spécifié dans l'*EndpointDescription*. Un *Client* ne peut pas se connecter à un *Endpoint* qui ne prend pas en charge une *SecurityPolicy* qu'il connaît.

Une *EndpointDescription* peut spécifier que le mode de sécurité de message est NONE. Cette configuration n'est pas recommandée, sauf si les applications communiquent sur un réseau physiquement isolé présentant un risque d'intrusion extrêmement faible. Si la sécurité de message est NONE, il est possible pour les *Clients* de pirater, délibérément ou accidentellement, des *Sessions* créées par d'autres *Clients*.

Un *Serveur* peut avoir plusieurs *HostNames*. Pour cette raison, le *Client* doit transmettre à ce *Service* l'URL qu'il a utilisée pour se connecter à l'*Endpoint*. L'implémentation de ce *Service* doit utiliser cette information pour retourner des réponses qui sont accessibles au *Client* via l'URL fournie.

Ce *Service* peut être utilisé sans sécurité. Il est donc vulnérable aux attaques par déni de service (DOS). Il convient qu'un *Serveur* réduise le plus possible la durée de traitement exigée pour envoyer la réponse à ce *Service*. Cela peut être fait en préparant le résultat à l'avance. Il convient que le *Serveur* ajoute également un faible retard avant de commencer le traitement d'une demande si les demandes sont nombreuses.

Certaines des *EndpointDescriptions* retournées dans une réponse doivent spécifier les informations de l'*Endpoint* pour un *Serveur passerelle* qui peut être utilisé pour accéder à un

autre *Serveur*. Dans ce cas, le *gatewayServerUri* est spécifié dans l'*EndpointDescription*, et tous les contrôles de sécurité utilisés pour vérifier les *Certificats* doivent utiliser le *gatewayServerUri* (voir 6.1.3) au lieu du *serverUri*.

Pour se connecter à un *Serveur* via la passerelle, le *Client* doit d'abord établir un *SecureChannel* avec le *Serveur passerelle*. Le *Client* doit appeler le service *CreateSession* et transmettre au *Serveur passerelle* le *serverUri* spécifié dans l'*EndpointDescription*. Le *Serveur passerelle* doit alors se connecter au *Serveur* sous-jacent pour le compte du *Client*. Le processus de connexion au *Serveur* via un *Serveur passerelle* est représenté à la Figure 10.




Figure 10 – Utilisation d'un Serveur passerelle

5.4.4.2 Paramètres

Le Tableau 5 définit les paramètres du *Service*.

Tableau 5 – Paramètres du service GetEndpoints

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. S'il est fourni, l' <i>authenticationToken</i> doit être ignoré. Le type <i>RequestHeader</i> est défini en 7.28.
endpointUrl	String	Adresse réseau que le <i>Client</i> a utilisée pour accéder au <i>DiscoveryEndpoint</i> . Le <i>Serveur</i> utilise cette information pour des diagnostics et pour déterminer les URL à retourner dans la réponse. Il convient que le <i>Serveur</i> retourne une URL adéquate par défaut s'il ne reconnaît pas le <i>HostName</i> dans l'URL.
localeIds []	LocaleId	Liste des paramètres de lieu à utiliser. Spécifie le paramètre de lieu à utiliser pour retourner des chaînes lisibles par l'homme. Ce paramètre est décrit en 5.4.2.2.
profileUris []	String	Liste de <i>Profils de transport</i> que doivent prendre en charge les <i>Endpoints</i> retournés. L'IEC 62541-7 définit les URI pour les <i>Profils de transport</i> . Tous les <i>Endpoints</i> sont retournés si la liste est vide. Si l'URI est un URL, cet URL peut avoir une chaîne de requête accolée. Les <i>Profils de transport</i> qui prennent en charge les chaînes de requête sont définis dans l'IEC 62541-7.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs. Le type <i>ResponseHeader</i> est défini en 7.29.
Endpoints []	EndpointDescription	Liste des <i>Endpoints</i> qui répondent aux critères spécifiés dans la demande. Cette liste est vide si aucun <i>Endpoint</i> ne répond aux critères. Le type <i>EndpointDescription</i> est défini en 7.10.

5.4.4.3 Résultats de service

Les *StatusCodes* communs sont définis dans le Tableau 177.

5.4.5 RegisterServer

5.4.5.1 Description

Ce *Service* est implémenté par les *serveurs Discovery*.

Ce *Service* enregistre un *Serveur* auprès d'un *Serveur Découverte*. Ce *Service* est appelé par un *Serveur* ou par un utilitaire de configuration distinct. Les *Clients* n'utilisent pas ce *Service*.

Un *Serveur* doit établir un *SecureChannel* avec le *Serveur Découverte* avant d'appeler ce *Service*. Le *SecureChannel* est décrit en 5.5. L'*Administrateur* du *serveur* doit fournir au *Serveur* une *EndpointDescription* pour le *serveur Discovery*, ce qui fait partie intégrante du processus de configuration. Les *serveurs Discovery* doivent rejeter les enregistrements si le *serverUri* fourni ne correspond pas à l'*applicationUri* contenue dans le *Certificat du serveur* utilisé pour créer le *SecureChannel*.

Ce *Service* ne peut être invoqué qu'à l'aide des *SecureChannels* qui prennent en charge l'authentification *Client* (c'est-à-dire que HTTPS ne peut être utilisé pour appeler ce *Service*).

Un *Serveur* fournit uniquement son *serverUri* et les URL des *DiscoveryEndpoints* au *Serveur Découverte*. Les *Clients* doivent utiliser le *Service GetEndpoints* pour rechercher directement auprès du *Serveur* les informations de configuration les plus récentes.

Le *Serveur* doit fournir un nom localisé pour lui-même dans tous paramètres de lieu qu'il prend en charge.

Les *Serveurs* doivent être capables de s'enregistrer auprès d'un *Serveur Découverte* s'exécutant sur la même machine. Les mécanismes exacts dépendent de l'implémentation du *Serveur Découverte* et sont décrits dans l'IEC 62541-6.

Il existe deux types d'applications *Serveur*: celles qui sont lancées manuellement, y compris par un démarrage par le système d'exploitation à l'amorçage et celles qui sont lancées automatiquement lorsqu'un *Client* tente de se connecter. Le processus d'enregistrement qu'un *Serveur* doit utiliser dépend de la catégorie dans laquelle il s'inscrit.

Le processus d'enregistrement pour les *Serveurs* lancés manuellement est représenté à la Figure 11.




Figure 11 – Processus d'enregistrement – Serveurs lancés manuellement

Le processus d'enregistrement pour les *Serveurs* lancés automatiquement est représenté à la Figure 12.




Figure 12 – Processus d'enregistrement – Serveurs lancés automatiquement

Le processus d'enregistrement est conçu pour être indépendant de la plateforme, être robuste et capable de réduire le plus possible les problèmes créés par des erreurs de configuration. Pour cette raison, les *Serveurs* doivent s'enregistrer plusieurs fois.

Dans des conditions normales, les *Serveurs* lancés manuellement doivent s'enregistrer régulièrement auprès du *serveur Discovery* tant qu'ils peuvent recevoir des connexions des *Clients*. Si un *Serveur* passe hors ligne, il doit s'enregistrer une fois de plus et indiquer s'il passe hors ligne. Il convient que la période d'enregistrement récurrent soit configurable; cependant, la valeur maximale est de 10 minutes. Si une erreur se produit au cours de l'enregistrement (par exemple, le *serveur Discovery* n'est pas en cours d'exécution), le *Serveur* doit tenter de se réenregistrer régulièrement. Il convient que la fréquence de ces tentatives commence à 1 s, puis augmente progressivement jusqu'à ce que la fréquence d'enregistrement atteigne la valeur qu'elle aurait eue si aucune erreur ne s'était produite. L'approche recommandée est de doubler la période de chaque tentative jusqu'à ce que le maximum soit atteint.

Lorsqu'un *Serveur* lancé automatiquement (ou son programme d'installation) s'enregistre auprès du *Serveur Découverte*, il doit fournir un chemin vers un fichier sémaphore que le *Serveur Découverte* peut utiliser pour déterminer si le *Serveur* a été désinstallé de la machine. Le *Serveur Découverte* doit avoir eu accès au système de fichiers qui contient le fichier.

5.4.5.2 Paramètres

Le Tableau 6 définit les paramètres du *Service*.

Tableau 6 – Paramètres du service RegisterServer

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. Le type <i>RequestHeader</i> est défini en 7.28.
Server	RegisteredServer	Serveur à enregistrer. Le type <i>RegisteredServer</i> est défini en 7.27.
Réponse		
ResponseHeader	ResponseHeader	Paramètres de réponse communs. Le type <i>ResponseHeader</i> est défini en 7.29.

5.4.5.3 Résultats de service

Le Tableau 7 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 7 – Codes de résultats du service RegisterServer

Id Symbolique	Description
Bad_InvalidArgument	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ServerUriInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ServerNameMissing	Aucun <i>ServerName</i> n'a été spécifié.
Bad_DiscoveryUriMissing	Aucune URL de découverte n'a été spécifiée.
Bad_SemaphoreFileMissing	Le fichier sémaphore spécifié n'est pas valide.

5.4.6 RegisterServer2

5.4.6.1 Description

Ce *Service* est implémenté par les *serveurs Discovery*.

Ce *Service* permet à un *serveur* d'enregistrer les *DiscoveryUrls* et les capacités avec un *serveur Discovery*. Il complète les informations relatives à l'enregistrement du *RegisterServer* avec les informations nécessaires à *FindServersOnNetwork*. Ce *Service* est appelé par un *Serveur* ou par un utilitaire de configuration distinct. Les *Clients* n'utilisent pas ce *Service*.

Les *Serveurs* qui prennent en charge *RegisterServer2* doivent tenter de s'enregistrer avec le *serveur Discovery* à l'aide de ce *Service* et doivent revenir au *RegisterServer* si *RegisterServer2* échoue avec le statut *Bad_ServiceUnsupported*.

Il est nécessaire qu'un *serveur Discovery* mettant en œuvre ce *Service* attribue des identifiants d'enregistrement uniques chaque fois que ce *Service* est appelé. Voir 5.4.3 pour plus d'informations.

Ce *Service* ne peut être invoqué qu'à l'aide des *SecureChannels* qui prennent en charge l'authentification *Client* (c'est-à-dire que HTTPS ne peut être utilisé pour appeler ce *Service*).

5.4.6.2 Paramètres

Le Tableau 8 définit les paramètres du *Service*.

Tableau 8 – RegisterServer2

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. Le type <i>RequestHeader</i> est défini en 7.28.
Server	RegisteredServer	Serveur à enregistrer. Le type <i>RegisteredServer</i> est défini en 7.27.
discoveryConfiguration []	ExtensibleParameter DiscoveryConfiguration	Réglages de configuration supplémentaires pour le <i>Serveur</i> à enregistrer. Le <i>discoveryConfiguration</i> est un type de paramètre extensible défini en 7.9. Les <i>serveurs Discovery</i> qui ne comprennent pas une configuration doivent retourner <i>Bad_NotSupported</i> pour cette configuration.
Réponse		
ResponseHeader	ResponseHeader	Paramètres de réponse communs. Le type <i>ResponseHeader</i> est défini en 7.29.
configurationResults []	StatusCode	Liste de résultats pour les paramètres <i>discoveryConfiguration</i> .
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les paramètres <i>discoveryConfiguration</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.4.6.3 Résultats de service

Le Tableau 9 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 9 – Codes de résultats du service RegisterServer2

ID symbolique	Description
Bad_InvalidArgument	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ServerUriInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ServerNameMissing	Aucun <i>ServerName</i> n'a été spécifié.
Bad_DiscoveryUriMissing	Aucune URL de découverte n'a été spécifiée.
Bad_SemaphoreFileMissing	Le fichier sémaphore spécifié n'est pas valide.
Bad_ServiceUnsupported	Voir Tableau 177 pour la description de ce code de résultat.

5.4.6.4 StatusCodes

Le Tableau 10 définit les valeurs de niveau opération pour les paramètres *configurationResults* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 10 – Codes de résultats de niveau opération RegisterServer2

ID symbolique	Description
Bad_NotSupported	Voir Tableau 178 pour la description de ce code de résultat.

5.5 Jeu de services SecureChannel

5.5.1 Vue d'ensemble

Ce *Jeu de services* définit les *Services* utilisés pour ouvrir un canal de communication qui assure la confidentialité et l'*Intégrité* de tous les *Messages* échangés avec le *Serveur*. Les concepts de base pour la sécurité OPC UA sont définis dans l'IEC TR 62541-2.

Les *Services SecureChannel* sont différents des autres *Services*, car ils ne sont pas mis en œuvre directement par l'*Application OPC UA*. Ils sont en effet fournis par la *Pile de communication* sur laquelle repose l'*Application OPC UA*. Par exemple, un *Serveur OPC UA* peut reposer sur une pile qui permet aux applications d'établir un *SecureChannel* en utilisant HTTPS. Dans ce cas, l'*Application OPC UA* doit vérifier qu'elle a bien reçu le *Message* dans le cadre d'une connexion HTTPS. L'IEC 62541-6 décrit comment les *Services SecureChannel* sont mis en œuvre.

Un *SecureChannel* est une connexion logique de longue durée établie entre un seul *Client* et un seul *Serveur*. Ce canal gère un jeu de clés connues seulement du *Client* et du *Serveur*, qui sont utilisées pour authentifier et chiffrer les *Messages* au sein du réseau. Les *Services SecureChannel* permettent au *Client* et au *Serveur* de négocier en toute sécurité les clés à utiliser.

Les connexions logiques peuvent être initialisées par le *Client* ou par le *Serveur* comme décrit dans l'IEC 62541-6. Après l'initialisation de la connexion, le *SecureChannel* est ouvert et fermé par le *Client* en utilisant les *Services SecureChannel*.

Une *EndpointDescription* indique au *Client* comment établir un *SecureChannel* avec un *Endpoint* donné. Un *Client* peut obtenir l'*EndpointDescription* provenant d'un *serveur Découverte*, soit par le biais d'un serveur non défini par l'UA, soit à partir de sa propre configuration.

Les algorithmes exacts utilisés pour authentifier et chiffrer les *Messages* sont décrits dans le champ *SecurityPolicy* de l'*EndpointDescription*. Un *Client* doit utiliser ces algorithmes lorsqu'il crée un *SecureChannel*.

Il convient de noter que certaines *SecurityPolicies* définies dans l'IEC 62541-7 désactivent l'authentification et le chiffrement, ce qui compromet la sécurité du *SecureChannel*.

Lorsqu'un *Client* et un *Serveur* sont en communication via un *SecureChannel*, ils doivent vérifier que tous les *Messages* entrants ont été signés et chiffrés conformément aux exigences spécifiées dans l'*EndpointDescription*. Une *Application OPC UA* ne doit traiter aucun *Message* qui ne serait pas conforme à ces exigences.

La relation entre le *SecureChannel* et l'*Application OPC UA* dépend de la technologie d'implémentation. L'IEC 62541-6 définit toutes les exigences induites par la technologie utilisée.

La corrélation entre une *Session d'Application OPC UA* et le *SecureChannel* est représentée à la Figure 13. La *Pile de communication* est utilisée par les *Applications OPC UA* pour échanger des *Messages*. Dans un premier temps, les *services SecureChannel* sont utilisés pour établir un *SecureChannel* entre les deux *Piles de communication*, ce qui permet l'échange sécurisé de *Messages*. Dans un second temps, les *Applications OPC UA* utilisent le *Jeu de services Session* pour établir une *Session d'Application OPC UA*.




Figure 13 – Services SecureChannel et Session

Lorsqu'un *Client* a établi une *Session*, il peut souhaiter accéder à cette *Session* à partir d'un *SecureChannel* différent. Le *Client* peut le faire en validant le nouveau *SecureChannel* avec le *Service ActivateSession* décrit en 5.6.3.

Si un *Serveur* agit comme *Client* auprès d'autres *Serveurs* (couramment appelé chaînage de *Serveurs*), le *Serveur* doit être capable de maintenir la sécurité au niveau de l'utilisateur. Cela signifie qu'il convient que l'identité de l'utilisateur soit transmise au *Serveur* sous-jacent, ou qu'elle soit mise en correspondance avec une identité d'utilisateur appropriée dans le serveur sous-jacent. Il est inacceptable d'ignorer la sécurité de niveau utilisateur. Elle est exigée pour s'assurer que la sécurité est maintenue et qu'un utilisateur n'obtienne pas d'informations auxquelles il convient qu'il n'ait pas accès. Autant que possible, il convient qu'un *Serveur* se substitue au *Client* d'origine en transmettant l'identité d'utilisateur du *Client* d'origine au *serveur* sous-jacent lorsqu'il appelle le *service ActivateSession*. Si la substitution d'identité n'est pas une option, le *Serveur* doit faire correspondre l'identité d'utilisateur du *Client* avec une nouvelle identité d'utilisateur que le *Serveur* sous-jacent reconnaît effectivement.

5.5.2 OpenSecureChannel

5.5.2.1 Description

Ce *Service* est utilisé pour ouvrir ou renouveler un *SecureChannel* qui peut être utilisé pour assurer la *Confidentialité* et l'*Intégrité* pour l'échange de *Messages* pendant cette *Session*. Ce *Service* exige de la *Pile de communication* qu'elle applique les différents algorithmes de sécurité aux *Messages* lorsqu'ils sont envoyés et reçus. Des mises en œuvre spécifiques de ce *Service* pour différentes *Piles de communication* sont décrites dans l'IEC 62541-6.

Chaque *SecureChannel* a un identificateur globalement unique et est valide pour une combinaison spécifique d'instances d'applications *Client* et *Serveur*. Chaque canal contient un ou plusieurs *SecurityTokens* qui identifient un jeu de clés cryptographiques qui sont utilisées pour chiffrer et authentifier des *Messages*. Les *SecurityTokens* ont aussi des identificateurs globalement uniques qui sont joints à chaque *Message* sécurisé avec le jeton. Cela permet à un récepteur autorisé de savoir comment déchiffrer et vérifier les *Messages*.

Les *SecurityTokens* ont une durée de vie finie négociée avec ce *Service*. Cependant, les différences entre les horloges système de plusieurs machines et les temps de latence réseau impliquent que des *Messages* valides peuvent arriver après l'expiration du jeton. Pour éviter que des *Messages* valides soient rejetés, il convient que les applications exécutent les actions suivantes:

- il convient que les *Clients* demandent un nouveau *SecurityToken* dès que 75 % de leur durée de vie se sont écoulés; il convient que cela garantisse que les *Clients* reçoivent le nouveau *SecurityToken* avant que l'ancien n'expire effectivement;

- b) les *Serveurs* doivent utiliser le *SecurityToken* existant pour sécuriser les *Messages* sortants jusqu'à ce que le *SecurityToken* expire ou que le *Serveur* reçoive un *message* sécurisé avec un nouveau *SecurityToken*; il convient que cela garantisse que les *Clients* ne rejettent pas des *Messages* sécurisés reçus avec le nouveau *SecurityToken* avant que le *Client* ne reçoive le nouveau *SecurityToken*;
- c) il convient que les *Clients* acceptent des *Messages* sécurisés par un *SecurityToken* expiré pendant une durée pouvant atteindre 25 % de la durée de vie du jeton; il convient que cela garantisse que les *Messages* envoyés par le *Serveur* avant l'expiration du jeton ne soient pas rejetés en raison des retards du réseau.

Chaque *SecureChannel* existe jusqu'à ce qu'il soit explicitement fermé ou jusqu'à ce que le dernier jeton ait expiré et que la période de chevauchement ait expiré. Il convient qu'une application *Serveur* limite le nombre de *SecureChannels*. Afin d'assurer une protection contre des *Clients* mal intentionnés et des attaques de déni de service, le *Serveur* doit fermer le plus ancien *SecureChannel* qui n'a pas de *Session* assignée avant d'atteindre le nombre maximal de *SecureChannels* pris en charge.

Les *Messages* de demande de l'*OpenSecureChannel* et de réponse doivent être signés avec la clé privée de l'expéditeur. Ces *Messages* doivent toujours être chiffrés. Si la couche transport ne procure pas de chiffrement, ces *Messages* doivent être chiffrés avec la clé privée du récepteur. Ces exigences pour *OpenSecureChannel* ne s'appliquent que si le *securityPolicyUri* n'est pas None.

Si le protocole défini dans l'IEC 62541-6 exige que des *Certificats d'instance d'application* soient utilisés dans le *Service OpenSecureChannel*, les *Clients* et *Serveurs* doivent vérifier que les mêmes *Certificats* ont été utilisés dans les services *CreateSession* et *ActivateSession*. Les *Certificats* ne sont pas fournis et ne doivent pas être vérifiés si le *securityPolicyUri* est None.

Si le *securityPolicyUri* n'est pas None, un *Client* doit vérifier que le *HostName* spécifié dans le *certificat du Serveur* est le même que le *HostName* contenu dans l'*endpointUrl*. En cas de différence, le *Client* doit alors signaler la différence et peut choisir de ne pas ouvrir le *SecureChannel*. Les *Serveurs* doivent ajouter tous les *HostNames* possibles, comme *MyHost* et *MyHost.local*, dans le *Certificat du Serveur*. Ceci comprend les adresses IP de l'hôte ou l'*HostName* présenté par un routeur NAT utilisé pour établir la connexion avec le *Serveur*.

Il convient que les *Clients* soient prêts à remplacer le *HostName* retourné dans l'*EndPointDescription* par le *HostName* ou les adresses IP qu'ils utilisent pour appeler les *GetEndpoints*.

5.5.2.2 Paramètres

Le Tableau 11 définit les paramètres du *Service*.

Contrairement aux autres *Services*, les paramètres applicables à ce *Service* ne fournissent qu'une définition abstraite. La représentation concrète sur le réseau dépend des mappings définis dans l'IEC 62541-6.

Tableau 11 – Paramètres du service OpenSecureChannel

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. Le type <i>RequestHeader</i> est défini en 7.28.
clientCertificate	ApplicationInstanceCertificate	<i>Certificat</i> qui identifie le <i>Client</i> . La demande d' <i>OpenSecureChannel</i> doit être signée avec la clé privée pour ce <i>Certificat</i> . Le type <i>ApplicationInstanceCertificate</i> est défini en 7.2. Si le <i>securityPolicyUri</i> est None, le <i>Serveur</i> doit ignorer l' <i>ApplicationInstanceCertificate</i> .
requestType	Enum RequestType "SecurityToken"	Type de demande de <i>SecurityToken</i> : Une énumération qui doit être l'une des suivantes: ISSUE_0 crée un nouveau <i>SecurityToken</i> pour un nouveau <i>SecureChannel</i> . RENEW_1 crée un nouveau <i>SecurityToken</i> pour un <i>SecureChannel</i> existant.
secureChannelId	BaseDataType	Identificateur du <i>SecureChannel</i> auquel il convient que le nouveau jeton appartienne. Ce paramètre doit être Null lors de la création d'un nouveau <i>SecureChannel</i> . La définition du protocole de sécurité concret dans l'IEC 62541-6 choisit le <i>Data Type</i> concret.
securityMode	Enum MessageSecurityMode	Type de sécurité à appliquer aux messages. Le type <i>MessageSecurityMode</i> est défini en 7.15. Il peut être nécessaire de créer un <i>SecureChannel</i> même si le <i>securityMode</i> est NONE. Le comportement exact dépend du mapping utilisé et est décrit dans l'IEC 62541-6.
securityPolicyUri	String	URI pour la <i>SecurityPolicy</i> à utiliser lors de la sécurisation des messages envoyés sur le <i>SecureChannel</i> . Le jeu d'URI connu et les <i>SecurityPolicies</i> qui leur sont associées sont définis dans l'IEC 62541-7.
clientNonce	ByteString	Nombre aléatoire ne devant pas être utilisé par une autre demande. Un nouveau <i>clientNonce</i> doit être généré chaque fois qu'un <i>SecureChannel</i> est renouvelé. Ce paramètre doit avoir une longueur égale à la <i>SecureChannelNonceLength</i> définie pour la <i>SecurityPolicy</i> dans l'IEC 62541-7. La <i>SecurityPolicy</i> est identifiée par le <i>SecurityPolicyUri</i> .
requestedLifetime	Duration	Durée de vie demandée, en millisecondes, pour le nouveau <i>SecurityToken</i> . Il spécifie le moment, où le <i>Client</i> compte renouveler le <i>SecureChannel</i> en appelant de nouveau le service <i>OpenSecureChannel</i> . Si un <i>SecureChannel</i> n'est pas renouvelé, tous les <i>Messages</i> envoyés en utilisant les <i>SecurityTokens</i> courants doivent être rejetés par le récepteur. L'augmentation du matériel disponible chiffré avec une clé spécifique facilite diverses attaques cryptanalytiques. En limitant le nombre de données traitées en utilisant une clé spécifique, ces attaques sont rendues plus difficiles. Par conséquent, le volume de données échangées entre le <i>Client</i> et le <i>Serveur</i> doit être limité en établissant un nouveau <i>SecurityToken</i> après la durée de vie. Le réglage de la durée de vie demandée dépend du nombre prévu de messages échangés, ainsi que de leur taille durant la durée de vie. Un volume plus important de données exige une durée de vie plus courte.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition du type <i>ResponseHeader</i>).
securityToken	ChannelSecurityToken	Décrit le nouveau <i>SecurityToken</i> émis par le <i>Serveur</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
channelId	BaseDataType	Identificateur unique pour le <i>SecureChannel</i> . Il s'agit de l'identificateur qui doit être fourni chaque fois que le <i>SecureChannel</i> est renouvelé. La définition du protocole de sécurité concret dans l'IEC 62541-6 choisit le <i>DataType</i> concret.
tokenId	ByteString	Identificateur unique pour un seul <i>SecurityToken</i> dans le canal. Il s'agit de l'identificateur qui doit être transmis avec chaque <i>Message</i> sécurisé avec le <i>SecurityToken</i> .
createdAt	UtcTime	Heure à laquelle le <i>SecurityToken</i> a été créé.
revisedLifetime	Duration	Durée de vie du <i>SecurityToken</i> en millisecondes. L'heure d'expiration TUC pour le jeton peut être calculée en ajoutant la durée de vie au temps <i>createdAt</i> .
serverNonce	ByteString	Nombre aléatoire ne devant pas être utilisé par une autre demande. Un nouveau <i>serverNonce</i> doit être généré chaque fois qu'un <i>SecureChannel</i> est renouvelé. Ce paramètre doit avoir une longueur égale à la <i>SecureChannelNonceLength</i> définie pour la <i>SecurityPolicy</i> dans l'IEC 62541-7. La <i>SecurityPolicy</i> est identifiée par le <i>SecurityPolicyUri</i> .

5.5.2.3 Résultats de service

Le Tableau 12 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 12 – Codes de résultats du service OpenSecureChannel

ID symbolique	Description
Bad_SecurityChecksFailed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateTimeInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerTimeInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateHostNameInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateUriInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateUseNotAllowed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerUseNotAllowed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateUntrusted	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateRevocationUnknown	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerRevocationUnknown	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateRevoked	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerRevoked	Voir Tableau 177 pour la description de ce code de résultat.
Bad_RequestTypeInvalid	Le type de la demande de jeton de sécurité n'est pas valide.
Bad_SecurityModeRejected	Le mode de sécurité ne répond pas aux exigences établies par le <i>Serveur</i> .
Bad_SecurityPolicyRejected	La politique de sécurité ne répond pas aux exigences établies par le <i>Serveur</i> .
Bad_SecureChannelIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_NonceInvalid	Voir Tableau 177 pour la description de ce code de résultat. Un <i>Serveur</i> doit vérifier la longueur minimale du nonce du <i>Client</i> et retourner ce statut si la longueur est inférieure à 32 octets. Une vérification des doublons de nonces ne peut être faite que dans les appels <i>OpenSecureChannel</i> avec le type de demande RENEW_1.

5.5.3 CloseSecureChannel

5.5.3.1 Description

Ce *Service* est utilisé pour mettre fin à un *SecureChannel*.

Les *Messages* de demande doivent être signés avec la clé appropriée associée au jeton courant pour le *SecureChannel*.

5.5.3.2 Paramètres

Le Tableau 13 définit les paramètres du *Service*.

Les mappings des protocoles spécifiques définis dans l'IEC 62541-6 peuvent choisir d'exclure la réponse.

Tableau 13 – Paramètres du service CloseSecureChannel

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. Le type <i>RequestHeader</i> est défini en 7.28.
secureChannelId	BaseDataType	Identificateur pour le <i>SecureChannel</i> à fermer. La définition du protocole de sécurité concret dans l'IEC 62541-6 choisit le <i>DataType</i> concret.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).

5.5.3.3 Résultats de service

Le Tableau 14 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 14 – Codes de résultats du service CloseSecureChannel

ID symbolique	Description
Bad_SecureChannelIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.

5.6 Jeu de services Session

5.6.1 Vue d'ensemble

Ce *Jeu de services* définit des *Services* pour l'établissement de connexions de couche application dans le contexte d'une *Session*.

5.6.2 CreateSession

5.6.2.1 Description

Ce *Service* est utilisé par un *Client* OPC UA pour créer une *Session* et le *Serveur* retourne deux valeurs qui identifient de façon unique cette *Session*. La première valeur est le *sessionId* qui est utilisé pour identifier cette *Session* dans les journaux d'audit et dans l'*AddressSpace* du *Serveur*. La seconde est l'*authenticationToken* qui est utilisé pour associer une demande entrante à une *Session*.

Avant d'appeler ce *Service*, le *Client* doit créer un *SecureChannel* avec le *Service OpenSecureChannel* pour garantir l'*Intégrité* de tous les *Messages* échangés au cours de cette *Session*. Ce *SecureChannel* a un identificateur unique, que le *Serveur* doit associer à l'*authenticationToken*. Le *Serveur* ne peut accepter des demandes avec l'*authenticationToken* que si elles sont associées au *SecureChannel* qui a été utilisé pour créer cette *Session*. Le *Client* peut associer un nouveau *SecureChannel* à la *Session* en appelant la méthode *ActivateSession*.

Le *SecureChannel* est toujours géré par la *Pile de communication*, ce qui signifie qu'il doit fournir les API que le *Serveur* peut utiliser pour trouver des informations du *SecureChannel* utilisé pour n'importe quelle demande donnée. La *Pile de communication* doit au minimum fournir la *SecurityPolicy* et le *SecurityMode* utilisés par le *SecureChannel*. Elle doit aussi fournir un *SecureChannelId* qui identifie de façon unique le *SecureChannel* ou le *Certificat Client* utilisé pour établir le *SecureChannel*. Le *Serveur* utilise l'un de ceux-ci pour identifier le

SecureChannel utilisé pour envoyer une demande. Le 7.31 décrit comment créer l'*authenticationToken* pour différents types de *piles de communication*.

Selon la *SecurityPolicy* et le *SecurityMode* du *SecureChannel*, l'échange des *ApplicationInstanceCertificates* et des *Nonces* peut être facultatif et les signatures peuvent être vides. Voir l'IEC 62541-7 pour la définition des *SecurityPolicies* et la gestion de ces paramètres.

Le *Serveur* retourne ses *EndpointDescriptions* dans la réponse. Les *Clients* utilisent ces informations pour déterminer si la liste des *EndpointDescriptions* retournées par le *DiscoveryEndpoint* correspond aux *Endpoints* que le *Serveur* possède. S'il y a une différence, le *Client* doit alors fermer cette *Session* et rapporter l'erreur. Le *Serveur* retourne toutes les *EndpointDescriptions* pour le *serverUri* spécifié par le *Client* dans la demande. Le *Client* vérifie uniquement les *EndpointDescriptions* ayant un *transportProfileUri* qui correspond au *profileUri* spécifié dans la demande de *GetEndpoints* initiale. Un *Client* peut ignorer cette vérification si les *EndpointDescriptions* ont été fournies par une source de confiance, comme l'*Administrateur*.

Une *Session* créée avec ce *Service* ne doit pas être utilisée tant que le *Client* n'a pas appelé le *Service ActivateSession*, fourni ses *SoftwareCertificates* et prouvé la possession de son *Certificat d'instance d'application* et de tout jeton d'identité utilisateur qu'il fournit.

Il convient qu'une application *Serveur* limite le nombre de *Sessions*. Afin d'assurer une protection contre des *Clients* mal intentionnés et des attaques de déni de service, le *Serveur* doit fermer la *Session* la plus ancienne qui n'est pas activée avant d'atteindre le nombre maximal de *Sessions* prises en charge.

Le paramètre *SoftwareCertificates* dans la réponse du *Serveur* est dévalorisé de façon à réduire la taille du message pour les Applications OPC UA aux ressources limitées. Les *SoftwareCertificates* sont fournis dans l'*AddressSpace* du *Serveur* comme défini dans l'IEC 62541-5. Un *SoftwareCertificate* identifie les capacités du *Serveur* et contient également la liste des *Profils* OPC UA pris en charge par le *Serveur*. Les *Profils* OPC UA sont définis dans l'IEC 62541-7.

Les *Certificats* supplémentaires émis par d'autres organisations peuvent être inclus pour identifier des fonctions supplémentaires du *Serveur*. Les exemples de ces *Profils* incluent la prise en charge pour des modèles spécifiques d'informations et la prise en charge pour l'accès aux types spécifiques d'appareils.

Lorsqu'une *Session* est créée, le *Serveur* ajoute une entrée pour le *Client* dans sa *Variable SessionDiagnosticsArray*. Voir l'IEC 62541-5 pour une description de cette *Variable*.

Les *Sessions* sont créées de façon à être indépendantes de la connexion de communication sous-jacente. Par conséquent, une défaillance de la connexion de communication n'affecte pas immédiatement la *Session*. Le mécanisme exact de reprise sur une erreur de la connexion de communication sous-jacente dépend du mapping de *SecureChannel* décrite dans l'IEC 62541-6.

Les *Sessions* sont automatiquement fermées par le *Serveur* si le *Client* oublie d'émettre une demande de *Service* sur la *Session* dans la période d'attente de temporisation négociée avec le *Serveur* dans la réponse à la demande de *Service CreateSession*. Le *Serveur* est ainsi protégé des défaillances du *Client* et des situations dans lesquelles la connexion sous-jacente défectueuse ne peut pas être rétablie. Les *Clients* doivent être prêts à transmettre des demandes dans le temps prévu afin d'empêcher la fermeture automatique de cette *Session*. Les *Clients* peuvent explicitement mettre fin à des sessions au moyen du service *CloseSession*.

Lorsqu'une *Session* est terminée, toutes les demandes en cours sur cette *Session* sont abandonnées et des *StatusCodes Bad_SessionClosed* sont retournés au *Client*. De plus, le *Serveur* supprime de sa *Variable SessionDiagnosticsArray* l'entrée pour le *Client* et en notifie les autres éventuels *Clients* qui s'étaient abonnés à cette entrée.

Si un *Client* invoque le service *CloseSession*, tous les *abonnements* associés à la *session* sont également supprimés si le fanion *deleteSubscriptions* est mis à TRUE. Si un *Serveur* met fin à une *Session* pour une quelconque raison, les *abonnements* associés à cette *Session* ne sont pas supprimés. Chaque *abonnement* a sa propre durée de vie pour se protéger des pertes de données en cas de fermeture d'une *Session*. Dans ces cas, l'*Abonnement* peut être réaffecté à un autre *Client* avant que sa durée de vie n'expire.

Certains *Serveurs*, comme les *Serveurs* de regroupement, jouent le rôle de *Clients* auprès d'autres *Serveurs*. Les *Serveurs* prennent généralement en charge plus d'un utilisateur système, en jouant le rôle de leurs agents auprès des *Serveurs* qu'ils représentent. La sécurité pour ces *Serveurs* est prise en charge à deux niveaux.

En premier lieu, chaque demande de *Service* OPC UA contient un paramètre de chaîne qui est utilisé pour porter un ID d'enregistrement d'audit. Un *Client* ou tout *Serveur* opérant en tant que *Client*, comme un *Serveur* de regroupement, peut créer une entrée locale de journal d'audit pour une demande qu'il soumet. Ce paramètre permet au *Client* de transmettre l'identificateur pour cette entrée avec la demande. Si le *Serveur* gère également un journal d'audits, il peut alors inclure cet ID dans l'entrée de journal d'audit qu'il écrit. Lorsque le journal est examiné et que l'entrée est trouvée, l'examineur peut la relier directement à l'entrée de journal d'audit créée par le *Client*. Cette fonction garantit la traçabilité des informations dans les journaux d'audit d'un système. Voir l'IEC TR 62541-2 pour des informations complémentaires relatives aux audits. Un *Serveur* qui gère un journal d'audit doit fournir les informations dans les entrées du journal d'audit via les *Messages* d'événements définis dans le présent document. Le *Serveur* peut choisir de fournir uniquement les informations d'*audit* via des *Messages* d'événements. L'*EventType Audit* est défini dans l'IEC 62541-3.

En second lieu, ces *Serveurs* de regroupement peuvent ouvrir des *Sessions* indépendantes sur les *Serveurs* sous-jacents pour chaque *Client* qui accède aux données auprès d'eux. La Figure 14 représente ce concept.




Figure 14 – Multiplexage d'utilisateurs sur une Session

5.6.2.2 Paramètres

Le Tableau 15 définit les paramètres du *Service*.

Tableau 15 – Paramètres du service CreateSession

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. L' <i>authenticationToken</i> est toujours Null. Le type <i>RequestHeader</i> est défini en 7.28.
clientDescription	Application Description	Informations qui décrivent l'application <i>Client</i> . Le type <i>ApplicationDescription</i> est défini en 7.1.
serverUri	String	Cette valeur est seulement spécifiée si l' <i>EndpointDescription</i> a un <i>gatewayServerUri</i> . Cette valeur est l' <i>applicationUri</i> issue de l' <i>EndpointDescription</i> qui est l' <i>applicationUri</i> pour le <i>Serveur</i> sous-jacent. Le type <i>EndpointDescription</i> est défini en 7.10.
endpointUrl	String	Adresse réseau que le <i>Client</i> a utilisée pour accéder à l' <i>Endpoint</i> de <i>Session</i> . Il convient que la partie <i>HostName</i> de l'URL soit l'un des <i>HostNames</i> pour l'application qui sont spécifiés dans l' <i>ApplicationInstanceCertificate</i> du <i>Serveur</i> (voir 7.2). Le <i>Serveur</i> doit générer un événement <i>AuditUrlMismatchEventType</i> si l'URL ne correspond pas aux <i>HostNames</i> du <i>Serveur</i> . Le type d'événement <i>AuditUrlMismatchEventType</i> est défini dans l'IEC 62541-5. Le <i>Serveur</i> utilise cette information pour des diagnostics et pour déterminer le jeu d' <i>EndpointDescriptions</i> à retourner dans la réponse.
sessionName	String	Chaîne interprétable par l'utilisateur qui identifie une <i>session</i> . Le <i>Serveur</i> rend ce nom et le <i>sessionId</i> visibles dans son <i>AddressSpace</i> à des fins de diagnostic. Il convient que le <i>Client</i> fournisse un nom qui est unique pour l'instance du <i>Client</i> . Si ce paramètre n'est pas spécifié, le <i>Serveur</i> doit attribuer une valeur.
clientNonce	ByteString	Nombre aléatoire qu'il convient de ne jamais utiliser dans aucune autre demande. Ce nombre doit avoir une longueur minimale de 32 octets. Les profils peuvent augmenter la longueur exigée. Le <i>Serveur</i> doit utiliser cette valeur pour prouver la possession de son <i>Certificat d'instance d'application</i> dans la réponse.
clientCertificate	ApplicationInstance Certificate	<i>Certificat d'instance d'application</i> transmis au <i>Client</i> . Le type <i>ApplicationInstanceCertificate</i> est défini en 7.2. Si le <i>securityPolicyUri</i> est None, le <i>Serveur</i> doit ignorer l' <i>ApplicationInstanceCertificate</i> .
Requested SessionTimeout	Duration	Nombre maximal demandé de millisecondes pendant lesquelles il convient qu'une <i>Session</i> reste ouverte sans activité. Si le <i>Client</i> oublie d'émettre une demande de <i>Service</i> dans la limite de cet intervalle, le <i>Serveur</i> doit automatiquement mettre fin à la <i>Session</i> du <i>Client</i> .
maxResponse MessageSize	UInt32	Taille maximale, en octets, pour le corps de tout message de réponse. Il convient que le <i>Serveur</i> retourne un défaut de service <i>Bad_ResponseTooLarge</i> si un message de réponse dépasse cette limite. La valeur 0 indique que ce paramètre n'est pas utilisé. Les protocoles de transport définis dans l'IEC 62541-6 peuvent impliquer des tailles de messages minimales. Le 5.3 fournit plus d'informations sur l'utilisation de ce paramètre.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition du <i>typeResponseHeader</i>).
sessionId	NodeId	<i>NodeId</i> unique affecté par le <i>Serveur</i> à une <i>Session</i> . Cet identificateur est utilisé pour accéder aux informations de diagnostic pour la <i>Session</i> dans l' <i>AddressSpace</i> du <i>Serveur</i> . Il est également utilisé dans les journaux d'audit et les éventuels événements qui rapportent des informations relatives à cette <i>Session</i> . Les informations de diagnostic d'une <i>Session</i> sont décrites dans l'IEC 62541-5. Les journaux d'audit et les événements associés sont décrits en 6.5.
authentication Token	Session AuthenticationToken	Identificateur unique affecté par le <i>Serveur</i> à la <i>Session</i> . Cet identificateur doit être transmis dans le <i>RequestHeader</i> de chaque demande et il est utilisé avec le <i>SecureChannelId</i> pour déterminer si un <i>Client</i> a accès à la <i>Session</i> . Cet identificateur ne doit pas être réutilisé d'une manière induisant le risque pour le <i>Client</i> ou le <i>Serveur</i> de le confondre avec une <i>Session</i> précédente ou existante. Le type <i>SessionAuthenticationToken</i> est décrit en 7.31.
revisedSession Temporisation	Duration	Nombre maximal effectif de millisecondes pendant lesquelles une <i>Session</i> doit rester ouverte sans activité. Il convient que le <i>Serveur</i> essaie d'honorer la demande du <i>Client</i> pour ce paramètre, mais il peut négocier cette valeur à la hausse ou à la baisse pour répondre à ses propres contraintes.
serverNonce	ByteString	Nombre aléatoire qu'il convient de ne jamais utiliser dans aucune autre demande. Ce nombre doit avoir une longueur minimale de 32 octets. Le <i>Client</i> doit utiliser cette valeur pour prouver la possession de son <i>Certificat d'instance d'application</i> dans la demande <i>ActivateSession</i> . Cette valeur peut aussi être utilisée pour prouver la possession du <i>userIdentityToken</i> qu'il a spécifié dans la demande <i>ActivateSession</i> .
serverCertificate	ApplicationInstance Certificate	<i>Certificat d'instance d'application</i> émis vers le <i>Serveur</i> . Un <i>Serveur</i> doit prouver la possession en utilisant la clé privée pour signer le <i>Nonce</i> fourni par le <i>Client</i> dans la demande. Le <i>Client</i> doit vérifier que ce <i>Certificat</i> est le même que celui qu'il a utilisé pour créer le <i>SecureChannel</i> . Le type <i>ApplicationInstanceCertificate</i> est défini en 7.2. Si le <i>securityPolicyUri</i> est NONE et qu'aucune <i>UserTokenPolicy</i> n'exige de chiffrement, le <i>Client</i> doit ignorer l' <i>ApplicationInstanceCertificate</i> .
serverEndpoints []	EndpointDescription	Liste des <i>Endpoints</i> que le <i>Serveur</i> prend en charge. Le <i>Serveur</i> doit retourner un jeu d' <i>EndpointDescriptions</i> disponibles pour le <i>serverUri</i> spécifié dans la demande. Le type <i>EndpointDescription</i> est défini en 7.10. Le <i>Client</i> doit vérifier cette liste avec la liste issue d'un <i>DiscoveryEndpoint</i> s'il a utilisé un <i>DiscoveryEndpoint</i> pour rechercher des <i>EndpointDescriptions</i> . Il est recommandé que les <i>Serveurs</i> ne comprennent que le <i>server.applicationUri</i> , l' <i>endpointUrl</i> , le <i>securityMode</i> , le <i>securityPolicyUri</i> , les <i>userIdentityTokens</i> , le <i>transportProfileUri</i> et le <i>securityLevel</i> avec tous les autres paramètres définis sur Null. Seuls les paramètres recommandés doivent être vérifiés par le <i>Client</i> .
serverSoftware Certificates []	Certificat SignedSoftware	Ce paramètre est dévalorisé et la matrice doit être vide. Les <i>SoftwareCertificates</i> sont fournis dans l' <i>AddressSpace</i> du <i>Serveur</i> comme défini dans l'IEC 62541-5.
serverSignature	SignatureData	Signature générée par la clé privée associée au <i>serverCertificate</i> . Ce paramètre est calculé en accolant le <i>clientNonce</i> au <i>clientCertificate</i> et en signant la séquence d'octets obtenue. Si le <i>clientCertificate</i> contient une chaîne, le calcul de la signature doit uniquement être effectué avec le <i>Certificat</i> feuille. Pour la rétrocomptabilité, un <i>Client</i> doit vérifier la signature avec la chaîne complète si la vérification avec le <i>Certificat</i> feuille échoue. Le <i>SignatureAlgorithm</i> doit être l' <i>AsymmetricSignatureAlgorithm</i> spécifié dans la <i>SecurityPolicy</i> pour l' <i>Endpoint</i> . Le type <i>SignatureData</i> est défini en 7.32.

Nom	Type	Description
maxRequest MessageSize	UInt32	Taille maximale, en octets, pour le corps de tout message de demande. Il convient que la <i>Pile de communication</i> du <i>Client</i> retourne une erreur <i>Bad_RequestTooLarge</i> à l'application si un message de demande dépasse cette limite. La valeur 0 indique que ce paramètre n'est pas utilisé. Voir l'IEC 62541-6 pour les valeurs par défaut ou minimales spécifiques au protocole. Le 5.3 fournit plus d'informations sur l'utilisation de ce paramètre.

5.6.2.3 Résultats de service

Le Tableau 16 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 16 – Codes de résultats du service CreateSession

ID symbolique	Description
Bad_SecureChannelIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_NonceInvalid	Voir Tableau 177 pour la description de ce code de résultat. Un <i>Serveur</i> doit vérifier la longueur minimale du nonce du <i>Client</i> et retourner ce statut si la longueur est inférieure à 32 octets. La vérification des doublons de nonces est facultative et nécessite un accès au nonce utilisé pour créer le canal sécurisé.
Bad_SecurityChecksFailed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateTimeInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerTimeInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateHostNameInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateUriInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateUseNotAllowed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerUseNotAllowed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateUntrusted	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateRevocationUnknown	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerRevocationUnknown	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateRevoked	Voir Tableau 177 pour la description de ce code de résultat.
Bad_CertificateIssuerRevoked	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManySessions	Le <i>Serveur</i> a atteint le nombre maximal de sessions.
Bad_ServerUriInvalid	Voir Tableau 177 pour la description de ce code de résultat.

5.6.3 ActivateSession

5.6.3.1 Description

Ce *Service* est utilisé par le *Client* pour spécifier l'identité de l'utilisateur associé à la *Session*. Cette demande de *Service* doit être émise par le *Client* avant qu'il n'émette une quelconque demande de *Service* autre que *CloseSession* après *CreateSession*. Si tel n'est pas le cas, le *Serveur* doit fermer la *Session*.

Chaque fois que le *Client* appelle ce *Service*, le *Client* doit prouver qu'il s'agit de la même application qui a appelé le *Service CreateSession*. Pour ce faire, le *Client* crée une signature avec la clé privée associée au *clientCertificate* spécifié dans la demande *CreateSession*.

Cette signature est créée en accolant le dernier *serverNonce* fourni par le *Serveur* au *serverCertificate* et en calculant la signature de la séquence d'octets obtenue.

Une fois qu'il a été utilisé, un *serverNonce* ne peut plus être réutilisé par la suite. C'est pourquoi le *Serveur* retourne un nouveau *serverNonce* chaque fois que le service *ActivateSession* est appelé.

Lorsque le Service *ActivateSession* est appelé la première fois, le *Serveur* doit rejeter la demande si le *SecureChannel* n'est pas le même que celui qui est associé à la demande *CreateSession*. Les appels consécutifs à *ActivateSession* peuvent être associés à des *SecureChannels* différents. Si tel est le cas, le *Serveur* doit vérifier que le *Certificat* utilisé par le *Client* pour créer le nouveau *SecureChannel* est le même que le *Certificat* utilisé pour créer le *SecureChannel* d'origine. En outre, le *Serveur* doit vérifier que le *Client* a fourni un *UserIdentityToken* identique au jeton réellement associé à la *Session*. Lorsque le *Serveur* a accepté le nouveau *SecureChannel*, il doit rejeter les demandes envoyées par l'ancien *SecureChannel*.

Le service *ActivateSession* est utilisé pour associer une identité d'utilisateur à une *Session*. Lorsqu'un *Client* fournit une identité d'utilisateur, il doit fournir la preuve qu'il est habilité à utiliser l'identité d'utilisateur en question. Le mécanisme exact utilisé pour fournir cette preuve dépend du type de *UserIdentityToken*. Si le jeton est un *UserNameIdentityToken*, la preuve est le *password* qui a été inclus dans le jeton. Si le jeton est un *X509IdentityToken*, la preuve est une signature générée avec la clé privée associée au *Certificat*. Les données à signer sont créées en accolant le dernier *serverNonce* au *serverCertificate* spécifié dans la réponse *CreateSession*. Si un jeton inclut un secret, il convient de le déchiffrer à l'aide de la clé publique contenue dans le *serverCertificate*.

Les *Serveurs* doivent prendre les mesures de protection appropriées contre les attaques à l'encontre des jetons d'identité d'utilisateur. L'hypothèse d'une telle attaque est retenue en cas de tentatives de connexion répétées avec des jetons d'identité d'utilisateur invalides. Une option consiste à bloquer un *Client* OPC UA pendant un certain temps, si la validation du jeton d'identité d'utilisateur échoue à plusieurs reprises. Le *Client* OPC UA est détecté par son adresse IP pour les connexions non sécurisées ou par l'*ApplicationInstanceUri* pour les connexions sécurisées. Une autre option consiste à retarder la réponse du Service lorsque la validation d'une identité d'utilisateur échoue. Ce temps de retard peut être supérieur en cas de défaillances répétées. Des défaillances sporadiques ne doivent pas retarder les connexions avec des jetons valides.

Les *Clients* peuvent modifier l'identité d'un utilisateur associé à une *Session* en appelant le Service *ActivateSession*. Le *Serveur* valide les signatures fournies avec la demande puis valide la nouvelle identité d'utilisateur. En l'absence d'erreurs, le *Serveur* remplace l'identité d'utilisateur pour la *Session*. La modification de l'identité utilisateur pour une *Session* peut induire des discontinuités des *Abonnements* actifs, car il peut être nécessaire que le *Serveur* casse des connexions à un système sous-jacent et les rétablir en utilisant les nouveaux authentifiants.

Lorsqu'un *Client* fournit une liste d'identificateurs de paramètres de lieu dans la demande, chacun d'eux doit nécessairement intégrer la composante langue. Il peut contenir la composante <pays/région> (facultative). Lorsque le *Serveur* retourne un *LocalizedText* dans le contexte d'une *Session*, il peut aussi retourner les deux composantes langue et pays/région ou juste la composante langue comme la valeur par défaut de son identificateur de paramètre de lieu.

Lorsqu'un *Serveur* retourne une chaîne au *Client*, il détermine d'abord si des traductions sont disponibles. Si elles le sont, le *Serveur* retourne la chaîne dont le paramètre de lieu correspond exactement à celui ayant la priorité la plus haute dans la liste fournie par le *Client*.

S'il n'existe pas de mappings exacts, le *Serveur* ignore alors la composante <pays/région> de l'identificateur de paramètre de lieu et retourne la chaîne dont la composante <langue> correspond à la composante <langue> de l'identificateur de paramètre de lieu ayant la priorité la plus haute dans la liste fournie par le *Client*.

S'il n'y a pas de mappings, le *Serveur* retourne alors la chaîne qu'il a avec l'identificateur de paramètre de lieu.

Un *Serveur passerelle* est supposé se substituer à l'utilisateur donné par le *Client* lorsqu'il se connecte au *Serveur* sous-jacent. Cela signifie qu'il doit recalculer les signatures sur le *UserIdentityToken* en utilisant le nonce fourni par le *Serveur* sous-jacent. Il sera nécessaire que le *Serveur passerelle* utilise ses propres authentifiants d'utilisateur, si le *UserIdentityToken* fourni par le *Client* ne prend pas en charge la substitution d'identité.

5.6.3.2 Paramètres

Le Tableau 17 définit les paramètres du *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 REV

Tableau 17 – Paramètres du service ActivateSession

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs. Le type <i>RequestHeader</i> est défini en 7.28.
clientSignature	SignatureData	Signature générée par la clé privée associée au <i>clientCertificate</i> . Ce paramètre est calculé en accolant le <i>serverNonce</i> au <i>serverCertificate</i> et en signant la séquence d'octets obtenue. Si le <i>serverCertificate</i> contient une chaîne, le calcul de la signature doit uniquement être effectué avec le <i>Certificat</i> feuille. Pour la rétrocomptabilité, un <i>Serveur</i> doit vérifier la signature avec la chaîne complète si la vérification avec le <i>Certificat</i> feuille échoue. Le <i>SignatureAlgorithm</i> doit être l' <i>AsymmetricSignatureAlgorithm</i> spécifié dans la <i>SecurityPolicy</i> pour l' <i>Endpoint</i> . Le type <i>SignatureData</i> est défini en 7.32.
clientSoftwareCertificates []	SignedSoftwareCertificate	Réservé pour utilisation ultérieure. Le type <i>SignedSoftwareCertificate</i> est défini en 7.33.
localeIds []	LocaleId	Liste des identificateurs de paramètres de lieu dans l'ordre des priorités pour les chaînes localisées. Le premier <i>LocaleId</i> de la liste a la priorité la plus élevée. Si le <i>Serveur</i> retourne une chaîne localisée au <i>Client</i> , le <i>Serveur</i> doit retourner la traduction avec la plus haute priorité qu'il peut. S'il n'a de traduction pour aucun des paramètres de lieu identifiés dans cette liste, il doit retourner la valeur de chaîne qu'il a et inclure l'identificateur de paramètre de lieu avec la chaîne. Voir l'IEC 62541-3 pour plus d'informations sur les identificateurs de paramètres de lieu. Si le <i>Client</i> ne réussit pas à spécifier au moins un identificateur de paramètre de lieu, le <i>Serveur</i> doit utiliser n'importe lequel en sa possession. Il n'est nécessaire de spécifier ce paramètre que pendant le premier appel au service <i>ActivateSession</i> au cours d'une seule <i>session</i> d'application. S'il n'est pas spécifié, le <i>Serveur</i> doit continuer d'utiliser les <i>localeIds</i> courants pour la <i>session</i> .
userIdentityToken	Paramètre Extensible UserIdentityToken	Authentifiants de l'utilisateur associé à l'application <i>Client</i> . Le <i>Serveur</i> utilise ces authentifiants pour déterminer s'il convient d'autoriser le <i>Client</i> à activer une <i>Session</i> et pour déterminer également les ressources auxquelles le <i>Client</i> a accès au cours de la <i>Session</i> . Le <i>UserIdentityToken</i> est un type de paramètre extensible défini en 7.36. L' <i>EndpointDescription</i> spécifie les <i>UserIdentityTokens</i> que le <i>Serveur</i> doit accepter. Les jetons d'utilisateur Null ou vides doivent toujours être interprétés comme anonymes.
userTokenSignature	SignatureData	Si le <i>Client</i> a spécifié un jeton d'identité d'utilisateur qui prend en charge les signatures numériques, il doit créer une signature et la transmettre comme ce paramètre. Dans le cas contraire, le paramètre est Null. Le <i>SignatureAlgorithm</i> dépend du type de jeton d'identité. Le type <i>SignatureData</i> est défini en 7.32.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
serverNonce	ByteString	Nombre aléatoire qu'il convient de ne jamais utiliser dans aucune autre demande. Ce nombre doit avoir une longueur minimale de 32 octets. Le <i>Client</i> doit utiliser cette valeur pour prouver la possession de son <i>Certificat d'instance d'application</i> lors de la prochaine demande <i>ActivateSession</i> .
results []	StatusCode	Liste de résultats de validation pour les <i>SoftwareCertificates</i> (voir 7.34 pour la définition de <i>StatusCode</i>).
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic associées à des erreurs de validation de <i>SoftwareCertificate</i> (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.6.3.3 Résultats de service

Le Tableau 18 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 18 – Codes de résultats du service ActivateSession

ID symbolique	Description
Bad_IdentityTokenInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_IdentityTokenRejected	Voir Tableau 177 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ApplicationSignatureInvalid	La signature fournie par l'application <i>Client</i> est absente ou non valide.
Bad_UserSignatureInvalid	La signature du jeton d'utilisateur est absente ou non valide.
Bad_NoValidCertificates	Le <i>Client</i> n'a pas fourni au moins un <i>SoftwareCertificate</i> qui est valide et répond aux exigences relatives aux profils pour le <i>Serveur</i> .
Bad_IdentityChangeNotSupported	Le <i>Serveur</i> ne prend pas en charge la modification de l'identité d'utilisateur affectée à la session.

5.6.4 CloseSession

5.6.4.1 Description

Ce *Service* est utilisé pour mettre fin à une *Session*. Le *Serveur* entreprend les actions suivantes lorsqu'il reçoit une demande *CloseSession*:

- Il arrête d'accepter des demandes pour cette *Session*. Toutes les demandes ultérieures reçues pour la *Session* sont rejetées.
- Il retourne des réponses négatives avec le *StatusCode* *Bad_SessionClosed* à toutes les demandes qui sont actuellement en cours pour permettre le retour en temps utile de la réponse *CloseSession*. Les *Clients* sont vivement encouragés à attendre que toutes les demandes en cours se terminent avant de présenter une demande *CloseSession*.
- Il retire l'entrée pour le *Client* dans sa *Variable SessionDiagnosticsArray*.

Lorsque le *Service CloseSession* est appelé avant que la *Session* ne soit activée avec succès, le *Serveur* doit rejeter la demande si le *SecureChannel* n'est pas le même que celui associé à la demande *CreateSession*.

5.6.4.2 Paramètres

Le Tableau 19 définit les paramètres du *Service*.

Tableau 19 – Paramètres du service CloseSession

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
deleteSubscriptions	Boolean	Si la valeur est TRUE, le <i>Serveur</i> supprime tous les Abonnements associés à la <i>Session</i> . Si la valeur est FALSE, le <i>Serveur</i> conserve les Abonnements associés à la <i>Session</i> jusqu'à ce que leur durée de vie expire.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).

5.6.4.3 Résultats de service

Le Tableau 20 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 20 – Codes de résultats du service CloseSession

ID symbolique	Description
Bad_SessionIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.

5.6.5 Cancel

5.6.5.1 Description

Ce *Service* est utilisé pour annuler des demandes de *Services* en cours. Les demandes de services qui ont été annulées avec succès doivent répondre avec *Bad_RequestCancelledByClient*.

5.6.5.2 Paramètres

Le Tableau 21 définit les paramètres du *Service*.

Tableau 21 – Paramètres du service Cancel

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
requestHandle	IntegerId	Le <i>requestHandle</i> affecté à une ou plusieurs demandes qu'il convient d'annuler. Toutes les demandes en cours avec le <i>requestHandle</i> correspondant doivent être annulées.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
cancelCount	UInt32	Nombre de demandes annulées.

5.6.5.3 Résultats de service

Les *StatusCodes* communs sont définis dans le Tableau 177.

5.7 Jeu de services NodeManagement

5.7.1 Vue d'ensemble

Ce *Jeu de services* définit les *Services* pour ajouter et supprimer des *Nœuds* d'*AddressSpace* et les *Références* entre eux. Tous les *Nœuds* ajoutés continuent d'exister dans l'*AddressSpace* même si le *Client* qui les a créés se déconnecte du *Serveur*.

5.7.2 AddNodes

5.7.2.1 Description

Ce *Service* est utilisé pour ajouter un ou plusieurs *Nœuds* dans la hiérarchie d'*AddressSpace*. En utilisant ce *Service*, chaque *Nœud* est ajouté comme étant le *TargetNode* d'une

HierarchicalReference pour assurer que l'*AddressSpace* est pleinement connecté et que le *Nœud* est ajouté comme enfant dans la hiérarchie d'*AddressSpace* (voir IEC 62541-3).

Lorsqu'un *Serveur* crée une instance de *TypeDefinitionNode*, il doit créer la même hiérarchie de *Nœuds* sous le nouvel *Objet* ou la nouvelle *Variable*, en fonction de la *ModellingRule* de chaque *InstanceDeclaration*. Tous les *Nœuds* d'une *ModellingRule* obligatoire doivent être créés ou un *Nœud* existant doit être référencé afin de se conformer à l'*InstanceDeclaration*. La création de *Nœuds* avec d'autres *ModellingRules* est spécifique au *Serveur*.

5.7.2.2 Paramètres

Le Tableau 22 définit les paramètres du *Service*.

Tableau 22 – Paramètres du service AddNodes

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
nodesToAdd []	AddNodesItem	Liste des <i>Nœuds</i> à ajouter. Tous les <i>Nœuds</i> sont ajoutés comme <i>Référence</i> à un <i>Nœud</i> existant à l'aide d'un <i>ReferenceType</i> hiérarchique. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
parentNodeid	Expanded Nodeid	<i>ExpandedNodeid</i> du <i>Nœud</i> parent pour la <i>Référence</i> . Le type <i>ExpandedNodeid</i> est défini en 7.11.
referenceTypeid	Nodeid	<i>Nodeid</i> du <i>ReferenceType</i> hiérarchique à utiliser pour la <i>Référence</i> allant du <i>Nœud</i> parent au nouveau <i>Nœud</i> .
requestedNewNodeid	Expanded Nodeid	Le <i>Client</i> a demandé le <i>Nodeid</i> étendu du <i>Nœud</i> à ajouter. Le <i>serverIndex</i> dans le <i>Nodeid</i> étendu doit être 0. Si le <i>Serveur</i> ne peut pas utiliser ce <i>Nodeid</i> , il rejette ce <i>Nœud</i> et renvoie un code d'erreur approprié. Si le <i>Client</i> ne souhaite pas demander un <i>Nodeid</i> , il met la valeur de ce paramètre au <i>Nodeid</i> étendu Null. Si le <i>Nœud</i> à ajouter est un <i>Nœud</i> de <i>ReferenceType</i> , il convient que son <i>Nodeid</i> soit un id numérique. Voir l'IEC 62541-3 pour une description des <i>Nodeids</i> de <i>ReferenceType</i> .
browseName	QualifiedName	Le nom de navigation du <i>Nœud</i> à ajouter.
nodeClass	NodeClass	<i>NodeClass</i> du <i>Nœud</i> à ajouter.
nodeAttributes	Paramètre Extensible NodeAttributes	Les <i>Attributs</i> qui sont spécifiques à la <i>NodeClass</i> . Le type de paramètre <i>NodeAttributes</i> est un type de paramètre extensible spécifié en 7.19. Un <i>Client</i> peut exclure les valeurs pour une partie ou l'ensemble des <i>Attributs</i> . Si une valeur d' <i>attribut</i> est Null, le <i>Serveur</i> doit utiliser les valeurs par défaut provenant du <i>TypeDefinitionNode</i> . Si un <i>TypeDefinitionNode</i> n'a pas été fourni, le <i>Serveur</i> doit choisir une valeur par défaut appropriée. Le <i>Serveur</i> peut encore ajouter un <i>Attribut</i> facultatif au <i>Nœud</i> avec une valeur par défaut adaptée, même si le <i>Client</i> ne spécifie pas une valeur.
typeDefinition	Expanded Nodeid	<i>Nodeid</i> du <i>TypeDefinitionNode</i> pour le <i>Nœud</i> à ajouter. Ce paramètre doit être Null pour toutes les <i>NodeClasses</i> autres qu' <i>Objet</i> et <i>Variable</i> , auquel cas il doit être fourni.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	AddNodesResult	Liste des résultats pour les <i>Nœuds</i> à ajouter. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToAdd</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> pour le <i>Nœud</i> à ajouter (voir 7.34 pour la définition de <i>StatusCode</i>).
addedNodeId	NodeId	<i>NodeId</i> affecté par le <i>Serveur</i> au <i>Nœud</i> ajouté. <i>NodeId</i> Null si l'opération a échoué.
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>Nœuds</i> à ajouter (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToAdd</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.7.2.3 Résultats de service

Le Tableau 23 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 23 – Codes de résultats du service AddNodes

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.7.2.4 StatusCodes

Le Tableau 24 définit les valeurs de niveau opération pour le paramètre *statusCode* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 24– Codes de résultats de niveau opération AddNodes

ID symbolique	Description
Bad_ParentNodeIdInvalid	L'identificateur de <i>Nœud</i> parent ne renvoie pas à un <i>Nœud</i> valide.
Bad_ReferenceTypeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ReferenceNotAllowed	La référence n'a pas pu être créée, car elle enfreint les contraintes imposées par le modèle de données.
Bad_NodeIdRejected	L'identificateur de nœud demandé a été rejeté soit parce qu'il était invalide, soit parce que le <i>Serveur</i> n'autorise pas la spécification des identificateurs de nœuds par le <i>Client</i> .
Bad_NodeIdExists	L'identificateur de nœuds demandé est déjà utilisé par un autre nœud.
Bad_NodeClassInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_BrowseNameInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_BrowseNameDuplicated	Le nom de navigation n'est pas unique parmi les nœuds qui partagent la même relation avec le parent.
Bad_NodeAttributesInvalid	Les <i>Attributs</i> de nœuds ne sont pas valides pour la classe de nœuds.
Bad_TypeDefinitionInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.

5.7.3 AddReferences

5.7.3.1 Description

Ce *Service* est utilisé pour ajouter une ou plusieurs *Références* à un ou plusieurs *Nœuds*. La *NodeClass* est un paramètre d'entrée qui est utilisé pour valider le fait que la *Référence* à ajouter corresponde à la *NodeClass* du *TargetNode*. Ce paramètre n'est pas validé si la *Référence* renvoie à un *TargetNode* dans un *Serveur* distant.

Dans certains cas, l'ajout de nouvelles *Références* à l'*AddressSpace* doit exiger que le *Serveur* ajoute de nouveaux identificateurs de *Serveur* à la *Variable ServerArray* du *Serveur*. C'est pourquoi les *Serveurs* distants sont identifiés par leur URI et non par leur indice de *ServerArray*. Le *Serveur* peut ainsi ajouter les URI de *Serveurs* distants à sa *ServerArray*.

5.7.3.2 Paramètres

Le Tableau 25 définit les paramètres du *Service*.

Tableau 25 – Paramètres du service AddReferences

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
referencesToAdd []	AddReferences Item	Liste d'instances <i>Référence</i> à ajouter au <i>SourceNode</i> . La <i>targetNodeClass</i> de chaque <i>Référence</i> dans la liste doit correspondre à la <i>NodeClass</i> du <i>TargetNode</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
sourceNodeid	Nodeid	<i>Nodeid</i> du <i>Nœud</i> auquel la <i>Référence</i> doit être ajoutée. Le <i>Nœud</i> source doit toujours exister dans le <i>Serveur</i> pour ajouter la <i>Référence</i> . Le paramètre <i>isForward</i> peut être mis à FALSE si le <i>Nœud</i> cible est sur le <i>Serveur</i> local et le <i>Nœud</i> source sur le <i>Serveur</i> distant.
referenceTypeid	Nodeid	<i>Nodeid</i> du <i>ReferenceType</i> qui définit la <i>Référence</i> .
isForward	Boolean	Si la valeur est TRUE, le <i>Serveur</i> crée une <i>Référence</i> directe. Si la valeur est FALSE, le <i>Serveur</i> crée une <i>Référence</i> inverse.
targetServerUri	String	URI du <i>Serveur</i> distant. Si ce paramètre n'est pas Null, il se substitue au <i>serverIndex</i> dans le <i>targetNodeid</i> .
targetNodeid	Expanded Nodeid	<i>Nodeid</i> étendu du <i>TargetNode</i> . Le type <i>ExpandedNodeid</i> est défini en 7.11.
targetNodeClass	NodeClass	<i>NodeClass</i> du <i>TargetNode</i> . Le <i>Client</i> doit spécifier cela, car le <i>TargetNode</i> peut ne pas être directement accessible pour le <i>Serveur</i> .
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	StatusCode	Liste des <i>StatusCodes</i> pour les <i>Références</i> à ajouter (voir 7.34 pour la définition des <i>StatusCodes</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>referencesToAdd</i> .
diagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic pour les <i>Références</i> à ajouter (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>referencesToAdd</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.7.3.3 Résultats de service

Le Tableau 26 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 26 – Codes de résultats du service AddReferences

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.7.3.4 StatusCodes

Le Tableau 27 définit les valeurs pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 27 – Codes de résultats de niveau opération AddReferences

ID symbolique	Description
Bad_SourceNodeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ReferenceTypeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ServerUriInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TargetNodeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeClassInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ReferenceNotAllowed	La référence n'a pas pu être créée, car elle enfreint les contraintes imposées par le modèle de données à ce <i>Serveur</i> .
Bad_ReferenceLocalOnly	Le type de référence n'est pas valide pour une référence à un <i>Serveur</i> distant.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_DuplicateReferenceNotAllowed	Le type de référence entre les nœuds est déjà défini.
Bad_InvalidSelfReference	Le <i>Serveur</i> n'autorise pas ce type d'autoréférence sur ce nœud.

5.7.4 DeleteNodes

5.7.4.1 Description

Ce *Service* est utilisé pour supprimer un ou plusieurs *Nœuds* de l'*AddressSpace*.

Lorsque l'un quelconque des *Nœuds* supprimés par une invocation de ce *Service* est le *TargetNode* d'une *Référence*, ces *Références* sont laissées non résolues sur la base du paramètre *deleteTargetReferences*.

Lorsque l'un quelconque des *Nœuds* supprimés par une invocation de ce *Service* est sous surveillance, une *Notification* contenant le code de statut *Bad_NodeIdUnknown* est envoyée au *Client* sous surveillance indiquant que le *Nœud* a été supprimé.

5.7.4.2 Paramètres

Le Tableau 28 définit les paramètres du *Service*.

Tableau 28 – Paramètres du service DeleteNodes

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
nodesToDelete []	DeleteNodes Item	Liste des <i>Nœuds</i> à supprimer. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
nodeId	NodeId	<i>NodeId</i> du <i>Nœud</i> à supprimer.
deleteTargetReferences	Boolean	Paramètre <i>booléen</i> ayant les valeurs suivantes: TRUE supprimer les <i>Références</i> dans les <i>TargetNodes</i> qui <i>réfèrent</i> le <i>Nœud</i> à supprimer. FALSE supprimer seulement les <i>Références</i> dont le <i>Nœud</i> à supprimer est la source. Le <i>Serveur</i> ne peut pas garantir qu'il est capable de supprimer toutes les <i>Références</i> des <i>TargetNodes</i> si ce paramètre est TRUE.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	StatusCode	Liste des <i>StatusCodes</i> pour les <i>Nœuds</i> à supprimer (voir 7.34 pour la définition des <i>StatusCode</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre de la liste du paramètre de demande <i>nodesToDelete</i> .
diagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic pour les <i>Nœuds</i> à supprimer (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToDelete</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.7.4.3 Résultats de service

Le Tableau 29 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 29 – Codes de résultats du service DeleteNodes

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.7.4.4 StatusCodes

Le Tableau 30 définit les valeurs pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 30 – Codes de résultat de niveau opération DeleteNodes

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_NoDeleteRights	Voir Tableau 178 pour la description de ce code de résultat.
Uncertain_ReferenceNotDeleted	Le <i>Serveur</i> n'était pas capable de supprimer toutes les références cibles.

5.7.5 DeleteReferences

5.7.5.1 Description

Ce *Service* est utilisé pour supprimer une ou plusieurs *Références* d'un *Nœud*.

Lorsqu'un nombre quelconque de *Références* supprimées par une invocation de ce *Service* est contenu dans une *Vue*, la *Propriété ViewVersion* est mise à jour si cette *Propriété* est prise en charge.

Le Tableau 31 définit les paramètres du *Service*.

IECNORM.COM : Click to view the full PDF of IEC 62541-4:2020 RLV

Tableau 31 – Paramètres du service DeleteReferences

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
referencesToDelete []	DeleteReferences Item	Liste des <i>Références</i> à supprimer. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
sourceNodeid	Nodeid	<i>Nodeid</i> du <i>Nœud</i> qui contient la <i>Référence</i> à supprimer.
referenceTypeid	Nodeid	<i>Nodeid</i> du <i>ReferenceType</i> qui définit la <i>Référence</i> à supprimer.
isForward	Boolean	Si la valeur est TRUE, le Serveur supprime une <i>Référence</i> directe. Si la valeur est FALSE, le Serveur supprime une <i>Référence</i> inverse.
targetNodeid	ExpandedNodeid	<i>Nodeid</i> du <i>TargetNode</i> de la <i>Référence</i> . Si l'indice du <i>Serveur</i> indique que le <i>TargetNode</i> est un <i>Nœud</i> distant, le <i>nodeid</i> doit contenir l'URI d'espace de noms absolu. Si le <i>TargetNode</i> est un <i>Nœud</i> local, le <i>nodeid</i> doit contenir l'indice d'espace de noms.
deleteBidirectional	Boolean	Paramètre <i>booléen</i> ayant les valeurs suivantes: TRUE supprimer la <i>Référence</i> spécifiée et la <i>Référence</i> opposée partant du <i>TargetNode</i> . Si le <i>TargetNode</i> réside sur un <i>Serveur</i> distant, le <i>Serveur</i> peut supprimer la <i>Référence</i> spécifiée uniquement. FALSE supprimer la <i>Référence</i> spécifiée seulement.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	StatusCode	Liste des <i>StatusCodes</i> pour les <i>Références</i> à supprimer (voir 7.34 pour la définition des <i>StatusCode</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>referencesToDelete</i> .
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>Références</i> à supprimer (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>referencesToDelete</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.7.5.2 Résultats de service

Le Tableau 32 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 32 – Codes de résultats du service DeleteReferences

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.7.5.3 StatusCodes

Le Tableau 33 définit les valeurs pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 33 – Codes de résultats de niveau opération DeleteReferences

ID symbolique	Description
Bad_SourceNodeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ReferenceTypeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ServerIndexInvalid	L'indice de <i>Serveur</i> n'est pas valide.
Bad_TargetNodeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_NoDeleteRights	Voir Tableau 178 pour la description de ce code de résultat.

5.8 Jeu de services View

5.8.1 Vue d'ensemble

Les *Clients* utilisent les *Services* navigation du *Jeu de services View* pour naviguer à travers l'*AddressSpace* ou à travers une *Vue* qui est un sous-ensemble de l'*AddressSpace*.

Une *Vue* est un sous-ensemble de l'*AddressSpace* créé par le *Serveur*. Les versions ultérieures de l'IEC 62541-4 peuvent aussi définir des services pour créer des *vues* définies par le *Client*. Voir l'IEC 62541-5 pour une description de l'organisation des vues dans l'*AddressSpace*.

5.8.2 Browse

5.8.2.1 Description

Ce *Service* est utilisé pour découvrir les *Références* d'un *Nœud* spécifié. La navigation peut en outre être limitée par l'utilisation d'une *Vue*. Ce *service Browse* prend aussi en charge une fonction de filtrage primitive.

Dans certains cas, le traitement de tous les nœuds à parcourir peut prendre plus de temps que la suggestion de temporisation du *Client* pour traiter tous les nœuds à parcourir. Dans ce cas, le *Serveur* peut retourner des résultats nuls avec un point de continuation pour les nœuds affectés avant que la temporisation n'expire.

5.8.2.2 Paramètres

Le Tableau 34 définit les paramètres du *Service*.

Tableau 34 – Paramètres du service Browse

Nom	Type	Description																		
Demande																				
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).																		
Vue	ViewDescription	Description de la <i>Vue</i> à parcourir (voir 7.39 pour la définition de <i>ViewDescription</i>). Une valeur <i>ViewDescription</i> vide représente l' <i>AddressSpace</i> dans son intégralité. L'utilisation de la valeur vide de <i>ViewDescription</i> implique que toutes les <i>Références</i> du <i>nodesToBrowse</i> sont retournées. L'utilisation d'une autre <i>Vue</i> implique que seules les <i>Références</i> du <i>nodesToBrowse</i> qui sont définies pour cette <i>Vue</i> sont retournées.																		
requestedMax ReferencesPerNode	Counter	Nombre maximal de références à retourner pour chaque Nœud de départ spécifié dans la demande. La valeur 0 indique que le <i>Client</i> n'impose aucune limitation (voir 7.5 pour la définition de <i>Counter</i>).																		
nodesToBrowse []	BrowseDescription	Liste de nœuds à parcourir. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.																		
nodeId	Nodeld	<i>Nodeld</i> du Nœud à parcourir. Si une <i>Vue</i> est fournie, elle doit inclure ce Nœud.																		
browseDirection	Enum BrowseDirection	Enumération qui spécifie le sens des <i>Références</i> à suivre. Ce paramètre a les valeurs suivantes: FORWARD_0 Choisir des <i>Références</i> directes seulement. INVERSE_1 Choisir des <i>Références</i> inverses seulement. BOTH_2 Choisir des <i>Références</i> directes et inverses. Les <i>Références</i> retournées indiquent effectivement le sens que le <i>Serveur</i> a suivi dans le paramètre <i>isForward</i> de la <i>ReferenceDescription</i> . Les <i>Références</i> symétriques sont toujours considérées être dans le sens direct; par conséquent, le fanion <i>isForward</i> est toujours mis à TRUE et les <i>Références</i> symétriques ne sont pas retournées si <i>browseDirection</i> est défini sur <i>INVERSE_1</i> .																		
referenceTypeId	Nodeld	Spécifie le <i>Nodeld</i> du <i>ReferenceType</i> à suivre. Seules des instances de ce <i>ReferenceType</i> ou de ses sous-types sont retournées. S'il n'est pas spécifié, toutes les <i>Références</i> sont retournées et <i>includeSubtypes</i> est ignoré.																		
includeSubtypes	Boolean	Indique s'il convient d'inclure dans la navigation les sous-types du <i>ReferenceType</i> . Si TRUE, les instances de <i>referenceTypeId</i> et de tous ses sous-types sont retournées.																		
nodeClassMask	UInt32	Spécifie les <i>NodeClasses</i> des <i>TargetNodes</i> . Seuls les <i>TargetNodes</i> avec les <i>NodeClasses</i> choisis sont retournés. Il est affecté aux <i>NodeClasses</i> les bits suivants: <table><tr><th>Bit</th><th>NodeClass</th></tr><tr><td>0</td><td>Objet</td></tr><tr><td>1</td><td>Variable</td></tr><tr><td>2</td><td>Méthode</td></tr><tr><td>3</td><td>ObjectType</td></tr><tr><td>4</td><td>VariableType</td></tr><tr><td>5</td><td>ReferenceType</td></tr><tr><td>6</td><td>DataType</td></tr><tr><td>7</td><td>Vue</td></tr></table> S'il est mis à zéro, toutes les <i>NodeClasses</i> sont retournées. Si le <i>NodeClass</i> n'est pas connu pour un Nœud distant, le <i>nodeClassMask</i> est ignoré.	Bit	NodeClass	0	Objet	1	Variable	2	Méthode	3	ObjectType	4	VariableType	5	ReferenceType	6	DataType	7	Vue
Bit	NodeClass																			
0	Objet																			
1	Variable																			
2	Méthode																			
3	ObjectType																			
4	VariableType																			
5	ReferenceType																			
6	DataType																			
7	Vue																			

Nom	Type	Description														
resultMask	UInt32	Spécifie les champs dans la structure <i>ReferenceDescription</i> qu'il convient de retourner. Il est affecté aux champs les bits suivants: <table><tr><th>Bit</th><th>Résultat</th></tr><tr><td>0</td><td>ReferenceType</td></tr><tr><td>1</td><td>IsForward</td></tr><tr><td>2</td><td>NodeClass</td></tr><tr><td>3</td><td>BrowseName</td></tr><tr><td>4</td><td>DisplayName</td></tr><tr><td>5</td><td>TypeDefinition</td></tr></table> Le type <i>ReferenceDescription</i> est défini en 7.25.	Bit	Résultat	0	ReferenceType	1	IsForward	2	NodeClass	3	BrowseName	4	DisplayName	5	TypeDefinition
Bit	Résultat															
0	ReferenceType															
1	IsForward															
2	NodeClass															
3	BrowseName															
4	DisplayName															
5	TypeDefinition															
Réponse																
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).														
results []	BrowseResult	Liste de <i>BrowseResults</i> La taille et l'ordre de la liste correspondent à la taille et à l'ordre du <i>nodesToBrowse</i> spécifié dans la demande. Le type <i>BrowseResult</i> est défini en 7.3.														
diagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic pour les <i>résultats</i> (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de réponse <i>results</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.														

5.8.2.3 Résultats de service

Le Tableau 35 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 35 – Codes de résultats du service Browse

ID symbolique	Description
Bad_ViewIdUnknown	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ViewTimestampInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ViewParameterMismatchInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ViewVersionInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.8.2.4 StatusCodes

Le Tableau 36 définit les valeurs pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 36 – Codes de résultats de niveau opération Browse

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ReferenceTypeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_BrowseDirectionInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeNotInView	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NoContinuationPoints	Voir Tableau 178 pour la description de ce code de résultat.
Uncertain_NotAllNodesAvailable	Les résultats de la navigation peuvent être incomplets en raison de l'indisponibilité d'un sous-système.

5.8.3 BrowseNext

5.8.3.1 Description

Ce service est utilisé pour demander le prochain jeu d'informations de réponse *Browse* ou *BrowseNext* qui sont trop volumineuses pour être envoyées dans une seule réponse. "Trop volumineuses" signifie dans ce contexte que le *Serveur* n'est pas capable de retourner une réponse plus volumineuse ou que le nombre de résultats à retourner dépasse le nombre maximal de résultats à retourner qui a été spécifié par le *Client* dans la demande *Browse* d'origine. Le *BrowseNext* doit être soumis dans la même *session* qui a été utilisée pour soumettre la demande *Browse* ou *BrowseNext* qui se poursuit.

5.8.3.2 Paramètres

Le Tableau 37 définit les paramètres du Service.

Tableau 37 – Paramètres du service BrowseNext

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
releaseContinuationPoints	Boolean	Paramètre <i>booléen</i> ayant les valeurs suivantes: TRUE les <i>continuationPoints</i> transmis doivent être réinitialisés pour libérer des ressources sur le <i>Serveur</i>. Les points de continuation sont libérés et les résultats et les matrices de <i>diagnosticInfos</i> sont vides. FALSE les <i>continuationPoints</i> transmis doivent être utilisés pour récupérer le prochain jeu d'informations de navigation. Un <i>Client</i> doit toujours utiliser le point de continuation retourné par une réponse <i>Browse</i> ou <i>BrowseNext</i> pour libérer les ressources pour le point de continuation dans le <i>Serveur</i>. Si le <i>Client</i> ne souhaite pas récupérer le prochain jeu d'informations de navigation, <i>BrowseNext</i> doit être appelé avec ce paramètre mis à TRUE.
continuationPoints []	Continuation Point	Liste de valeurs opaques définies par le <i>Serveur</i> qui représentent des points de continuation. La valeur pour un point de continuation a été retournée au <i>Client</i> dans une réponse <i>Browse</i> ou <i>BrowseNext</i> précédente. Ces valeurs sont utilisées pour identifier la demande <i>Browse</i> ou <i>BrowseNext</i> précédemment traitée qui se poursuit et le point dans le jeu de résultats à partir duquel la réponse de navigation doit se poursuivre. Les <i>Clients</i> peuvent mélanger les points de continuation provenant de réponses <i>Browse</i> ou <i>BrowseNext</i> différentes. Le type <i>ContinuationPoint</i> est décrit en 7.6.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	BrowseResult	Liste de références qui satisfont aux critères spécifiés dans la demande <i>Browse</i> d'origine. La taille et l'ordre de cette liste correspondent à la taille et à l'ordre du paramètre de demande <i>continuationPoints</i>. Le type <i>BrowseResult</i> est défini en 7.3.
diagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic pour les <i>résultats</i> (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de réponse <i>results</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.8.3.3 Résultats de service

Le Tableau 38 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 38 – Codes de résultats du service BrowseNext

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.8.3.4 StatusCodes

Le Tableau 39 définit les valeurs pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 39 – Codes de résultats de niveau opération BrowseNext

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ReferenceTypeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_BrowseDirectionInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeNotInView	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ContinuationPointInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.8.4 TranslateBrowsePathsToNodeIds

5.8.4.1 Description

Ce *Service* est utilisé pour demander au *Serveur* de traduire un ou plusieurs chemins de navigation en des *NodeIds*. Chaque chemin de navigation est composé d'un *Nœud* de départ et d'un *RelativePath*. Le *Nœud* de départ spécifié identifie le *Nœud* sur lequel repose le *RelativePath*. Le *RelativePath* contient une séquence de *ReferenceTypes* et de *BrowseNames*.

Un objectif de ce *Service* est de permettre la programmation en fonction des définitions de type. Sachant que les *BrowseNames* doivent être uniques dans le contexte des définitions de type, un *Client* peut créer un chemin de navigation qui est valide pour une définition de type et utiliser ce chemin sur des instances du type. Par exemple, un *ObjectType* "Boiler" peut avoir une *Variable* "HeatSensor" comme *InstanceDeclaration*. Il peut être nécessaire qu'un élément graphique programmé en fonction du "Boiler" affiche la *Valeur* du "HeatSensor". Si l'élément graphique était appelé sur "Boiler1", une instance de "Boiler", il serait nécessaire qu'il appelle ce *Service* en spécifiant le *NodeId* de "Boiler1" comme *Nœud* de départ et le *BrowseName* du "HeatSensor" comme chemin de navigation. Le *Service* retourne le *NodeId* du "HeatSensor" de "Boiler1" et l'élément graphique peut s'abonner à son *attribut Valeur*.

Si un *Nœud* a plusieurs cibles ayant le même *BrowseName*, le *Serveur* doit retourner une liste de *NodeIds*. Cependant, comme l'un des principaux objectifs de ce *Service* est de prendre en charge la programmation en fonction des définitions de type, le *NodeId* du *Nœud* reposant sur la définition de type du *Nœud* de départ est retourné comme le premier *NodeId* dans la liste.

5.8.4.2 Paramètres

Le Tableau 40 définit les paramètres du *Service*.

Tableau 40 – Paramètres du service TranslateBrowsePathsToNodeIds

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
browsePaths []	BrowsePath	Liste des chemins de navigation pour lesquels les <i>NodeIds</i> sont demandés. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
startingNode	NodeId	<i>NodeId</i> du <i>Nœud</i> de départ pour le chemin de navigation.
relativePath	RelativePath	Chemin à suivre à partir du <i>startingNode</i> . Le dernier élément dans le <i>relativePath</i> doit toujours avoir un <i>targetName</i> spécifié. Cela restreint encore plus la définition du type <i>RelativePath</i> . Le <i>Serveur</i> doit retourner <i>Bad_BrowseNameInvalid</i> si le <i>targetName</i> est absent. La structure <i>RelativePath</i> est définie en 7.26.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	BrowsePathResult	Liste de résultats pour la liste de chemins de navigation. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>browsePaths</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> pour le chemin de navigation (voir 7.34 pour la définition de <i>StatusCode</i>).
targets []	BrowsePathTarget	Liste de cibles pour le <i>relativePath</i> à partir du <i>startingNode</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants. Un <i>Serveur</i> peut rencontrer une <i>Référence</i> à un <i>Nœud</i> dans un autre <i>Serveur</i> qu'il ne peut pas suivre pendant qu'il traite le <i>RelativePath</i> . Si cela se produit, le <i>Serveur</i> retourne le <i>NodeId</i> du <i>Nœud</i> externe et définit le paramètre <i>remainingPathIndex</i> pour indiquer les éléments <i>RelativePath</i> qu'il est encore nécessaire de traiter. Pour terminer l'opération, le <i>Client</i> doit se connecter à l'autre <i>Serveur</i> et appeler de nouveau ce service en utilisant la cible comme le <i>startingNode</i> et les éléments non traités comme le <i>relativePath</i> .
targetId	ExpandedNodeId	Identificateur pour une cible du <i>RelativePath</i> .
remainingPathIndex	Index	Indice du premier élément non traité dans le <i>RelativePath</i> . Cette valeur doit être égale à la valeur maximale du type de données <i>Index</i> si tous les éléments ont été traités (voir 7.13 pour la définition d' <i>Index</i>).
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour la liste des chemins de navigation (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>browsePaths</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.8.4.3 Résultats de service

Le Tableau 41 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le 7.34.

Tableau 41 – Codes de résultats du service TranslateBrowsePathsToNodeIds

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.8.4.4 StatusCodes

Le Tableau 42 définit les valeurs de niveau opération pour les paramètres *statusCode* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 42 – Codes de résultats de niveau opération TranslateBrowsePathsToNodeIds

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat. Ce code indique que le <i>relativePath</i> contenait une liste vide.
Bad_BrowseNameInvalid	Voir Tableau 178 pour la description de ce code de résultat. Ce code indique qu'un <i>TargetName</i> était absent dans un <i>RelativePath</i> .
Uncertain_ReferenceOutOfServer	L'élément de chemin a des cibles qui résident sur un autre serveur.
Bad_TooManyMatches	L'opération demandée a de trop nombreux mappings à retourner. Il convient que les utilisateurs emploient des interrogations pour les volumineux jeux de résultats. Il convient que les <i>Serveurs</i> autorisent au moins 10 mappings avant de retourner ce code d'erreur.
Bad_QueryTooComplex	L'opération demandée exige un trop grand nombre de ressources sur le serveur.
Bad_NoMatch	Le <i>relativePath</i> demandé ne peut être résolu en une cible à retourner.

5.8.5 RegisterNodes

5.8.5.1 Description

Un *Serveur* n'a souvent pas d'accès direct aux informations qu'il gère. Les variables ou services peuvent résider dans des systèmes sous-jacents, où l'établissement d'une connexion à ces systèmes doit nécessiter des efforts supplémentaires. Le *Service RegisterNodes* peut être utilisé par des *Clients* pour enregistrer les *Nœuds* auxquels ils savent qu'ils vont accéder de façon répétée (par exemple *Write*, *Call*). Il permet aux *Serveurs* d'installer les éléments nécessaires pour optimiser les opérations d'accès. Les *Clients* peuvent espérer des améliorations de performance lorsqu'ils utilisent des *NodeIds* enregistrés, mais les mesures d'optimisation sont spécifiques au fournisseur. Pour les *Nœuds de variable*, les *Serveurs* doivent concentrer leurs efforts d'optimisation sur l'*attribut Valeur*.

La validité des *NodeIds* enregistrés ne peut être garantie que dans la *Session* courante. Les *Clients* doivent immédiatement supprimer les enregistrements des identificateurs inutiles pour libérer des ressources.

RegisterNodes ne valide pas les *NodeIds* issus de la demande. Les *Serveurs* copient simplement les *NodeIds* inconnus dans la réponse. Les erreurs de structure des *NodeIds* (violations de taille, types d'identificateurs non valides) entraînent la défaillance du *service* complet.

Pour les besoins de l'*audit*, les *Serveurs* ne doivent pas utiliser les *NodeIds* enregistrés, mais seulement les *NodeIds* canoniques qui sont la valeur de l'*attribut "NodeId"*.

5.8.5.2 Paramètres

Le Tableau 43 définit les paramètres du *Service*.

Tableau 43 – Paramètres du service RegisterNodes

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
nodesToRegister []	NodeId	Liste des <i>NodeIds</i> à enregistrer que le <i>Client</i> a récupérés par navigation, une requête ou de quelque autre manière.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
registeredNodeIds []	NodeId	Liste de <i>NodeIds</i> que le <i>Client</i> doit utiliser pour des opérations d'accès ultérieures. La taille et l'ordre de cette liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToRegister</i>. Le <i>Serveur</i> peut retourner le <i>NodeId</i> de la demande ou un nouveau <i>NodeId</i>. Il est recommandé que le <i>Serveur</i> retourne un <i>NodeId</i> numérique pour la dénomination. Si aucune optimisation n'est prise en charge pour un <i>Nœud</i>, le <i>Serveur</i> doit retourner le <i>NodeId</i> issu de la demande.

5.8.5.3 Résultats de service

Le Tableau 44 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 44 – Codes de résultats du service RegisterNodes

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat. Les <i>Serveurs</i> doivent rejeter complètement la demande <i>RegisterNodes</i> si l'un quelconque des <i>NodeIds</i> dans le paramètre <i>nodesToRegister</i> est structurellement non valide.

5.8.6 UnregisterNodes

5.8.6.1 Description

Ce *Service* est utilisé pour supprimer l'enregistrement des *NodeIds* qui ont été obtenus via le service *RegisterNodes*.

UnregisterNodes ne valide pas les *NodeIds* issus de la demande. Les *Serveurs* doivent simplement supprimer l'enregistrement des *NodeIds* qui sont connus comme des *NodeIds* enregistrés. Tous les *NodeIds* qui sont dans la liste, mais ne sont pas des *NodeIds* enregistrés, sont simplement ignorés.

5.8.6.2 Paramètres

Le Tableau 50 définit les paramètres du *Service*.

Tableau 45 – Paramètres du service UnregisterNodes

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
nodesToUnregister []	NodeId	Une liste des <i>NodeIds</i> qui ont été obtenus via le service <i>RegisterNodes</i> .
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).

5.8.6.3 Résultats de service

Le Tableau 51 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 46 – Codes de résultats du service UnregisterNodes

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.9 Jeu de services Query

5.9.1 Vue d'ensemble

Ce *Jeu de services* est utilisé pour émettre une *requête* vers un *Serveur*. La *requête* OPC UA est générique dans le sens, où elle fournit un mécanisme de stockage sous-jacent indépendant de la fonction *Query* qui peut être utilisé pour accéder à différents magasins de données et systèmes de gestion d'informations d'OPC UA. La *requête* OPC UA permet à un *Client* d'accéder aux données gérées par un *Serveur* sans connaître le plan logique utilisé pour le stockage interne des données. La connaissance de l'*AddressSpace* est suffisante.

Une *application OPC UA* est supposée utiliser les *services Query* d'OPC UA dans le cadre d'un processus d'initialisation ou d'une étape de synchronisation d'informations. Par exemple, la *requête* OPC UA serait utilisée pour l'accès aux données en vrac d'un magasin de données persistant afin d'initialiser une application d'analyse avec l'état courant d'une configuration système. Une *requête* peut aussi être utilisée pour initialiser ou alimenter des données aux fins d'un rapport.

Une *requête* définit les instances d'un ou de plusieurs *TypeDefinitionNodes* dans l'*AddressSpace* qu'il convient d'appliquer à un jeu d'*attributs*. Les résultats retournés par un *Serveur* sont sous la forme d'une matrice de *QueryDataSets*. Les valeurs d'*attributs* choisies dans chaque *QueryDataSet* proviennent de la définition des *TypeDefinitionNodes* ou *TypeDefinitionNodes* connexes choisis et les résultats apparaissent dans le même ordre que les *attributs* qui ont été transmis dans la *requête*. La *requête* prend également en charge le filtrage des *nœuds* en fonction des valeurs d'*attributs*, ainsi que les relations entre *TypeDefinitionNodes*.

Voir Annexe B pour des exemples d'interrogations.

5.9.2 Vues de requêtes

Une *Vue* est un sous-ensemble de l'*AddressSpace* disponible dans le *Serveur*. Voir l'IEC 62541-5 pour une description de l'organisation des *Vues* dans l'*AddressSpace*.

Pour toute *vue* existante, une *requête* peut être utilisée pour retourner un sous-ensemble de données de la *vue*. Lorsqu'une application émet une *requête* en fonction d'une *vue*, seules les données définies par la *vue* sont retournées. Les données non incluses dans la *Vue*, mais incluses dans l'*AddressSpace* d'origine ne sont pas retournées.

Les *services Query* prennent en charge l'accès aux données courantes et historiques. Le *Service* prend en charge un *Client* qui interroge une version antérieure de l'*AddressSpace*. Les *Clients* peuvent spécifier une *ViewVersion* ou un *Timestamp* dans une *requête* pour accéder à des versions antérieures de l'*AddressSpace*. Le service Query OPC UA est complémentaire à l'accès à l'historique, dans le sens où le premier est utilisé pour la *requête* d'un *AddressSpace* qui existait à un moment donné et que le second est utilisé pour une *requête* afin d'obtenir la valeur des *attributs* dans le temps. Ainsi, une *requête* peut être utilisée pour récupérer une partie d'un *AddressSpace* antérieur et, de ce fait, un historique des valeurs de l'*attribut* peut être accessible à l'aide de l'accès à l'historique même si le *nœud* ne se trouve plus dans l'*AddressSpace* courant.

Les *Serveurs* qui prennent en charge la *Requête* sont supposés être capables d'accéder à l'*AddressSpace* qui est associé au *Serveur* local et à n'importe quelle *Vue* disponible sur le *Serveur* local. Si une *Vue* ou un *AddressSpace* référence également un *Serveur* distant, la requête peut être capable d'accéder à l'*AddressSpace* du *Serveur* distant, mais cela n'est pas exigé. Si un *Serveur* accède à un *Server* distant, l'accès doit être réalisé avec l'identité utilisateur du *Client* décrite en 5.5.1.

5.9.3 QueryFirst

5.9.3.1 Description

Ce *Service* est utilisé pour émettre une demande *Query* vers le *Serveur*. La complexité de la *requête* peut être très simple ou extrêmement sophistiquée. La *Requête* peut simplement demander des données issues d'instances d'un *TypeDefinitionNode* ou *TypeDefinitionNode* soumis à des restrictions spécifiées par le filtre. Par ailleurs, la *requête* peut demander des données issues d'instances de type *Nœuds* connexes en spécifiant un *RelativePath* partant d'un *TypeDefinitionNode* d'origine. Dans le filtre, un jeu séparé de chemins peut être construit pour limiter les instances qui fournissent les données. Un chemin filtrant peut inclure plusieurs opérateurs *RelatedTo* pour définir un chemin à sauts multiples entre des instances sources et des instances cibles. Par exemple, il est possible de filtrer les étudiants d'une école particulière, mais de retourner des informations relatives aux étudiants et à leurs familles. Dans ce cas, la relation étudiant-école est parcourue pour le filtrage, mais seule la relation étudiant-famille est parcourue pour le tri des données. Pour une description complète de *ContentFilter*, voir 7.4; voir aussi B.1 pour des exemples simples et B.2 pour de plus complexes exemples de filtre de contenu et d'interrogations.

Le *Client* fournit une matrice de *NodeTypeDescription* qui spécifie le *NodeId* d'un *TypeDefinitionNode* et choisit les *Attributs* à retourner dans la réponse. Un *Client* peut aussi fournir un jeu de *RelativePaths* à travers le système de type commençant à partir d'un *TypeDefinitionNode* de départ. En utilisant ces chemins, le *Client* choisit un jeu d'*attributs* à partir des *nœuds* qui sont reliés à des instances du *TypeDefinitionNode* de départ. En outre, le *Client* peut demander au *Serveur* de retourner des instances des sous-types de *TypeDefinitionNodes*. Si un *Attribut* choisi n'existe pas dans un *TypeDefinitionNode*, mais qu'il existe dans un sous-type, l'hypothèse retenue est qu'il a une valeur Null dans ce *TypeDefinitionNode*. Par conséquent, ceci ne constitue pas une condition d'erreur et une valeur Null est retournée pour l'*Attribut*.

Le *Client* peut utiliser le paramètre filtre pour limiter le jeu de résultats en limitant les *Attributs* et les *Propriétés* à certaines valeurs. Une autre manière dont le *Client* peut utiliser un filtre

pour limiter le jeu de résultats consiste à spécifier comment il convient de relier les instances, en utilisant des opérateurs *RelatedTo*. Dans ce cas, si une instance au sommet du chemin *RelatedTo* ne peut pas être suivie jusqu'au bas du chemin via les sauts spécifiés, aucun *QueryDataSet* n'est retourné pour l'instance de départ ou l'une des instances intermédiaires.

Lorsqu'il effectue une requête relative à des instances connexes dans le *RelativePath*, le *Client* peut demander des *Références*. Une *Référence* est demandée via un *RelativePath* qui inclut seulement un *ReferenceType*. Si toutes les *Références* sont souhaitées, le *ReferenceType* racine est alors énuméré. Ces *Références* sont retournées comme partie des *QueryDataSets*.

Les services Query permettent une gestion particulière du champ *targetName* dans le *RelativePath*. Dans plusieurs cas d'utilisation des requêtes, un *NodeId* de type est nécessaire dans le chemin, à la place d'un *QualifiedName*. Le *Client* peut donc spécifier un *NodeId* dans le *QualifiedName*. Cela est réalisé en mettant le *namespaceIndex* du *targetName* à zéro et la partie nom du *targetName* à la représentation en XML du *NodeId*. La représentation XML est définie dans l'IEC 62541-6. Lorsque les instances de mapping sont retournées comme le nœud cible, le nœud cible doit être une instance du type spécifié ou de son sous-type.

Le Tableau 47 définit les paramètres de demande et le Tableau 48 les paramètres de réponse pour le *Service QueryFirst*.

Tableau 47 – Paramètres de demande QueryFirst

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
Vue	ViewDescription	Spécifie une <i>Vue</i> et un contexte temporel à un <i>Serveur</i> (voir 7.39 pour la définition de <i>ViewDescription</i>).
nodeTypes []	NodeTypeDescription	Description du type de <i>Nœud</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
typeDefinitionNode	ExpandedNodeId	<i>NodeId</i> du <i>TypeDefinitionNode</i> de départ des instances pour lesquelles les données doivent être retournées.
includeSubtypes	Boolean	Un fanion qui indique s'il convient que le <i>Serveur</i> inclue des instances des sous-types du <i>TypeDefinitionNode</i> dans la liste des instances du type de <i>Nœud</i> .
dataToReturn []	QueryDataDescription	Spécifie un <i>Attribut</i> ou une <i>Référence</i> partant du <i>typeDefinitionNode</i> de départ le long d'un <i>relativePath</i> donné pour lequel retourner des données. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
relativePath	RelativePath	Chemin de navigation relatif au <i>Nœud</i> de départ qui identifie le <i>Nœud</i> contenant les données demandées, le <i>Nœud</i> de départ étant un <i>Nœud</i> d'instance du type défini par le <i>Nœud</i> de définition de type. Les <i>Nœuds</i> d'instance font l'objet d'une limitation supplémentaire par le filtre fourni comme partie de cet appel. Pour une définition de <i>relativePath</i> , voir 7.26. Ce chemin relatif peut se terminer sur une <i>Référence</i> et, dans ce cas, la <i>ReferenceDescription</i> de la <i>Référence</i> est retournée comme étant sa valeur. Le champ <i>targetName</i> du <i>relativePath</i> peut contenir un type de <i>NodeId</i> . Cela est réalisé en mettant le <i>namespaceIndex</i> du <i>targetName</i> à zéro et la partie <i>nom</i> du <i>targetName</i> à la représentation en XML du <i>NodeId</i> . La représentation XML est définie dans l'IEC 62541-6. Lorsque les instances de mapping sont retournées comme le nœud cible, le nœud cible doit être une instance du type spécifié ou de son sous-type.
AttributeID	IntegerId	Identificateur de l' <i>Attribut</i> . Il doit s'agir d'un identificateur d' <i>attribut</i> valide. L' <i>IntegerId</i> est défini en 7.14. Les <i>IntegerIds</i> pour les <i>Attributs</i> sont définis dans l'IEC 62541-6. Si le <i>RelativePath</i> a abouti à une <i>Référence</i> , ce paramètre est 0 et est ignoré par le serveur.

Nom	Type	Description
indexRange	NumericRange	Ce paramètre est utilisé pour identifier un élément simple d'une structure ou d'une matrice ou bien une plage simple d'indices pour les matrices. Si une plage d'éléments est spécifiée, les valeurs sont retournées sous forme de composé. Le premier élément est identifié par l'indice 0 (zéro). Le type <i>NumericRange</i> est défini en 7.22. Ce paramètre est Null si l'<i>Attribut</i> spécifié n'est pas une matrice ou une structure. Cependant, si l'<i>Attribut</i> spécifié est une matrice ou une structure, et que ce paramètre est Null, tous les éléments doivent être inclus dans la plage.
Filtre	ContentFilter	Les <i>Nœuds</i> obtenus doivent être limités aux <i>Nœuds</i> qui correspondent aux critères définis par le filtre. <i>ContentFilter</i> est l'objet d'une discussion en 7.4. Si un filtre vide est fourni, l'<i>AddressSpace</i> doit être examiné dans son intégralité et tous les <i>Nœuds</i> contenant un <i>Attribut</i> ou une <i>Référence</i> demandé(e) correspondant aux critères sont retournés.
maxDataSetsToReturn	Counter	Nombre de <i>QueryDataSets</i> que le <i>Client</i> souhaite que le <i>Serveur</i> retourne dans la réponse et sur chaque réponse ultérieure d'appel de continuation. Le <i>Serveur</i> peut restreindre encore la réponse, mais ne doit pas dépasser cette limite. Une valeur 0 indique que le <i>Client</i> n'impose aucune limitation.
maxReferencesToReturn	Counter	Nombre de <i>Références</i> que le <i>Client</i> souhaite que le <i>Serveur</i> retourne dans la réponse pour chaque <i>QueryDataSet</i> et sur chaque réponse ultérieure d'appel de continuation. Le <i>Serveur</i> peut restreindre encore la réponse, mais ne doit pas dépasser cette limite. Une valeur 0 indique que le <i>Client</i> n'impose aucune limitation. Par exemple, dans le cas d'un résultat où 4 <i>Nœuds</i> sont retournés, avec chacun 100 <i>Références</i> et une limite établie à 50, seules les 50 premières <i>Références</i> de chaque <i>Nœud</i> sont retournées sur l'appel initial, et un point de continuation est établi pour indiquer l'existence de données supplémentaires.

Tableau 48 – Paramètres de réponse QueryFirst

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
queryDataSets []	QueryDataSet	Matrice de <i>QueryDataSets</i>. Cette matrice est vide s'il n'y a pas de <i>Nœuds</i> ou de <i>Références</i> répondant aux critères des <i>nodeTypes</i>. Dans ce cas, le paramètre <i>continuationPoint</i> doit être vide. Le type <i>QueryDataSet</i> est défini en 7.23.
continuationPoint	ContinuationPoint	Valeur opaque définie par le serveur qui identifie le point de continuation. Le point de continuation est utilisé seulement lorsque les résultats de la <i>requête</i> sont trop volumineux pour être retournés dans une seule réponse. "Trop volumineux" signifie dans ce contexte que le <i>Serveur</i> n'est pas capable de retourner une réponse plus volumineuse ou que le nombre de <i>QueryDataSets</i> à retourner dépasse le nombre maximal de <i>QueryDataSets</i> à retourner qui a été spécifié par le <i>Client</i> dans la demande. Le point de continuation est utilisé dans le <i>Service QueryNext</i>. Lorsqu'il n'est pas utilisé, la valeur de ce paramètre est Null. Si un point de continuation est retourné, le <i>Client</i> doit appeler <i>QueryNext</i> pour récupérer le prochain jeu de <i>QueryDataSets</i> ou pour libérer des ressources pour le point de continuation dans le <i>Serveur</i>. Un point de continuation doit rester actif jusqu'à ce que le <i>Client</i> transmette le point de continuation à <i>QueryNext</i> ou jusqu'à ce que la session soit fermée. Si le nombre maximal de points de continuation a été atteint, le point de continuation le plus ancien doit être réinitialisé. Le type <i>ContinuationPoint</i> est décrit en 7.6.

Nom	Type	Description
parsingResults[]	ParsingResult	Liste de résultats d'analyse pour <i>QueryFirst</i> . La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>NodeTypes</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants. Cette liste est alimentée par tous les codes de statut qui se rapportent au traitement des types de nœuds qui relèvent de la requête. La matrice peut être vide s'il ne s'est pas produit d'erreurs. Si un quelconque type de nœud a rencontré une erreur, tous les types de nœuds doivent avoir un code de statut associé.
statusCode	StatusCode	Résultat d'analyse pour la <i>NodeTypeDescription</i> demandée.
dataStatusCodes []	StatusCode	Liste de résultats pour <i>dataToReturn</i> . La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>dataToReturn</i> . La matrice peut être vide s'il ne s'est pas produit d'erreurs.
dataDiagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>dataToReturn</i> (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>dataToReturn</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été rencontrée dans le traitement de la demande de requête.
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour la <i>NodeTypeDescription</i> demandée. Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été rencontrée dans le traitement de la demande de requête.
filterResult	ContentFilter Résultat	Structure contenant toutes les erreurs éventuelles associées au filtre. Cette structure doit être vide s'il ne s'est pas produit d'erreurs. Le type <i>ContentFilterResult</i> est défini en 7.4.2.

5.9.3.2 Résultats de service

Si la *requête* n'est pas valide ou ne peut pas être traitée, les *QueryDataSets* ne sont pas retournés et seul un résultat de *service* (*filterResult*, *parsingResults* et *DiagnosticInfo* facultatif) est retourné. Le Tableau 49 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 49 – Codes de résultats du service QueryFirst

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ContentFilterInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_ViewIdUnknown	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ViewTimestampInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ViewParameterMismatchInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ViewVersionInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_InvalidFilter	Le filtre fourni n'est pas valide (voir <i>filterResult</i> pour les erreurs spécifiques).
Bad_NodelistError	Les <i>NodeTypes</i> fournis contiennent une erreur (voir <i>parsingResults</i> pour les erreurs spécifiques).
Bad_InvalidView	La <i>ViewDescription</i> fournie n'est pas une <i>ViewDescription</i> valide.
Good_ResultsMayBeIncomplete	Il convient que le <i>Serveur</i> ait suivi une référence à un nœud dans un <i>Serveur</i> distant, mais il ne l'a pas fait. Le jeu de résultats peut être incomplet.

5.9.3.3 StatusCodes

Le Tableau 50 définit les valeurs pour le paramètre *statusCode* de *parsingResults* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 50 – Codes de résultats de niveau opération QueryFirst

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NotTypeDefinition	Le NodeId fourni n'était pas un NodeId de définition de type.
Bad_AttributeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.9.4 QueryNext

5.9.4.1 Descriptions

Ce *Service* est utilisé pour demander le prochain jeu d'informations de réponse *QueryFirst* ou *QueryNext*, qui sont trop volumineuses pour être envoyées dans une seule réponse. "Trop volumineuses" signifie dans ce contexte que le *Serveur* n'est pas capable de retourner une réponse plus volumineuse ou que le nombre de *QueryDataSets* à retourner dépasse le nombre maximal de *QueryDataSets* à retourner qui a été spécifié par le *Client* dans la demande d'origine. Le *QueryNext* doit être présenté dans la même session qui a été utilisée pour présenter le *QueryFirst* ou *QueryNext* qui se poursuit.

5.9.4.2 Paramètres

Le Tableau 51 définit les paramètres du *Service*.

Tableau 51 – Paramètres du service QueryNext

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
releaseContinuationPoint	Boolean	Paramètre <i>booléen</i> ayant les valeurs suivantes: TRUE Le <i>continuationPoint</i> transmis doit être réinitialisé pour libérer des ressources pour le point de continuation dans le <i>Serveur</i>. FALSE Le <i>continuationPoint</i> transmis doit être utilisé pour récupérer le prochain jeu de <i>QueryDataSets</i>. Un <i>Client</i> doit toujours utiliser le point de continuation retourné par une réponse <i>QueryFirst</i> ou <i>QueryNext</i> pour libérer les ressources pour le point de continuation dans le <i>Serveur</i>. Si le <i>Client</i> ne souhaite pas récupérer le prochain jeu d'informations <i>Query</i>, <i>QueryNext</i> doit être appelé avec ce paramètre mis à TRUE. Si le paramètre est mis à TRUE, tous les paramètres de matrices dans la réponse doivent contenir des matrices vides.
continuationPoint	ContinuationPoint	Valeur opaque définie par le serveur qui représente le point de continuation. La valeur pour le point de continuation a été retournée au <i>Client</i> dans une réponse <i>QueryFirst</i> ou <i>QueryNext</i> précédente. Cette valeur est utilisée pour identifier la demande <i>QueryFirst</i> ou <i>QueryNext</i> précédemment traitée qui se poursuit, ainsi que le point dans le jeu de résultats à partir duquel la réponse de navigation doit se poursuivre. Le type <i>ContinuationPoint</i> est décrit en 7.6.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
queryDataSets []	QueryDataSet	Matrice de <i>QueryDataSets</i>. Le type <i>QueryDataSet</i> est défini en 7.23.
revisedContinuationPoint	ContinuationPoint	Valeur opaque définie par le serveur qui représente le point de continuation. Elle est utilisée seulement si les informations à retourner sont trop volumineuses pour être contenues dans une seule réponse. Lorsqu'elle n'est pas utilisée ou lorsque le <i>releaseContinuationPoint</i> est défini, la valeur de ce paramètre est Null. Le type <i>ContinuationPoint</i> est décrit en 7.6.

5.9.4.3 Résultats de service

Le Tableau 52 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 52 – Codes de résultats du service QueryNext

ID symbolique	Description
Bad_ContinuationPointInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.10 Jeu de services Attribute

5.10.1 Vue d'ensemble

Ce *Jeu de services* fournit les *Services* pour accéder à des *Attributs* qui sont partie intégrante des *Nœuds*.

5.10.2 Read

5.10.2.1 Description

Ce *Service* est utilisé pour lire un ou plusieurs *Attributs* d'un ou plusieurs *Nœuds*. Pour les valeurs d'*attributs* construits dont les éléments ont des indices (comme une matrice), ce *Service* permet aux *Clients* de lire la totalité du jeu de valeurs à indices comme un composé, de lire des éléments individuels ou de lire des plages d'éléments du composé.

Le paramètre *maxAge* est utilisé pour ordonner au *Serveur* d'accéder à la valeur à partir de la source de données sous-jacente (comme un appareil) si sa copie des données est plus ancienne que celle que le *maxAge* spécifie. Si le *Serveur* ne peut pas répondre à l'âge maximal demandé, il retourne sa meilleure valeur possible au lieu de rejeter la demande.

5.10.2.2 Paramètres

Le Tableau 53 définit les paramètres du *Service*.

Tableau 53 – Paramètres du service Read

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
maxAge	Duration	Age maximal de la valeur à lire, en millisecondes. L'âge de la valeur repose sur la différence entre le <i>ServerTimestamp</i> et l'heure à laquelle le <i>Serveur</i> commence à traiter la demande. Par exemple, si le <i>Client</i> spécifie un <i>maxAge</i> de 500 ms et que la durée d'attente est de 100 ms avant que le <i>Serveur</i> ne commence à traiter la demande, l'âge de la valeur retournée peut être 600 ms avant l'heure à laquelle elle a été demandée. Si le <i>Serveur</i> a une ou plusieurs valeurs d'un <i>Attribut</i> dans la limite de l'âge maximal, il peut retourner n'importe laquelle de ces valeurs ou il peut lire une nouvelle valeur dans la source de données. Le nombre de valeurs d'un <i>Attribut</i> que possède un <i>Serveur</i> dépend du nombre de <i>MonitoredItems</i> qui sont définis pour cet <i>Attribut</i>. Dans tous les cas, le <i>Client</i> ne peut présupposer de la copie des données qui est retournée. Si le <i>Serveur</i> n'a pas de valeur s'inscrivant dans la limite de l'âge maximal, il doit tenter de lire une nouvelle valeur dans la source de données. Si le <i>Serveur</i> ne peut pas répondre au <i>maxAge</i> demandé, il retourne sa meilleure valeur possible au lieu de rejeter la demande. Ce cas de figure peut se produire lorsque le temps mis par le <i>Serveur</i> pour traiter et retourner la nouvelle valeur de donnée après accès est supérieur à l'âge maximal spécifié. Si <i>maxAge</i> est mis à 0, le <i>Serveur</i> doit tenter de lire une nouvelle valeur dans la source de données. Si <i>maxAge</i> est mis à une valeur supérieure ou égale à la valeur <i>Int32</i> maximale, le <i>Serveur</i> doit tenter de récupérer une valeur mise en cache. Les valeurs négatives ne sont pas valides pour <i>maxAge</i>.
timestampsToReturn	Enum TimestampsToReturn	Énumération qui spécifie les <i>horodatages</i> à retourner pour chaque <i>Attribut</i> de valeur de <i>variable</i> . L'énumération <i>TimestampsToReturn</i> est définie en 7.35.
nodesToRead []	ReadValueld	Liste des <i>Nœuds</i> et de leurs <i>Attributs</i> à lire. Pour chaque entrée de la liste, un <i>StatusCode</i> est retourné, et s'il indique un succès, la <i>valeur d'attribut</i> est également retournée. Le type de paramètre <i>ReadValueld</i> est défini en 7.24.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	DataValue	Liste de valeurs d' <i>attribut</i> (voir 7.7 pour la définition <i>DataValue</i>). La taille et l'ordre de cette liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToRead</i> . Cette liste comporte une entrée pour chaque <i>Nœud</i> contenu dans le paramètre <i>nodesToRead</i> .
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de cette liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToRead</i> . Cette liste comporte une entrée pour chaque <i>Nœud</i> contenu dans le paramètre <i>nodesToRead</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.10.2.3 Résultats de service

Le Tableau 54 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 54 – Codes de résultats du service Read

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_MaxAgeInvalid	Le paramètre d'âge maximal n'est pas valide.
Bad_TimestampsToReturnInvalid	Voir Tableau 177 pour la description de ce code de résultat.

5.10.2.4 StatusCodes

Le Tableau 55 définit les valeurs pour le *statusCode* de niveau opération contenu dans la structure *DataValue* de chaque élément *résultats*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 55 – Codes de résultats de niveau opération Read

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_AttributeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeNoData	Voir Tableau 178 pour la description de ce code de résultat.
Bad_DataEncodingInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_DataEncodingUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NotReadable	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_SecurityModeInsufficient	Voir Tableau 178 pour la description de ce code de résultat.

5.10.3 HistoryRead

5.10.3.1 Description

Ce *Service* est utilisé pour lire des valeurs ou *Événements* historiques d'un ou plusieurs *Nœuds*. Pour les valeurs d'*attributs* construits dont les éléments ont des indices (comme une matrice), ce *Service* permet aux *Clients* de lire la totalité du jeu de valeurs à indices comme un composé, de lire des éléments individuels ou de lire des plages d'éléments du composé. Les *Serveurs* peuvent rendre les valeurs historiques disponibles aux *Clients* en utilisant ce *Service*, bien que les valeurs historiques elles-mêmes ne soient pas visibles dans l'*AddressSpace*.

L'*Attribut AccessLevel* défini dans l'IEC 62541-3 indique une prise en charge des valeurs historiques par des *Nœuds*. Plusieurs paramètres de demande indiquent comment le *Serveur* doit accéder aux valeurs de la source sous-jacente de données d'historique. L'*Attribut EventNotifier* défini dans l'IEC 62541-3 indique la prise en charge d'*Événements* historiques par des *Nœuds*.

Le paramètre *continuationPoint* dans l'*HistoryRead* est utilisé pour marquer le point à partir duquel poursuivre la lecture si toutes les valeurs n'ont pas pu être retournées dans une seule réponse. La valeur n'est pas connue du *Client* et elle est utilisée seulement pour gérer les informations d'état pour le point de continuation du *Serveur*. Un *Serveur* peut utiliser l'horodatage du dernier élément de données retourné si l'horodatage est unique. Cela peut réduire la nécessité de stocker des informations d'état dans le *Serveur* pour le point de continuation.

Dans certains cas, le traitement de tous les nœuds à parcourir peut prendre plus de temps que la suggestion de temporisation du *Client* pour tous les nœuds à lire. Le *Serveur* peut alors retourner des résultats nuls avec un point de continuation pour les nœuds affectés avant l'expiration de la temporisation. Ceci permet au *Serveur* de poursuivre l'acquisition des données sur l'appel de lecture du *Client* suivant.

Pour des informations complémentaires relatives à la lecture de données historiques et d'*Événements* historiques, voir l'IEC 62541-11.

5.10.3.2 Paramètres

Le Tableau 56 définit les paramètres du *Service*.

Tableau 56 – Paramètres du service HistoryRead

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
historyReadDetails	Paramètre extensible HistoryReadDetails	Les informations définissent les types de lectures d'historiques qui peuvent être effectués. Le type de paramètre <i>HistoryReadDetails</i> est un type de paramètre extensible qui est défini formellement dans l'IEC 62541-11. Le type <i>ExtensibleParameter</i> est défini en 7.12.
timestampsToReturn	Enum TimestampsToReturn	Énumération qui spécifie les horodatages à retourner pour chaque <i>Attribut de valeur de variable</i> . L'énumération <i>TimestampsToReturn</i> est définie en 7.35. Il n'est pas valide de spécifier un <i>TimestampsToReturn</i> de NEITHER. Dans ce cas, le <i>Serveur</i> doit retourner un <i>StatusCode Bad_InvalidTimestampArgument</i> . L'IEC 62541-11 définit les exceptions pour lesquelles ce paramètre doit être ignoré.

Nom	Type	Description
releaseContinuation Points	Boolean	Paramètre <i>booléen</i> ayant les valeurs suivantes: TRUE les <i>continuationPoints</i> transmis doivent être réinitialisés pour libérer des ressources sur le <i>Serveur</i>. FALSE les <i>continuationPoints</i> transmis doivent être utilisés pour récupérer le prochain jeu d'informations historiques. Un <i>Client</i> doit toujours utiliser le point de continuation retourné par une réponse <i>HistoryRead</i> pour libérer les ressources pour le point de continuation dans le <i>Serveur</i>. Si le <i>Client</i> ne souhaite pas récupérer le prochain jeu d'informations historiques, <i>HistoryRead</i> doit être appelé avec ce paramètre mis à TRUE.
nodesToRead []	HistoryReadValueId	Ce paramètre contient la liste d'éléments sur lesquels la récupération d'historique doit être réalisée. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
nodeId	NodeId	Si l'<i>HistoryReadDetails</i> est RAW, PROCESSED, MODIFIED ou ATTIME: <i>NodeId</i> des <i>Nœuds</i> dont les valeurs historiques doivent être lues. La valeur retournée doit toujours comprendre un horodatage. Si l'<i>HistoryReadDetails</i> est EVENTS: Le <i>NodeId</i> du <i>Nœud</i> dont l'historique des <i>Événements</i> doit être lu. Si le <i>Nœud</i> ne prend pas en charge l'accès demandé pour des valeurs historiques ou des <i>Événements</i> historiques, une réponse d'erreur appropriée pour le <i>Nœud</i> donné doit être générée.
indexRange	NumericRange	Ce paramètre est utilisé pour identifier un élément simple d'une matrice ou une plage simple d'indices pour les matrices. Si une plage d'éléments est spécifiée, les valeurs sont retournées sous forme de composé. Le premier élément est identifié par l'indice 0 (zéro). Le type <i>NumericRange</i> est défini en 7.22. Ce paramètre est Null si la valeur n'est pas une matrice. Cependant, si la valeur est une matrice et que ce paramètre est Null, tous les éléments doivent être inclus dans la plage.
dataEncoding	QualifiedName	<i>QualifiedName</i> qui spécifie le codage de données à retourner pour la <i>Valeur</i> à lire (voir 7.24 pour la définition de la façon de spécifier le codage de données). Ce paramètre ne s'applique que si le <i>DataType</i> de la <i>Variable</i> est un sous-type de <i>Structure</i>. Il est erroné de spécifier ce paramètre si le <i>DataType</i> de la <i>Variable</i> n'est pas un sous-type de <i>Structure</i>. Le paramètre est ignoré lors de la lecture de l'historique des <i>Événements</i>.
continuationPoint	ContinuationPoint	Pour chaque <i>NodesToRead</i>, ce paramètre spécifie un point de continuation retourné à partir d'un précédent appel <i>HistoryRead</i>, permettant au <i>Client</i> de poursuivre cette lecture à partir de la dernière valeur reçue. Le <i>HistoryRead</i> est utilisé pour choisir une séquence ordonnée de valeurs ou événements historiques. Un point de continuation marque un point dans cette séquence ordonnée et, ainsi, le <i>Serveur</i> retourne le sous-ensemble de la séquence qui suit le point en question. Une valeur Null indique que ce paramètre n'est pas utilisé. Voir 7.6 pour une description générale des points de continuation. Ce point de continuation est décrit de façon plus précise dans l'IEC 62541-11.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition du type <i>ResponseHeader</i>).
results []	HistoryReadResult	Liste de résultats lus. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToRead</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> pour l'élément <i>NodesToRead</i> (voir 7.34 pour la définition de <i>StatusCode</i>).

Nom	Type	Description
continuationPoint	ContinuationPoint	Ce paramètre est utilisé uniquement si le nombre de valeurs à retourner est trop important pour être retourné dans une seule réponse ou si la temporisation fournie à titre indicatif par le <i>Client</i> est sur le point d'expirer et que tous les nœuds n'ont pas été traités. Lorsque ce paramètre n'est pas utilisé, sa valeur est Null. Les <i>Serveurs</i> doivent prendre en charge au moins un point de continuation par <i>Session</i>. Les <i>Serveurs</i> spécifient un nombre maximal de points de continuation historiques dans l'<i>Objet</i> des capacités du <i>Serveur</i> défini dans l'IEC 62541-5. Un point de continuation doit rester actif jusqu'à ce que le <i>Client</i> transmette le point de continuation à <i>HistoryRead</i> ou jusqu'à ce que la session soit fermée. Si le nombre maximal de points de continuation a été atteint, le point de continuation le plus ancien doit être réinitialisé.
historyData	Paramètre extensible HistoryData	Données d'historique retournées par le <i>Nœud</i>. Le type de paramètre <i>HistoryData</i> est un type de paramètre extensible qui est défini formellement dans l'IEC 62541-11. Il spécifie les types de données d'historique qui peuvent être retournées. Le type de base <i>ExtensibleParameter</i> est défini en 7.12.
diagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToRead</i>. Cette liste comporte une entrée pour chaque <i>Nœud</i> contenu dans le paramètre <i>nodesToRead</i>. Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.10.3.3 Résultats de service

Le Tableau 57 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 57 – Codes de résultats du service HistoryRead

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TimestampsToReturnInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_HistoryOperationInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_HistoryOperationUnsupported	Voir Tableau 178 pour la description de ce code de résultat. L'opération demandée sur l'historique n'est pas prise en charge par le <i>Serveur</i>.

5.10.3.4 StatusCodes

Le Tableau 58 définit les valeurs de niveau opération pour le paramètre *statusCode* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178. Les *StatusCodes* spécifiques à l'accès aux historiques sont définis dans l'IEC 62541-11.

Tableau 58 – Codes de résultats de niveau opération HistoryRead

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_DataEncodingInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_DataEncodingUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_ContinuationPointInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_InvalidTimestampArgument	L'horodatage défini à retourner n'était pas valide.
Bad_HistoryOperationUnsupported	Voir Tableau 178 pour la description de ce code de résultat. L'opération demandée sur l'historique n'est pas prise en charge pour le Nœud demandé.
Bad_NoContinuationPoints	Voir Tableau 178 pour la description de ce code de résultat. Voir 7.6 pour les règles à appliquer à ce code de statut.

5.10.4 Write

5.10.4.1 Description

Ce *Service* est utilisé pour écrire des valeurs dans un ou plusieurs *Attributs* d'un ou plusieurs *Nœuds*. Pour les valeurs d'*attributs* construits dont les éléments ont des indices (comme une matrice), ce *Service* permet aux *Clients* d'écrire la totalité du jeu de valeurs à indices comme un composé, d'écrire des éléments individuels ou d'écrire des plages d'éléments du composé.

Les valeurs sont écrites dans la source de données (comme un appareil) et le *Service* ne retourne rien jusqu'à ce qu'il écrive les valeurs ou détermine que la valeur ne peut pas être écrite. Dans certains cas, le *Serveur* écrit avec succès dans un système ou *Serveur* intermédiaire et ne sait pas si la source de données a été correctement mise à jour. Dans ces cas, il convient que le *Serveur* rapporte un code de succès qui indique que l'écriture n'a pas été vérifiée. Dans les cas où le *Serveur* peut vérifier s'il a réussi à écrire dans la source de données, il indique un succès inconditionnel.

L'ordre dans lequel les opérations sont traitées dans le *Serveur* n'est pas défini et dépend des différentes sources de données et de la logique interne du *Serveur*. Si une combinaison *Attribut* et *Nœud* est contenue dans plus d'une opération, l'ordre de traitement n'est pas défini. Si un *Client* exige un traitement séquentiel, le *Client* nécessite des appels de *Services* distincts.

Il est possible que le *Serveur* puisse réussir à écrire certains *Attributs* et pas d'autres. Le retour en arrière est du ressort du *Client*.

Si un *Serveur* autorise l'écriture d'*attributs* ayant le *DataType LocalizedText*, le *Client* peut ajouter ou écraser le texte pour un paramètre de lieu en écrivant le texte avec le *LocaleId* associé. Le fait d'écrire une *String* Null pour le texte pour un paramètre de lieu doit supprimer la *String* pour le paramètre de lieu en question. Le fait d'écrire une *String* Null pour le paramètre de lieu et une *String* non Null pour le texte met le texte à un paramètre de lieu invariable. Le fait d'écrire une *String* Null pour le texte et une *String* non Null pour le paramètre de lieu doit supprimer les entrées pour tous les paramètres de lieu. Si un *Client* tente d'écrire un paramètre de lieu qui est soit non valide syntaxiquement, soit non pris en charge, le *Serveur* renvoie *Bad_LocaleNotSupported*. Le comportement *Write* pour les *Attributs de valeur* avec un *LocalizedText DataType* est spécifique au *Serveur*, mais il est recommandé de suivre les mêmes règles.

5.10.4.2 Paramètres

Le Tableau 59 définit les paramètres du *Service*.

Tableau 59 – Paramètres du service Write

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
nodesToWrite []	WriteValue	Liste des <i>Nœuds</i> et de leurs <i>Attributs</i> à écrire. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
nodeId	NodeId	<i>NodeId</i> du <i>Nœud</i> qui contient les <i>Attributs</i> .
AttributeId	IntegerId	Identificateur de l' <i>Attribut</i> . Il doit s'agir d'un identificateur d' <i>attribut</i> valide. L' <i>IntegerId</i> est défini en 7.14. Les <i>IntegerIds</i> pour les <i>Attributs</i> sont définis dans l'IEC 62541-6.
indexRange	NumericRange	Ce paramètre est utilisé pour identifier un élément simple d'une matrice ou une plage simple d'indices pour les matrices. Le premier élément est identifié par l'indice 0 (zéro). Le type <i>NumericRange</i> est défini en 7.22. Ce paramètre n'est pas utilisé si l'<i>Attribut</i> spécifié n'est pas une matrice. Cependant, si l'<i>Attribut</i> spécifié est une matrice et que ce paramètre n'est pas utilisé, tous les éléments doivent être inclus dans la plage. Ce paramètre est Null s'il n'est pas utilisé. Un <i>Serveur</i> doit retourner une erreur <i>Bad_WriteNotSupported</i> si une <i>indexRange</i> est fournie et si l'écriture d'<i>indexRange</i> n'est pas possible pour le <i>Nœud</i>.
value	DataValue	Valeur d'<i>attribut</i> du <i>nœud</i> (voir 7.7 pour la définition de <i>DataValue</i>). Si le paramètre <i>indexRange</i> est spécifié, la <i>Value</i> doit être une matrice même si un seul élément est en cours d'écriture. Si le <i>SourceTimestamp</i> ou le <i>ServerTimestamp</i> est spécifié, le <i>Serveur</i> doit utiliser ces valeurs. Le <i>Serveur</i> retourne une erreur <i>Bad_WriteNotSupported</i> s'il ne prend pas en charge l'écriture d'horodatages. Un <i>Serveur</i> doit retourner une erreur <i>Bad_TypeMismatch</i> si le type de données de la valeur écrite n'est pas le même que le type ou le sous-type du <i>DataType</i> de l'<i>Attribut</i>. Sur la base de la hiérarchie du <i>DataType</i>, les sous-types du <i>DataType</i> de l'<i>Attribut</i> doivent être acceptés par le <i>Serveur</i>. Les <i>Serveurs</i> peuvent rejeter les sous-types définis dans les versions de spécifications plus récentes que celles prises en charge par le <i>Serveur</i> avec <i>Bad_TypeMismatch</i>. Pour l'<i>attribut Valeur</i>, le <i>DataType</i> est défini au moyen de l'<i>attribut "DataType"</i>. La structure d'une <i>ByteString</i> est la même que celle d'une matrice <i>Byte</i> à une dimension. Un <i>Serveur</i> doit accepter une <i>ByteString</i> si une matrice <i>Byte</i> est prévue. Le <i>Serveur</i> retourne une erreur <i>Bad_DataEncodingUnsupported</i> s'il ne prend pas en charge le codage de données fourni. Les <i>DataTypes Simples</i> (voir IEC 62541-3) utilisent la même représentation à la volée que leurs super types, donc écrire une valeur d'un <i>DataType</i> simple ne peut être distingué du fait d'écrire une valeur de son super type. Lorsqu'il reçoit la représentation filaire adéquate pour un <i>DataType</i> simple, le <i>Serveur</i> doit admettre par hypothèse que le type correct a été choisi. Les <i>Serveurs</i> peuvent imposer des validations supplémentaires de données sur la valeur indépendante du codage (par exemple, une image au format GIF dans une <i>ByteString</i>). Dans ce cas, le <i>Serveur</i> doit retourner une erreur <i>Bad_TypeMismatch</i> si la validation échoue.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	StatusCode	Liste des résultats pour les <i>Nœuds</i> à écrire (voir 7.34 pour la définition des <i>StatusCode</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToWrite</i> . Cette liste comporte une entrée pour chaque <i>Nœud</i> contenu dans le paramètre <i>nodesToWrite</i> .
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>Nœuds</i> à écrire (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>nodesToWrite</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.10.4.3 Résultats de service

Le Tableau 60 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 60 – Codes de résultats du service Write

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.10.4.4 StatusCodes

Le Tableau 61 définit les valeurs pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 61 – Codes de résultats de niveau opération Write

ID symbolique	Description
Good_CompletesAsynchronously	Voir Tableau 177 pour la description de ce code de résultat. La valeur a été écrite avec succès dans un système intermédiaire, mais le <i>Serveur</i> ne sait pas si la source de données a été correctement mise à jour.
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_AttributeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeNoData	Voir Tableau 178 pour la description de ce code de résultat.
Bad_WriteNotSupported	L'opération d'écriture demandée n'est pas prise en charge. Si un <i>Client</i> tente d'écrire une combinaison de valeur, code de statut ou horodatage alors que le <i>Serveur</i> ne prend pas en charge la combinaison demandée (qui peut être une grandeur unique, comme un simple horodatage), le <i>Serveur</i> ne doit pas effectuer d'écriture sur ce <i>Nœud</i> et doit retourner ce <i>StatusCode</i> pour ce <i>Nœud</i> . Il est aussi utilisé si l'écriture d'une <i>IndexRange</i> n'est pas prise en charge pour un <i>Nœud</i> .
Bad_NotWritable	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat. L'utilisateur courant n'a pas la permission d'écrire l'attribut.
Bad_OutOfRange	Voir Tableau 178 pour la description de ce code de résultat. Si un <i>Client</i> tente d'écrire une valeur hors de la plage valide, telle qu'une valeur non contenue dans le type de données d'énumération du <i>Nœud</i> , le <i>Serveur</i> doit retourner ce <i>StatusCode</i> pour ce <i>Nœud</i> .
Bad_TypeMismatch	Voir Tableau 178 pour la description de ce code de résultat.
Bad_DataEncodingUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NoCommunication	Voir Tableau 178 pour la description de ce code de résultat.
Bad_LocaleNotSupported	Le paramètre de lieu dans l'opération d'écriture demandée n'est pas pris en charge.

5.10.5 HistoryUpdate

5.10.5.1 Description

Ce *Service* est utilisé pour mettre à jour des valeurs ou *Événements* historiques d'un ou plusieurs *Nœuds*. Plusieurs paramètres de demande indiquent comment le *Serveur* doit mettre à jour une valeur ou un *événement* historique. Les actions valides sont Insérer, Remplacer ou Supprimer.

5.10.5.2 Paramètres

Le Tableau 62 définit les paramètres du *Service*.

Tableau 62 – Paramètres du service HistoryUpdate

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
historyUpdateDetails []	Paramètre extensible HistoryUpdate Details	Les informations définies pour la mise à jour. Le type de paramètre <i>HistoryUpdateDetails</i> est un type de paramètre extensible qui est défini formellement dans l'IEC 62541-11. Il spécifie les types des mises à jour d'historique qui peuvent être accomplies. Le type <i>ExtensibleParameter</i> est défini en 7.12.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	HistoryUpdate Result	Liste des résultats de mise à jour pour les informations de mise à jour d'historiques. La taille et l'ordre de la liste correspondent à la taille et à l'ordre des éléments d'informations du paramètre <i>historyUpdateDetails</i> spécifié dans la demande. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> pour la mise à jour du <i>Nœud</i> (voir 7.34 pour la définition de <i>StatusCode</i>).
operationResults []	StatusCode	Liste des <i>StatusCodes</i> pour les opérations à exécuter sur un <i>Nœud</i> . La taille et l'ordre de la liste correspondent à la taille et à l'ordre de toute liste définie par l'élément "details" communiqué dans un rapport par l'entrée de résultat.
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les opérations à exécuter sur un <i>Nœud</i> (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre de toute liste définie par l'élément "details" communiqué dans un rapport par l'entrée de <i>résultats</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les informations de mise à jour d'historiques. La taille et l'ordre de la liste correspondent à la taille et à l'ordre des éléments d'informations du paramètre <i>historyUpdateDetails</i> spécifié dans la demande. Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.10.5.3 Résultats de service

Le Tableau 63 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 63 – Codes de résultats du service HistoryUpdate

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.10.5.4 StatusCodes

Le Tableau 64 définit les valeurs pour les paramètres *statusCode* et *operationResults* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178. Les *StatusCodes* spécifiques à l'accès aux historiques sont définis dans l'IEC 62541-11.

Tableau 64 – Codes de résultat de niveau d'opération HistoryUpdate

ID symbolique	Description
Bad_NotWritable	Voir Tableau 178 pour la description de ce code de résultat.
Bad_HistoryOperationInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_HistoryOperationUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat. L'utilisateur courant n'a pas la permission de mettre à jour l'historique.

5.11 Jeu de services Method

5.11.1 Vue d'ensemble

Les *Méthodes* représentent les appels de fonction des *Objets*. Elles sont définies dans l'IEC 62541-3. Les *Méthodes* sont invoquées et retournent un résultat (avec ou sans succès) seulement après leur achèvement. Les temps d'exécution pour les *Méthodes* peuvent varier, selon la fonction qu'elles accomplissent.

Le *Jeu de services Method* définit les moyens pour invoquer des *Méthodes*. Une *Méthode* doit être un *composant* d'un *Objet*. La découverte est fournie par le biais des *services Browse* et *Query*. Les *Clients* trouvent les *Méthodes* prises en charge par un *Serveur* en explorant l'appartenance des *Références "Objets"* qui identifient les *Méthodes* prises en charge.

Parce que des *Méthodes* peuvent commander un certain aspect des opérations d'une installation, l'invocation de *Méthodes* peut dépendre des conditions d'environnement ou autres. Cela peut être particulièrement vrai lorsqu'une *Méthode* qui a terminé son exécution fait l'objet d'une nouvelle tentative d'invocation. Les conditions nécessaires pour invoquer la *Méthode* peuvent ne pas encore être retournées à l'état qui permet à la *Méthode* de redémarrer.

5.11.2 Call

5.11.2.1 Description

Ce *service* est utilisé pour appeler (invoquer) une liste de *Méthodes*.

Ce *Service* permet le passage d'arguments d'entrée et de sortie en direction/provenance d'une *Méthode*. Ces arguments sont définis par les *propriétés* de la *Méthode*.

Si la *Méthode* est invoquée dans le contexte d'une *Session* et que cette *Session* est terminée, les résultats de l'exécution de la *Méthode* ne peuvent pas être retournés au *Client* et sont éliminés. Cela est indépendant de la tâche réellement exécutée au niveau du *Serveur*.

L'ordre dans lequel les opérations sont traitées dans le *Serveur* n'est pas défini et dépend des différentes tâches et de la logique interne du *Serveur*. Si une *Méthode* est contenue dans plusieurs opérations, l'ordre de traitement n'est pas défini. Si un *Client* exige un traitement séquentiel, le *Client* nécessite des appels de *Services* distincts.

5.11.2.2 Paramètres

Le Tableau 65 définit les paramètres du *Service*.

Tableau 65 – Paramètres du service Call

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
methodsToCall []	CallMethodRequest	Liste des <i>Méthodes</i> à appeler. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
objectId	NodeId	Le <i>NodeId</i> doit être celui de l'<i>Objet</i> ou <i>ObjectType</i> sur lequel la <i>Méthode</i> est invoquée. Dans le cas d'un <i>ObjectType</i>, l'<i>ObjectType</i> ou le super type de l'<i>ObjectType</i> doit être à la source d'une Référence "HasComponent" (ou sous-type de Référence <i>HasComponent</i>) à la <i>Méthode</i> spécifiée dans <i>methodId</i>. Dans le cas d'un <i>Objet</i>, l'<i>Objet</i> ou l'<i>ObjectType</i> de l'<i>Objet</i> ou un super type de cet <i>ObjectType</i> doit être à la source d'une Référence "HasComponent" (ou sous-type de Référence <i>HasComponent</i>) à la <i>Méthode</i> spécifiée dans <i>methodId</i>. Voir l'IEC 62541-3 pour une description des <i>Objets</i> et de leurs <i>Méthodes</i>.
methodId	NodeId	<i>NodeId</i> de la <i>Méthode</i> à invoquer. Si l'<i>objectId</i> est le <i>NodeId</i> d'un <i>Objet</i>, il peut utiliser le <i>NodeId</i> d'une <i>Méthode</i> qui est la cible d'une Référence <i>HasComponent</i> de l'<i>ObjectType</i> de l'<i>Objet</i>.
inputArguments []	BaseDataType	Liste de valeurs des arguments d'entrée. Une liste vide indique qu'il n'y a pas d'arguments d'entrée. La taille et l'ordre de cette liste correspondent à la taille et à l'ordre des arguments d'entrée définis par la Propriété "InputArguments" de la <i>Méthode</i>. Le nom, une description et le type de données de chaque argument sont définis par la structure de l'<i>Argument</i> dans chaque élément de la Propriété "InputArguments" de la méthode.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	CallMethodResult	Résultat pour les appels de la <i>Méthode</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> de la <i>Méthode</i> exécutée dans le serveur. Ce <i>StatusCode</i> est mis au <i>StatusCode</i> Bad_InvalidArgument si au moins un argument d'entrée a enfreint une contrainte (par exemple type de données erroné, valeur hors plage). Ce <i>StatusCode</i> est mis à un <i>StatusCode</i> Bad si l'exécution de la <i>Méthode</i> a échoué dans le <i>Serveur</i> , par exemple sur la base d'une exception.
inputArgumentResults []	StatusCode	Liste de <i>StatusCodes</i> correspondant aux <i>inputArguments</i> . Cette liste est vide sauf si le résultat de niveau opération est Bad_InvalidArgument Si cette liste est remplie, elle a la même longueur que la liste des <i>inputArguments</i> .
inputArgumentDiagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic correspondant aux <i>inputArguments</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.
outputArguments []	BaseDataType	Liste de valeurs des arguments de sortie. Une liste vide indique qu'il n'y a pas d'arguments de sortie. La taille et l'ordre de cette liste correspondent à la taille et à l'ordre des arguments de sortie définis par la <i>Propriété</i> "OutputArguments" de la <i>Méthode</i> . Le nom, une description et le type de données de chaque argument sont définis par la structure de l' <i>Argument</i> dans chaque élément de la <i>Propriété</i> "OutputArguments" de la méthode
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour le <i>statusCode</i> des <i>résultats</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.11.2.3 Résultats de service

Le Tableau 66 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 66 – Codes de résultats du service Call

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.

5.11.2.4 StatusCodes

Le Tableau 67 définit les valeurs pour le paramètre *statusCode* et le Tableau 68 définit les valeurs pour le paramètre *inputArgumentResults* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Les fournisseurs de serveurs ou les spécifications d'accompagnement OPC UA peuvent réutiliser les *StatusCodes* existants pour les informations d'erreurs spécifiques à l'application. Cela est valable tant que la description canonique du *StatusCode* n'a pas une signification différente de celle de la description spécifique de l'application. Afin d'éliminer toute

imprécision, il convient que le *Serveur* inclue la description spécifique de l'application dans le *DiagnosticInfo*.

Les *StatusCodes* Good avec sous-états ne doivent pas être utilisés comme des *statusCodes*, car les liaisons de langage de programmation trop nombreuses risquent d'entraîner la génération d'une exception par de tels codes.

Tableau 67 – Codes de résultats de niveau opération Call

ID symbolique	Description
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat. Sert à indiquer que l' <i>Objet</i> spécifié n'est pas valide.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat. Sert à indiquer que l' <i>Objet</i> spécifié n'est pas valide.
Bad_NotExecutable	L' <i>Attribut</i> exécutable ne permet pas l'exécution de la <i>Méthode</i> .
Bad_ArgumentsMissing	Le <i>Client</i> n'a pas spécifié la totalité des arguments d'entrée pour la <i>Méthode</i> .
Bad_TooManyArguments	Le <i>Client</i> a spécifié plus d'arguments d'entrée que le nombre défini pour la <i>Méthode</i> .
Bad_InvalidArgument	Voir Tableau 177 pour la description de ce code de résultat. Sert à indiquer dans les résultats de niveau opération qu'un ou plusieurs des arguments d'entrée ne sont pas valides. Les <i>inputArgumentResults</i> contiennent le code de statut spécifique pour chaque argument non valide.
Bad_UserAccessDenied	Voir Tableau 177 pour la description de ce code de résultat.
Bad_SecurityModelInsufficient	Voir Tableau 178 pour la description de ce code de résultat.
Bad_MethodInvalid	L'identificateur de la méthode ne désigne pas une <i>Méthode</i> pour l' <i>Objet</i> spécifié.
Bad_NoCommunication	Voir Tableau 178 pour la description de ce code de résultat.

Tableau 68 – Codes de résultats des arguments d'entrée "Call Input"

ID symbolique	Description
Bad_OutOfRange	Voir Tableau 178 pour la description de ce code de résultat. Sert à indiquer qu'un argument d'entrée se trouve à l'extérieur de la plage acceptable.
Bad_TypeMismatch	Voir Tableau 178 pour la description de ce code de résultat. Sert à indiquer qu'un argument d'entrée n'a pas le type de données correct. La structure d'une <i>ByteString</i> est la même que celle d'une matrice Byte à une dimension. Un <i>Serveur</i> doit accepter une <i>ByteString</i> si une matrice Byte est prévue.

5.12 Jeu de services MonitoredItem


5.12.1 Modèle MonitoredItem

5.12.1.1 Vue d'ensemble

Les *Clients* définissent des *MonitoredItems* pour s'abonner à des données et à des *Événements*. Chaque *MonitoredItem* identifie l'élément à surveiller et l'*Abonnement* à utiliser pour envoyer des *Notifications*. L'élément à surveiller peut être n'importe quel *Attribut Nœud*.

Les *Notifications* sont des structures de données qui décrivent l'apparition des modifications de données et d'*Événements*. Elles sont empaquetées dans des *NotificationMessages* pour le transfert au *Client*. L'*Abonnement* envoie régulièrement des *NotificationMessages* selon un intervalle d'édition spécifié par l'utilisateur, et le cycle durant lequel ces messages sont envoyés est appelé cycle d'édition.

Quatre paramètres primaires sont définis pour les *MonitoredItems* qui indiquent au *Serveur* comment l'élément doit être échantillonné, évalués et rapportés. Ces paramètres sont l'intervalle d'échantillonnage, le mode de surveillance, le filtre et le paramètre de mise en file d'attente. La Figure 15 représente ces concepts.



IEC

Figure 15 – Modèle MonitoredItem

Les *Attributs* autres que l'*attribut Valeur* ne sont surveillés que pour détecter une modification de valeur. Le filtre n'est pas utilisé pour ces *Attributs*. Toute modification de valeur pour ces *Attributs* entraîne la création d'une *Notification*.

L'*attribut Valeur* est utilisé pour surveiller des *Variables*. Les valeurs des *Variables* sont surveillées pour détecter une modification de valeur ou une modification de leur statut. Les filtres définis dans le présent document (voir 7.17.2) et dans l'IEC 62541-8 sont utilisés pour déterminer si la modification de la valeur est suffisamment forte pour entraîner la création d'une *Notification* pour cette *Variable*.

Les *Objets* et les vues peuvent être utilisés pour surveiller les *Événements*. Les *événements* sont disponibles uniquement à partir des *Nœuds* pour lesquels le bit *SubscribeToEvents* de l'*Attribut "EventNotifier"* est défini. Le filtre défini dans le présent document (voir 7.17.3) est utilisé pour déterminer si un *Événement* reçu du *Nœud* est envoyé au *Client*. Le filtre permet aussi de choisir les champs de l'*EventType* qui sont contenus dans l'*Événement*, comme *EventId*, *EventType*, *SourceNode*, *Time* et *Description*.

L'IEC 62541-3 décrit le modèle *Événement* et les *EventTypes* de base.

Les *Propriétés* des *EventTypes* de base et la représentation des *EventTypes* de base dans l'*AddressSpace* sont spécifiées dans l'IEC 62541-5.

5.12.1.2 Intervalle d'échantillonnage

Chaque *MonitoredItem* créé par le *Client* reçoit un intervalle d'échantillonnage qui lui est affecté et qui est soit hérité de l'intervalle d'édition de l'*Abonnement*, soit défini spécifiquement pour se substituer à cette fréquence. Un nombre négatif indique que l'intervalle d'échantillonnage par défaut défini par l'intervalle d'édition de l'*Abonnement* est demandé. L'intervalle d'échantillonnage indique la fréquence la plus rapide à laquelle il convient que le *Serveur* échantillonne sa source sous-jacente pour les modifications de données.

Un *Client* doit définir un intervalle d'échantillonnage de 0 s'il s'abonne à des *Événements*.

L'intervalle d'échantillonnage affecté définit une fréquence cyclique de "meilleur effort" que le *Serveur* utilise pour échantillonner l'élément à partir de sa source. "Meilleur effort" signifie

dans ce contexte que le *Serveur* fait de son mieux pour échantillonner à cette fréquence. L'échantillonnage à des fréquences plus rapides que celle-ci est acceptable, mais n'est pas nécessaire pour répondre aux besoins du *Client*. La façon dont le *Serveur* traite la fréquence d'échantillonnage et la fréquence à laquelle il sonde effectivement ses sources de données en interne sont un détail d'implémentation du *Serveur*. Cependant, le temps entre les valeurs retournées au *Client* doit être supérieur ou égal à l'intervalle d'échantillonnage.

Le *Client* peut aussi spécifier 0 pour l'intervalle d'échantillonnage, qui indique qu'il convient pour le *Serveur* d'utiliser la fréquence pratique la plus rapide. Les *Serveurs* sont supposés prendre en charge uniquement un jeu limité d'intervalles d'échantillonnage pour optimiser leur fonctionnement. Si l'intervalle exact demandé par le *Client* n'est pas pris en charge par le *Serveur*, le *Serveur* affecte au *MonitoredItem* l'intervalle qu'il a déterminé comme étant le plus approprié. Il retourne au *Client* cet intervalle affecté. L'*Objet* "Capacités" *Serveur* défini dans l'IEC 62541-5 identifie les intervalles d'échantillonnage pris en charge par le *Serveur*.

Le *Serveur* peut prendre en charge les données qui sont recueillies sur la base d'un modèle d'échantillonnage ou générées sur la base d'un modèle piloté par des exceptions. L'intervalle d'échantillonnage le plus rapide pris en charge peut être égal à 0, ce qui indique que l'élément de donnée est piloté par des exceptions au lieu d'être échantillonné régulièrement. Un modèle piloté par des exceptions signifie que le système sous-jacent ne nécessite pas d'échantillonnage et rapporte les modifications apportées aux données.

Le *Client* peut utiliser des valeurs d'intervalle d'échantillonnage révisées comme suggestion pour l'établissement de l'intervalle d'édition ainsi que le compte keep-alive d'un *Abonnement*. Si, par exemple, le plus faible intervalle d'échantillonnage révisé pour les *MonitoredItems* est de cinq secondes, il convient que le temps avant l'émission d'un message keep-alive soit supérieur à cinq secondes.

Noter que, dans de nombreux cas, le *Serveur* OPC UA fournit l'accès à un système découplé et n'a donc aucune connaissance de la logique de mise à jour de données. Dans ce cas, même si le *Serveur* OPC UA effectue l'échantillonnage à la fréquence négociée, les données peuvent être mises à jour par le système sous-jacent à une fréquence nettement inférieure. Dans ce cas, les modifications ne peuvent être détectées qu'à cette fréquence plus lente.

Si le comportement par lequel le système sous-jacent met à jour l'élément est connu, il est disponible via l'*attribut* "MinimumSamplingInterval" défini dans l'IEC 62541-3. Si le *Serveur* spécifie une valeur pour l'*attribut* *MinimumSamplingInterval*, il doit toujours retourner un *revisedSamplingInterval* qui est supérieur ou égal au *MinimumSamplingInterval*, si le *Client* souscrit à l'*attribut* *Valeur*.

Il convient que les *Clients* sachent que l'échantillonnage par le *Serveur* OPC UA et le cycle de mise à jour du système sous-jacent ne sont pas, en général, synchronisés. Cela peut introduire des retards supplémentaires dans la détection des modifications, comme représenté à la Figure 16.




Figure 16 – Retard type pour détecter des modifications

5.12.1.3 Mode de surveillance

Le paramètre mode de surveillance est utilisé pour activer et désactiver l'échantillonnage d'un *MonitoredItem*, et aussi pour permettre d'activer et de désactiver de façon indépendante la production de rapports de *Notifications*. Cette fonction permet de configurer un *MonitoredItem* pour échantillonner, échantillonner et rapporter, ou ni l'un ni l'autre. La désactivation de l'échantillonnage ne modifie les valeurs d'aucun des autres paramètres de *MonitoredItem*, pas même son intervalle d'échantillonnage.

Lorsqu'un *MonitoredItem* est activé (c'est-à-dire lorsque le *MonitoringMode* est passé de *DISABLED* à *SAMPLING* ou à *REPORTING*) ou qu'il est créé à l'état activé, le *Serveur* doit signaler le premier échantillon dès que possible dans le rapport et l'heure de cet échantillon devient le point de départ pour le prochain intervalle d'échantillonnage.

5.12.1.4 Filtre

Chaque fois qu'un *MonitoredItem* est échantillonné, le *Serveur* évalue l'échantillon en utilisant le filtre défini pour le *MonitoredItem*. Le paramètre de filtrage définit les critères que le *Serveur* utilise pour déterminer s'il convient de générer une *Notification* pour l'échantillon. Le type de filtre dépend du type d'élément qui est surveillé. Par exemple, le *DataChangeFilter* et l'*AggregateFilter* sont utilisés pour surveiller les *Valeurs de variables* et l'*EventFilter* est utilisé pour surveiller les *Événements*. L'échantillonnage et l'évaluation, y compris l'utilisation de filtres, sont décrits dans le présent document. Des filtres supplémentaires peuvent être définis dans d'autres parties de l'IEC 62541.

5.12.1.5 Paramètres de mise en file d'attente

Si l'échantillon répond aux critères de filtrage, une *Notification* est générée et mise en file d'attente pour le transfert par l'*Abonnement*. La taille de la file d'attente est définie lorsque *MonitoredItem* est créé. Lorsque la file d'attente est pleine et qu'une nouvelle *Notification* est reçue, soit le *Serveur* élimine la plus ancienne *Notification* et place la nouvelle dans la file d'attente, soit il remplace la dernière valeur ajoutée à la file d'attente par la nouvelle valeur. A sa création, le *MonitoredItem* est configuré pour l'une de ces politiques de rejet. Si une *Notification* est rejetée pour une *DataValue* et que la taille de la file d'attente est supérieure à un, le bit (fanion) *Overflow* dans la partie *InfoBits* du *statusCode* de la *DataValue* est établi. Si *discardOldest* est TRUE, la valeur la plus ancienne est supprimée de la file d'attente et la

valeur suivante dans la file d'attente établit le fanion. Si *discardOldest* est FALSE, la dernière valeur ajoutée à la file d'attente est remplacée par la nouvelle valeur. La nouvelle valeur établit le fanion pour indiquer les valeurs perdues dans le *NotificationMessage* suivant. La Figure 17 représente la gestion du dépassement de capacité de la file d'attente.




Figure 17 – Gestion du dépassement de capacité de la file d'attente

Si la taille de la file d'attente est 1, la file d'attente devient un tampon qui contient toujours la *Notification* la plus récente. Dans ce cas, si l'intervalle d'échantillonnage du *MonitoredItem* est inférieur à l'intervalle d'édition de l'*Abonnement*, le *MonitoredItem* est en suréchantillonnage et le *Client* reçoit toujours la valeur la plus récente. La politique de rejet est ignorée si la taille de la file d'attente est un.

Par ailleurs, le *Client* peut souhaiter s'abonner à un flux continu de *Notifications* sans le moindre trou, mais ne souhaite pas que celles-ci soient rapportées à l'intervalle d'échantillonnage. Dans ce cas, le *MonitoredItem* met en place une file d'attente d'une longueur suffisante pour contenir toutes les *Notifications* générées entre deux cycles d'édition consécutifs. Puis, à chaque cycle d'édition, l'*Abonnement* envoie au *Client* toutes les *Notifications* mises en file d'attente pour le *MonitoredItem*. Le *Serveur* doit retourner les *Notifications* pour tout élément particulier dans l'ordre dans lequel elles sont stockées dans la file d'attente.

Le *Serveur* peut échantillonner à une fréquence plus rapide que l'intervalle d'échantillonnage pour prendre en charge d'autres *Clients*; il convient que le *Client* attende seulement les valeurs à l'intervalle d'échantillonnage négocié. Le *Serveur* peut délivrer moins de valeurs que ne le dicte l'intervalle d'échantillonnage, sur la base des contraintes du filtre et de l'implémentation. Si un *DataChangeFilter* est configuré pour un *MonitoredItem*, il est toujours appliqué à la valeur la plus récente dans la file d'attente en comparaison à l'échantillon courant.

Si, par exemple, l'*AbsoluteDeadband* dans le *DataChangeFilter* est "10", la file d'attente peut être constituée de valeurs dans l'ordre suivant:

- 100
- 111
- 100

- 89
- 100

La mise en file d'attente des données peut conduire à un comportement inattendu lorsqu'un filtre *Deadband* est utilisé et que le nombre de modifications rencontrées est supérieur au nombre de valeurs pouvant être gérées. La nouvelle première valeur de la file d'attente peut ne pas dépasser la limite de *Deadband* de la valeur précédente envoyée au *Client*.

La taille de la file d'attente est la valeur maximale prise en charge par le *Serveur* lors de la surveillance d'*Événements*. Dans ce cas, le *Serveur* a la responsabilité du tampon d'*Événements*. Si des *Événements* sont perdus, un *Événement* du type *EventQueueOverflowEventType* est placé dans la file d'attente. Cet *Événement* est généré lorsque le premier *Événement* est éliminé sur un *MonitoredItem* s'abonnant à des *Événements*. Il est mis dans la File d'attente du *MonitoredItem* en plus de la taille de la File d'attente définie pour ce *MonitoredItem* sans rejeter aucun autre *Événement*. Si *discardOldest* est mis à TRUE, il est placé au début de la file d'attente et n'est jamais rejeté, sinon à la fin. Un *Serveur* de regroupement ne doit pas transmettre un tel *événement*. Il doit être géré comme les autres scénarios d'erreur de connexion.

5.12.1.6 Modèle de déclenchement

Le *Service MonitoredItems* permet d'ajouter des éléments qui sont signalés dans un rapport seulement lorsqu'un autre élément est déclenché (l'élément de déclenchement). Cela est réalisé en créant des liens entre les éléments déclenchés et les éléments à rapporter. Le mode de surveillance des éléments à rapporter est mis en échantillonnage seul. De ce fait, il échantillonne et met les *Notifications* en file d'attente sans les signaler dans un rapport. La Figure 18 représente ce concept.




Figure 18 – Modèle de déclenchement

Le mécanisme de déclenchement est une caractéristique utile qui permet aux *Clients* de réduire le volume de données à la volée en configurant certains éléments pour qu'ils échantillonnent fréquemment, mais se limitent seulement à produire un rapport lorsqu'un autre *Événement* se produit.

Les comportements de déclenchement suivants sont spécifiés.

- Si le mode de surveillance de l'élément déclencheur est *SAMPLING*, lorsque l'élément déclencheur déclenche les éléments à rapporter, cela ne fait pas l'objet d'un rapport.
- Si le mode de surveillance de l'élément déclencheur est *REPORTING*, lorsque l'élément déclencheur déclenche les éléments à rapporter, cela fait l'objet d'un rapport.
- Si le mode de surveillance de l'élément déclencheur est *DISABLED*, alors l'élément déclencheur ne déclenche pas les éléments à rapporter.
- Si le mode de surveillance de l'élément à rapporter est *SAMPLING*, lorsque l'élément déclencheur déclenche les éléments à rapporter, cela fait l'objet d'un rapport.

- e) Si le mode de surveillance de l'élément à rapporter est *REPORTING*, l'élément déclencheur est ignoré. Toutes les notifications des éléments à rapporter sont envoyées à l'expiration de l'intervalle d'édition.
- f) Si le mode de surveillance de l'élément à rapporter est *DISABLED*, alors il n'y a pas d'échantillonnage de l'élément à rapporter et donc aucune notification à rapporter.
- g) Le premier déclenchement doit se produire lorsque la première notification est mise en file d'attente pour l'élément déclencheur après la création de la liaison.

Les *Clients* créent et suppriment des liaisons de déclenchement entre un élément déclencheur et un ensemble d'éléments à rapporter. Si le *MonitoredItem* qui représente un élément à rapporter est supprimé avant que sa liaison associée ne soit supprimée, la liaison de déclenchement est également supprimée, mais l'élément déclencheur n'est pas altéré.

Il convient de ne pas confondre la suppression d'un *MonitoredItem* avec le retrait de l'*Attribut* qu'il surveille. Si le *Nœud* qui contient l'*attribut* surveillé est supprimé, le *MonitoredItem* génère une *Notification* avec un *StatusCode* *Bad_NodeIdUnknown* qui indique la suppression, mais le *MonitoredItem* n'est pas supprimé.

5.12.2 CreateMonitoredItems

5.12.2.1 Description

Ce *Service* est utilisé pour créer et ajouter un ou plusieurs *MonitoredItems* à un *Abonnement*. Un *MonitoredItem* est supprimé automatiquement par le *Serveur* lorsque l'*Abonnement* est supprimé. La suppression d'un *MonitoredItem* entraîne celle de l'ensemble complet des liaisons d'éléments déclenchés, mais n'a pas d'effet sur les *MonitoredItems* référencés par les éléments déclenchés.

L'appel à répétitions du *service* *CreateMonitoredItems* pour ajouter chaque fois une faible quantité de *MonitoredItems* peut affecter défavorablement la performance du *Serveur*. Il convient que les *Clients* ajoutent plutôt un jeu complet de *MonitoredItems* à un *Abonnement* chaque fois que c'est possible.

Lorsqu'un *MonitoredItem* est ajouté, le *Serveur* effectue le processus d'initialisation à cet effet. Le processus d'initialisation est défini par le type de *Notification* de l'élément surveillé. Les types de *Notifications* sont spécifiés dans le présent document et dans les parties de l'IEC 62541 qui spécifient les types d'accès, comme l'IEC 62541-8. Voir IEC TR 62541-1 pour une description des types d'accès.

Lorsqu'un utilisateur ajoute un élément surveillé dont l'accès en lecture est refusé à l'utilisateur, l'opération d'ajout de l'élément doit réussir et le mauvais statut *Bad_NotReadable* ou *Bad_UserAccessDenied* doit être retourné dans la réponse *Publish*. Il s'agit du même comportement que dans le cas où les droits d'accès sont modifiés après l'appel à *CreateMonitoredItems*. Si les droits d'accès passent aux droits de lecture, le *Serveur* doit commencer à envoyer des données pour le *MonitoredItem*. La même procédure doit être appliquée pour un *IndexRange* qui ne fournit pas de données pour la valeur courante, mais qui peut fournir des données plus tard.

Les *Nœuds* surveillés peuvent être retirés de l'*AddressSpace* après la création d'un *MonitoredItem*. Cela n'affecte pas la validité du *MonitoredItem*, mais un *Bad_NodeIdUnknown* doit être retourné dans la réponse *Publish*. Il est possible que le *MonitoredItem* redevienne valable si le *Nœud* est à nouveau ajouté à l'*AddressSpace* et que le *MonitoredItem* existe encore.

Si un *NodeId* est considéré valide par un *Serveur*, mais que les *Attributs de nœuds* correspondants sont indisponibles actuellement, le *serveur* peut permettre la création d'un *MonitoredItem* et retourner un *StatusCode* *Bad* dans la réponse *Publish*.

Le réglage des informations de diagnostic de retour dans l'en-tête de demande de *CreateMonitoredItems* ou le dernier *service ModifyMonitoredItems* est appliqué aux *éléments surveillés* et sert de réglages d'informations de diagnostic lors de l'envoi de Notifications dans la réponse *Publish*.

Les valeurs de demande illégales pour les paramètres pouvant être révisés ne génèrent pas d'erreur. En revanche, le *Serveur* choisit des valeurs par défaut et les indique dans le paramètre révisé correspondant.

Il est fortement recommandé par l'OPC UA qu'un *Client* réutilise un *Abonnement* après une courte interruption du réseau en activant la *Session* existante sur un nouveau *SecureChannel* comme décrit en 6.7. Si un *Client* a appelé *CreateMonitoredItems* pendant l'interruption du réseau et si l'appel a réussi dans le *Serveur*, mais n'est pas retourné au *Client*, alors le *Client* ne sait pas si l'appel a réussi. Le *Client* peut recevoir des modifications de données pour ces éléments surveillés, mais ne peut pas les retirer, car il ne connaît pas l'identifiant du *Serveur* pour chaque élément surveillé. Il n'est pas non plus possible pour le *Client* de détecter si la création a réussi. Pour supprimer et recréer, utiliser l'*Abonnement* n'est pas non plus possible, car il peut y avoir plusieurs éléments surveillés fonctionnant normalement qu'il convient de ne pas interrompre. Pour résoudre cette situation, l'*Objet "Serveur"* fournit une *Méthode GetMonitoredItems* qui retourne la liste des identifiants du *Client* et du *Serveur* pour les éléments surveillés dans un *Abonnement*. Cette *Méthode* est définie dans l'IEC 62541-5. Le *Serveur* doit vérifier que la *Méthode* est appelée dans le contexte de *Session* de la *Session* qui détient l'*Abonnement*.

5.12.2.2 Paramètres

Le Tableau 69 définit les paramètres du *Service*.

Tableau 69 – Paramètres du service CreateMonitoredItems

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
subscriptionId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour l' <i>Abonnement</i> qui signale dans un rapport les <i>Notifications</i> pour ce <i>MonitoredItem</i> (voir 7.14 pour la définition d' <i>IntegerId</i>).
timestampsToReturn	Enum TimestampsToReturn	Enumération qui spécifie les <i>Attributs</i> d'horodatage à transmettre pour chaque <i>MonitoredItem</i> . L'énumération <i>TimestampsToReturn</i> est définie en 7.35. Lors de la surveillance d' <i>Evénements</i> , cela s'applique seulement aux champs <i>Evénement</i> qui sont du type <i>DataValue</i> .
itemsToCreate []	MonitoredItem "CreateRequest"	Liste des <i>MonitoredItems</i> à créer et à affecter à l' <i>Abonnement</i> spécifié. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
itemToMonitor	ReadValueId	Identifie un élément dans l' <i>AddressSpace</i> à surveiller. Pour surveiller en vue de détecter des <i>Evénements</i> , l'élément <i>attributId</i> de la structure <i>ReadValueId</i> est l'identificateur de l' <i>Attribut EventNotifier</i> . Le type <i>ReadValueId</i> est défini en 7.24.
monitoringMode	Enum MonitoringMode	Mode de surveillance à établir pour le <i>MonitoredItem</i> . L'énumération <i>MonitoringMode</i> est définie en 7.18.
requestedParameters	Paramètres de surveillance	Paramètres de surveillance demandés. Les <i>Serveurs</i> négocient les valeurs de ces paramètres sur la base de l' <i>Abonnement</i> et des fonctions du <i>Serveur</i> . Le type <i>MonitoringParameters</i> est défini en 7.16.

Nom	Type	Description
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	MonitoredItem "CreateResult"	Liste de résultats pour les <i>MonitoredItems</i> à créer. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>itemsToCreate</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> pour le <i>MonitoredItem</i> à créer (voir 7.34 pour la définition de <i>StatusCode</i>).
monitoredItemId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour le <i>MonitoredItem</i> (voir 7.14 pour la définition d' <i>IntegerId</i>). Cet identificateur est unique dans l' <i>Abonnement</i> , mais peut ne pas être à l'échelle du <i>Serveur</i> ou de la <i>Session</i> . Ce paramètre est présent seulement si le <i>statusCode</i> indique que le <i>MonitoredItem</i> a été créé avec succès.
revisedSampling Intervalle	Duration	Intervalle d'échantillonnage réel utilisé par le <i>Serveur</i> . Cette valeur est fondée sur plusieurs facteurs, y compris les fonctions du système sous-jacent. Le <i>Serveur</i> doit toujours retourner un <i>revisedSamplingInterval</i> supérieur ou égal au <i>samplingInterval</i> demandé. Si le <i>samplingInterval</i> demandé est supérieur à l'intervalle d'échantillonnage maximal pris en charge par le <i>Serveur</i> , l'intervalle d'échantillonnage maximal est retourné.
revisedQueueSize	Counter	Taille réelle de file d'attente que le <i>Serveur</i> utilise.
filterResult	Paramètre extensible MonitoringFilterRe sult	Contient les éventuelles valeurs de paramètre révisées ou tous les éventuels résultats d'erreur associés au <i>MonitoringFilter</i> spécifié dans <i>requestedParameters</i> . Ce paramètre peut être nul s'il ne s'est pas produit d'erreurs. Le type de paramètre <i>MonitoringFilterResult</i> est un type de paramètre extensible spécifié en 7.17.
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>MonitoredItems</i> à créer (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>itemsToCreate</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.12.2.3 Résultats de service

Le Tableau 70 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 70 – Codes de résultats du service CreateMonitoredItems

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TimestampsToReturnInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_SubscriptionIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.

5.12.2.4 StatusCodes

Le Tableau 71 définit les valeurs de niveau opération pour le paramètre *statusCode* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 71 – Codes de résultats de niveau opération CreateMonitoredItems

ID symbolique	Description
Bad_MonitoringModelInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_NodeIdUnknown	Voir Tableau 178 pour la description de ce code de résultat.
Bad_AttributeIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_IndexRangeNoData	Voir Tableau 178 pour la description de ce code de résultat. Si les <i>ArrayDimensions</i> ont une longueur fixe qui ne peut être modifiée et qu'aucune donnée n'existe dans la plage d'indices spécifiée, <i>Bad_IndexRangeNoData</i> est retourné dans <i>CreateMonitoredItems</i> . Sinon, si la longueur de la matrice est dynamique, le <i>Serveur</i> doit retourner ce statut dans une réponse <i>Publish</i> pour le <i>MonitoredItem</i> si aucune donnée n'existe dans la plage.
Bad_DataEncodingInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_DataEncodingUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_MonitoredItemFilterInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_MonitoredItemFilterUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_FilterNotAllowed	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyMonitoredItems	Le <i>Serveur</i> a atteint son nombre maximal d'éléments surveillés.

5.12.3 ModifyMonitoredItems

5.12.3.1 Description

Ce *Service* est utilisé pour modifier les *MonitoredItems* d'un *Abonnement*. Les modifications apportées aux réglages de *MonitoredItem* doivent être appliquées immédiatement par le *Serveur*. Elles prennent effet le plus tôt possible, sans toutefois dépasser une durée égale au double du nouveau *revisedSamplingInterval*.

Le réglage des informations de diagnostic de retour dans l'en-tête de demande de *CreateMonitoredItems* ou le dernier *service ModifyMonitoredItems* est appliqué aux *éléments surveillés* et sert de réglages d'informations de diagnostic lors de l'envoi de Notifications dans la réponse *Publish*.

Les valeurs de demande illégales pour les paramètres pouvant être révisés ne génèrent pas d'erreur. En revanche, le *Serveur* choisit des valeurs par défaut et les indique dans le paramètre révisé correspondant.

5.12.3.2 Paramètres

Le Tableau 72 définit les paramètres du *Service*.

Tableau 72 – Paramètres du service ModifyMonitoredItems

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
subscriptionId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour l' <i>Abonnement</i> et qui est utilisé pour qualifier le <i>monitoredItemId</i> (voir 7.14 pour la définition d' <i>IntegerId</i>).
timestampsToReturn	Enum TimestampsToReturn	Énumération qui spécifie les <i>Attributs</i> d'horodatage à transmettre pour chaque <i>MonitoredItem</i> à modifier. L'énumération <i>TimestampsToReturn</i> est définie en 7.35. Lors de la surveillance d' <i>Événements</i> , cela s'applique seulement aux champs <i>Événement</i> qui sont du type <i>DataValue</i> .
itemsToModify []	MonitoredItemModifyRequest	Liste des <i>MonitoredItems</i> à modifier. Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
monitoredItemId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour le <i>MonitoredItem</i> .
requestedParameters	Paramètres de surveillance	Valeurs demandées pour les paramètres de surveillance. Le type <i>MonitoringParameters</i> est défini en 7.16. Si le nombre de notifications dans la file d'attente dépasse la taille de la nouvelle file d'attente, les notifications en excédent doivent être rejetées selon la politique de rejet configurée.
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	MonitoredItemModifyResult	Liste de résultats pour les <i>MonitoredItems</i> à modifier. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>itemsToModify</i> . Cette structure est définie de façon cohérente par rapport aux éléments prévus suivants.
statusCode	StatusCode	<i>StatusCode</i> pour le <i>MonitoredItem</i> à modifier (voir 7.34 pour la définition de <i>StatusCode</i>).
revisedSampling Intervalle	Duration	Intervalle d'échantillonnage réel utilisé par le <i>Serveur</i> . Le <i>Serveur</i> retourne la valeur qu'il utilisera réellement pour l'intervalle d'échantillonnage. Cette valeur est fondée sur plusieurs facteurs, y compris les fonctions du système sous-jacent. Le <i>Serveur</i> doit toujours retourner un <i>revisedSamplingInterval</i> supérieur ou égal au <i>samplingInterval</i> demandé. Si le <i>samplingInterval</i> demandé est supérieur à l'intervalle d'échantillonnage maximal pris en charge par le <i>Serveur</i> , l'intervalle d'échantillonnage maximal est retourné.
revisedQueueSize	Counter	Taille réelle de file d'attente que le <i>Serveur</i> utilise.
filterResult	Paramètre extensible Résultat du MonitoringFilter	Contient les éventuelles valeurs de paramètre révisées ou tous les éventuels résultats d'erreurs associés au <i>MonitoringFilter</i> spécifié dans la demande. Ce paramètre peut être nul s'il ne s'est pas produit d'erreurs. Le type de paramètre <i>MonitoringFilterResult</i> est un type de paramètre extensible spécifié en 7.17.
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>MonitoredItems</i> à modifier (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>itemsToModify</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.12.3.3 Résultats de service

Le Tableau 73 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 73 – Codes de résultats du service ModifyMonitoredItems

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TimestampsToReturnInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_SubscriptionIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.

5.12.3.4 StatusCodes

Le Tableau 74 définit les valeurs de niveau opération pour le paramètre *statusCode* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 74 – Codes de résultats de niveau opération ModifyMonitoredItems

ID symbolique	Description
Bad_MonitoredItemIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_MonitoredItemFilterInvalid	Voir Tableau 178 pour la description de ce code de résultat.
Bad_MonitoredItemFilterUnsupported	Voir Tableau 178 pour la description de ce code de résultat.
Bad_FilterNotAllowed	Voir Tableau 177 pour la description de ce code de résultat.

5.12.4 SetMonitoringMode

5.12.4.1 Description

Ce *Service* est utilisé pour établir le mode de surveillance pour un ou plusieurs *MonitoredItems* d'un *Abonnement*. Le fait de mettre le mode à DISABLED entraîne la suppression de toutes les *Notifications* placées en file d'attente.

5.12.4.2 Paramètres

Le Tableau 75 définit les paramètres du *Service*.

Tableau 75 – Paramètres du service SetMonitoringMode

Nom	Type	Description
Demande		
requestHeader	RequestHeader	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
subscriptionId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour l' <i>Abonnement</i> et qui est utilisé pour qualifier les <i>monitoredItemIds</i> (voir 7.14 pour la définition d' <i>IntegerId</i>).
monitoringMode	Enum MonitoringMode	Mode de surveillance à établir pour les <i>MonitoredItems</i> . L'énumération <i>MonitoringMode</i> est définie en 7.18.
monitoredItemIds []	IntegerId	Liste des identificateurs affectés par le <i>Serveur</i> pour les <i>MonitoredItems</i> .
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
results []	StatusCode	Liste des <i>StatusCodes</i> pour les <i>MonitoredItems</i> à activer/désactiver (voir 7.34 pour la définition des <i>StatusCode</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>monitoredItemIds</i> .
diagnosticInfos []	DiagnosticInfo	Liste d'informations de diagnostic pour les <i>MonitoredItems</i> à activer/désactiver (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>monitoredItemIds</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.12.4.3 Résultats de service

Le Tableau 76 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 177.

Tableau 76 – Codes de résultats du service SetMonitoringMode

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_SubscriptionIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_MonitoringModeInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.12.4.4 StatusCodes

Le Tableau 77 définit les valeurs de niveau opération pour le paramètre *results* qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 77 – Codes de résultats de niveau opération SetMonitoringMode

Valeur	Description
Bad_MonitoredItemIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.12.5 SetTriggering

5.12.5.1 Description

Ce *Service* est utilisé pour créer et supprimer des liaisons de déclenchement pour un élément déclencheur. L'élément déclencheur et les éléments à rapporter doivent appartenir au même *Abonnement*.

Chaque liaison de déclenchement relie un élément déclencheur à un élément à rapporter. Chaque liaison est représentée par l'identificateur de *MonitoredItem* pour l'élément à rapporter. Un code d'erreur est retourné si cet identificateur n'est pas valide.

Voir 5.12.1.6 pour une description du modèle de déclenchement.

5.12.5.2 Paramètres

Le Tableau 78 définit les paramètres du *Service*.

Tableau 78 – Paramètres du service SetTriggering

Nom	Type	Description
Demande		
requestHeader	Request Header	Paramètres de demande communs (voir 7.28 pour la définition de <i>RequestHeader</i>).
subscriptionId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour l' <i>Abonnement</i> qui contient l'élément déclencheur et les éléments à rapporter (voir 7.14 pour la définition d' <i>IntegerId</i>).
triggeringItemId	IntegerId	Identificateur affecté par le <i>Serveur</i> pour le <i>MonitoredItem</i> utilisé comme l'élément déclencheur.
linksToAdd []	IntegerId	Liste des identificateurs affectés par le <i>Serveur</i> aux éléments à rapporter et qui doivent être ajoutés comme liaisons de déclenchement. La liste de <i>linksToRemove</i> est traitée avant le <i>linksToAdd</i> .
linksToRemove []	IntegerId	Liste des identificateurs affectés par le <i>Serveur</i> des éléments à rapporter pour les liaisons de déclenchement à supprimer. La liste de <i>linksToRemove</i> est traitée avant le <i>linksToAdd</i> .
Réponse		
responseHeader	ResponseHeader	Paramètres de réponse communs (voir 7.29 pour la définition de <i>ResponseHeader</i>).
addResults []	StatusCode	Liste des <i>StatusCodes</i> pour les éléments à ajouter (voir 7.34 pour la définition des <i>StatusCode</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre <i>linksToAdd</i> spécifié dans la demande.
addDiagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic pour les liaisons à ajouter (voir 7.8 pour la définition de <i>DiagnosticInfo</i>). La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>linksToAdd</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.
removeResults []	StatusCode	Liste des <i>StatusCodes</i> pour les éléments à supprimer. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre <i>linksToRemove</i> spécifié dans la demande.
removeDiagnosticInfos []	Informations de diagnostic	Liste d'informations de diagnostic pour les liaisons à supprimer. La taille et l'ordre de la liste correspondent à la taille et à l'ordre du paramètre de demande <i>linksToRemove</i> . Cette liste est vide si les informations de diagnostic n'ont pas été demandées dans l'en-tête de demande ou si aucune information de diagnostic n'a été trouvée dans le traitement de la demande.

5.12.5.3 Résultats de service

Le Tableau 79 définit les résultats de *Service* spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le 7.34.

Tableau 79 – Codes de résultats du service SetTriggering

ID symbolique	Description
Bad_NothingToDo	Voir Tableau 177 pour la description de ce code de résultat.
Bad_TooManyOperations	Voir Tableau 177 pour la description de ce code de résultat.
Bad_SubscriptionIdInvalid	Voir Tableau 177 pour la description de ce code de résultat.
Bad_MonitoredItemIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.12.5.4 StatusCodes

Le Tableau 80 définit les valeurs pour les paramètres de résultats qui sont spécifiques à ce *Service*. Les *StatusCodes* communs sont définis dans le Tableau 178.

Tableau 80 – Codes de résultats de niveau opération SetTriggering

ID symbolique	Description
Bad_MonitoredItemIdInvalid	Voir Tableau 178 pour la description de ce code de résultat.

5.12.6 DeleteMonitoredItems

5.12.6.1 Description

Ce *Service* est utilisé pour retirer un ou plusieurs *MonitoredItems* d'un *Abonnement*. Lorsqu'un *MonitoredItem* est supprimé, ses liaisons d'éléments déclenchés sont supprimées aussi.

Le retrait réussi d'un *MonitoredItem* peut toutefois ne pas retirer les *Notifications* du *MonitoredItem* qui sont en cours d'envoi par l'*Abonnement*. Par conséquent, les *Clients* peuvent recevoir des *Notifications* pour le *MonitoredItem* après avoir reçu une réponse positive indiquant que le *MonitoredItem* a été supprimé.

5.12.6.2 Paramètres

Le Tableau 81 définit les paramètres du *Service*.