

INTERNATIONAL STANDARD



**OPC unified architecture –
Part 6: Mappings**

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 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 It also gives information on projects, replaced and withdrawn publications.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 60384-6:2020 HV



IEC 62541-6

Edition 3.0 2020-07
REDLINE VERSION

INTERNATIONAL STANDARD



OPC unified architecture –
Part 6: Mappings

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 25.040.40; 35.100.05

ISBN 978-2-8322-8665-4

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

CONTENTS

FOREWORD	8
1 Scope	11
2 Normative references	11
3 Terms, definitions, abbreviated terms and symbols	14
3.1 Terms and definitions	14
3.2 Abbreviated terms and symbols	15
4 Overview	16
5 Data encoding	17
5.1 General	17
5.1.1 Overview	17
5.1.2 Built-in Types	17
5.1.3 Guid	18
5.1.4 ByteString	19
5.1.5 ExtensionObject	19
5.1.6 Variant	19
5.1.7 Decimal	20
5.2 OPC UA Binary	21
5.2.1 General	21
5.2.2 Built-in Types	21
5.2.3 Decimal	32
5.2.4 Enumerations	32
5.2.5 Arrays	32
5.2.6 Structures	33
5.2.7 Structures with optional fields	35
5.2.8 Unions	37
5.2.9 Messages	38
5.3 OPC UA XML	39
5.3.1 Built-in Types	39
5.3.2 Decimal	45
5.3.3 Enumerations	45
5.3.4 Arrays	46
5.3.5 Structures	46
5.3.6 Structures with optional fields	47
5.3.7 Unions	47
5.3.8 Messages	48
5.4 OPC UA JSON	48
5.4.1 General	48
5.4.2 Built-in Types	49
5.4.3 Decimal	54
5.4.4 Enumerations	54
5.4.5 Arrays	54
5.4.6 Structures	55
5.4.7 Structures with optional fields	55
5.4.8 Unions	56
5.4.9 Messages	56
6 Message Security Protocols	57

6.1	Security handshake	57
6.2	Certificates	59
6.2.1	General	59
6.2.2	Application Instance Certificate.....	59
6.2.3	Signed Software Certificate	61
6.2.3	Certificate Chains	61
6.3	Time synchronization	61
6.4	UTC and International Atomic Time (TAI).....	62
6.5	Issued User Identity Tokens.....	62
6.5.1	Kerberos.....	62
6.5.2	JSON Web Token (JWT).....	63
6.5.3	OAuth2	63
6.6	WS Secure Conversation	65
6.7	OPC UA Secure Conversation	70
6.7.1	Overview	70
6.7.2	MessageChunk structure	70
6.7.3	MessageChunks and error handling.....	75
6.7.4	Establishing a SecureChannel	75
6.7.5	Deriving keys.....	77
6.7.6	Verifying Message security	79
7	TransportProtocols	80
7.1	OPC UA TCP Connection Protocol.....	80
7.1.1	Overview	80
7.1.2	Message structure	80
7.1.3	Establishing a connection	84
7.1.4	Closing a connection	87
7.1.5	Error handling.....	87
7.2	OPC UA TCP	91
7.3	SOAP/HTTP.....	91
7.4	OPC UA HTTPS.....	93
7.4.1	Overview	93
7.4.2	Session-less Services.....	95
7.4.3	XML Encoding	95
7.4.4	OPC UA Binary Encoding	96
7.4.5	JSON Encoding	97
7.5	WebSockets.....	97
7.5.1	Overview	97
7.5.2	Protocol Mapping.....	98
7.5.3	Security	98
7.6	Well known addresses	99
8	Normative Contracts	100
8.1	OPC Binary Schema	100
8.2	XML Schema and WSDL.....	100
8.3	Information Model Schema.....	100
8.4	Formal definition of UA Information Model.....	100
8.5	Constants	100
8.6	DataType encoding.....	100
8.7	Security configuration	100
Annex A	(normative) Constants.....	101

A.1	Attribute Ids	101
A.2	Status Codes	101
A.3	Numeric Node Ids	102
Annex B	(normative) OPC UA Nodeset	103
Annex C	(normative) Type declarations for the OPC UA native Mapping	104
Annex D	(normative) WSDL for the XML Mapping	105
D.1	XML Schema	105
D.2	WDSL Port Types	105
D.3	WSDL Bindings	105
Annex E	(normative) Security settings management	106
E.1	Overview	106
E.2	SecuredApplication	107
E.3	CertificateIdentifier	110
E.4	CertificateStoreIdentifier	112
E.5	CertificateList	113
E.6	CertificateValidationOptions	113
Annex F	(normative) Information Model XML Schema	115
F.1	Overview	115
F.2	UANodeSet	115
F.3	UANode	117
F.4	Reference	118
F.5	RolePermission	118
F.6	UAType	118
F.7	UAInstance	119
F.8	UAVariable	119
F.9	UAMethod	120
F.10	TranslationType	121
F.11	UADatatype	122
F.12	DataTypeDefinition	122
F.13	DataTypeField	123
F.14	Variant	124
F.15	Example	125
F.16	UANodeSetChanges	127
F.17	NodesToAdd	128
F.18	ReferencesToChange	128
F.19	ReferenceToChange	129
F.20	NodesToDelete	129
F.21	NodeToDelete	129
F.22	UANodeSetChangesStatus	130
F.23	NodeSetStatusList	130
F.24	NodeSetStatus	131
Bibliography	132
Figure 1	– The OPC UA Stack Overview	17
Figure 2	– Encoding Integers in a binary stream	21
Figure 3	– Encoding Floating Points in a binary stream	22
Figure 4	– Encoding Strings in a binary stream	22

Figure 5 – Encoding Guid in a binary stream	23
Figure 6 – Encoding XmlElement in a binary stream	24
Figure 7 – A String NodeId	25
Figure 8 – A Two Byte NodeId	26
Figure 9 – A Four Byte NodeId	26
Figure 10 – Security handshake	57
Figure 11 – OPC UA Secure Conversation MessageChunk	70
Figure 12 – OPC UA TCP Connection Protocol Message structure	80
Figure 13 – Client initiated OPC UA Connection Protocol connection	86
Figure 14 – Server initiated OPC UA Connection Protocol connection	86
Figure 15 – Closing a OPC UA TCP Connection Protocol connection	87
Figure 16 – Scenarios for the HTTPS Transport	94
Figure 17 – Setting up Communication over a WebSocket	98
Table 1 – Built-in Data Types	18
Table 2 – Guid structure	18
Table 3 – Layout of Decimal	20
Table 4 – Supported Floating Point Types	22
Table 5 – NodeId components	24
Table 6 – NodeId DataEncoding values	25
Table 7 – Standard NodeId Binary DataEncoding	25
Table 8 – Two Byte NodeId Binary DataEncoding	26
Table 9 – Four Byte NodeId Binary DataEncoding	26
Table 10 – ExpandedNodeId Binary DataEncoding	27
Table 11 – DiagnosticInfo Binary DataEncoding	28
Table 12 – QualifiedName Binary DataEncoding	28
Table 13 – LocalizedText Binary DataEncoding	29
Table 14 – Extension Object Binary DataEncoding	30
Table 15 – Variant Binary DataEncoding	31
Table 16 – Data Value Binary DataEncoding	32
Table 17 – Sample OPC UA Binary Encoded structure	34
Table 18 – Sample OPC UA Binary Encoded Structure with optional fields	36
Table 19 – Sample OPC UA Binary Encoded Structure	37
Table 20 – XML Data Type Mappings for Integers	39
Table 21 – XML Data Type Mappings for Floating Points	39
Table 22 – Components of NodeId	41
Table 23 – Components of ExpandedNodeId	42
Table 24 – Components of Enumeration	46
Table 25 – JSON Object Definition for a NodeId	50
Table 26 – JSON Object Definition for an ExpandedNodeId	51
Table 27 – JSON Object Definition for a StatusCode	51
Table 28 – JSON Object Definition for a DiagnosticInfo	52
Table 29 – JSON Object Definition for a QualifiedName	52

Table 30 – JSON Object Definition for a LocalizedText	52
Table 31 – JSON Object Definition for an ExtensionObject	53
Table 32 – JSON Object Definition for a Variant	53
Table 33 – JSON Object Definition for a DataValue	54
Table 34 – JSON Object Definition for a Decimal	54
Table 35 – JSON Object Definition for a <i>Structure</i> with Optional Fields	55
Table 36 – JSON Object Definition for a Union	56
Table 37 – SecurityPolicy	58
Table 38 – Application Instance Certificate	60
Table 39 – Kerberos UserTokenPolicy	62
Table 40 – JWT UserTokenPolicy	63
Table 41 – JWT IssuerEndpointUrl Definition	63
Table 42 – Access Token Claims	64
Table 43 – OPC UA Secure Conversation Message header	71
Table 44 – Asymmetric algorithm Security header	72
Table 45 – Symmetric algorithm Security header	73
Table 46 – Sequence header	73
Table 47 – OPC UA Secure Conversation Message footer	74
Table 48 – OPC UA Secure Conversation Message abort body	75
Table 49 – OPC UA Secure Conversation OpenSecureChannel Service	76
Table 50 – PRF inputs for RSA based SecurityPolicies	78
Table 51 – Cryptography key generation parameters	78
Table 52 – OPC UA TCP Connection Protocol Message header	81
Table 53 – OPC UA TCP Connection Protocol Hello Message	82
Table 54 – OPC UA TCP Connection Protocol Acknowledge Message	83
Table 55 – OPC UA TCP Connection Protocol Error Message	83
Table 56 – OPC UA Connection Protocol ReverseHello Message	84
Table 57 – OPC UA Connection Protocol error codes	89
Table 58 – WebSocket Protocols Mappings	98
Table 59 – Well known addresses for Local Discovery Servers	99
Table A.1 – Identifiers assigned to Attributes	101
Table E.1 – SecuredApplication	108
Table E.2 – CertificateIdentifier	111
Table E.3 – Structured directory store	112
Table E.4 – CertificateStoreIdentifier	113
Table E.5 – CertificateList	113
Table E.6 – CertificateValidationOptions	114
Table F.1 – UANodeSet	116
Table F.2 – UANode	117
Table F.3 – Reference	118
Table F.4 – RolePermission	118
Table F.5 – UANodeSet Type Nodes	118
Table F.6 – UANodeSet Instance Nodes	119

Table F.7 – UAInstance	119
Table F.8 – UAVariable	120
Table F.9 – UAMethod	120
Table F.10 – TranslationType	121
Table F.11 – UADatatype	122
Table F.12 – DataTypeDefinition	122
Table F.13 – DataTypeField	123
Table F.14 – UANodeSetChanges	127
Table F.15 – NodesToAdd	128
Table F.16 – ReferencesToChange	129
Table F.17 – ReferencesToChange	129
Table F.18 – NodesToDelete	129
Table F.19 – ReferencesToChange	130
Table F.20 – UANodeSetChangesStatus	130
Table F.21 – NodeSetStatusList	131
Table F.22 – NodeSetStatus	131

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

INTERNATIONAL ELECTROTECHNICAL COMMISSION

OPC UNIFIED ARCHITECTURE –

Part 6: Mappings

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 62541-6 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) Encodings:

- added JSON encoding for PubSub (non-reversible);
- added JSON encoding for Client/Server (reversible);
- added support for optional fields in structures;
- added support for Unions.

b) Transport mappings:

- added WebSocket secure connection – WSS;
- added support for reverse connectivity;
- added support for session-less service invocation in HTTPS.

c) Deprecated Transport (missing support on most platforms):

- SOAP/HTTP with WS-SecureConversation (all encodings).

d) Added mapping for JSON Web Token.

e) Added support for Unions to NodeSet Schema.

f) Added batch operations to add/delete nodes to/from NodeSet Schema.

g) Added support for multi-dimensional arrays outside of Variants.

h) Added binary representation for Decimal data types.

i) Added mapping for an OAuth2 Authorization Framework.

The text of this International Standard is based on the following documents:

FDIS	Report on voting
65E/718/FDIS	65E/734/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 IEC 62541, 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 62541-6:2020 RLV

OPC UNIFIED ARCHITECTURE –

Part 6: Mappings

1 Scope

This part of IEC 62541 specifies the OPC Unified Architecture (OPC UA) mapping between the security model described in IEC TR 62541-2, the abstract service definitions, ~~described~~ ~~specified~~ in IEC 62541-4, the data structures defined in IEC 62541-5 and the physical network protocols that can be used to implement the OPC UA specification.

2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. 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 62541-4, *OPC Unified Architecture – Part 4: Services*

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

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

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

ISO 8601-1:2019, *Date and time – Representations for information interchange – Part 1: Basic rules*

~~XML Schema Part 1: XML Schema Part 1: Structures~~

~~<http://www.w3.org/TR/xmlschema-1/>~~

~~XML Schema Part 2: XML Schema Part 2: Datatypes~~

~~<http://www.w3.org/TR/xmlschema-2/>~~

~~SOAP Part 1: SOAP Version 1.2 Part 1: Messaging Framework~~

~~<http://www.w3.org/TR/soap12-part1/>~~

~~SOAP Part 2: SOAP Version 1.2 Part 2: Adjuncts~~

~~<http://www.w3.org/TR/soap12-part2/>~~

~~XML Encryption: XML Encryption Syntax and Processing~~

~~<http://www.w3.org/TR/xmlenc-core/>~~

~~XML Signature: XML Signature Syntax and Processing~~

~~<http://www.w3.org/TR/xmlsig-core/>~~

~~WS Security: SOAP Message Security 1.1~~

~~<http://www.oasis-open.org/committees/download.php/16790/wss-v1.1-spec-os-SOAPMessageSecurity.pdf>~~

~~WS Addressing: Web Services Addressing (WS-Addressing)~~

~~<http://www.w3.org/Submission/ws-addressing/>~~

~~WS Trust: WS Trust 1.3~~

~~<http://docs.oasis-open.org/ws-sx/ws-trust/v1.3/ws-trust.html>~~

~~WS Secure Conversation: WS Secure Conversation 1.3~~

~~<http://docs.oasis-open.org/ws-sx/ws-secureconversation/v1.3/ws-secureconversation.html>~~

~~WS Security Policy: WS Security Policy 1.2~~

~~<http://docs.oasis-open.org/ws-sx/ws-securitypolicy/200702/ws-securitypolicy-1.2-spec-os.html>~~

~~SSL/TLS: RFC 5246 — The TLS Protocol Version 1.2~~

~~<http://tools.ietf.org/html/rfc5246.txt>~~

~~X509: X.509 Public Key Certificate Infrastructure~~

~~<http://www.itu.int/rec/T-REC-X.509-200003-I/e>~~

~~WS-I Basic Profile 1.1: WS-I Basic Profile Version 1.1~~

~~<http://www.ws-i.org/Profiles/BasicProfile-1.1.html>~~

~~WS-I Basic Security Profile 1.1: WS-I Basic Security Profile Version 1.1~~

~~<http://www.ws-i.org/Profiles/BasicSecurityProfile-1.1.html>~~

~~HTTP: RFC 2616 — Hypertext Transfer Protocol — HTTP/1.1~~

~~<http://www.ietf.org/rfc/rfc2616.txt>~~

~~Base64: RFC 3548 — The Base16, Base32, and Base64 Data Encodings~~

~~<http://www.ietf.org/rfc/rfc3548.txt>~~

~~X.690: ITU-T X.690 — Basic (BER), Canonical (CER) and Distinguished (DER) Encoding Rules~~

~~<http://www.itu.int/ITU-T/studygroups/com17/languages/X.690-0207.pdf>~~

~~IEEE 754: Standard for Binary Floating-Point Arithmetic~~

~~<http://grouper.ieee.org/groups/754/>~~

~~HMAC: HMAC — Keyed-Hashing for Message Authentication~~

~~<http://www.ietf.org/rfc/rfc2104.txt>~~

~~PKCS #1: PKCS #1 — RSA Cryptography Specifications Version 2.0~~

~~<http://www.ietf.org/rfc/rfc2437.txt>~~

~~FIPS 180-2: Secure Hash Standard (SHA)~~

~~<http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>~~

~~FIPS 197: Advanced Encryption Standard (AES)~~

~~<http://www.csrc.nist.gov/publications/fips/fips197/fips-197.pdf>~~

~~UTF8: UTF-8, a transformation format of ISO 10646~~

~~<http://tools.ietf.org/html/rfc3629>~~

~~RFC 3280: RFC 3280 – X.509 Public Key Infrastructure Certificate and CRL Profile~~

~~<http://www.ietf.org/rfc/rfc3280.txt>~~

~~RFC 4514: RFC 4514 – LDAP: String Representation of Distinguished Names~~

~~<http://www.ietf.org/rfc/rfc4514.txt>~~

~~NTP: RFC 1305 – Network Time Protocol (Version 3)~~

~~<http://www.ietf.org/rfc/rfc1305.txt>~~

~~Kerberos: WS Security Kerberos Token Profile 1.1~~

~~<http://docs.oasis-open.org/wss/v1.1/wss-v1.1-spec-os-KerberosTokenProfile.pdf>~~

XML Schema Part 2: XML Schema Part 2: Datatypes

<http://www.w3.org/TR/xmlschema-2/>

SOAP Part 1: SOAP Version 1.2 Part 1: Messaging Framework

<http://www.w3.org/TR/soap12-part1/>

SSL/TLS: RFC 5246 – The TLS Protocol Version 1.2

<http://tools.ietf.org/html/rfc5246.txt>

X.509 v3: ISO/IEC 9594-8 (ITU-T Rec. X.509), *Information technology – Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks*

HTTP: RFC 2616 – Hypertext Transfer Protocol – HTTP/1.1

<http://www.ietf.org/rfc/rfc2616.txt>

HTTPS: RFC 2818 – HTTP Over TLS

<http://www.ietf.org/rfc/rfc2818.txt>

Base64: RFC 3548 – The Base16, Base32, and Base64 Data Encodings

<http://www.ietf.org/rfc/rfc3548.txt>

X690: ISO/IEC 8825-1 (ITU-T Rec. X.690), *Information technology – ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)*

IEEE-754: Standard for Floating-Point Arithmetic

HMAC: HMAC – Keyed-Hashing for Message Authentication

<http://www.ietf.org/rfc/rfc2104.txt>

PKCS #1: PKCS #1 – RSA Cryptography Specifications Version 2.0

<http://www.ietf.org/rfc/rfc2437.txt>

PKCS #12: PKCS #12 – Personal Information Exchange Syntax v1.1

<http://www.ietf.org/rfc/rfc7292.txt>

FIPS 180-4: Secure Hash Standard (SHS)

<https://csrc.nist.gov/publications/detail/fips/180/4/final>

FIPS 197: Advanced Encryption Standard (AES)

<https://csrc.nist.gov/publications/detail/fips/197/final>

UTF-8: UTF-8, a transformation format of ISO 10646

<http://www.ietf.org/rfc/rfc3629.txt>

RFC 3280: RFC 3280 – X.509 Public Key Infrastructure Certificate and CRL Profile

<http://www.ietf.org/rfc/rfc3280.txt>

RFC 4514: RFC 4514 – LDAP: String Representation of Distinguished Names

<http://www.ietf.org/rfc/rfc4514.txt>

NTP: RFC 1305 – Network Time Protocol (Version 3) Specification, Implementation and Analysis

<http://www.ietf.org/rfc/rfc1305.txt>

Kerberos: Web Services Security – Kerberos Token Profile 1.1

<http://docs.oasis-open.org/wss/v1.1/wss-v1.1-spec-os-KerberosTokenProfile.pdf>

RFC 1738: RFC 1738 – Uniform Resource Locators (URL)

<http://www.ietf.org/rfc/rfc1738.txt>

RFC 2141: RFC 2141 – URN Syntax

<http://www.ietf.org/rfc/rfc2141.txt>

RFC 6455: RFC 6455 – The WebSocket Protocol

<http://www.ietf.org/rfc/rfc6455.txt>

RFC 7159: The JavaScript Object Notation (JSON) Data Interchange Format

<http://www.ietf.org/rfc/rfc7159.txt>

RFC 7523: JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants

<https://tools.ietf.org/rfc/rfc7523.txt>

RFC 6749: The OAuth 2.0 Authorization Framework

<http://www.ietf.org/rfc/rfc6749.txt>

OpenID-Core: OpenID Connect Core 1.0

http://openid.net/specs/openid-connect-core-1_0.html

OpenID-Discovery: OpenID Connect Discovery 1.0

https://openid.net/specs/openid-connect-discovery-1_0.html

RFC 6960: RFC 6960 – X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP

<https://www.ietf.org/rfc/rfc6960.txt>

3 Terms, definitions, abbreviated terms and symbols

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

CertificateDigest

short identifier used to uniquely identify an X.509 v3 *Certificate*

Note 1 to entry: This is the SHA1 hash of DER encoded form of the *Certificate*.

3.1.2

DataEncoding

way to serialize OPC UA *Messages* and data structures

3.1.3

DevelopmentPlatform

suite of tools and/or programming languages used to create software

3.1.4

Mapping

~~specifies~~ specification on how to implement an OPC UA feature with a specific technology

Note 1 to entry: For example, the OPC UA Binary Encoding is a *Mapping* that specifies how to serialize OPC UA data structures as sequences of bytes.

3.1.5

SecurityProtocol

protocol which ensures the integrity and privacy of UA *Messages* that are exchanged between OPC UA applications

3.1.6

StackProfile

combination of *DataEncodings*, *SecurityProtocol* and *TransportProtocol Mappings*

Note 1 to entry: OPC UA applications implement one or more *StackProfiles* and can only communicate with OPC UA applications that support a *StackProfile* that they support.

3.1.7

TransportConnection

full-duplex communication link established between OPC UA applications

Note 1 to entry: A TCP/IP socket is an example of a *TransportConnection*.

3.1.8

TransportProtocol

way to exchange serialized OPC UA *Messages* between OPC UA applications

3.2 Abbreviated terms and symbols

API	application programming interface
ASN.1	Abstract Syntax Notation #1 (used in X690)
BP	WS-I Basic Profile Version
BSP	WS-I Basic Security Profile
CSV	comma separated value (file format)
ECC	elliptic curve cryptography
HTTP	Hypertext Transfer Protocol
HTTPS	Secure Hypertext Transfer Protocol
IPSec	Internet Protocol Security
RST	Request Security Token

OID	object identifier (used with ASN.1)
RSTR	Request Security Token Response
SCT	Security Context Token
PRF	pseudo random function
RSA	Rivest, Shamir and Adleman [public key encryption system]
SHA1	secure hash algorithm
SOAP	Simple Object Access Protocol
SSL	Secure Sockets Layer (defined in SSL/TLS)
TCP	Transmission Control Protocol
TLS	Transport Layer Security (defined in SSL/TLS)
UTF8	Unicode Transformation Format (8-bit) (Defined in UTF8)
UA	Unified Architecture
UACP	OPC UA Connection Protocol
UASC	OPC UA Secure Conversation
WS-*	XML Web Services specifications
WSS	WS Security
WS-SC	WS Secure Conversation
XML	eXtensible Markup Language

4 Overview

Other parts of the IEC 62541 series are written to be independent of the technology used for implementation. This approach means OPC UA is a flexible specification that will continue to be applicable as technology evolves. On the other hand, this approach means that it is not possible to build an OPC UA application with the information contained in IEC TR 62541-1 through to IEC 62541-5 because important implementation details have been left out.

This document defines *Mappings* between the abstract specifications and technologies that can be used to implement them. The *Mappings* are organized into three groups: *DataEncodings*, *SecurityProtocols* and *TransportProtocols*. Different *Mappings* are combined together to create *StackProfiles*. All OPC UA applications shall implement at least one *StackProfile* and can only communicate with other OPC UA applications that implement the same *StackProfile*.

This document defines the *DataEncodings* in Clause 5, the *SecurityProtocols* in ~~Clause 6~~ 5.4 and the *TransportProtocols* in 6.7.6. The *StackProfiles* are defined in IEC 62541-7.

All communication between OPC UA applications is based on the exchange of *Messages*. The parameters contained in the *Messages* are defined in IEC 62541-4; however, their format is specified by the *DataEncoding* and *TransportProtocol*. For this reason, each *Message* defined in IEC 62541-4 shall have a normative description which specifies exactly what shall be put on the wire. The normative descriptions are defined in the annexes.

A *Stack* is a collection of software libraries that implement one or more *StackProfiles*. The interface between an OPC UA application and the *Stack* is a non-normative API which hides the details of the *Stack* implementation. An API depends on a specific *DevelopmentPlatform*. Note that the datatypes exposed in the API for a *DevelopmentPlatform* may not match the datatypes defined by the specification because of limitations of the *DevelopmentPlatform*. For example, the Java programming language does not support an unsigned integer which means that any Java API will need to map unsigned integers onto a signed integer type.

Figure 1 illustrates the relationships between the different concepts defined in this document.

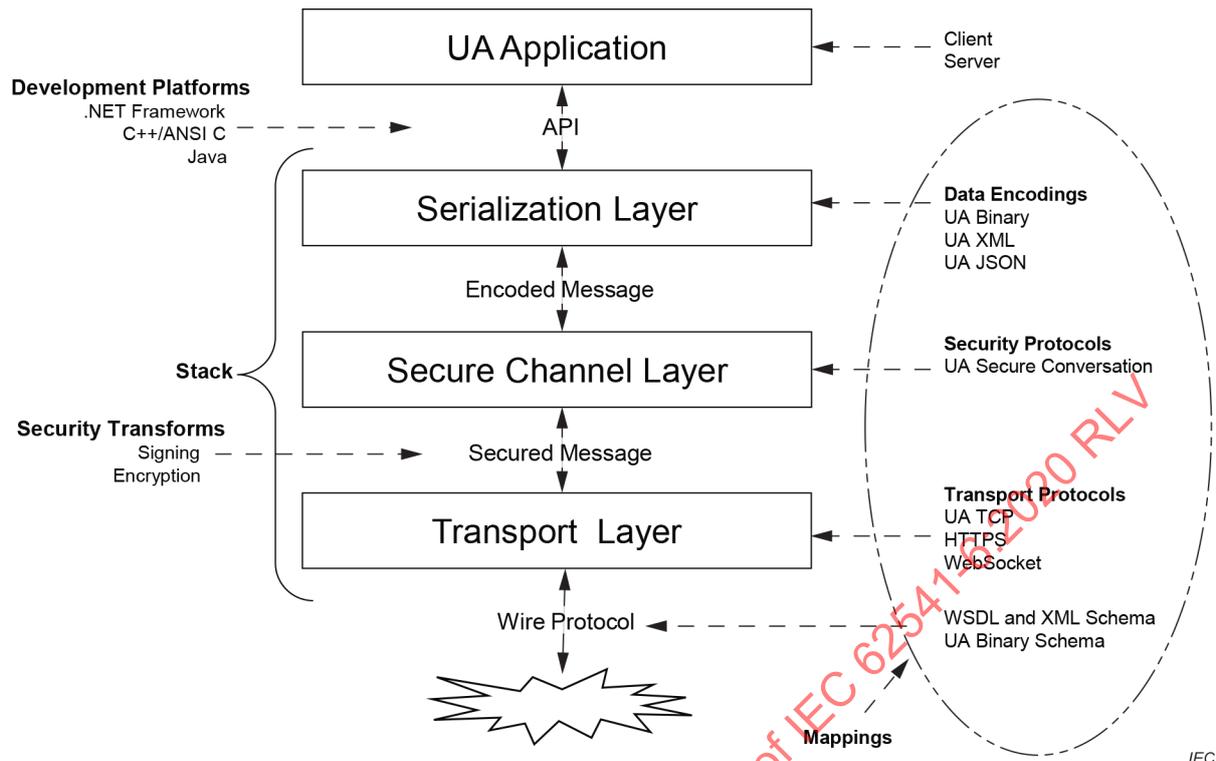


Figure 1 – The OPC UA Stack Overview

The layers described in this document do not correspond to layers in the OSI 7-layer model [X200]. Each OPC UA *StackProfile* should be treated as a single Layer 7 (application) protocol that is built on an existing Layer 5, 6 or 7 protocol such as TCP/IP, TLS or HTTP. The *SecureChannel* layer is always present even if the *SecurityMode* is *None*. In this situation, no security is applied but the *SecurityProtocol* implementation shall maintain a logical channel with a unique identifier. Users and administrators are expected to understand that a *SecureChannel* with *SecurityMode* set to *None* cannot be trusted unless the application is operating on a physically secure network or a low-level protocol such as IPsec is being used.

5 Data encoding

5.1 General

5.1.1 Overview

This document defines ~~two~~ three data encodings: OPC UA Binary, OPC UA XML and OPC UA ~~XML~~ JSON. It describes how to construct *Messages* using each of these encodings.

5.1.2 Built-in Types

All OPC UA *DataEncodings* are based on rules that are defined for a standard set of built-in types. These built-in types are then used to construct structures, arrays and *Messages*. The built-in types are described in Table 1.

Table 1 – Built-in Data Types

ID	Name	Description
1	Boolean	A two-state logical value (true or false).
2	SByte	An integer value between –128 and 127 inclusive.
3	Byte	An integer value between 0 and 256 255 inclusive.
4	Int16	An integer value between –32 768 and 32 767 inclusive.
5	UInt16	An integer value between 0 and 65 535 inclusive.
6	Int32	An integer value between –2 147 483 648 and 2 147 483 647 inclusive.
7	UInt32	An integer value between 0 and 4 294 967 295 inclusive.
8	Int64	An integer value between –9 223 372 036 854 775 808 and 9 223 372 036 854 775 807 inclusive.
9	UInt64	An integer value between 0 and 18 446 744 073 709 551 615 inclusive.
10	Float	An IEEE single precision (32 bit) floating point value.
11	Double	An IEEE double precision (64 bit) floating point value.
12	String	A sequence of Unicode characters.
13	DateTime	An instance in time.
14	Guid	A 16-byte value that can be used as a globally unique identifier.
15	ByteString	A sequence of octets.
16	XmlElement	An XML element.
17	NodeId	An identifier for a node in the address space of an OPC UA Server.
18	ExpandedNodeId	A NodeId that allows the namespace URI to be specified instead of an index.
19	StatusCode	A numeric identifier for an error or condition that is associated with a value or an operation.
20	QualifiedName	A name qualified by a namespace.
21	LocalizedText	Human readable text with an optional locale identifier.
22	ExtensionObject	A structure that contains an application specific data type that may not be recognized by the receiver.
23	DataValue	A data value with an associated status code and timestamps.
24	Variant	A union of all of the types specified above.
25	DiagnosticInfo	A structure that contains detailed error and diagnostic information associated with a StatusCode.

Most of these data types are the same as the abstract types defined in IEC 62541-3 and IEC 62541-4. However, the *ExtensionObject* and *Variant* types are defined in this document. In addition, this document defines a representation for the *Guid* type defined in IEC 62541-3.

5.1.3 Guid

A *Guid* is a 16-byte globally unique identifier with the layout shown in Table 2.

Table 2 – Guid structure

Component	Data Type
Data1	UInt32
Data2	UInt16
Data3	UInt16
Data4	Byte [8]

Guid values may be represented as a string in this form:

<Data1>–<Data2>–<Data3>–<Data4[0:1]>–<Data4[2:7]>

where Data1 is 8 characters wide, Data2 and Data3 are 4 characters wide and each Byte in Data4 is 2 characters wide. Each value is formatted as a hexadecimal number with padded zeros. A typical *Guid* value would look like this when formatted as a string:

C496578A-0DFE-4B8F-870A-745238C6AEAE

5.1.4 ByteString

A *ByteString* is structurally the same as a one-dimensional array of *Byte*. It is represented as a distinct built-in data type because it allows encoders to optimize the transmission of the value. However, some *DevelopmentPlatforms* will not be able to preserve the distinction between a *ByteString* and a one-dimensional array of *Byte*.

If a decoder for *DevelopmentPlatform* cannot preserve the distinction it shall convert all one-dimensional arrays of *Byte* to *ByteStrings*.

Each element in a one-dimensional array of *ByteString* can have a different length which means is structurally different from a two-dimensional array of *Byte* where the length of each dimension is the same. This means decoders shall preserve the distinction between two or more dimension arrays of *Byte* and one or more dimension arrays of *ByteString*.

If a *DevelopmentPlatform* does not support unsigned integers, then it will ~~have~~ need to represent *ByteStrings* as arrays of *SByte*. In this case, the requirements for *Byte* would then apply to *SByte*.

5.1.5 ExtensionObject

An *ExtensionObject* is a container for any ~~Complex Data types~~ *Structured DataTypes* which cannot be encoded as one of the other built-in data types. The *ExtensionObject* contains a complex value serialized as a sequence of bytes or as an XML element. It also contains an identifier which indicates what data it contains and how it is encoded.

~~Complex Data types~~ *Structured DataTypes* are represented in a *Server* address space as subtypes of the *Structure DataType*. The *DataEncodings* available for any given ~~Complex Data type~~ *Structured DataTypes* are represented as a *DataTypeEncoding Object* in the *Server AddressSpace*. The *NodeId* for the *DataTypeEncoding Object* is the identifier stored in the *ExtensionObject*. IEC 62541-3 describes how *DataTypeEncoding Nodes* are related to other *Nodes* of the *AddressSpace*.

Server implementers should use namespace qualified numeric *NodeIds* for any *DataTypeEncoding Objects* they define. This will minimize the overhead introduced by packing ~~Complex Data~~ *Structured DataType* values into ~~ExtensionObjects~~ an *ExtensionObject*.

ExtensionObjects and Variants allow unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

5.1.6 Variant

A *Variant* is a union of all built-in data types including an *ExtensionObject*. *Variants* can also contain arrays of any of these built-in types. *Variants* are used to store any value or parameter with a data type of *BaseDataType* or one of its subtypes.

Variants can be empty. An empty *Variant* is described as having a null value and should be treated like a null column in a SQL database. A null value in a *Variant* may not be the same

as a null value for data types that support nulls such as *Strings*. Some *DevelopmentPlatforms* may not be able to preserve the distinction between a null for a *DataType* and a null for a *Variant*; therefore, applications shall not rely on this distinction. This requirement also means that if an *Attribute* supports the writing of a null value it shall also support writing of an empty *Variant* and vice versa.

Variants can contain arrays of *Variants* but they cannot directly contain another *Variant*.

~~*DataValue*~~ and *DiagnosticInfo* types only have meaning when returned in a response message with an associated *StatusCode* and table of strings. As a result, *Variants* cannot contain instances of ~~*DataValue*~~ or *DiagnosticInfo*.

Values of *Attributes* are always returned in instances of *DataValues*. Therefore, the *DataType* of an *Attribute* cannot be a *DataValue*. *Variants* can contain *DataValue* when used in other contexts such as *Method Arguments* or *PubSub Messages*. The *Variant* in a *DataValue* cannot, directly or indirectly, contain another *DataValue*.

Variables with a *DataType* of *BaseDataType* are mapped to a *Variant*, however, the *ValueRank* and *ArrayDimensions Attributes* place restrictions on what is allowed in the *Variant*. For example, if the *ValueRank* is *Scalar* then the *Variant* may only contain scalar values.

ExtensionObjects and *Variants* allow unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

5.1.7 Decimal

A *Decimal* is a high-precision signed decimal number. It consists of an arbitrary precision integer unscaled value and an integer scale. The scale is the power of ten that is applied to the unscaled value.

A *Decimal* has the fields described in Table 3.

Table 3 – Layout of Decimal

Field	Type	Description
Typeld	NodeId	The identifier for the <i>Decimal DataType</i> .
Encoding	Byte	This value is always 1.
Length	Int32	The length of the <i>Decimal</i> . If the length is less than or equal to 0 then the <i>Decimal</i> value is 0.
Scale	Int16	A signed integer representing the power of ten used to scale the value. i.e. the decimal number of the value multiplied by 10^{-scale} The integer is encoded starting with the least significant bit.
Value	Byte [*]	A 2-complement signed integer representing the unscaled value. The number of bits is inferred from the length of the <i>length</i> field. If the number of bits is 0 then the value is 0. The integer is encoded with the least significant byte first.

When a *Decimal* is encoded in a *Variant*, the built-in type is set to *ExtensionObject*. Decoders that do not understand the *Decimal* type shall treat it like any other unknown *Structure* and pass it on to the application. Decoders that do understand the *Decimal* can parse the value and use any construct that is suitable for the *DevelopmentPlatform*.

If a *Decimal* is embedded in another *Structure* then the *DataTypeDefinition* for the field shall specify the *NodeId* of the *Decimal Node* as the *DataType*. If a *Server* publishes an OPC Binary type description for the *Structure* then the type description shall set the *DataType* for the field to *ExtensionObject*.

5.2 OPC UA Binary

5.2.1 General

The OPC UA *Binary DataEncoding* is a data format developed to meet the performance needs of OPC UA applications. This format is designed primarily for fast encoding and decoding, however, the size of the encoded data on the wire was also a consideration.

The OPC UA *Binary DataEncoding* relies on several primitive data types with clearly defined encoding rules that can be sequentially written to or read from a binary stream. A structure is encoded by sequentially writing the encoded form of each field. If a given field is also a structure, then the values of its fields are written sequentially before writing the next field in the containing structure. All fields shall be written to the stream even if they contain null values. The encodings for each primitive type specify how to encode either a null or a default value for the type.

The OPC UA *Binary DataEncoding* does not include any type or field name information because all OPC UA applications are expected to have advance knowledge of the services and structures that they support. An exception is an *ExtensionObject* which provides an identifier and a size for the ~~Complex Data~~ *Structured DataType* structure it represents. This allows a decoder to skip over types that it does not recognize.

5.2.2 Built-in Types

5.2.2.1 Boolean

A *Boolean* value shall be encoded as a single byte where a value of 0 (zero) is false and any non-zero value is true.

Encoders shall use the value of 1 to indicate a true value; however, decoders shall treat any non-zero value as true.

5.2.2.2 Integer

All integer types shall be encoded as little endian values where the least significant byte appears first in the stream.

Figure 2 illustrates how value 1 000 000 000 (Hex: 3B9ACA00) ~~should be~~ is encoded as a 32-bit integer in the stream.

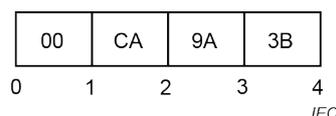


Figure 2 – Encoding Integers in a binary stream

5.2.2.3 Floating Point

All floating-point values shall be encoded with the appropriate IEEE-754 binary representation which has three basic components: the sign, the exponent, and the fraction. The bit ranges assigned to each component depend on the width of the type. Table 4 lists the bit ranges for the supported floating point types.

Table 4 – Supported Floating Point Types

Name	Width (bits)	Fraction	Exponent	Sign
Float	32	0 to 22	23 to 30	31
Double	64	0 to 51	52 to 62	63

In addition, the order of bytes in the stream is significant. All floating point values shall be encoded with the least significant byte appearing first (i.e. little endian).

Figure 3 illustrates how the value -6,5 (Hex: C0D00000) ~~should be~~ is encoded as a *Float*.

The floating-point type supports positive and negative infinity and not-a-number (NaN). The IEEE specification allows for multiple NaN variants; however, the encoders/decoders may not preserve the distinction. Encoders shall encode a NaN value as an IEEE quiet-NAN (000000000000F8FF) or (0000C0FF). Any unsupported types such as denormalized numbers shall also be encoded as an IEEE quiet-NAN. Any test for equality between NaN values always fails.

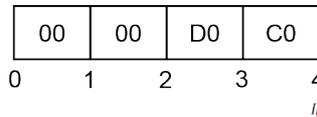


Figure 3 – Encoding Floating Points in a binary stream

5.2.2.4 String

All *String* values are encoded as a sequence of UTF-8 characters without a null terminator and preceded by the length in bytes.

The length in bytes is encoded as *Int32*. A value of -1 is used to indicate a 'null' string.

Figure 4 illustrates how the multilingual string "水Boy" ~~should be~~ is encoded in a byte stream.



Figure 4 – Encoding Strings in a binary stream

5.2.2.5 DateTime

A *DateTime* value shall be encoded as a 64-bit signed integer (see 5.2.2.2) which represents the number of 100 nanosecond intervals since January 1, 1601 (UTC).

Not all *DevelopmentPlatforms* will be able to represent the full range of dates and times that can be represented with this *DataEncoding*. For example, the UNIX time_t structure only has a 1 second resolution and cannot represent dates prior to 1970. For this reason, a number of rules shall be applied when dealing with date/time values that exceed the dynamic range of a *DevelopmentPlatform*. These rules are:

- a) A date/time value is encoded as 0 if either
 - 1) The value is equal to or earlier than 1601-01-01 12:00AM UTC.

- 2) The value is the earliest date that can be represented with the *DevelopmentPlatform's* encoding.
- a) A date/time is encoded as the maximum value for an *Int64* if either
 - 1) The value is equal to or greater than 9999-01-01 12:31 11:59:59PM UTC,
 - 2) The value is the latest date that can be represented with the *DevelopmentPlatform's* encoding.
 - b) A date/time is decoded as the earliest time that can be represented on the platform if either
 - 1) The encoded value is 0,
 - 2) The encoded value represents a time earlier than the earliest time that can be represented with the *DevelopmentPlatform's* encoding.
 - c) A date/time is decoded as the latest time that can be represented on the platform if either
 - 1) The encoded value is the maximum value for an *Int64*,
 - 2) The encoded value represents a time later than the latest time that can be represented with the *DevelopmentPlatform's* encoding.

These rules imply that the earliest and latest times that can be represented on a given platform are invalid date/time values and should be treated that way by applications.

A decoder shall truncate the value if a decoder encounters a *DateTime* value with a resolution that is greater than the resolution supported on the *DevelopmentPlatform*.

5.2.2.6 Guid

A *Guid* is encoded in a structure as shown in Table 2. Fields are encoded sequentially according to the data type for field.

Figure 5 illustrates how the *Guid* "72962B91-FA75-4AE6-8D28-B404DC7DAF63" ~~should be~~ is encoded in a byte stream.

Data1				Data2		Data3		Data4							
91	2B	96	72	75	FA	E6	4A	8D	28	B4	04	DC	7D	AF	63
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

IEC

Figure 5 – Encoding Guids in a binary stream

5.2.2.7 ByteString

A *ByteString* is encoded as sequence of bytes preceded by its length in bytes. The length is encoded as a 32-bit signed integer as described above.

If the length of the byte string is -1 then the byte string is 'null'.

5.2.2.8 XmlElement

An *XmlElement* is an XML fragment serialized as UTF-8 string and then encoded as *ByteString*.

Figure 6 illustrates how the *XmlElement* "<A>Hot水" ~~should be~~ is encoded in a byte stream.

Length				<A>			Hot			水							
0D	00	00	00	3C	41	3E	72	6F	74	E6	B0	B4	3C	3F	41	3E	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17

IEC

Figure 6 – Encoding XmlElement in a binary stream

A decoder may choose to parse the XML after decoding; if an unrecoverable parsing error occurs then the decoder should try to continue processing the stream. For example, if the *XmlElement* is the body of a *Variant* or an element in an array which is the body of a *Variant* then this error can be reported by setting the value of the *Variant* to the *StatusCode Bad_DecodingError*.

5.2.2.9 NodeId

The components of a *NodeId* are described the Table 5.

Table 5 – NodeId components

Name	Data Type	Description
Namespace	UInt16	The index for a namespace URI. An index of 0 is used for OPC UA defined <i>NodeIds</i> .
IdentifierType	Enumeration	The format and data type of the identifier. The value may be one of the following: NUMERIC - the value is an <i>UInteger</i> ; STRING - the value is <i>String</i> ; GUID - the value is a <i>Guid</i> ; OPAQUE - the value is a <i>ByteString</i> ;
Value	*	The identifier for a node in the address space of an OPC UA <i>Server</i> .

The *DataEncoding* of a *NodeId* varies according to the contents of the instance. For that reason, the first byte of the encoded form indicates the format of the rest of the encoded *NodeId*. The possible *DataEncoding* formats are shown in Table 6. Table 6 through Table 9 describe the structure of each possible format (they exclude the byte which indicates the format).

Table 6 – NodeId DataEncoding values

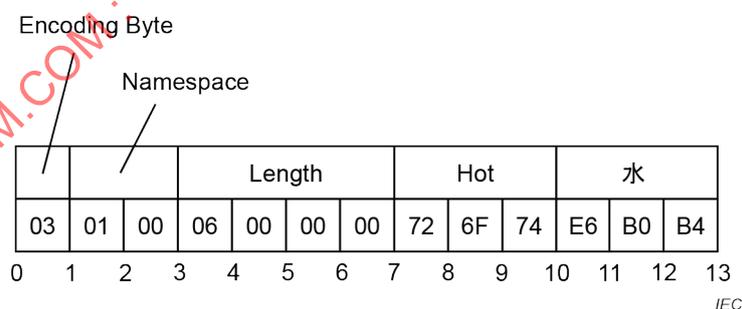
Name	Value	Description
Two Byte	0x00	A numeric value that fits into the two-byte representation.
Four Byte	0x01	A numeric value that fits into the four-byte representation.
Numeric	0x02	A numeric value that does not fit into the two or four byte representations.
String	0x03	A String value.
Guid	0x04	A Guid value.
ByteString	0x05	An opaque (ByteString) value.
NamespaceUri Flag	0x80	See discussion of <i>ExpandedNodeId</i> in 5.2.2.10.
ServerIndex Flag	0x40	See discussion of <i>ExpandedNodeId</i> in 5.2.2.10.

The standard *NodeId DataEncoding* has the structure shown in Table 7. The standard *DataEncoding* is used for all formats that do not have an explicit format defined.

Table 7 – Standard NodeId Binary DataEncoding

Name	Data Type	Description
Namespace	UInt16	The <i>NamespaceIndex</i> .
Identifier	*	The identifier which is encoded according to the following rules: NUMERIC UInt32 STRING String GUID Guid OPAQUE ByteString

An example of a String *NodeId* with Namespace = 1 and Identifier = "Hot水" is shown in Figure 7.

**Figure 7 – A String NodeId**

The Two Byte *NodeId DataEncoding* has the structure shown in Table 8.

Table 8 – Two Byte NodeId Binary DataEncoding

Name	Data Type	Description
Identifier	Byte	The <i>Namespace</i> is the default OPC UA namespace (i.e. 0). The <i>Identifier</i> Type is 'Numeric'. The <i>Identifier</i> shall be in the range 0 to 255.

An example of a Two Byte *NodeId* with Identifier = 72 is shown in Figure 8.

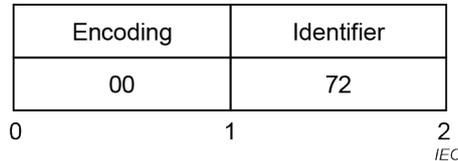


Figure 8 – A Two Byte NodeId

The Four Byte *NodeId DataEncoding* has the structure shown in Table 9.

Table 9 – Four Byte NodeId Binary DataEncoding

Name	Data Type	Description
Namespace	Byte	The <i>Namespace</i> shall be in the range 0 to 255.
Identifier	UInt16	The <i>Identifier</i> Type is 'Numeric'. The <i>Identifier</i> shall be an integer in the range 0 to 65 535.

An example of a Four Byte *NodeId* with Namespace = 5 and Identifier = 1 025 is shown in Figure 9.

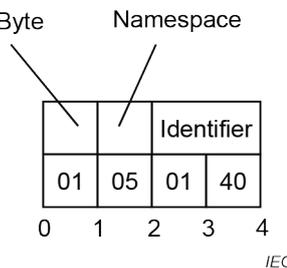


Figure 9 – A Four Byte NodeId

5.2.2.10 ExpandedNodeId

An *ExpandedNodeId* extends the *NodeId* structure by allowing the *NamespaceUri* to be explicitly specified instead of using the *NamespaceIndex*. The *NamespaceUri* is optional. If it is specified, then the *NamespaceIndex* inside the *NodeId* shall be ignored.

The *ExpandedNodeId* is encoded by first encoding a *NodeId* as described in 5.2.2.9 and then encoding *NamespaceUri* as a *String*.

An instance of an *ExpandedNodeId* may still use the *NamespaceIndex* instead of the *NamespaceUri*. In this case, the *NamespaceUri* is not encoded in the stream. The presence of the *NamespaceUri* in the stream is indicated by setting the *NamespaceUri* flag in the encoding format byte for the *NodeId*.

If the *NamespaceUri* is present, then the encoder shall encode the *NamespaceIndex* as 0 in the stream when the *NodeId* portion is encoded. The unused *NamespaceIndex* is included in the stream for consistency.

An *ExpandedNodeId* may also have a *ServerIndex* which is encoded as a *UInt32* after the *NamespaceUri*. The *ServerIndex* flag in the *NodeId* encoding byte indicates whether the *ServerIndex* is present in the stream. The *ServerIndex* is omitted if it is equal to zero.

The *ExpandedNodeId* encoding has the structure shown in Table 10.

Table 10 – ExpandedNodeId Binary DataEncoding

Name	Data Type	Description
NodeId	NodeId	The NamespaceUri and ServerIndex flags in the NodeId encoding indicate whether those fields are present in the stream.
NamespaceUri	String	Not present if null or Empty.
ServerIndex	UInt32	Not present if 0.

5.2.2.11 StatusCode

A *StatusCode* is encoded as a *UInt32*.

5.2.2.12 DiagnosticInfo

A *DiagnosticInfo* structure is described in IEC 62541-4. It specifies a number of fields that could be missing. For that reason, the encoding uses a bit mask to indicate which fields are actually present in the encoded form.

As described in IEC 62541-4, the *SymbolicId*, *NamespaceUri*, *LocalizedText* and *Locale* fields are indexes in a string table which is returned in the response header. Only the index of the corresponding string in the string table is encoded. An index of -1 indicates that there is no value for the string.

DiagnosticInfo allows unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

Table 11 – DiagnosticInfo Binary DataEncoding

Name	Data Type	Description
Encoding Mask	Byte	A bit mask that indicates which fields are present in the stream. The mask has the following bits: 0x01 Symbolic Id 0x02 Namespace 0x04 LocalizedText 0x08 Locale 0x10 Additional Info 0x20 InnerStatusCode 0x40 InnerDiagnosticInfo
SymbolicId	Int32	A symbolic name for the status code.
NamespaceUri	Int32	A namespace that qualifies the symbolic id.
Locale	Int32	The locale used for the localized text.
LocalizedText	Int32	A human readable summary of the status code.
Additional Info	String	Detailed application specific diagnostic information.
Inner StatusCode	StatusCode	A status code provided by an underlying system.
Inner DiagnosticInfo	DiagnosticInfo	Diagnostic info associated with the inner status code.

5.2.2.13 QualifiedName

A *QualifiedName* structure is encoded as shown in Table 12.

The abstract *QualifiedName* structure is defined in IEC 62541-3.

Table 12 – QualifiedName Binary DataEncoding

Name	Data Type	Description
NamespaceIndex	UInt16	The namespace index.
Name	String	The name.

5.2.2.14 LocalizedText

A *LocalizedText* structure contains two fields that could be missing. For that reason, the encoding uses a bit mask to indicate which fields are actually present in the encoded form.

The abstract *LocalizedText* structure is defined in IEC 62541-3.

Table 13 – LocalizedText Binary DataEncoding

Name	Data Type	Description
EncodingMask	Byte	A bit mask that indicates which fields are present in the stream. The mask has the following bits: 0x01 Locale 0x02 Text
Locale	String	The locale. Omitted is null or empty.
Text	String	The text in the specified locale. Omitted is null or empty.

5.2.2.15 ExtensionObject

An *ExtensionObject* is encoded as sequence of bytes prefixed by the *NodeId* of its *DataTypeEncoding* and the number of bytes encoded.

An *ExtensionObject* may be encoded by the application which means it is passed as a *ByteString* or an *XmlElement* to the encoder. In this case, the encoder will be able to write the number of bytes in the object before it encodes the bytes. However, an *ExtensionObject* may know how to encode/decode itself which means the encoder shall calculate the number of bytes before it encodes the object or it shall be able to seek backwards in the stream and update the length after encoding the body.

When a decoder encounters an *ExtensionObject* it shall check if it recognizes the *DataTypeEncoding* identifier. If it does, then it can call the appropriate function to decode the object body. If the decoder does not recognize the type it shall use the *EncodingMask* to determine if the body is a *ByteString* or an *XmlElement* and then decode the object body or treat it as opaque data and skip over it.

The serialized form of an *ExtensionObject* is shown in Table 14.

Table 14 – Extension Object Binary DataEncoding

Name	Data Type	Description
TypeId	NodeId	The identifier for the <i>DataTypeEncoding</i> node in the <i>Server's AddressSpace</i> . <i>ExtensionObjects</i> defined by the OPC UA specification have a numeric node identifier assigned to them with a <i>NamespaceIndex</i> of 0. The numeric identifiers are defined in A.3. Decoders use this field to determine the syntax of the <i>Body</i> . For example, if this field is the <i>NodeId</i> of the <i>JSON Encoding Object</i> for a <i>DataType</i> then the <i>Body</i> is a <i>ByteString</i> containing a JSON document encoded as a UTF-8 string.
Encoding	Byte	An enumeration that indicates how the body is encoded. The parameter may have the following values: 0x00 No body is encoded. 0x01 The body is encoded as a <i>ByteString</i> . 0x02 The body is encoded as a <i>XmlElement</i> .
Length	Int32	The length of the object body. The length shall be specified if the body is encoded.
Body	Byte [*]	The object body. This field contains the raw bytes for <i>ByteString</i> bodies. For <i>XmlElement</i> bodies this field contains the XML encoded as a UTF-8 string without any null terminator. Some binary encoded structures may have a serialized length that is not a multiple of 8 bits. Encoders shall append 0 bits to ensure the serialized length is a multiple of 8 bits. Decoders that understand the serialized format shall ignore the padding bits.

ExtensionObjects are used in two contexts: as values contained in *Variant* structures or as parameters in OPC UA Messages.

A decoder may choose to parse an *XmlElement* body after decoding; if an unrecoverable parsing error occurs then the decoder should try to continue processing the stream. For example, if the *ExtensionObject* is the body of a *Variant* or an element in an array that is the body of *Variant* then this error can be reported by setting the value of the *Variant* to the *StatusCode Bad_DecodingError*.

5.2.2.16 Variant

A *Variant* is a union of the built-in types.

The structure of a *Variant* is shown in Table 15.

Table 15 – Variant Binary DataEncoding

Name	Data Type	Description
EncodingMask	Byte	<p>The type of data encoded in the stream. A value of 0 specifies a NULL and that no other fields are encoded. The mask has the following bits assigned:</p> <p>0:5 Built-in Type Id (see Table 1). 6 True if the Array Dimensions field is encoded. 7 True if an array of values is encoded.</p> <p>The Built-in Type Ids 26 through 31 are not currently assigned but may be used in the future. Decoders shall accept these IDs, assume the <i>Value</i> contains a <i>ByteString</i> and pass both onto the application. Encoders shall not use these IDs.</p>
ArrayLength	Int32	<p>The number of elements in the array. This field is only present if the array bit is set in the encoding mask.</p> <p>Multi-dimensional arrays are encoded as a one-dimensional array and this field specifies the total number of elements. The original array can be reconstructed from the dimensions that are encoded after the value field.</p> <p>Higher rank dimensions are serialized first. For example, an array with dimensions [2,2,2] is written in this order: [0,0,0], [0,0,1], [0,1,0], [0,1,1], [1,0,0], [1,0,1], [1,1,0], [1,1,1]</p>
Value	*	<p>The value encoded according to its built-in data type.</p> <p>If the array bit is set in the encoding mask, then each element in the array is encoded sequentially. Since many types have variable length encoding each element shall be decoded in order.</p> <p>The value shall not be a <i>Variant</i> but it could be an array of <i>Variants</i>.</p> <p>Many implementation platforms do not distinguish between one dimensional Arrays of <i>Bytes</i> and <i>ByteStrings</i>. For this reason, decoders are allowed to automatically convert an Array of <i>Bytes</i> to a <i>ByteString</i>.</p>
ArrayDimensions Length	Int32	<p>The number of dimensions. This field is only present if the ArrayDimensions flag is set in the encoding mask.</p>
ArrayDimensions	Int32[*]	<p>The length of each dimension encoded as a sequence of Int32 values This field is only present if the ArrayDimensions flag is set in the encoding mask. The lower rank dimensions appear first in the array.</p> <p>All dimensions shall be specified and shall be greater than zero.</p> <p>If <i>ArrayDimensions</i> are inconsistent with the <i>ArrayLength</i> then the decoder shall stop and raise a <i>Bad_DecodingError</i>.</p>

The types and their identifiers that can be encoded in a *Variant* are shown in Table 1.

5.2.2.17 DataValue

A *DataValue* is always preceded by a mask that indicates which fields are present in the stream.

The fields of a *DataValue* are described in Table 16.

Table 16 – Data Value Binary DataEncoding

Name	Data Type	Description
Encoding Mask	Byte	A bit mask that indicates which fields are present in the stream. The mask has the following bits: 0x01 False if the Value is <i>Null</i> . 0x02 False if the StatusCode is Good. 0x04 False if the Source Timestamp is <i>DateTime.MinValue</i> . 0x08 False if the Server Timestamp is <i>DateTime.MinValue</i> . 0x10 False if the Source PicoSeconds is 0. 0x20 False if the Server PicoSeconds is 0.
Value	VARIANT	The value. Not present if the Value bit in the EncodingMask is False.
Status	StatusCode	The status associated with the value. Not present if the StatusCode bit in the EncodingMask is False.
SourceTimestamp	DateTime	The source timestamp associated with the value. Not present if the SourceTimestamp bit in the EncodingMask is False.
SourcePicoSeconds	UInt16	The number of 10 picosecond intervals for the SourceTimestamp. Not present if the SourcePicoSeconds bit in the EncodingMask is False. If the source timestamp is missing the picoseconds are ignored.
ServerTimestamp	DateTime	The Server timestamp associated with the value. Not present if the ServerTimestamp bit in the EncodingMask is False.
ServerPicoSeconds	UInt16	The number of 10 picosecond intervals for the ServerTimestamp. Not present if the ServerPicoSeconds bit in the EncodingMask is False. If the Server timestamp is missing the picoseconds are ignored.

The *PicoSeconds* fields store the difference between a high-resolution timestamp with a resolution of 10 picoseconds and the *Timestamp* field value which only has a 100 ns resolution. The *PicoSeconds* fields shall contain values less than 10 000. The decoder shall treat values greater than or equal to 10 000 as the value '9999'.

5.2.3 Decimal

Decimals are encoded as described in 5.1.7.

A *Decimal* does not have a NULL value.

5.2.4 Enumerations

Enumerations are encoded as *Int32* values.

An *Enumeration* does not have a NULL value.

5.2.5 Arrays

One dimensional *Arrays* ~~that occur outside of a Variant~~ are encoded as a sequence of elements preceded by the number of elements encoded as an *Int32* value. If an *Array* is null, then its length is encoded as -1. An *Array* of zero length is different from an *Array* that is null so encoders and decoders shall preserve this distinction.

~~Multi-dimensional arrays can only be encoded within a Variant.~~

Multi-dimensional *Arrays* are encoded as an *Int32 Array* containing the dimensions followed by a list of all the values in the *Array*. The total number of values is equal to the product of the dimensions. The number of values is 0 if one or more dimensions are less than or equal to 0. The process for reconstructing the multi-dimensional array is described in 5.2.2.16.

5.2.6 Structures

Structures are encoded as a sequence of fields in the order that they appear in the definition. The encoding for each field is determined by the built-in type for the field.

All fields specified in the ~~complex type~~ structure shall be encoded. If optional fields exist in the structure then see 5.2.7.

Structures do not have a null value. If an encoder is written in a programming language that allows structures to have null values, then the encoder shall create a new instance with default values for all fields and serialize that. Encoders shall not generate an encoding error in this situation.

The following is an example of a structure using C/C++ syntax:

```

class Type2
{
    int A;
    int B;
}

class Type1
{
    int X;
    int NoOfY;
    Type2* Y;
    int Z;
}

struct Type2
{
    Int32 A;
    Int32 B;
};

struct Type1
{
    Int32 X;
    Byte NoOfY;
    Type2* Y;
    Int32 Z;
};

```

In the C/C++ example above, the Y field is a pointer to an array with a length stored in NoOfY. When encoding an array, the length is part of the array encoding so the NoOfY field is not encoded. That said, encoders and decoders use NoOfY during encoding.

An instance of *Type1* which contains an array of two *Type2* instances would be encoded as ~~37~~ 28-byte sequence. If the instance of *Type1* was encoded in an *ExtensionObject* it would have ~~the encoded form~~ an additional prefix shown in Table 17 which would make the total length 37 bytes. The *TypeId*, Encoding and the *Length* are fields defined by the *ExtensionObject*. The encoding of the *Type2* instances do not include any type identifier because it is explicitly defined in *Type1*.

Table 17 – Sample OPC UA Binary Encoded structure

Field	Bytes	Value
Type Id	4	The identifier for the Type1 Binary Encoding Node
Encoding	1	0x1 for ByteString
Length	4	28
X	4	The value of field 'X'
NoOf Y.Length	4	2
Y.A	4	The value of field 'Y[0].A'
Y.B	4	The value of field 'Y[0].B'
Y.A	4	The value of field 'Y[1].A'
Y.B	4	The value of field 'Y[1].B'
Z	4	The value of field 'Z'

The Value of the *DataTypeDefinition Attribute* for a DataType Node describing Type1 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type1_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure without optional fields]
fields [0]	StructureField	
name	String	"X"
description	LocalizedText	Description of X
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"Y"
description	LocalizedText	Description of Y-Array
dataType	NodeId	NodeId of the Type2 DataType Node (e.g. "ns=3; s=MyType2")
valueRank	Int32	1 (OneDimension)
isOptional	Boolean	false
fields [2]	StructureField	
name	String	"Z"
description	LocalizedText	Description of Z
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false

The *Value* of the *DataTypeDefinition Attribute* for a *DataType Node* describing Type2 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type2_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure without optional fields]
fields [0]	StructureField	
name	String	"A"
description	LocalizedText	Description of A
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"B"
description	LocalizedText	Description of B
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false

5.2.7 Structures with optional fields

Structures with optional fields are encoded with an encoding mask preceding a sequence of fields in the order that they appear in the definition. The encoding for each field is determined by the data type for the field.

The *EncodingMask* is a 32-bit unsigned integer. Each optional field is assigned exactly one bit. The first optional field is assigned bit '0', the second optional field is assigned bit '1' and so on until all optional fields are assigned bits. A maximum of 32 optional fields can appear within a single *Structure*. Unassigned bits are set to 0 by encoders. Decoders shall report an error if unassigned bits are not 0.

The following is an example of a structure with optional fields using C++ syntax:

```
struct TypeA
{
    Int32 X;
    Int32* O1;
    SByte Y;
    Int32* O2;
};
```

O1 and O2 are optional fields which are NULL if not present.

An instance of *TypeA* which contains two mandatory (X and Y) and two optional (O1 and O2) fields would be encoded as a byte sequence. The length of the byte sequence depends on the available optional fields. An encoding mask field determines the available optional fields.

An instance of *TypeA* where field O2 is available and field O1 is not available would be encoded as a 13-byte sequence. If the instance of *TypeA* was encoded in an *ExtensionObject* it would have the encoded form shown in Table 18 and have a total length of 22 bytes. The length of the *TypeId*, *Encoding* and the *Length* are fields defined by the *ExtensionObject*.

Table 18 – Sample OPC UA Binary Encoded Structure with optional fields

Field	Bytes	Value
Type Id	4	The identifier for the TypeA Binary Encoding Node
Encoding	1	0x1 for ByteString
Length	4	13
EncodingMask	4	0x02 for O2
X	4	The value of X
Y	1	The value of Y
O2	4	The value of O2

If a *Structure* with optional fields is subtyped, the subtypes extend the *EncodingMask* defined for the parent.

The Value of the *DataTypeDefinition* Attribute for a *DataType* Node describing *TypeA* is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "TypeA_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	StructureWithOptionalFields_1 [Structure without optional fields]
fields [0]	StructureField	
name	String	"X"
description	LocalizedText	Description of X
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"O1"
description	LocalizedText	Description of O1
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true
fields [2]	StructureField	
name	String	"Y"
description	LocalizedText	Description of Z
dataType	NodeId	"i=2" [SByte]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [3]	StructureField	
name	String	"O2"
description	LocalizedText	Description of O2
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true

5.2.8 Unions

Unions are encoded as a switch field preceding one of the possible fields. The encoding for the selected field is determined by the data type for the field.

The switch field is encoded as a UInt32.

The switch field is the index of the available union fields starting with 1. If the switch field is 0 then no field is present. For any value greater than the number of defined union fields the encoders or decoders shall report an error.

A *Union* with no fields present has the same meaning as a NULL value. A *Union* with any field present is not a NULL value even if the value of the field itself is NULL.

The following is an example of a union using C/C++ syntax:

```

struct Type2
{
    Int32 A;
    Int32 B;
};

struct Type1
{
    Byte Selector;

    union
    {
        Int32 Field1;
        Type2 Field2;
    }
    Value;
};

```

In the C/C++ example above, the Selector and Value are semantically coupled to form a union. The order of the fields does not matter.

An instance of *Type1* would be encoded as a byte sequence. The length of the byte sequence depends on the selected field.

An instance of *Type1* where field *Field1* is available would be encoded as 8-byte sequence. If the instance of *Type1* was encoded in an *ExtensionObject* it would have the encoded form shown in Table 19 and it would have a total length of 17 bytes. The *TypeId*, *Encoding* and the *Length* are fields defined by the *ExtensionObject*.

Table 19 – Sample OPC UA Binary Encoded Structure

Field	Bytes	Value
Type Id	4	The identifier for Type1
Encoding	1	0x1 for ByteString
Length	4	8
SwitchValue	4	1 for Field1
Field1	4	The value of Field1

The Value of the *DataTypeDefinition Attribute* for a *DataType Node* describing Type1 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type1_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Union]
structureType	StructureType	Union_2 [Union]
fields [0]	StructureField	
name	String	"Field1"
description	LocalizedText	Description of Field1
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true
fields [1]	StructureField	
name	String	"Field2"
description	LocalizedText	Description of Field2
dataType	NodeId	NodeId of the Type2 DataType Node (e.g. "ns=3; s=MyType2")
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true

The Value of the *DataTypeDefinition Attribute* for a *DataType Node* describing Type2 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type2_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure without optional fields]
fields [0]	StructureField	
name	String	"A"
description	LocalizedText	Description of A
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"B"
description	LocalizedText	Description of B
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false

5.2.9 Messages

~~Messages are encoded as ExtensionObjects. The parameters in each Message are serialized in the same way the fields of a Structure are serialized. The TypeId field contains the DataTypeEncoding identifier for the Message. The Length field is omitted since the Messages are defined by this series of OPC UA standards.~~

~~Each OPC UA Service described in IEC 62541-4 has a request and response Message. The DataTypeEncoding IDs assigned to each Service are given in A.3.~~

Messages are Structures encoded as sequence of bytes prefixed by the *NodeId* of for the OPC UA Binary *DataTypeEncoding* defined for the Message.

Each OPC UA *Service* described in IEC 62541-4 has a request and response *Message*. The *DataTypeEncoding* IDs assigned to each *Service* are specified in Clause A.3.

5.3 OPC UA XML

5.3.1 Built-in Types

5.3.1.1 General

Most built-in types are encoded in XML using the formats defined in XML Schema Part 2 specification. Any special restrictions or usages are discussed below. Some of the built-in types have an XML Schema defined for them using the syntax defined in XML Schema Part 42.

The prefix *xs:* is used to denote a symbol defined by the XML Schema specification.

5.3.1.2 Boolean

A Boolean value is encoded as an *xs:boolean* value.

5.3.1.3 Integer

Integer values are encoded using one of the subtypes of the *xs:decimal* type. The mappings between the OPC UA integer types and XML schema data types are shown in Table 20.

Table 20 – XML Data Type Mappings for Integers

Name	XML Type
SByte	xs:byte
Byte	xs:unsignedByte
Int16	xs:short
UInt16	xs:unsignedShort
Int32	xs:int
UInt32	xs:unsignedInt
Int64	xs:long
UInt64	xs:unsignedLong

5.3.1.4 Floating Point

Floating point values are encoded using one of the XML floating point types. The mappings between the OPC UA floating point types and XML schema data types are shown in Table 21.

Table 21 – XML Data Type Mappings for Floating Points

Name	XML Type
Float	xs:float
Double	xs:double

The XML floating point type supports positive infinity (INF), negative infinity (-INF) and not-a-number (NaN).

5.3.1.5 String

A *String* value is encoded as an *xs:string* value.

5.3.1.6 DateTime

A *DateTime* value is encoded as an *xs:dateTime* value.

All *DateTime* values shall be encoded as UTC times or with the time zone explicitly specified.

Correct:

```
2002-10-10T00:00:00+05:00
2002-10-09T19:00:00Z
```

Incorrect:

```
2002-10-09T19:00:00
```

It is recommended that all *xs:dateTime* values be represented in UTC format.

The earliest and latest date/time values that can be represented on a *DevelopmentPlatform* have special meaning and shall not be literally encoded in XML.

The earliest date/time value on a *DevelopmentPlatform* shall be encoded in XML as '0001-01-01T00:00:00Z'.

The latest date/time value on a *DevelopmentPlatform* shall be encoded in XML as '9999-12-31T23:59:59Z'

If a decoder encounters a *xs:dateTime* value that cannot be represented on the *DevelopmentPlatform* it should convert the value to either the earliest or latest date/time that can be represented on the *DevelopmentPlatform*. The XML decoder should not generate an error if it encounters an out of range date value.

The earliest date/time value on a *DevelopmentPlatform* is equivalent to a null date/time value.

5.3.1.7 Guid

A *Guid* is encoded using the string representation defined in 5.1.3.

The XML schema for a *Guid* is:

```
<xs:complexType name="Guid">
  <xs:sequence>
    <xs:element name="String" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.8 ByteString

A *ByteString* value is encoded as an *xs:base64Binary* value (see Base64).

The XML schema for a *ByteString* is:

```
<xs:element name="ByteString" type="xs:base64Binary" nillable="true"/>
```

5.3.1.9 XmlElement

An *XmlElement* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="XmlElement">
  <xs:sequence>
    <xs:any minOccurs="0" maxOccurs="1" processContents="lax" />
  </xs:sequence>
</xs:complexType>
```

XmlElements may only be used inside *Variant* or *ExtensionObject* values.

5.3.1.10 NodeId

A *NodeId* value is encoded as an *xs:string* with the syntax:

```
ns=<namespaceindex>;<type>=<value>
```

The elements of the syntax are described in Table 22.

Table 22 – Components of NodeId

Field	Data Type	Description
<namespaceindex>	UInt16	The <i>NamespaceIndex</i> formatted as a base 10 number. If the index is 0 then the entire 'ns=0;' clause shall be omitted.
<type>	Enumeration	A flag that specifies the <i>IdentifierType</i> . The flag has the following values: i NUMERIC (<i>UInteger</i> UInt32) s STRING (String) g GUID (Guid) b OPAQUE (ByteString)
<value>	*	The <i>Identifier</i> encoded as string. The <i>Identifier</i> is formatted using the XML data type mapping for the <i>IdentifierType</i> . Note that the <i>Identifier</i> may contain any non-null UTF-8 character including whitespace.

Examples of *NodeIds*:

```
i=13
ns=10;i=-1
ns=10;s>Hello:World
g=09087e75-8e5e-499b-954f-f2a9603db28a
ns=1;b=M/RbKBSRVkePCePcx24oRA==
```

The XML schema for a *NodeId* is:

```
<xs:complexType name="NodeId">
  <xs:sequence>
    <xs:element name="Identifier" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.11 ExpandedNodeId

An *ExpandedNodeId* value is encoded as an *xs:string* with the syntax:

```
svr=<serverindex>;ns=<namespaceindex>;<type>=<value>
or
svr=<serverindex>;nsu=<uri>;<type>=<value>
```

The possible fields are shown in Table 23.

Table 23 – Components of ExpandedNodeId

Field	Data Type	Description
<serverindex>	UInt32	The <i>ServerIndex</i> formatted as a base 10 number. If the <i>ServerIndex</i> is 0 then the entire 'svr=0;' clause shall be omitted.
<namespaceindex>	UInt16	The <i>NamespaceIndex</i> formatted as a base 10 number. If the <i>NamespaceIndex</i> is 0 then the entire 'ns=0;' clause shall be omitted. The <i>NamespaceIndex</i> shall not be present if the URI is present.
<uri>	String	The <i>NamespaceUri</i> formatted as a string. Any reserved characters in the URI shall be replaced with a '%' followed by its 8 bit ANSI value encoded as two hexadecimal digits (case insensitive). For example, the character ';' would be replaced by '%3B'. The reserved characters are ';' and '%'. If the <i>NamespaceUri</i> is null or empty, then 'nsu=' clause shall be omitted.
<type>	Enumeration	A flag that specifies the <i>IdentifierType</i> . This field is described in Table 22.
<value>	*	The <i>Identifier</i> encoded as string. This field is described in Table 22.

The XML schema for an *ExpandedNodeId* is:

```
<xs:complexType name="ExpandedNodeId">
  <xs:sequence>
    <xs:element name="Identifier" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.12 StatusCode

A *StatusCode* is encoded as an *xs:unsignedInt* with the following XML schema:

```
<xs:complexType name="StatusCode">
  <xs:sequence>
    <xs:element name="Code" type="xs:unsignedInt" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.13 DiagnosticInfo

An *DiagnosticInfo* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="DiagnosticInfo">
  <xs:sequence>
    <xs:element name="SymbolicId" type="xs:int" minOccurs="0" />
    <xs:element name="NamespaceUri" type="xs:int" minOccurs="0" />
<xs:element name="LocalizedText" type="xs:int" minOccurs="0"/>
    <xs:element name="Locale" type="xs:int" minOccurs="0"/>
    <xs:element name="LocalizedText" type="xs:int" minOccurs="0"/>
    <xs:element name="AdditionalInfo" type="xs:string" minOccurs="0"/>
    <xs:element name="InnerStatusCode" type="tns:StatusCode"
      minOccurs="0" />
    <xs:element name="InnerDiagnosticInfo" type="tns:DiagnosticInfo"
      minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

DiagnosticInfo allows unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

5.3.1.14 QualifiedName

A *QualifiedName* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="QualifiedName">
  <xs:sequence>
    <xs:element name="NamespaceIndex" type="xs:int" minOccurs="0" />
    <xs:element name="Name" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.15 LocalizedText

A *LocalizedText* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="LocalizedText">
  <xs:sequence>
    <xs:element name="Locale" type="xs:string" minOccurs="0" />
    <xs:element name="Text" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.16 ExtensionObject

An *ExtensionObject* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="ExtensionObject">
  <xs:sequence>
    <xs:element name="TypeId" type="tns:NodeId" minOccurs="0" />
    <xs:element name="Body" minOccurs="0">
      <xs:complexType>
        <xs:sequence>
          <xs:any minOccurs="0" processContents="lax"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

The body of the *ExtensionObject* contains a single element which is either a *ByteString* or XML encoded *Structure*. A decoder can distinguish between the two by inspecting the top-level element. An element with the name `tns:ByteString` contains an OPC UA Binary encoded body. Any other name shall contain an OPC UA XML encoded body. The *TypeId* specifies the syntax of a *ByteString* body which could be UTF-8 encoded JSON, UA Binary or some other format.

The *TypeId* is the *NodeId* for the *DataTypeEncoding Object*.

5.3.1.17 Variant

A *Variant* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="Variant">
  <xs:sequence>
    <xs:element name="Value" minOccurs="0" nillable="true">
      <xs:complexType>
        <xs:sequence>
          <xs:any minOccurs="0" processContents="lax"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

If the *Variant* represents a scalar value, then it shall contain a single child element with the name of the built-in type. For example, the single precision floating point value 3,141 5 would be encoded as:

```
<tns:Float>3.1415</tns:Float>
```

If the *Variant* represents a single dimensional array, then it shall contain a single child element with the prefix 'Listof' and the name built-in type. For example, an *Array* of strings would be encoded as:

```
<tns:ListofString>
  <tns:String>Hello</tns:String>
  <tns:String>World</tns:String>
</tns:ListofString>
```

If the *Variant* represents a multidimensional *Array*, then it shall contain a child element with the name '*Matrix*' with the two sub-elements shown in this example:

```
<tns:Matrix>
  <tns:Dimensions>
    <tns:Int32>2</tns:Int32>
    <tns:Int32>2</tns:Int32>
  </tns:Dimensions>
  <tns:Elements>
    <tns:String>A</tns:String>
    <tns:String>B</tns:String>
    <tns:String>C</tns:String>
    <tns:String>D</tns:String>
  </tns:Elements>
</tns:Matrix>
```

In this example, the array has the following elements:

```
[0,0] = "A"; [0,1] = "B"; [1,0] = "C"; [1,1] = "D"
```

The elements of a multi-dimensional *Array* are always flattened into a single dimensional *Array* where the higher rank dimensions are serialized first. This single dimensional *Array* is encoded as a child of the 'Elements' element. The 'Dimensions' element is an *Array* of *Int32* values that specify the dimensions of the array starting with the lowest rank dimension. The multi-dimensional *Array* can be reconstructed by using the dimensions encoded. All dimensions shall be specified and shall be greater than zero. If the dimensions are inconsistent with the number of elements in the array, then the decoder shall stop and raise a *Bad_DecodingError*.

The complete set of built-in type names is found in Table 1.

5.3.1.18 DataValue

A *DataValue* value is encoded as a *xs:complexType* with the following XML schema:

```
<xs:complexType name="DataValue">
  <xs:sequence>
    <xs:element name="Value" type="tns:Variant" minOccurs="0"
      nillable="true" />
    <xs:element name="StatusCode" type="tns:StatusCode"
      minOccurs="0" />
    <xs:element name="SourceTimestamp" type="xs:dateTime"
      minOccurs="0" />
    <xs:element name="SourcePicoseconds" type="xs:unsignedShort"
      minOccurs="0"/>
    <xs:element name="ServerTimestamp" type="xs:dateTime"
      minOccurs="0" />
    <xs:element name="ServerPicoseconds" type="xs:unsignedShort"
      minOccurs="0"/>
  </xs:sequence>
</xs:complexType>
```

5.3.2 Decimal

A *Decimal Value* is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="Decimal">
  <xs:sequence>
    <xs:element name="TypeId" type="tns:NodeId" minOccurs="0" />
    <xs:element name="Body" minOccurs="0">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="Scale" type="xs:unsignedShort" />
          <xs:element name="Value" type="xs:string" />
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

The *NodeId* is always the *NodeId* of the *Decimal DataType*. When encoded in a *Variant* the *Decimal* is encoded as an *ExtensionObject*. Arrays of *Decimals* are *Arrays* of *ExtensionObjects*.

The *Value* is a base-10 signed integer with no limit on size. See 5.1.7 for a description of the *Scale* and *Value* fields.

5.3.3 Enumerations

Enumerations that are used as parameters in the *Messages* defined in IEC 62541-4 are encoded as *xs:string* with the following syntax:

<symbol>_<value>

The elements of the syntax are described in Table 24.

Table 24 – Components of Enumeration

Field	Type	Description
<symbol>	String	The symbolic name for the enumerated value.
<value>	UInt32	The numeric value associated with enumerated value.

For example, the XML schema for the *NodeClass* enumeration is:

```
<xs:simpleType name="NodeClass">
  <xs:restriction base="xs:string">
    <xs:enumeration value="Unspecified_0" />
    <xs:enumeration value="Object_1" />
    <xs:enumeration value="Variable_2" />
    <xs:enumeration value="Method_4" />
    <xs:enumeration value="ObjectType_8" />
    <xs:enumeration value="VariableType_16" />
    <xs:enumeration value="ReferenceType_32" />
    <xs:enumeration value="DataType_64" />
    <xs:enumeration value="View_128" />
  </xs:restriction>
</xs:simpleType>
```

Enumerations that are stored in a *Variant* are encoded as an *Int32* value.

For example, any *Variable* could have a value with a *DataType* of *NodeClass*. In this case, the corresponding numeric value is placed in the *Variant* (e.g. *NodeClass Object* would be stored as a 1).

5.3.4 Arrays

One dimensional *Array* parameters are always encoded by wrapping the elements in a container element and inserting the container into the structure. The name of the container element should be the name of the parameter. The name of the element in the array shall be the type name.

For example, the *Read* service takes an array of *ReadValueIds*. The XML schema would look like:

```
<xs:complexType name="ListOfReadValueId">
  <xs:sequence>
    <xs:element name="ReadValueId" type="tns:ReadValueId"
      minOccurs="0" maxOccurs="unbounded" nillable="true" />
  </xs:sequence>
</xs:complexType>
```

The nillable attribute shall be specified because XML encoders will drop elements in arrays if those elements are empty.

Multi-dimensional *Array* parameters are encoded using the *Matrix* type defined in 5.3.1.17.

5.3.5 Structures

Structures are encoded as a *xs:complexType* with all of the fields appearing in a sequence. All fields are encoded as an *xs:element* and have *xs:maxOccurs* set to 1. All elements have *minOccurs* set 0 to allow for compact XML representations. If an element is missing the

default value for the field type is used. If the field type is a structure, the default value is an instance of the structure with default values for each contained field.

Types which have a NULL value defined shall have the `nillable="true"` flag set.

For example, the Read service has a *ReadValueId* structure in the request. The XML schema would look like:

```
<xs:complexType name="ReadValueId">
  <xs:sequence>
    <xs:element name="NodeId" type="tns:NodeId" minOccurs="1" />
      minOccurs="0" nillable="true" />
    <xs:element name="AttributeId" type="xs:int" minOccurs="1" />
    <xs:element name="IndexRange" type="xs:string"
      minOccurs="0" nillable="true" />
    <xs:element name="DataEncoding" type="tns:NodeId" minOccurs="1" />
      minOccurs="0" nillable="true" />
  </xs:sequence>
</xs:complexType>
```

5.3.6 Structures with optional fields

Structures with optional fields are encoded as an *xs:complexType* with all of the fields appearing in a sequence. The first element is a bit mask that specifies what fields are encoded. The bits in the mask are sequentially assigned to optional fields in the order they appear in the *Structure*.

To allow for compact XML, any field can be omitted from the XML so decoders shall assign default values based on the field type for any mandatory fields.

For example, the following *Structure* has one mandatory and two optional fields. The XML schema would look like:

```
<xs:complexType name="OptionalType">
  <xs:sequence>
    <xs:element name="EncodingMask" type="xs:unsignedLong" />
    <xs:element name="X" type="xs:int" minOccurs="0" />
    <xs:element name="O1" type="xs:int" minOccurs="0" />
    <xs:element name="Y" type="xs:byte" minOccurs="0" />
    <xs:element name="O2" type="xs:int" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

In the example above, the *EncodingMask* has a value of 3 if both *O1* and *O2* are encoded. Encoders shall set unused bits to 0 and decoders shall ignore unused bits.

If a *Structure* with optional fields is subtyped, the subtypes extend the *EncodingMask* defined for the parent.

5.3.7 Unions

Unions are encoded as an *xs:complexType* containing an *xs:sequence* with two entries.

The first entry in the sequence is the *SwitchField xs:element* and specifies a numeric value which identifies which element in the *xs:choice* is encoded. The name of the element may be any valid text.

The second entry in the sequence is an *xs:choice* which specifies the possible fields. The order in the *xs:choice* determines the value of the *SwitchField* when that choice is encoded.

The first element has a *SwitchField* value of 1 and the last value has a *SwitchField* equal to the number of choices.

No additional elements in the sequence are permitted. If the *SwitchField* is missing or 0 then the union has a NULL value. Encoders or decoders shall report an error for any *SwitchField* value greater than the number of defined union fields.

For example, the following union has two fields. The XML schema would look like:

```
<xs:complexType name="Type1">
  <xs:sequence>
    <xs:element name="SwitchField"
      type="xs:unsignedInt" minOccurs="0"/>
    <xs:choice>
      <xs:element name="Field1" type="xs:int" minOccurs="0"/>
      <xs:element name="Field2" type="tns:Field2" minOccurs="0"/>
    </xs:choice>
  </xs:sequence>
</xs:complexType>
```

5.3.8 Messages

Messages are encoded as an `xs:complexType`. The parameters in each *Message* are serialized in the same way the fields of a *Structure* are serialized.

5.4 OPC UA JSON

5.4.1 General

The JSON *DataEncoding* was developed to allow OPC UA applications to interoperate with web and enterprise software that use this format. The OPC UA JSON *DataEncoding* defines standard JSON representations for all OPC UA Built-In types.

The JSON format is defined in RFC 7159. It is partially self-describing because each field has a name encoded in addition to the value; however, JSON has no mechanism to qualify names with namespaces.

The JSON format does not have a published standard for a schema that can be used to describe the contents of a JSON document. However, the schema mechanisms defined in this document can be used to describe JSON documents. Specifically, the *DataTypeDescription* structure defined in IEC 62541-3 can define any JSON document that conforms to the rules described below.

Servers that support the JSON *DataEncoding* shall add *DataTypeEncoding Nodes* called "Default JSON" to all *DataTypes* which can be serialized with the JSON encoding. The *NodeIds* of these *Nodes* are defined by the information model which defines the *DataType*. These *NodeIds* are used in *ExtensionObjects* as described in 5.4.2.16.

There are two important use cases for the JSON encoding: Cloud applications which consume *PubSub* messages and JavaScript *Clients* (JSON is the preferred serialization format for JavaScript). For the Cloud application use case, the *PubSub* message needs to be self-contained which implies it cannot contain numeric references to an externally defined namespace table. Cloud applications also often rely on scripting languages to process the incoming messages, so artefacts in the *DataEncoding* that exist to ensure fidelity during decoding are not necessary. For this reason, this *DataEncoding* defines a 'non-reversible' form which is designed to meet the needs of Cloud applications. Applications, such as JavaScript Clients, which use the *DataEncoding* for communication with other OPC UA applications use the normal or 'reversible' form. The differences, if any, between the reversible and non-reversible forms are described for each type.

5.4.2 Built-in Types

5.4.2.1 General

Any value for a Built-In type that is NULL shall be encoded as the JSON literal 'null' if the value is an element of an array. If the NULL value is a field within a *Structure* or *Union*, the field shall not be encoded.

5.4.2.2 Boolean

A *Boolean* value shall be encoded as the JSON literal 'true' or 'false'.

5.4.2.3 Integer

Integer values other than *Int64* and *UInt64* shall be encoded as a JSON number.

Int64 and *UInt64* values shall be formatted as a decimal number encoded as a JSON string

(See the XML encoding of 64-bit values described in 5.3.1.3).

5.4.2.4 Floating point

Normal *Float* and *Double* values shall be encoded as a JSON number.

Special floating-point numbers such as positive infinity (INF), negative infinity (-INF) and not-a-number (NaN) shall be represented by the values "Infinity", "-Infinity" and "NaN" encoded as a JSON string. See 5.2.2.3 for more information on the different types of special floating-point numbers.

5.4.2.5 String

String values shall be encoded as JSON strings.

Any characters which are not allowed in JSON strings are escaped using the rules defined in RFC 7159.

5.4.2.6 DateTime

DateTime values shall be formatted as specified by ISO 8601-1:2019 and encoded as a JSON string.

DateTime values which exceed the minimum or maximum values supported on a platform shall be encoded as "0001-01-01T00:00:00Z" or "9999-12-31T23:59:59Z", respectively. During decoding, these values shall be converted to the minimum or maximum values supported on the platform.

DateTime values equal to "0001-01-01T00:00:00Z" are considered to be NULL values.

5.4.2.7 Guid

Guid values shall be formatted as described in 5.1.3 and encoded as a JSON string.

5.4.2.8 ByteString

ByteString values shall be formatted as a Base64 text and encoded as a JSON string.

Any characters which are not allowed in JSON strings are escaped using the rules defined in RFC 7159.

5.4.2.9 XmlElement

XmlElement value shall be encoded as a String as described in 5.3.1.9.

5.4.2.10 NodeId

NodeId values shall be encoded as a JSON object with the fields defined in Table 25.

The abstract *NodeId* structure is defined in IEC 62541-3 and has three fields: *Identifier*, *IdentifierType* and *NamespaceIndex*. The representation of these abstract fields is described in Table 25.

Table 25 – JSON Object Definition for a NodeId

Name	Description
IdType	The <i>IdentifierType</i> encoded as a JSON number. Allowed values are: 0 – <i>UInt32 Identifier</i> encoded as a JSON number. 1 – A <i>String Identifier</i> encoded as a JSON string. 2 – A <i>Guid Identifier</i> encoded as described in 5.4.2.7. 3 – A <i>ByteString Identifier</i> encoded as described in 5.4.2.8. This field is omitted for <i>UInt32</i> identifiers.
Id	The <i>Identifier</i> . The value of the <i>id</i> field specifies the encoding of this field.
Namespace	The <i>NamespaceIndex</i> for the <i>NodeId</i> . The field is encoded as a JSON number for the reversible encoding. The field is omitted if the <i>NamespaceIndex</i> equals 0. For the non-reversible encoding, the field is the <i>NamespaceUri</i> associated with the <i>NamespaceIndex</i> , encoded as a JSON string. A <i>NamespaceIndex</i> of 1 is always encoded as a JSON number.

5.4.2.11 ExpandedNodeId

ExpandedNodeId values shall be encoded as a JSON object with the fields defined in Table 26.

The abstract *ExpandedNodeId* structure is defined in IEC 62541-3 and has five fields: *Identifier*, *IdentifierType*, *NamespaceIndex*, *NamespaceUri* and *ServerIndex*. The representation of these abstract fields is described in Table 26.

Table 26 – JSON Object Definition for an ExpandedNodeId

Name	Description
IdType	The <i>IdentifierType</i> encoded as a JSON number. Allowed values are: 0 – <i>UInt32 Identifier</i> encoded as a JSON number. 1 – A <i>String Identifier</i> encoded as a JSON string. 2 – A <i>Guid Identifier</i> encoded as described in 5.4.2.7. 3 – A <i>ByteString Identifier</i> encoded as described in 5.4.2.8. This field is omitted for UInt32 identifiers.
Id	The <i>Identifier</i> . The value of the 't' field specifies the encoding of this field.
Namespace	The <i>NamespaceIndex</i> or the <i>NamespaceUri</i> for the <i>ExpandedNodeId</i> . If the <i>NamespaceUri</i> is not specified, the <i>NamespaceIndex</i> is encoded with these rules: The field is encoded as a JSON number for the reversible encoding. The field is omitted if the <i>NamespaceIndex</i> equals 0. For the non-reversible encoding the field is the <i>NamespaceUri</i> associated with the <i>NamespaceIndex</i> encoded as a JSON string. A <i>NamespaceIndex</i> of 1 is always encoded as a JSON number. If the <i>NamespaceUri</i> is specified it is encoded as a JSON string in this field.
ServerUri	The <i>ServerIndex</i> for the <i>ExpandedNodeId</i> . This field is encoded as a JSON number for the reversible encoding. This field is omitted if the <i>ServerIndex</i> equals 0. For the non-reversible encoding, this field is the <i>ServerUri</i> associated with the <i>ServerIndex</i> portion of the <i>ExpandedNodeId</i> , encoded as a JSON string.

5.4.2.12 StatusCode

StatusCode values shall be encoded as a JSON number for the reversible encoding.

For the non-reversible form, *StatusCode* values shall be encoded as a JSON object with the fields defined in Table 27.

Table 27 – JSON Object Definition for a StatusCode

Name	Description
Code	The numeric code encoded as a JSON number. The <i>Code</i> is omitted if the numeric code is 0 (Good).
Symbol	The string literal associated with the numeric code encoded as JSON string. e.g. 0x80AB0000 has the associated literal "BadInvalidArgument". The <i>Symbol</i> is omitted if the numeric code is 0 (Good).

A *StatusCode* of Good (0) is treated like a NULL and not encoded. If it is an element of an JSON array it is encoded as the JSON literal 'null'.

5.4.2.13 DiagnosticInfo

DiagnosticInfo values shall be encoded as a JSON object with the fields shown in Table 28.

Table 28 – JSON Object Definition for a DiagnosticInfo

Name	Data Type	Description
SymbolicId	Int32	A symbolic name for the status code.
NamespaceUri	Int32	A namespace that qualifies the symbolic id.
Locale	Int32	The locale used for the localized text.
LocalizedText	Int32	A human readable summary of the status code.
AdditionalInfo	String	Detailed application-specific diagnostic information.
InnerStatusCode	StatusCode	A status code provided by an underlying system.
InnerDiagnosticInfo	DiagnosticInfo	Diagnostic info associated with the inner status code.

Each field is encoded using the rules defined for the built-in type specified in the Data Type column.

The *SymbolicId*, *NamespaceUri*, *Locale* and *LocalizedText* fields are encoded as JSON numbers which reference the *StringTable* contained in the *ResponseHeader*.

5.4.2.14 QualifiedName

QualifiedName values shall be encoded as a JSON object with the fields shown in Table 29.

The abstract *QualifiedName* structure is defined in IEC 62541-3 and has two fields *Name* and *NamespaceIndex*. The *NamespaceIndex* is represented by the *Uri* field in the JSON object.

Table 29 – JSON Object Definition for a QualifiedName

Name	Description
Name	The Name component of the <i>QualifiedName</i> .
Uri	The <i>NamespaceIndex</i> component of the <i>QualifiedName</i> encoded as a JSON number. The <i>Uri</i> field is omitted if the <i>NamespaceIndex</i> equals 0. For the non-reversible form, the <i>NamespaceUri</i> associated with the <i>NamespaceIndex</i> portion of the <i>QualifiedName</i> is encoded as JSON string unless the <i>NamespaceIndex</i> is 1 or if <i>NamespaceUri</i> is unknown. In these cases, the <i>NamespaceIndex</i> is encoded as a JSON number.

5.4.2.15 LocalizedText

LocalizedText values shall be encoded as a JSON object with the fields shown in Table 30.

The abstract *LocalizedText* structure is defined in IEC 62541-3 and has two fields *Text* and *Locale*.

Table 30 – JSON Object Definition for a LocalizedText

Name	Description
Locale	The <i>Locale</i> portion of <i>LocalizedText</i> values shall be encoded as a JSON string
Text	The <i>Text</i> portion of <i>LocalizedText</i> values shall be encoded as a JSON string.

For the non-reversible form, *LocalizedText* value shall be encoded as a JSON string containing the *Text* component.

5.4.2.16 ExtensionObject

ExtensionObject values shall be encoded as a JSON object with the fields shown in Table 31.

Table 31 – JSON Object Definition for an ExtensionObject

Name	Description
TypeId	The <i>NodId</i> of a <i>DataTypeEncoding Node</i> formatted using the rules in 5.4.2.10.
Encoding	The format of the <i>Body</i> field encoded as a JSON number. This value is 0 if the body is <i>Structure</i> encoded as a JSON object (see 5.4.6). This value is 1 if the body is a <i>ByteString</i> value encoded as a JSON string (see 5.4.2.8). This value is 2 if the body is an <i>XmlElement</i> value encoded as a JSON string (see 5.4.2.9). This field is omitted if the value is 0.
Body	<i>Body</i> of the <i>ExtensionObject</i> . The type of this field is specified by the <i>Encoding</i> field. If the <i>Body</i> is empty, the <i>ExtensionObject</i> is NULL and is omitted or encoded as a JSON null.

For the non-reversible form, *ExtensionObject* values shall be encoded as a JSON object containing only the value of the *Body* field. The *TypeId* and *Encoding* fields are dropped.

5.4.2.17 Variant

Variant values shall be encoded as a JSON object with the fields shown in Table 32.

Table 32 – JSON Object Definition for a Variant

Name	Description
Type	The Built-in type for the value contained in the <i>Body</i> (see Table 1) encoded as a JSON number. If type is 0 (NULL) the <i>Variant</i> contains a NULL value and the containing JSON object shall be omitted or replaced by the JSON literal 'null' (when an element of a JSON array).
Body	If the value is a scalar it is encoded using the rules for type specified for the <i>Type</i> . If the value is a one-dimensional array it is encoded as a JSON array (see 5.4.5). Multi-dimensional arrays are encoded as a one dimensional JSON array which is reconstructed using the value of the <i>Dimensions</i> field (see 5.2.2.16).
Dimensions	The dimensions of the array encoded as a JSON array of JSON numbers. The <i>Dimensions</i> are omitted for scalar and one-dimensional array values.

For the non-reversible form, *Variant* values shall be encoded as a JSON object containing only the value of the *Body* field. The *Type* and *Dimensions* fields are dropped. Multi-dimensional arrays are encoded as a multi-dimensional JSON array as described in 5.4.5.

5.4.2.18 DataValue

DataValue values shall be encoded as a JSON object with the fields shown in Table 33.

Table 33 – JSON Object Definition for a DataValue

Name	Data Type	Description
Value	Variant	The value.
Status	StatusCode	The status associated with the value.
SourceTimestamp	DateTime	The source timestamp associated with the value.
SourcePicoSeconds	UInt16	The number of 10 picosecond intervals for the SourceTimestamp.
ServerTimestamp	DateTime	The <i>Server</i> timestamp associated with the value.
ServerPicoSeconds	UInt16	The number of 10 picosecond intervals for the ServerTimestamp.

If a field has a null or default value it is omitted. Each field is encoded using the rules defined for the built-in type specified in the Data Type column.

5.4.3 Decimal

Decimal values shall be encoded as a JSON object with the fields in Table 34.

Table 34 – JSON Object Definition for a Decimal

Name	Description
Scale	A JSON number with the scale applied to the Value.
Value	A JSON string with the Value encoded as a base-10 signed integer. (See the XML encoding of Integer values described in 5.3.1.3).

See 5.1.7 for a description of the *Scale* and *Value* fields.

5.4.4 Enumerations

Enumeration values shall be encoded as a JSON number for the reversible encoding.

For the non-reversible form, *Enumeration* values are encoded as literals with the value appended as a JSON string.

The format of the string literal is:

<name>_<value>

where the name is the enumeration literal and the value is the numeric value.

If the literal is not known to the encoder, the numeric value is encoded as a JSON string.

5.4.5 Arrays

One dimensional *Arrays* shall be encoded as JSON arrays.

If an element is NULL, the element shall be encoded as the JSON literal 'null'.

Otherwise, the element is encoded according to the rules defined for the type.

Multi-dimensional *Arrays* are encoded as nested JSON arrays. The outer array is the first dimension and the innermost array is the last dimension. For example, the following matrix

0 2 3
1 3 4

is encoded in JSON as

```
[ [0, 2, 3], [1, 3, 4] ]
```

5.4.6 Structures

Structures shall be encoded as JSON objects.

If the value of a field is NULL it shall be omitted from the encoding.

For example, instances of the structures:

```
struct Type2
{
  Int32 A;
  Int32 B;
  Char* C;
};

struct Type1
{
  Int32 X;
  Int32 NoOfY;
  Type2* Y;
  Int32 Z;
};
```

are represented in JSON as:

```
{
  "X":1234,
  "Y":[ { "A":1, "B":2, "C":"Hello" }, { "A":3, "B":4 } ],
  "Z":5678
}
```

where "C" is omitted from the second Type2 instance because it has a NULL value.

5.4.7 Structures with optional fields

Structures with optional fields shall be encoded as JSON objects as shown in Table 35.

Table 35 – JSON Object Definition for a *Structure* with Optional Fields

Name	Description
EncodingMask	A bit mask indicating what fields are encoded in the structure (see 5.2.7) This mask is encoded as a JSON number. The bits are sequentially assigned to optional fields in the order that they are defined.
<FieldName>	The fields' structure is encoded according to the rules defined for their <i>DataType</i> .

For the non-reversible form, *Structures* with optional fields are encoded like *Structures*.

If a *Structure* with optional fields is subtyped, the subtypes extend the *EncodingMask* defined for the parent.

The following is an example of a structure with optional fields using C++ syntax:

```

struct TypeA
{
  Int32 X;
  Int32* O1;
  SByte Y;
  Int32* O2;
};

```

O1 and O2 are optional fields where a NULL indicates that the field is not present.

Assume that O1 is not specified and the value of O2 is 0.

The reversible encoding would be:

```

{ "EncodingMask": 2, "X": 1, "Y": 2 }

```

where decoders would assign the default value of 0 to O2 since the mask bit is set, even though the field was omitted (this is the behaviour defined for the *Int32 DataType*). Decoders would mark O1 as 'not specified'.

5.4.8 Unions

Unions shall be encoded as JSON objects as shown in Table 36 for the reversible encoding.

Table 36 – JSON Object Definition for a Union

Name	Description
SwitchField	The identifier for the field in the Union which is encoded as a JSON number.
Value	The value of the field encoded using the rules that apply to the data type.

For the non-reversible form, *Union* values are encoded using the rule for the current value.

For example, instances of the union:

```

struct Union1
{
  Byte Selector;

  {
    Int32 A;
    Double B;
    Char* C;
  }
  Value;
};

```

would be represented in reversible form as:

```

{ "SwitchField":2, "Value":3.1415 }

```

In non- reversible form, it is represented as:

```

3.1415

```

5.4.9 Messages

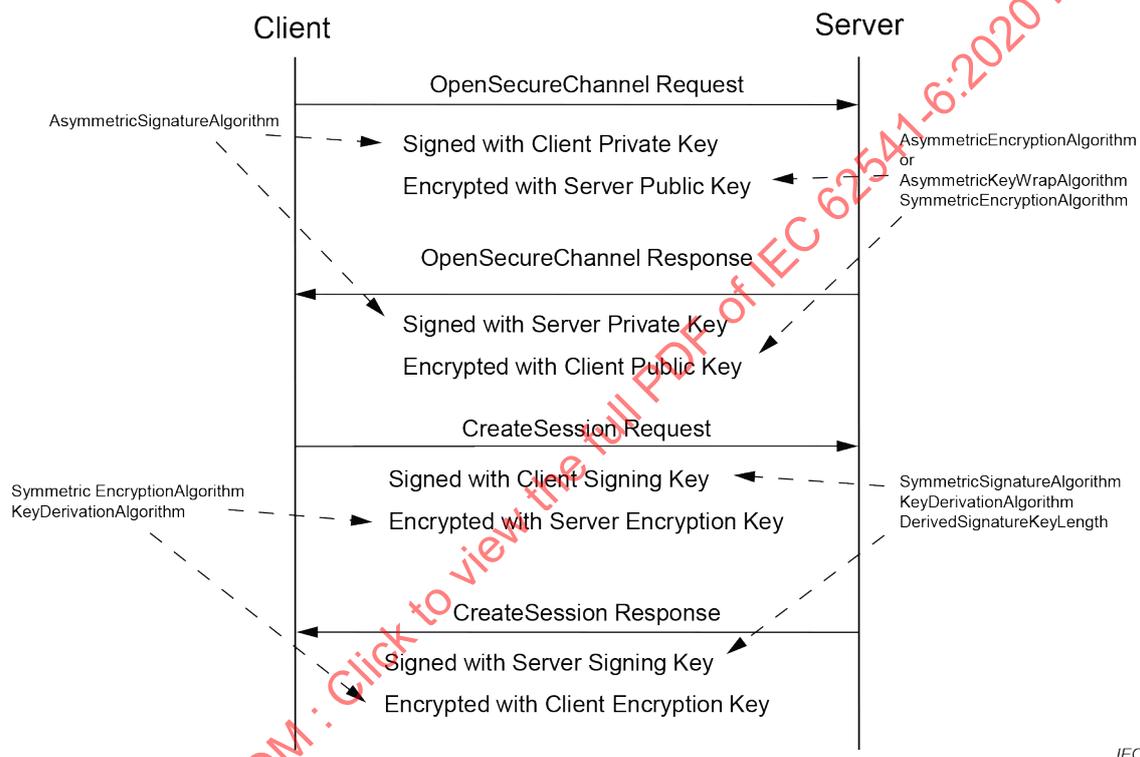
Messages are encoded *ExtensionObjects* (see 5.4.2.16).

6 Message SecurityProtocols

6.1 Security handshake

All *SecurityProtocols* shall implement the *OpenSecureChannel* and *CloseSecureChannel* services defined in IEC 62541-4. These *Services* specify how to establish a *SecureChannel* and how to apply security to *Messages* exchanged over that *SecureChannel*. The *Messages* exchanged and the security algorithms applied to them are shown in Figure 10.

SecurityProtocols shall support three *SecurityModes*: *None*, *Sign* and *SignAndEncrypt*. If the *SecurityMode* is *None* then no security is used and the security handshake shown in Figure 10 is not required. However, a *SecurityProtocol* implementation shall still maintain a logical channel and provide a unique identifier for the *SecureChannel*.



IEC

Figure 10 – Security handshake

Each *SecurityProtocol* mapping specifies exactly how to apply the security algorithms to the *Message*. A set of security algorithms that shall be used together during a security handshake is called a *SecurityPolicy*. IEC 62541-7 defines standard *SecurityPolicies* as parts of the standard *Profiles* which OPC UA applications are expected to support. IEC 62541-7 also defines a URI for each standard *SecurityPolicy*.

A *Stack* is expected to have built in knowledge of the *SecurityPolicies* that it supports. Applications specify the *SecurityPolicy* they wish to use by passing the URI to the *Stack*.

Table 37 defines the contents of a *SecurityPolicy*. Each *SecurityProtocol* mapping specifies how to use each of the parameters in the *SecurityPolicy*. A *SecurityProtocol* mapping may not make use of all of the parameters.

Table 37 – SecurityPolicy

Name	Description
PolicyUri	The URI assigned to the <i>SecurityPolicy</i> .
SymmetricSignatureAlgorithm	The URI of The symmetric signature algorithm to use.
SymmetricEncryptionAlgorithm	The URI of The symmetric key encryption algorithm to use.
AsymmetricSignatureAlgorithm	The URI of The asymmetric signature algorithm to use.
AsymmetricKeyWrapAlgorithm	The URI of the asymmetric key wrap algorithm to use.
AsymmetricEncryptionAlgorithm	The URI of The asymmetric key encryption algorithm to use.
MinAsymmetricKeyLength	The minimum length, in bits, for an asymmetric key.
MaxAsymmetricKeyLength	The maximum length, in bits, for an asymmetric key.
KeyDerivationAlgorithm	The key derivation algorithm to use.
DerivedSignatureKeyLength	The length in bits of the derived key used for <i>Message</i> authentication.
CertificateSignatureAlgorithm	The asymmetric signature algorithm used to sign certificates.
SecureChannelNonceLength	The length, in bytes, of the <i>Nonces</i> exchanged when creating a <i>SecureChannel</i> .

~~The *AsymmetricEncryptionAlgorithm* is used when encrypting the entire *Message* with an asymmetric key. Some *SecurityProtocols* do not encrypt the entire *Message* with an asymmetric key. Instead, they use the *AsymmetricKeyWrapAlgorithm* to encrypt a symmetric key and then use the *SymmetricEncryptionAlgorithm* to encrypt the *Message*.~~

~~The *AsymmetricSignatureAlgorithm* is used to sign a *Message* with an asymmetric key.~~

~~The *KeyDerivationAlgorithm* is used to create the keys used to secure *Messages* sent over the *SecureChannel*. The length of the keys used for encryption is implied by the *SymmetricEncryptionAlgorithm*. The length of the keys used for creating *Symmetric Signatures* depends on the *SymmetricSignatureAlgorithm* and may be different from the encryption key length.~~

The *KeyDerivationAlgorithm* is used to create the keys used to secure *Messages* sent over the *SecureChannel*. The length of the keys used for encryption is implied by the *SymmetricEncryptionAlgorithm*. The length of the keys used for creating *Signatures* is specified by the *DerivedSignatureKeyLength*.

The *MinAsymmetricKeyLength* and *MaxAsymmetricKeyLength* are constraints that apply to all *Certificates* (including *Issuers* in the chain). In addition, the key length of issued *Certificates* shall be less than or equal to the key length of the issuer *Certificate*. See 6.2.3 for information on *Certificate* chains.

The *CertificateKeyAlgorithm* and *EphemeralKeyAlgorithm* are used to generate new asymmetric key pairs used with *Certificates* and during the *SecureChannel* handshake. IEC 62541-7 specifies the bit lengths that need to be supported for each *SecurityPolicy*.

The *CertificateSignatureAlgorithm* applies to the *Certificate* and all *Issuer Certificates*. If a *CertificateSignatureAlgorithm* allows for more than one algorithm then the algorithms are listed in order of increasing priority. Each *Issuer* in a chain shall have an algorithm that is the same or higher priority than any *Certificate* it issues.

The *SecureChannelNonceLength* specifies the length of the *Nonces* exchanged when establishing a *SecureChannel* (see 6.7.4).

6.2 Certificates

6.2.1 General

OPC UA applications use *Certificates* to store the *Public Keys* needed for *Asymmetric Cryptography* operations. All *SecurityProtocols* use X.509 v3 *Certificates* (see X.509 v3) encoded using the DER format (see X690). *Certificates* used by OPC UA applications shall also conform to RFC 3280 which defines a profile for X.509 v3 *Certificates* when they are used as part of an Internet based application.

The *ServerCertificate* and *ClientCertificate* parameters used in the abstract *OpenSecureChannel* service are instances of the *Application Instance Certificate Data Type*. 6.2.2 describes how to create an X.509 v3 *Certificate* that can be used as an *Application Instance Certificate*.

~~The *ServerSoftwareCertificates* and *ClientSoftwareCertificates* parameters in the abstract *CreateSession* and *ActivateSession* Services are instances of the *SignedSoftwareCertificate Data Type*. Subclause 6.2.3 describes how to create an X.509 *Certificate* that can be used as a *SignedSoftwareCertificate*.~~

6.2.2 Application Instance Certificate

An *Application Instance Certificate* is a *ByteString* containing the DER encoded form (see X690) of an X.509 v3 *Certificate*. This *Certificate* is issued by certifying authority and identifies an instance of an application running on a single host. The X.509 v3 fields contained in an *Application Instance Certificate* are described in Table 38. The fields are defined completely in RFC 3280.

Table 38 also provides a mapping from the RFC 3280 terms to the terms used in the abstract definition of an *Application Instance Certificate* defined in IEC 62541-4.

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Table 38 – Application Instance Certificate

Name	IEC 62541-4 Parameter Name	Description
Application Instance Certificate		An X.509 v3 <i>Certificate</i> .
version	version	shall be "V3"
serialNumber	serialNumber	The serial number assigned by the issuer.
signatureAlgorithm	signatureAlgorithm	The algorithm used to sign the <i>Certificate</i> .
signature	signature	The signature created by the Issuer.
issuer	issuer	The distinguished name of the <i>Certificate</i> used to create the signature. The <i>issuer</i> field is completely described in RFC 3280.
validity	validTo, validFrom	When the <i>Certificate</i> becomes valid and when it expires.
subject	subject	The distinguished name of the application <i>Instance</i> . The Common Name attribute shall be specified and should be the <i>productName</i> or a suitable equivalent. The Organization Name attribute shall be the name of the Organization that executes the application instance. This organization is usually not the vendor of the application. Other attributes may be specified. The <i>subject</i> field is completely described in RFC 3280.
subjectAltName	applicationUri, hostnames	The alternate names for the application <i>Instance</i> . Shall include a uniformResourceIdentifier which is equal to the <i>applicationUri</i> . The URI shall be a valid URL (see RFC 1738) or a valid URN (see RFC 2141). Servers shall specify a partial or a fully qualified <i>dnsName</i> or a static <i>IPAddress</i> which identifies the machine where the application <i>Instance</i> runs. Additional <i>dnsNames</i> may be specified if the machine has multiple names. The <i>IPAddress</i> should not be specified if the Server has <i>dnsName</i>. The <i>subjectAltName</i> field is completely described in RFC 3280.
publicKey	publicKey	The public key associated with the <i>Certificate</i> .
keyUsage	keyUsage	Specifies how the <i>Certificate</i> key may be used. Shall include digitalSignature, nonRepudiation, keyEncipherment and dataEncipherment. Other key uses are allowed.
extendedKeyUsage	keyUsage	Specifies additional key uses for the <i>Certificate</i> . Shall specify 'serverAuth and/or clientAuth'. Other key uses are allowed.
authorityKeyIdentifier	(no mapping)	Provides more information about the key used to sign the <i>Certificate</i> . It shall be specified for <i>Certificates</i> signed by a CA. It should be specified for self-signed <i>Certificates</i> .

~~6.2.3 Signed Software Certificate~~

~~A *SignedSoftwareCertificate* is a *ByteString* containing the DER encoded form of an X509v3 *Certificate*. This *Certificate* is issued by a certifying authority and contains an X509v3 extension with the *SoftwareCertificate* which specifies the claims verified by the certifying authority. The X509v3 fields contained in a *SignedSoftwareCertificate* are described in Table 24. The fields are defined completely in RFC 3280.~~

Table 24 – SignedSoftwareCertificate

Name		Description
SignedSoftwareCertificate		An X509v3 <i>Certificate</i> .
—version	—version	Shall be “V3”
—serialNumber	—serialNumber	The serial number assigned by the issuer.
—signatureAlgorithm	—signatureAlgorithm	The algorithm used to sign the <i>Certificate</i> .
—signature	—signature	The signature created by the Issuer.
—issuer	—issuer	The distinguished name of the <i>Certificate</i> used to create the signature. The <i>issuer</i> field is completely described in RFC 3280.
—validity	—validTo, validFrom	When the <i>Certificate</i> becomes valid and when it expires.
—subject	—subject	The distinguished name of the product. The Common Name attribute shall be the same as the <i>productName</i> in the <i>SoftwareCertificate</i> and the Organization Name attribute shall be the <i>vendorName</i> in the <i>SoftwareCertificate</i> . Other attributes may be specified. The <i>subject</i> field is completely described in RFC 3280.
—subjectAltName	—productUri	The alternate names for the product. It shall include a ‘uniformResourceIdentifier’ which is equal to the <i>productUri</i> specified in the <i>SoftwareCertificate</i> . The <i>subjectAltName</i> field is completely described in RFC 3280.
—publicKey	—publicKey	The public key associated with the <i>Certificate</i> .
—keyUsage	—keyUsage	Specifies how the <i>Certificate</i> key may be used. shall be ‘digitalSignature’ and ‘nonRepudiation’ Other key uses are not allowed.
—extendedKeyUsage	—keyUsage	Specifies additional key uses for the <i>Certificate</i> . May specify ‘codeSigning’. Other key usages are not allowed.
—softwareCertificate	—softwareCertificate	The XML encoded form of the <i>SoftwareCertificate</i> stored as UTF8 text. Subclause 5.3.4 describes how to encode a <i>SoftwareCertificate in XML</i> . The ASN.1 Object Identifier (OID) for this extension is: 1.2.840.113556.1.8000.2264.1.6.1

6.2.3 Certificate Chains

Any X.509 v3 *Certificate* may be signed by CA which means that validating the signature requires access to the X.509 v3 *Certificate* belonging to the signing CA. Whenever an application validates a signature it shall recursively build a chain of *Certificates* by finding the issuer *Certificate*, validating the *Certificate* and then repeating the process for the issuer *Certificate*. The chain ends with a self-signed *Certificate*.

The number of CAs used in a system should be small so it is common to install the necessary CAs on each machine with an OPC UA application. However, applications have the option of including a partial or complete chain whenever they pass a *Certificate* to a peer during the *SecureChannel* negotiation and during the *CreateSession/ActivateSession* handshake. All OPC UA applications shall accept partial or complete chains in any field that contains a DER encoded *Certificate*.

Chains are stored in a *ByteString* by simply appending the DER encoded form of the *Certificates*. The first *Certificate* shall be the end *Certificate* followed by its issuer. If the root CA is sent as part of the chain, the last *Certificate* is appended to the *ByteString*.

Chains are parsed by extracting the length of each *Certificate* from the DER encoding. For *Certificates* with lengths less than 65 535 bytes, it is an MSB encoded UInt16 starting at the third byte.

6.3 Time synchronization

All *SecurityProtocols* require that system clocks on communicating machines be reasonably synchronized in order to check the expiry times for *Certificates* or *Messages*. The amount of clock skew that can be tolerated depends on the system security requirements and

applications shall allow administrators to configure the acceptable clock skew when verifying times. A suitable default value is 5 minutes.

The Network Time Protocol (NTP) provides a standard way to synchronize a machine clock with a time server on the network. Systems running on a machine with a full featured operating system like Windows or Linux will already support NTP or an equivalent. Devices running embedded operating systems should support NTP.

If a device operating system cannot practically support NTP then an OPC UA application can use the *Timestamps* in the *ResponseHeader* (see IEC 62541-4) to synchronize its clock. In this scenario, the OPC UA application will ~~have~~ need to know the URL for a *Discovery Server* on a machine known to have the correct time. The OPC UA application or a separate background utility would call the *FindServers Service* and set its clock to the time specified in the *ResponseHeader*. This process will need to be repeated periodically because clocks can drift over time.

6.4 UTC and International Atomic Time (TAI)

All times in OPC UA are in UTC, however, UTC can include discontinuities due to leap seconds or repeating seconds added to deal with variations in the earth's orbit and rotation. *Servers* that have access to source for International Atomic Time (TAI) may choose to use this instead of UTC. That said, *Clients* ~~must~~ always need to be prepared to deal with discontinuities due to the UTC or simply because the system clock is adjusted on the *Server* machine.

6.5 Issued User Identity Tokens

6.5.1 Kerberos

Kerberos *UserIdentityTokens* can be passed to the *Server* using the *IssuedIdentityToken*. The body of the token is an XML element that contains the WS-Security token as defined in the Kerberos Token Profile (Kerberos) specification.

Servers that support Kerberos authentication shall provide a *UserTokenPolicy* which specifies what version of the Kerberos Token Profile is being used, the Kerberos Realm and the Kerberos Principal Name for the *Server*. The Realm and Principal name are combined together with a simple syntax and placed in the *issuerEndpointUri* as shown in Table 39.

Table 39 – Kerberos UserTokenPolicy

Name	Description
tokenType	ISSUEDTOKEN_3
issuedTokenType	http://docs.oasis-open.org/wss/oasis-wss-kerberos-token-profile-1.1
issuerEndpointUri	A string with the form \\<realm>\<server principal name> where <realm> is the Kerberos realm name (e.g. Windows Domain); <server principal name> is the Kerberos principal name for the OPC UA Server.

The interface between the *Client* and *Server* applications and the Kerberos Authentication Service is application specific. The realm is the DomainName when using a Windows Domain controller as the Kerberos provider.

6.5.2 JSON Web Token (JWT)

JSON Web Token (JWT) *UserIdentityTokens* can be passed to the *Server* using the *IssuedIdentityToken*. The body of the token is a string that contains the JWT as defined in RFC 7159.

Servers that support JWT authentication shall provide a *UserTokenPolicy* which specifies the *Authorization Service* which provides the token and the parameters needed to access that service. The parameters are specified by a JSON object specified as the *issuerEndpointUrl*. The contents of this JSON object are described in Table 41. The general *UserTokenPolicy* settings for JWT are defined in Table 40.

Table 40 – JWT UserTokenPolicy

Name	Description
tokenType	ISSUEDTOKEN_3
issuedTokenType	http://opcfoundation.org/UA/UserToken#JWT
issuerEndpointUrl	For JWTs this is a JSON object with fields defined in Table 41.

Table 41 – JWT IssuerEndpointUrl Definition

Name	Type	Description
IssuerEndpointUrl	JSON object	Specifies the parameters for a JWT <i>UserIdentityToken</i> .
ua:resourceId	String	The URI identifying the <i>Server</i> to the <i>Authorization Service</i> . If not specified, the <i>Server's ApplicationUri</i> is used.
ua:authorityUrl	String	The base URL for the <i>Authorization Service</i> . This URL may be used to discover additional information about the authority. This field is equivalent to the "issuer" defined in OpenID-Discovery.
ua:authorityProfileUri	String	The profile that defines the interactions with the authority. If not specified, the URI is "http://opcfoundation.org/UA/Authorization#OAuth2".
ua:tokenEndpoint	String	A path relative to the base URL used to request <i>Access Tokens</i> . This field is equivalent to the "token_endpoint" defined in OpenID-Discovery.
ua:authorizationEndpoint	String	A path relative to the base URL used to validate user credentials. This field is equivalent to the "authorization_endpoint" defined in OpenID-Discovery.
ua:requestTypes	JSON array String	The list of request types supported by the authority. The possible values depend on the authorityProfileUri. IEC 62541-7 specifies the default for each authority profile defined.
ua:scopes	JSON array String	A list of Scopes that are understood by the <i>Server</i> . If not specified, the <i>Client</i> may be able to access any <i>Scope</i> supported by the <i>Authorization Service</i> . This field is equivalent to the "scopes_supported" defined in OpenID-Discovery.

6.5.3 OAuth2

6.5.3.1 General

The OAuth2 Authorization Framework (see RFC 6749) provides a web based mechanism to request claims based *Access Tokens* from an *Authorization Service* (AS) that is supported by

many major companies providing cloud infrastructure. These *Access Tokens* are passed to a *Server* by a *Client* in a *UserIdentityToken* as described in IEC 62541-4.

The OpenID Connect specification (see OpenID) builds on the OAuth2 specification by defining the contents of the *Access Tokens* more strictly.

The OAuth2 specification supports a number of use cases (called 'flows') to handle different application requirements. The use cases that are relevant to OPC UA are discussed below.

6.5.3.2 Access Tokens

The JSON Web Token is the *Access Token* format which this document requires when using OAuth2. The JWT supports signatures using asymmetric cryptography which implies that *Servers* which accept the *Access Token* need to have access to the *Certificate* used by the *Authorization Service* (AS). The OpenID Connect Discovery specification is implemented by many AS products and provides a mechanism to fetch the AS *Certificate* via an HTTP request. If the AS does not support the discovery specification, then the signing *Certificate* will need to be provided to the *Server* when the location of the AS is added to the *Server* configuration.

Access Tokens expire and all *Servers* should revoke any privileges granted to the *Session* when the *Access Token* expires. If the *Server* allows for anonymous users, the *Server* should allow the *Session* to stay open but treat it as an anonymous user. If the *Server* does not allow anonymous users, it should close the *Session* immediately.

Clients know when the *Access Token* will expire and should request a new *Access Token* and call *ActivateSession* before the old *Access Token* expires.

The JWT format allows the *Authorization Service* to insert any number of fields. The mandatory fields are defined in RFC 7159. Some additional fields are defined in Table 42 (see RFC 7523).

Table 42 – Access Token Claims

Field	Description
sub	The subject for the token. Usually the <i>client_id</i> which identifies the <i>Client</i> . If returned from an <i>Identity Provider</i> it may be a unique identifier for the user.
aud	The audience for the token. Usually the <i>resource_id</i> which identifies the <i>Server</i> or the <i>Server ApplicationUri</i> .
name	A human readable name for the <i>Client</i> application or user.
scp	A list of <i>Scopes</i> granted to the subject. Scopes apply to the <i>Access Token</i> and restrict how it may be used. Usually permissions or other restriction which limit access rights.
nonce	A nonce used to mitigate replay attacks. Shall be the value provided by the <i>Client</i> in the request.
groups	A list of groups which are assigned to the subject. Usually a list of unique identifiers for platform-specific security groups. For example, Azure AD user account groups may be returned in this claim.
roles	A list of roles which are assigned to the subject. Roles apply to the requestor and describe what the requestor can do with the resource. Usually a list of unique identifiers for roles known to the <i>Authorization Service</i> . These values are typically mapped to the <i>Roles</i> defined in IEC 62541-3.

6.5.3.3 Authorization Code

The authorization code flow is available to *Clients* which allow interaction with a human user. The *Client* application displays a window with a web browser which sends an HTTP GET to the *Identity Provider*. When the human user enters credentials that the *Identity Provider* validates, the *Identity Provider* returns an authorization code which is passed to the *Authorization Service*. The *Authorization Service* validates the code and returns an *Access Token* to the *Client*.

The complete flow is described in RFC 6749, 4.1.

A *requestType* of "authorization_code" in the *UserTokenPolicy* (see 6.5.2) means the *Authorization Service* supports the authorization code flow.

6.5.3.4 Refresh Token

The refresh token flow applies when a *Client* application has access to a refresh token returned in a previous response to an authorization code request. The refresh token allows applications to skip the step that requires human interaction with the *Identity Provider*. This flow is initiated when the *Client* sends the refresh token to *Authorization Service* which validates it and returns an *Access Token*. A *Client* that saves the refresh token for later use shall use encryption or other means to ensure the refresh token cannot be accessed by unauthorized parties.

The complete flow is described in RFC 6749, Clause 6.

A *requestType* is not defined since support for refresh token is determined by checking the response to an authorization code request.

6.5.3.5 Client Credentials

The client credentials flow applies when a *Client* application cannot prompt a human user for input. This flow requires a secret known to the *Authorization Service* which the *Client* application can protect. This flow is initiated when the *Client* sends the *client_secret* to *Authorization Service* which validates it and returns an *Access Token*.

The complete flow is described in RFC 6749, 4.4.

A *requestType* of "client_credentials" in the *UserTokenPolicy* (see 6.5.2) means the *Authorization Service* supports the client credentials flow.

6.6 WS Secure Conversation

6.6.1 Overview

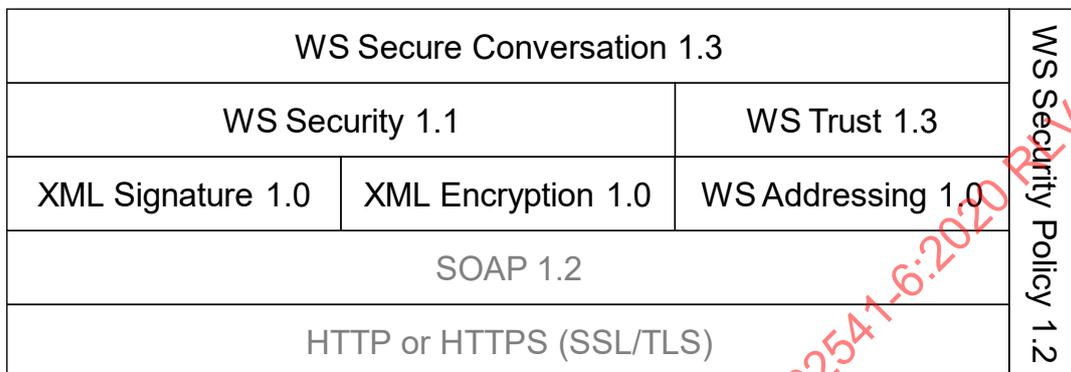
~~Any Message sent via SOAP may be secured with the WS Secure Conversation. This protocol specifies a way to negotiate shared secrets via WS Trust and then use these secrets to secure Messages exchanged with the mechanisms defined in WS Security.~~

~~The mechanisms for actually signing XML elements are described in the XML Signature specification. The mechanisms for encrypting XML elements are described in the XML Encryption specification.~~

~~WS Security Policy defines standard algorithm suites which can be used to secure SOAP Messages. These algorithm suites map directly onto the SecurityPolicies that are defined in IEC 62541-7. WS-I Basic Security Profile 1.1 defines best practices when using WS Security which will help ensure interoperability. All OPC UA implementations shall conform to this specification.~~

~~The *Timestamp* header defined by WS Security is used to prevent replay attacks and shall be present and signed in all Messages exchanged.~~

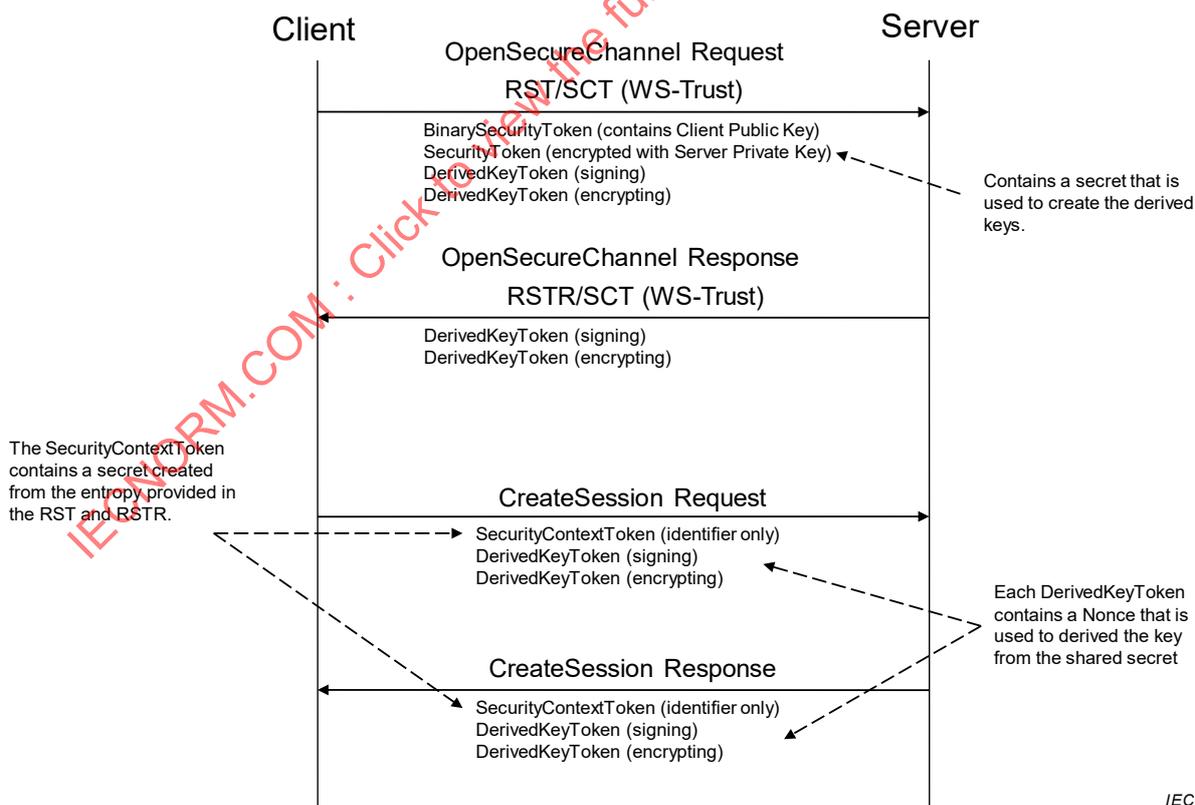
~~Figure 11 illustrates the relationship between the different WS-* specifications that are used by this mapping. The versions of the WS-* specifications shown in the diagram were the most current versions at the time of publication. IEC 62541-7 may define Profiles that require support for future versions of these specifications.~~



IEC

Figure 11 – Relevant XML Web Services specifications

~~Figure 12 illustrates how these WS-* specifications are used in the security handshake.~~



IEC

Figure 12 – The WS Secure Conversation handshake

~~The RST (Request Security Token) and RSTR (Request Security Token Response) Messages are defined by WS Trust. WS Secure Conversation defines new actions for these Messages~~

that tell the *Server* that the *Client* wants to create a SCT (Security Context Token). The SCT contains the shared keys that the *Applications* use to secure *Messages* sent over the *SecureChannel*.

Individual *Messages* are secured with keys derived from the SCT using the mechanism defined in WS Secure Conversation. The subclauses below specify the structure of the individual *Messages* and illustrate which features from the WS-* specifications are required to implement the OPC UA security handshake.

6.6.2 Notation

SOAP *Messages* use XML elements defined in a number of different specifications. This document uses the prefixes in Table 26 to identify the specification that defines an XML element.

Table 26 – WS-* Namespace prefixes

Prefix	Specification
wsu	WS-Security Utilities
wsse	WS-Security Extensions
wst	WS-Trust
wsc	WS-Secure Conversation
wsa	WS-Addressing
xenc	XML Encryption

6.6.3 Request Security Token (RST/SCT)

The Request Security Token *Messages* implements the abstract *OpenSecureChannel* request *Message* defined in IEC 62541-4. The syntax of this *Message* is defined by WS-Trust. The structure of the *Message* is described in detail in WS-Secure Conversation.

This *Message* shall have the following tokens:

- a) A wsse:BinarySecurityToken containing the *Client's Public Key*. The *Public Key* is sent in a DER encoded X509v3 *Certificate*.
- b) An encrypted wsse:SecurityToken containing *ClientNonce* used to derive keys. This *SecurityToken* shall be encrypted with the *AsymmetricKeyWrapAlgorithm* and the *Public Key* associated with the *Server's Application Instance Certificate*.
- c) A wsc:DerivedKeyToken which is used to sign the body, the WS-Addressing headers and the wsu:Timestamp header using the *SymmetricSignatureAlgorithm*. The signature element shall then be signed using the *AsymmetricSignatureAlgorithm* with the *Client's Private Key*. The wsc:DerivedKeyToken shall also specify a *Nonce*.
- d) A wsc:DerivedKeyToken which is used to encrypt the body of the *Message* using the *SymmetricEncryptionAlgorithm*.

This *Message* shall have the wsa:Action, wsa:MessageId, wsa:ReplyTo and wsa:To headers defined by WS-Addressing. The *Message* shall also have a wsu:Timestamp header defined by WS-Security. These headers shall also be signed with the derived key used to sign the *Message* body.

The signature shall be calculated before applying encryption and the signature shall be encrypted.

The mapping between the *OpenSecureChannel* request parameters and the elements of the RST/SCT *Message* are shown in Table 27.

Table 27 – RST/SCT Mapping to an OpenSecureChannel Request

OpenSecureChannel Parameter	RST/SCT Element	Description
clientCertificate	wsse:BinarySecurityToken	Passed in the SOAP header.
requestType	wst:RequestType	Shall be "http://schemas.xmlsoap.org/ws/2005/02/trust/Issue" when creating a new SCT. Shall be "http://schemas.xmlsoap.org/ws/2005/02/trust/Renew" when renewing a SCT.
secureChannelId	wsse:SecurityTokenReference	Passed in the SOAP header when renewing an SCT.
securityMode securityPolicyUri	wst:SignatureAlgorithm wst:EncryptionAlgorithm wst:KeySize	These elements describe the <i>SecurityPolicy</i> requested by the <i>Client</i> . These elements shall match the <i>SecurityPolicy</i> used by the <i>Endpoint</i> that the <i>Client</i> wishes to connect to. These elements are optional.
clientNonce	wst:Entropy	This contains the <i>Nonce</i> specified by the <i>Client</i> . The <i>Nonce</i> is specified with the wst:BinarySecret element.
requestedLifetime	wst:Lifetime	The requested lifetime for the SCT. This element is optional.

6.6.4 Request Security Token Response (RSTR/SCT)

The ~~Request Security Token Response Message~~ implements the abstract *OpenSecureChannel* response *Message* defined in IEC 62541-4. The syntax of this *Message* is defined by WS Trust. The use of the *Message* is described in detail in WS Secure Conversation. This *Message* not signed or encrypted with the asymmetric algorithms as described in IEC 62541-4. The symmetric algorithms and a key provided in the request *Message* are used instead.

This *Message* shall have the following tokens:

- a) A wsc:DerivedKeyToken which is used to sign the body, the WS Addressing headers and the wsu:Timestamp header using the *SymmetricSignatureAlgorithm*. This key is derived from the encrypted *SecurityToken* specified in the RST/SCT *Message*. The wsc:DerivedKeyToken shall also specify a *Nonce*.
- b) A wsc:DerivedKeyToken which is used to encrypt the body of the *Message* using the *SymmetricEncryptionAlgorithm*. This key is derived from the encrypted *SecurityToken* specified in the RST/SCT *Message*. The wsc:DerivedKeyToken shall also specify a *Nonce*.

This *Message* shall have the wsa:Action and wsa:RelatesTo headers defined by WS Addressing. The *Message* shall also have a wsu:Timestamp header defined by WS Security. These headers shall also be signed with the derived key used to sign the *Message* body.

The signature shall be calculated before applying encryption and the signature shall be encrypted.

The mapping between the *OpenSecureChannel* response parameters and the elements of the RSTR/SCT *Message* are shown Table 28.

Table 28 – RSTR/SCT Mapping to an OpenSecureChannel Response

OpenSecureChannel Parameter	RSTR/SCT Element	Description
---	wst:RequestedProofToken	This contains a wst:ComputedKey element which specifies the algorithm used to compute the shared secret key from the Nonces provided by the Client and the Server.
---	wst:TokenType	Specifies the type of SecurityToken issued.
securityToken	wst:RequestedSecurityToken	Specifies the new SCT (Security Context Token) or renewed SCT.
---channelId	wsc:Identifier	An absolute URI which identifies the SCT.
---tokenId	wsc:Instance	An identifier for a set of keys issued for a context. It shall be unique within the context.
---createdAt	wsu:Created	This is optional element in the wsc:SecurityContextToken returned in the header.
revisedLifetime	wst:Lifetime	The revised lifetime for the SCT.
serverNonce	wst:Entropy	This contains the Nonce specified by the Server. The Nonce is specified with the wst:BinarySecret element. The xenc:EncryptedData element is not used in OPC-UA because the Message body shall be encrypted.

The lifetime specifies the UTC expiration time for the security context token. The Client shall renew the SCT before that time by sending the RSTR/SCT Message again. The exact behaviour is described in IEC 62541-4.

6.6.5 Using the SCT

Once the Client receives the RSTR/SCT Message it can use the SCT to secure all other Messages.

An identifier for the SCT used shall be passed as an wsc:SecurityContextToken in each request Message. The response Message shall reference the SecurityContextToken used in the request.

If encryption is used it shall be applied before the signature is calculated.

Any Message secured with the SecurityContextToken shall have the following additional tokens:

- A wsc:DerivedKeyToken which is used to sign the body, the WS Addressing headers and the wsu:Timestamp header using the SymmetricSignatureAlgorithm. This key is derived from the SecurityContextToken. The wsc:DerivedKeyToken shall also specify a Nonce.
- A wsc:DerivedKeyToken which is used to encrypt the body of the Message using the SymmetricEncryptionAlgorithm. This key is derived from the SecurityContextToken. The wsc:DerivedKeyToken shall also specify a Nonce.

This Message shall have the wsa:Action and wsa:RelatesTo headers defined by WS Addressing. The Message shall also have a wsu:Timestamp header defined by WS Security.

6.6.6 Cancelling Security contexts

The Cancel Message defined by WS Trust implements the abstract CloseSecureChannel request Message defined in IEC 62541-4.

This Message shall be secured with the SCT.

NOTE Deprecated in Version 1.03 because WS-SecureConversation has not been widely adopted by industry.

6.7 OPC UA Secure Conversation

6.7.1 Overview

OPC UA Secure Conversation (UASC) ~~is a binary version of WS-Secure Conversation. It allows secure communication over transports that do not use SOAP or XML~~ using binary encoded *Messages*.

UASC is designed to operate with different *TransportProtocols* that may have limited buffer sizes. For this reason, OPC UA Secure Conversation will break OPC UA *Messages* into several pieces (called '*MessageChunks*') that are smaller than the buffer size allowed by the *TransportProtocol*. UASC requires a *TransportProtocol* buffer size that is at least ~~8196~~ 8 192 bytes.

All security is applied to individual *MessageChunks* and not the entire OPC UA *Message*. A *Stack* that implements UASC is responsible for verifying the security on each *MessageChunk* received and reconstructing the original OPC UA *Message*.

All *MessageChunks* will have a 4-byte sequence assigned to them. These sequence numbers are used to detect and prevent replay attacks.

UASC requires a *TransportProtocol* that will preserve the order of *MessageChunks*, however, a UASC implementation does not necessarily process the *Messages* in the order that they were received.

6.7.2 MessageChunk structure

6.7.2.1 Overview

Figure 11 shows the structure of a *MessageChunk* and how security is applied to the *Message*.

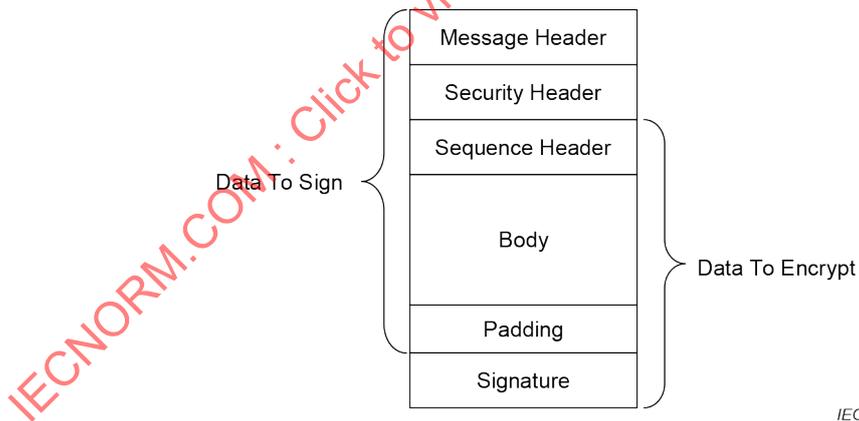


Figure 11 – OPC UA Secure Conversation MessageChunk

6.7.2.2 Message Header

Every *MessageChunk* has a *Message* header with the fields defined in Table 43.

Table 43 – OPC UA Secure Conversation Message header

Name	Data Type	Description
MessageType	Byte [3]	A three byte ASCII code that identifies the <i>Message</i> type. The following values are defined at this time: MSG A <i>Message</i> secured with the keys associated with a channel. OPN OpenSecureChannel <i>Message</i> . CLO CloseSecureChannel <i>Message</i> .
IsFinal	Byte	A one byte ASCII code that indicates whether the <i>MessageChunk</i> is the final chunk in a <i>Message</i> . The following values are defined at this time: C An intermediate chunk. F The final chunk. A The final chunk (used when an error occurred and the <i>Message</i> is aborted). This field is only meaningful for MessageType of 'MSG' This field is always 'F' for other MessageTypes.
MessageSize	UInt32	The length of the <i>MessageChunk</i> , in bytes. This value includes size of the <i>Message</i> header. The length of the <i>MessageChunk</i> , in bytes. The length starts from the beginning of the MessageType field.
SecureChannelId	UInt32	A unique identifier for the <i>SecureChannel</i> assigned by the <i>Server</i> . If a <i>Server</i> receives a SecureChannelId which it does not recognize it shall return an appropriate transport layer error. When a <i>Server</i> starts the first <i>SecureChannelId</i> used should be a value that is likely to be unique after each restart. This ensures that a <i>Server</i> restart does not cause previously connected <i>Clients</i> to accidentally 'reuse' <i>SecureChannels</i> that did not belong to them.

6.7.2.3 Security Header

The *Message* header is followed by a security header which specifies what cryptography operations have been applied to the *Message*. There are two versions of the security header which depend on the type of security applied to the *Message*. The security header used for asymmetric algorithms is defined in Table 44. Asymmetric algorithms are used to secure the *OpenSecureChannel Messages*. PKCS #1 defines a set of asymmetric algorithms that may be used by UASC implementations. The *AsymmetricKeyWrapAlgorithm* element of the *SecurityPolicy* structure defined in Table 37 is not used by UASC implementations.

Table 44 – Asymmetric algorithm Security header

Name	Data Type	Description
SecurityPolicyUriLength	Int32	The length of the <i>SecurityPolicyUri</i> in bytes. This value shall not exceed 255 bytes. If a URI is not specified this value may be 0 or -1. Other negative values are invalid.
SecurityPolicyUri	Byte [*]	The URI of the <i>Security Policy</i> used to secure the <i>Message</i> . This field is encoded as a UTF-8 string without a null terminator.
SenderCertificateLength	Int32	The length of the <i>SenderCertificate</i> in bytes. This value shall not exceed <i>MaxCertificateSize</i> <i>MaxSenderCertificateSize</i> bytes. If a certificate is not specified this value may be 0 or -1. Other negative values are invalid.
SenderCertificate	Byte [*]	The X.509 v3 <i>Certificate</i> assigned to the sending application <i>Instance</i> . This is a DER encoded blob. The structure of an X.509 v3 <i>Certificate</i> is defined in X.509 v3. The DER format for a <i>Certificate</i> is defined in X690 This indicates what <i>Private Key</i> was used to sign the <i>MessageChunk</i> . The <i>Stack</i> shall close the channel and report an error to the application if the <i>SenderCertificate</i> is too large for the buffer size supported by the transport layer. This field shall be null if the <i>Message</i> is not signed. If the <i>Certificate</i> is signed by a CA, the DER encoded CA <i>Certificate</i> may be appended after the <i>Certificate</i> in the byte array. If the CA <i>Certificate</i> is also signed by another CA this process is repeated until the entire <i>Certificate</i> chain is in the buffer or if <i>MaxSenderCertificateSize</i> limit is reached (the process stops after the last whole <i>Certificate</i> that can be added without exceeding the <i>MaxSenderCertificateSize</i> limit). Receivers can extract the <i>Certificates</i> from the byte array by using the <i>Certificate</i> size contained in DER header (see X.509 v3). Receivers that do not handle <i>Certificate</i> chains shall ignore the extra bytes.
ReceiverCertificateThumbprintLength	Int32	The length of the <i>ReceiverCertificateThumbprint</i> in bytes. If encrypted the length of this field is <i>always</i> 20 bytes. If not encrypted the value may be 0 or -1. Other negative values are invalid.
ReceiverCertificateThumbprint	Byte [*]	The thumbprint of the X.509 v3 <i>Certificate</i> assigned to the receiving application <i>Instance</i> . The thumbprint is the <i>SHA1-digest</i> <i>CertificateDigest</i> of the DER encoded form of the <i>Certificate</i> . This indicates what public key was used to encrypt the <i>MessageChunk</i> . This field shall be null if the <i>Message</i> is not encrypted.

The receiver shall close the communication channel if any of the fields in the security header have invalid lengths.

The *SenderCertificate*, including any chains, shall be small enough to fit into a single *MessageChunk* and leave room for at least one byte of body information. The maximum size for the *SenderCertificate* can be calculated with this formula:

```

MaxSenderCertificateSize =
  MessageChunkSize -
  12 - // Header size
  4 - // SecurityPolicyUriLength
  SecurityPolicyUri - // UTF-8 encoded string
  4 - // SenderCertificateLength
  4 - // ReceiverCertificateThumbprintLength
  20 - // ReceiverCertificateThumbprint
  8 - // SequenceHeader size
  1 - // Minimum body size
  1 - // PaddingSize if present
  Padding - // Padding if present
  ExtraPadding - // ExtraPadding if present
  AsymmetricSignatureSize // If present

```

The *MessageChunkSize* depends on the transport protocol but shall be at least ~~8196~~ 8 192 bytes. The *AsymmetricSignatureSize* depends on the number of bits in the public key for the *SenderCertificate*. The *Int32FieldLength* is the length of an encoded Int32 value and it is always 4 bytes.

The security header used for symmetric algorithms defined in Table 45. Symmetric algorithms are used to secure all *Messages* other than the *OpenSecureChannel Messages*. FIPS 197 define symmetric encryption algorithms that UASC implementations may use. FIPS 180-2-4 and HMAC define some symmetric signature algorithms.

Table 45 – Symmetric algorithm Security header

Name	Data Type	Description
TokenId	UInt32	A unique identifier for the <i>SecureChannel SecurityToken</i> used to secure the <i>Message</i> . This identifier is returned by the <i>Server</i> in an <i>OpenSecureChannel</i> response <i>Message</i> . If a <i>Server</i> receives a <i>TokenId</i> which it does not recognize it shall return an appropriate transport layer error.

6.7.2.4 Sequence header

The security header is always followed by the sequence header which is defined in Table 46. The sequence header ensures that the first encrypted block of every *Message* sent over a channel will start with different data.

Table 46 – Sequence header

Name	Data Type	Description
SequenceNumber	UInt32	A monotonically increasing sequence number assigned by the sender to each <i>MessageChunk</i> sent over the <i>SecureChannel</i> .
RequestId	UInt32	An identifier assigned by the <i>Client</i> to OPC UA request <i>Message</i> . All <i>MessageChunks</i> for the request and the associated response use the same identifier.

A *SequenceNumbers* may not be reused for any *TokenId*. The *SecurityToken* lifetime should be short enough to ensure that this never happens; however, if it does the receiver ~~should~~ shall treat it as a transport error and force a reconnect.

~~The *SequenceNumber* shall also monotonically increase for all *Messages* and shall not wrap around until it is greater than 4 294 966 271 (UInt32.MaxValue – 1 024). The first number after the wrap around shall be less than 1 024. Note that this requirement means that *SequenceNumbers* do not reset when a new *TokenId* is issued. The *SequenceNumber* shall be incremented by exactly one for each *MessageChunk* sent unless the communication channel was interrupted and re-established. Gaps are permitted between the *SequenceNumber* for the last *MessageChunk* received before the interruption and the *SequenceNumber* for first *MessageChunk* received after communication was reestablished.~~

~~Note that the first MessageChunk after a network interruption is always an OpenSecureChannel request or response.~~

The *SequenceNumber* shall start at 1 023 and monotonically increase for all *Messages* and shall wrap around when it is equal to 4 294 966 271 (UInt32.MaxValue - 1 024). The first number after the wrap around shall be less than 1 024. Note that this requirement means that a *SequenceNumber* does not reset when a new *TokenId* is issued. The *SequenceNumber* shall be incremented by exactly one for each *MessageChunk* sent. For backward compatibility, receivers shall accept *SequenceNumbers* less than 1 023 and greater than 4 294 966 271 provided they are in sequence. Administrators shall be able to disable this backward compatibility. Receivers shall log a warning when a rollover does not conform to the current specification.

The sequence header is followed by the *Message* body which is encoded with the OPC UA Binary encoding as described in 5.2.9. The body may be split across multiple *MessageChunks*.

6.7.2.5 Message footer

Each *MessageChunk* also has a footer with the fields defined in Table 47.

Table 47 – OPC UA Secure Conversation Message footer

Name	Data Type	Description
PaddingSize	Byte	The number of padding bytes (not including the byte for the PaddingSize).
Padding	Byte [*]	Padding added to the end of the <i>Message</i> to ensure length of the data to encrypt is an integer multiple of the encryption block size. The value of each byte of the padding is equal to PaddingSize.
ExtraPaddingSize	Byte	The most significant byte of a two-byte integer used to specify the padding size when the key used to encrypt the message chunk is larger than 2 048 bits. This field is omitted if the key length is less than or equal to 2 048 bits.
Signature	Byte [*]	The signature for the <i>MessageChunk</i> . The signature includes the all headers, all <i>Message</i> data, the PaddingSize and the Padding.

The formula to calculate the amount of padding depends on the amount of data that needs to be sent (called *BytesToWrite*). The sender shall first calculate the maximum amount of space available in the *MessageChunk* (called *MaxBodySize*) using the following formula:

$$\text{MaxBodySize} = \text{PlainTextBlockSize} * \text{Floor} ((\text{MessageChunkSize} - \text{HeaderSize} - \text{SignatureSize} - 1) / \text{CipherTextBlockSize}) - \text{SequenceHeaderSize} - \text{SignatureSize};$$

The *HeaderSize* includes the *MessageHeader* and the *SecurityHeader*. The *SequenceHeaderSize* is always 8 bytes.

During encryption a block with a size equal to *PlainTextBlockSize* is processed to produce a block with size equal to *CipherTextBlockSize*. These values depend on the encryption algorithm and may be the same.

The OPC UA *Message* can fit into a single chunk if *BytesToWrite* is less than or equal to the *MaxBodySize*. In this case the *PaddingSize* is calculated with this formula:

$$\text{PaddingSize} = \text{PlainTextBlockSize} - ((\text{BytesToWrite} + \text{SignatureSize} + 1) \% \text{PlainTextBlockSize});$$

If the *BytesToWrite* is greater than *MaxBodySize* the sender shall write *MaxBodySize* bytes with a *PaddingSize* of 0. The remaining *BytesToWrite* – *MaxBodySize* bytes shall be sent in subsequent *MessageChunks*.

The *PaddingSize* and *Padding* fields are not present if the *MessageChunk* is not encrypted.

The *Signature* field is not present if the *MessageChunk* is not signed.

6.7.3 MessageChunks and error handling

MessageChunks are sent as they are encoded. *MessageChunks* belonging to the same *Message* shall be sent sequentially. If an error occurs creating a *MessageChunk* then the sender shall send a final *MessageChunk* to the receiver that tells the receiver that an error occurred and that it should discard the previous chunks. The sender indicates that the *MessageChunk* contains an error by setting the *IsFinal* flag to 'A' (for Abort). Table 48 specifies the contents of the *Message* abort *MessageChunk*.

Table 48 – OPC UA Secure Conversation Message abort body

Name	Data Type	Description
Error	UInt32	The numeric code for the error. This shall be one of the values listed in Table 57.
Reason	String	A more verbose description of the error. This string shall not be more than 4 096 characters bytes. A <i>Client</i> shall ignore strings that are longer than this.

The receiver shall check the security on the abort *MessageChunk* before processing it. If everything is ok, then the receiver shall ignore the *Message* but shall not close the *SecureChannel*. The *Client* shall report the error back to the application as *StatusCode* for the request. If the *Client* is the sender, then it shall report the error without waiting for a response from the *Server*.

6.7.4 Establishing a SecureChannel

Most *Messages* require a *SecureChannel* to be established. A *Client* does this by sending an *OpenSecureChannel* request to the *Server*. The *Server* shall validate the *Message* and the *ClientCertificate* and return an *OpenSecureChannel* response. Some of the parameters defined for the *OpenSecureChannel* service are specified in the security header (see 6.7.2) instead of the body of the *Message*. ~~For this reason, the *OpenSecureChannel* Service is not the same as the one specified in IEC 62541-4.~~ Table 49 lists the parameters that appear in the body of the *Message*.

Note that IEC 62541-4 is an abstract specification which defines interfaces that can work with any protocol. This document provides a concrete implementation for specific protocols. This document is the normative reference for all protocols and takes precedence if there are differences with IEC 62541-4.

Table 49 – OPC UA Secure Conversation OpenSecureChannel Service

Name	Data Type
Request	
RequestHeader	RequestHeader
ClientProtocolVersion	UInt32
RequestType	SecurityTokenRequestType
SecurityMode	MessageSecurityMode
ClientNonce	ByteString
RequestedLifetime	Int32 UInt32
Response	
ResponseHeader	ResponseHeader
ServerProtocolVersion	UInt32
SecurityToken	ChannelSecurityToken
SecureChannelId	UInt32
TokenId	UInt32
CreatedAt	DateTime UtcTime
RevisedLifetime	Int32 UInt32
ServerNonce	ByteString

The *ClientProtocolVersion* and *ServerProtocolVersion* parameters are not defined in IEC 62541-4 and are added to the *Message* to allow backward compatibility if OPC UA-SecureConversation needs to be updated in the future. Receivers always accept numbers greater than the latest version that they support. The receiver with the higher version number is expected to ensure backward compatibility.

If OPC UA-SecureConversation is used with the OPC UA-TCP protocol (see 7.1) then the version numbers specified in the *OpenSecureChannel Messages* shall be the same as the version numbers specified in the OPC UA-TCP protocol *Hello/Acknowledge Messages*. The receiver shall close the channel and report a *Bad_ProtocolVersionUnsupported* error if there is a mismatch.

The *Server* shall return an error response as described in IEC 62541-4 if there are any errors with the parameters specified by the *Client*.

The *RevisedLifetime* tells the *Client* when it shall renew the *SecurityToken* by sending another *OpenSecureChannel* request. The *Client* shall continue to accept the old *SecurityToken* until it receives the *OpenSecureChannel* response. The *Server* ~~has to~~ shall accept requests secured with the old *SecurityToken* until that *SecurityToken* expires or until it receives a *Message* from the *Client* secured with the new *SecurityToken*. The *Server* shall reject renew requests if the *SenderCertificate* is not the same as the one used to create the *SecureChannel* or if there is a problem decrypting or verifying the signature. The *Client* shall abandon the *SecureChannel* if the *Certificate* used to sign the response is not the same as the *Certificate* used to encrypt the request. Note that datatype is a *UInt32* value representing the number of milliseconds instead of the *Double (Duration)* defined in IEC 62541-4. This optimization is possible because sub-millisecond timeouts are not supported.

The *OpenSecureChannel Messages* are signed and encrypted if the *SecurityMode* is not *None* (even if the *SecurityMode* is *Sign*).

The *Nonces* shall be cryptographic random numbers with a length specified by the *SecureChannelNonceLength* of the *SecurityPolicy*.

See IEC TR 62541-2 for more information on the requirements for random number generators. The *OpenSecureChannel Messages* are not signed or encrypted if the *SecurityMode* is *None*. The *Nonces* are ignored and should be set to null. The *SecureChannelId* and the *TokenId* are still assigned but no security is applied to *Messages* exchanged via the channel. The *SecurityToken* shall still be renewed before the *RevisedLifetime* expires. Receivers shall still ignore invalid or expired *TokenIds*.

If the communication channel breaks the *Server* shall maintain the *SecureChannel* long enough to allow the *Client* to reconnect. The ~~*RevisedLifetime*~~ *RevisedLifetime* parameter also tells the *Client* how long the *Server* will wait. If the *Client* cannot reconnect within that period it shall assume the *SecureChannel* has been closed.

The *AuthenticationToken* in the *RequestHeader* shall be set to null.

If an error occurs after the *Server* has verified *Message* security it shall return a *ServiceFault* instead of a *OpenSecureChannel* response. The *ServiceFault Message* is described in IEC 62541-4.

If the *SecurityMode* is not *None* then the *Server* shall verify that a *SenderCertificate* and a *ReceiverCertificateThumbprint* were specified in the *SecurityHeader*.

6.7.5 Deriving keys

Once the *SecureChannel* is established the *Messages* are signed and encrypted with keys derived from the *Nonces* exchanged in the *OpenSecureChannel* call. These keys are derived by passing the *Nonces* to a pseudo-random function which produces a sequence of bytes from a set of inputs. A pseudo-random function is represented by the following function declaration:

```
Byte[] PRF(  
    Byte[] secret,  
    Byte[] seed,  
    Int32 length,  
    Int32 offset)
```

Where *length* is the number of bytes to return and *offset* is a number of bytes from the beginning of the sequence.

The lengths of the keys that need to be generated depend on the *SecurityPolicy* used for the channel. The following information is specified by the *SecurityPolicy*:

- a) *SigningKeyLength* (from the *DerivedSignatureKeyLength*);
- b) *EncryptingKeyLength* (implied by the *SymmetricEncryptionAlgorithm*);
- c) *EncryptingBlockSize* (implied by the *SymmetricEncryptionAlgorithm*).

The pseudo random function requires a secret and a seed. These values are derived from the *Nonces* exchanged in the *OpenSecureChannel* request and response. Table 50 specifies how to derive the secrets and seeds when using RSA based *SecurityPolicies*.

Table 50 – PRF inputs for RSA based SecurityPolicies

Name	Derivation
ClientSecret	The value of the <i>ClientNonce</i> provided in the <i>OpenSecureChannel</i> request.
ClientSeed	The value of the <i>ClientNonce</i> provided in the <i>OpenSecureChannel</i> request.
ServerSecret	The value of the <i>ServerNonce</i> provided in the <i>OpenSecureChannel</i> response.
ServerSeed	The value of the <i>ServerNonce</i> provided in the <i>OpenSecureChannel</i> response.

The parameters passed to the pseudo random function are specified in Table 51.

Table 51 – Cryptography key generation parameters

Key	Secret	Seed	Length	Offset
ClientSigningKey	ServerNonce ServerSecret	ClientNonce ClientSeed	SigningKeyLength	0
ClientEncryptingKey	ServerNonce ServerSecret	ClientNonce ClientSeed	EncryptingKeyLength	SigningKeyLength
ClientInitializationVector	ServerNonce ServerSecret	ClientNonce ClientSeed	EncryptingBlockSize	SigningKeyLength+ EncryptingKeyLength
ServerSigningKey	ClientNonce ClientSecret	ServerNonce ServerSeed	SigningKeyLength	0
ServerEncryptingKey	ClientNonce ClientSecret	ServerNonce ServerSeed	EncryptingKeyLength	SigningKeyLength
ServerInitializationVector	ClientNonce ClientSecret	ServerNonce ServerSeed	EncryptingBlockSize	SigningKeyLength+ EncryptingKeyLength

The *Client* keys are used to secure *Messages* sent by the *Client*. The *Server* keys are used to secure *Messages* sent by the *Server*.

~~The SSL/TLS specification defines a pseudo random function called P_SHA1 which is used for some SecurityProfiles. The P_SHA1 algorithm is defined as follows:~~

~~$$P_SHA1(secret, seed) = HMAC_SHA1(secret, A(1) + seed) +$$~~
~~$$HMAC_SHA1(secret, A(2) + seed) +$$~~
~~$$HMAC_SHA1(secret, A(3) + seed) + \dots$$~~

~~Where A(n) is defined as:~~

~~$$A(0) = seed$$~~

~~$$A(n) = HMAC_SHA1(secret, A(n-1))$$~~

~~+ indicates that the results are appended to previous results.~~

The SSL/TLS specification defines a pseudo random function called P_HASH which is used for this purpose. The function is iterated until it produces enough data for all of the required keys. The Offset in Table 51 references to the offset from the start of the generated data.

The P_hash algorithm is defined as follows:

$$P_HASH(secret, seed) = HMAC_HASH(secret, A(1) + seed) +$$

$$HMAC_HASH(secret, A(2) + seed) +$$

$$HMAC_HASH(secret, A(3) + seed) + \dots$$

Where A(n) is defined as:

$$A(0) = seed$$

$$A(n) = HMAC_HASH(secret, A(n-1))$$

+ indicates that the results are appended to previous results.

where 'HASH' is a hash function such as SHA256. The hash function to use depends on the *SecurityPolicyUri*.

6.7.6 Verifying Message security

The contents of the *MessageChunk* shall not be interpreted until the *Message* is decrypted and the signature and sequence number verified.

If an error occurs during *Message* verification the receiver shall close the communication channel. If the receiver is the *Server*, it shall also send a transport error *Message* before closing the channel. Once the channel is closed the *Client* shall attempt to re-open the channel and request a new *SecurityToken* by sending an *OpenSecureChannel* request. The mechanism for sending transport errors to the *Client* depends on the communication channel.

The receiver shall first check the *SecureChannelId*. This value may be 0 if the *Message* is an *OpenSecureChannel* request. For other *Messages*, it shall report a *Bad_SecureChannelUnknown* error if the *SecureChannelId* is not recognized. If the *Message* is an *OpenSecureChannel* request and the *SecureChannelId* is not 0 then the *SenderCertificate* shall be the same as the *SenderCertificate* used to create the channel.

If the *Message* is secured with asymmetric algorithms, then the receiver shall verify that it supports the requested *SecurityPolicy*. If the *Message* is the response sent to the *Client*, then the *SecurityPolicy* shall be the same as the one specified in the request. In the *Server*, the *SecurityPolicy* shall be the same as the one used to originally create the *SecureChannel*. The receiver shall check that the *Certificate* is trusted first and return *Bad_CertificateUntrusted* on error. The receiver shall then verify the *SenderCertificate* using the rules defined in IEC 62541-4. The receiver shall report the appropriate error if *Certificate* validation fails. The receiver shall verify the *ReceiverCertificateThumbprint* and report a *Bad_CertificateUnknownCertificateInvalid* error if it does not recognize it.

If the *Message* is secured with symmetric algorithms, then a *Bad_SecureChannelTokenUnknown* error shall be reported if the *TokenId* refers to a *SecurityToken* that has expired or is not recognized.

If decryption or signature validation fails, then a *Bad_SecurityChecksFailed* error is reported. If an implementation allows multiple *SecurityModes* to be used the receiver shall also verify that the *Message* was secured properly as required by the *SecurityMode* specified in the *OpenSecureChannel* request.

After the security validation is complete the receiver shall verify the *RequestId* and the *SequenceNumber*. If these checks fail a *Bad_SecurityChecksFailed* error is reported. The *RequestId* only needs to be verified by the *Client* since only the *Client* knows if it is valid or not. If the *SequenceNumber* is not valid, the receiver shall log a *Bad_SequenceNumberInvalid* error.

At this point the *SecureChannel* knows it is dealing with an authenticated *Message* that was not tampered with or resent. This means the *SecureChannel* can return secured error responses if any further problems are encountered.

Stacks that implement UASC shall have a mechanism to log errors when invalid *Messages* are discarded. This mechanism is intended for developers, systems integrators and administrators to debug network system configuration issues and to detect attacks on the network.

7 TransportProtocols

7.1 OPC UA TCP Connection Protocol

7.1.1 Overview

~~OPC UA TCP is a simple TCP based protocol that establishes a full duplex channel between a Client and Server. This protocol has two key features that differentiate it from HTTP. First, this protocol allows responses to be returned in any order. Second, this protocol allows responses to be returned on a different TCP transport end-point if communication failures cause temporary TCP session interruption.~~

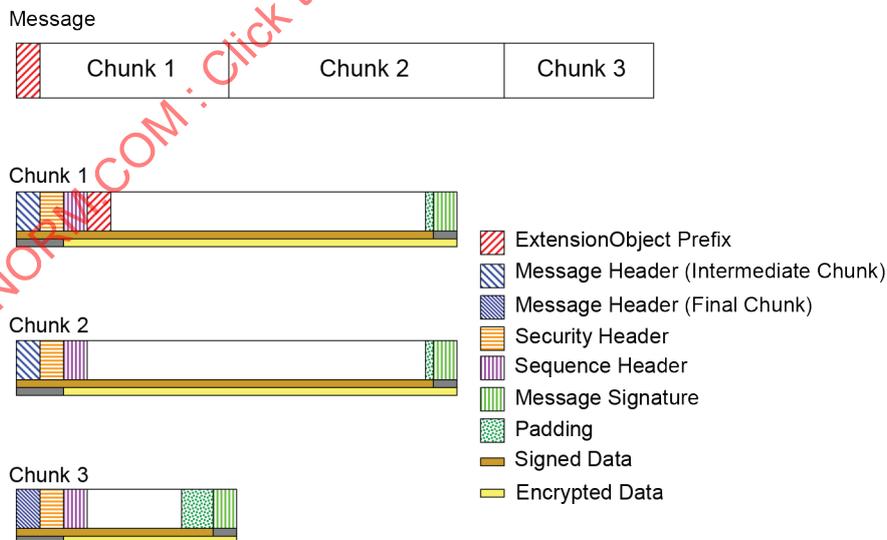
OPC UA Connection Protocol (UACP) is an abstract protocol that establishes a full duplex channel between a *Client* and *Server*. Concrete implementations of the UACP can be built with any middleware that supports full-duplex exchange of messages including TCP/IP and WebSockets. The term "*TransportConnection*" describes the middleware-specific connection used to exchange messages. For example, a socket is the *TransportConnection* for TCP/IP. *TransportConnections* allow responses to be returned in any order and allow responses to be returned on a different physical *TransportConnection* if communication failures cause temporary interruptions.

The OPC UA TCP Connection Protocol is designed to work with the *SecureChannel* implemented by a layer higher in the stack. For this reason, the OPC UA TCP Connection Protocol defines its interactions with the *SecureChannel* in addition to the wire protocol.

7.1.2 Message structure

7.1.2.1 Overview

Figure 12 illustrates the structure of a *Message* placed on the wire. This also illustrates how the *Message* elements defined by the OPC UA Binary Encoding mapping (see 5.2) and the OPC UA Secure Conversation mapping (see 6.7) relate to the OPC UA Connection Protocol *Messages*.



IEC

Figure 12 – OPC UA TCP Connection Protocol Message structure

7.1.2.2 Message Header

Every OPC UA TCP Connection Protocol *Message* has a header with the fields defined in Table 52.

Table 52 – OPC UA TCP Connection Protocol Message header

Name	Type	Description
MessageType	Byte [3]	A three byte ASCII code that identifies the <i>Message</i> type. The following values are defined at this time: HEL a <i>Hello Message</i> . ACK an <i>Acknowledge Message</i> . ERR an <i>Error Message</i> . RHE a <i>ReverseHello Message</i> . The <i>SecureChannel</i> layer defines additional values which the OPC UA TCP Connection Protocol layer shall accept.
Reserved	Byte [1]	Ignored. shall be set to the ASCII codes for 'F' if the <i>MessageType</i> is one of the values supported by the OPC UA TCP Connection Protocol.
MessageSize	UInt32	The length of the <i>Message</i> , in bytes. This value includes the 8 bytes for the <i>Message</i> header.

The layout of the OPC UA TCP Connection Protocol Message header is intentionally identical to the first 8 bytes of the OPC UA Secure Conversation Message header defined in Table 43. This allows the OPC UA TCP Connection Protocol layer to extract the *SecureChannel Messages* from the incoming stream even if it does not understand their contents.

The OPC UA TCP Connection Protocol layer shall verify the *MessageType* and make sure the *MessageSize* is less than the negotiated *ReceiveBufferSize* before passing any *Message* onto the *SecureChannel* layer.

7.1.2.3 Hello Message

The Hello Message has the additional fields shown in Table 53.

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Table 53 – OPC UA TCP Connection Protocol Hello Message

Name	Data Type	Description
ProtocolVersion	UInt32	The latest version of the OPC UA TCP UACP protocol supported by the <i>Client</i> . The <i>Server</i> may reject the <i>Client</i> by returning <i>Bad_ProtocolVersionUnsupported</i> . If the <i>Server</i> accepts the connection, it is responsible for ensuring that it returns <i>Messages</i> that conform to this version of the protocol. The <i>Server</i> shall always accept versions greater than what it supports.
ReceiveBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender can receive. This value shall be greater than 8 192 bytes.
SendBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender will send. This value shall be greater than 8 192 bytes.
MaxMessageSize	UInt32	The maximum size for any response <i>Message</i> . The <i>Server</i> shall abort the <i>Message</i> with a <i>Bad_ResponseTooLarge</i> <i>StatusCode</i> <i>Error Message</i> if a response <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. The <i>Message</i> size is calculated using the unencrypted <i>Message</i> body. A value of zero indicates that the <i>Client</i> has no limit.
MaxChunkCount	UInt32	The maximum number of chunks in any response <i>Message</i> . The <i>Server</i> shall abort the <i>Message</i> with a <i>Bad_ResponseTooLarge</i> <i>StatusCode</i> <i>Error Message</i> if a response <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. A value of zero indicates that the <i>Client</i> has no limit.
EndpointUrl	String	The URL of the <i>Endpoint</i> which the <i>Client</i> wished to connect to. The encoded value shall be less than 4 096 bytes. Servers shall return a <i>Bad_TcpUrlRejected</i> <i>error</i> <i>TcpEndpointUrlInvalid</i> <i>Error Message</i> and close the connection if the length exceeds 4 096 bytes or if it does not recognize the resource identified by the URL.

The *EndpointUrl* parameter is used to allow multiple *Servers* to share the same ~~port~~ endpoint on a machine. The process listening (also known as the proxy) on the ~~port~~ endpoint would connect to the *Server* identified by the *EndpointUrl* and would forward all *Messages* to the *Server* via this socket. If one socket closes, then the proxy shall close the other socket.

If the *Server* does not have sufficient resources to allow the establishment of a new *SecureChannel*, it shall immediately return a *Bad_TcpNotEnoughResources* *Error Message* and gracefully close the socket. *Clients* should not overload *Servers* that return this error by immediately trying to create a new *SecureChannel*.

7.1.2.4 Acknowledge Message

The *Acknowledge Message* has the additional fields shown in Table 54.

Table 54 – OPC UA TCP Connection Protocol Acknowledge Message

Name	Type	Description
ProtocolVersion	UInt32	The latest version of the OPC UA TCP UACP protocol supported by the <i>Server</i> . If the <i>Client</i> accepts the connection, it is responsible for ensuring that it sends <i>Messages</i> that conform to this version of the protocol. The <i>Client</i> shall always accept versions greater than what it supports.
ReceiveBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender can receive. This value shall not be larger than what the <i>Client</i> requested in the Hello <i>Message</i> . This value shall be greater than 8 192 bytes.
SendBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender will send. This value shall not be larger than what the <i>Client</i> requested in the Hello <i>Message</i> . This value shall be greater than 8 192 bytes.
MaxMessageSize	UInt32	The maximum size for any request <i>Message</i> . The <i>Client</i> shall abort the <i>Message</i> with a <i>Bad_RequestTooLarge</i> <i>StatusCode</i> if a request <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. The <i>Message</i> size is calculated using the unencrypted <i>Message</i> body. A value of zero indicates that the <i>Server</i> has no limit.
MaxChunkCount	UInt32	The maximum number of chunks in any request <i>Message</i> . The <i>Client</i> shall abort the <i>Message</i> with a <i>Bad_RequestTooLarge</i> <i>StatusCode</i> if a request <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. A value of zero indicates that the <i>Server</i> has no limit.

7.1.2.5 Error Message

The *Error Message* has the additional fields shown in Table 55.

Table 55 – OPC UA TCP Connection Protocol Error Message

Name	Type	Description
Error	UInt32	The numeric code for the error. This shall be one of the values listed in Table 57.
Reason	String	A more verbose description of the error. This string shall not be more than 4 096 characters bytes. A <i>Client</i> shall ignore strings that are longer than this.

The socket is always closed gracefully by the ~~Server~~ *Client* after it ~~sends~~ receives an *Error Message*.

7.1.2.6 ReverseHello Message

The *ReverseHello Message* has the additional fields shown in Table 56.

Table 56 – OPC UA Connection Protocol ReverseHello Message

Name	Data Type	Description
ServerUri	String	The <i>ApplicationUri</i> of the <i>Server</i> which sent the <i>Message</i> . The encoded value shall be less than 4 096 bytes. <i>Client</i> shall return a <i>Bad_TcpEndpointUrlInvalid</i> error and close the connection if the length exceeds 4 096 bytes or if it does not recognize the <i>Server</i> identified by the URI.
EndpointUrl	String	The URL of the <i>Endpoint</i> which the <i>Client</i> uses when establishing the <i>SecureChannel</i> . This value shall be passed back to the <i>Server</i> in the <i>Hello Message</i> . The encoded value shall be less than 4 096 bytes. <i>Clients</i> shall return a <i>Bad_TcpEndpointUrlInvalid</i> error and close the connection if the length exceeds 4 096 bytes or if it does not recognize the resource identified by the URL. This value is a unique identifier for the <i>Server</i> which the <i>Client</i> may use to look up configuration information. It should be one of the URLs returned by the <i>GetEndpoints Service</i> .

For connection based protocols, such as TCP, the *ReverseHello Message* allows *Servers* behind firewalls with no open ports to connect to a *Client* and request that the *Client* establish a *SecureChannel* using the socket created by the *Server*.

For message based protocols the *ReverseHello Message* allows *Servers* to announce their presence to a *Client*. In this scenario, the *EndpointUrl* specifies the *Server's* protocol-specific address and any tokens required to access it.

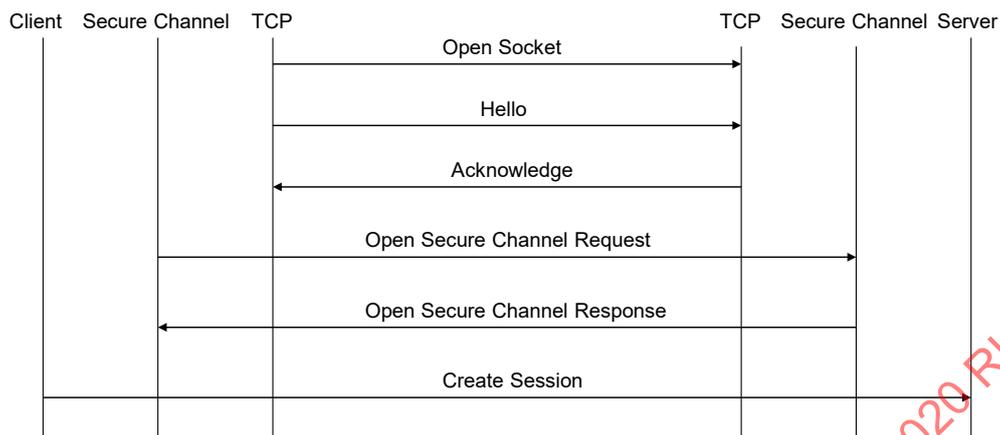
7.1.3 Establishing a connection

~~Connections are always initiated by the *Client* which creates the socket before it sends the first *OpenSecureChannel* request. After creating the socket the first *Message* sent shall be a *Hello* which specifies the buffer sizes that the *Client* supports. The *Server* shall respond with an *Acknowledge Message* which completes the buffer negotiation. The negotiated buffer size shall be reported to the *SecureChannel* layer. The negotiated *SendBufferSize* specifies the size of the *MessageChunks* to use for *Messages* sent over the connection.~~

~~The *Hello/Acknowledge Messages* may only be sent once. If they are received again the receiver shall report an error and close the socket. *Servers* shall close any socket after a period of time if it does not receive a *Hello Message*. This period of time shall be configurable and have a default value which does not exceed two minutes.~~

~~The *Client* sends the *OpenSecureChannel* request once it receives the *Acknowledge* back from the *Server*. If the *Server* accepts the new channel it shall associate the socket with the *SecureChannelId*. The *Server* uses this association to determine which socket to use when it has to send a response to the *Client*. The *Client* does the same when it receives the *OpenSecureChannel* response.~~

~~The sequence of *Messages* when establishing a OPC UA TCP connection are shown in Figure 15.~~



IEC

Figure 15 – Establishing a OPC UA TCP connection

~~The Server Application does not do any processing while the SecureChannel is negotiated; however, the Server Application shall to provide the Stack with the list of trusted Certificates. The Stack shall provide notifications to the Server Application whenever it receives an OpenSecureChannel request. These notifications shall include the OpenSecureChannel or Error response returned to the Client.~~

Connections may be initiated by the *Client* or by the *Server* when they create a *TransportConnection* and establish a communication with their peer. If the *Client* creates the *TransportConnection*, the first *Message* sent shall be a *Hello* which specifies the buffer sizes that the *Client* supports. The *Server* shall respond with an *Acknowledge Message* which completes the buffer negotiation. The negotiated buffer size shall be reported to the *SecureChannel* layer. The negotiated *SendBufferSize* specifies the size of the *MessageChunks* to use for *Messages* sent over the connection.

If the *Server* creates the *TransportConnection* the first *Message* shall be a *ReverseHello* sent to the *Client*. If the *Client* accepts the connection, it sends a *Hello* message back to the *Server* which starts the buffer negotiation described for the *Client* initiated connection.

The *Hello/Acknowledge Messages* may only be sent once. If they are received again the receiver shall report an error and close the *TransportConnection*. Applications accepting incoming connections shall close any *TransportConnection* after a period of time if it does not receive a *Hello* or *ReverseHello Message*. This period of time shall be configurable and have a default value which does not exceed two minutes.

The *Client* sends the *OpenSecureChannel* request once it receives the *Acknowledge* back from the *Server*. If the *Server* accepts the new channel it shall associate the *TransportConnection* with the *SecureChannelId*. The *Server* uses this association to determine which *TransportConnection* to use when it needs to send a response to the *Client*. The *Client* does the same when it receives the *OpenSecureChannel* response.

The sequence of *Messages* when establishing a *Client* initiated OPC UA Connection Protocol connection is shown in Figure 13.

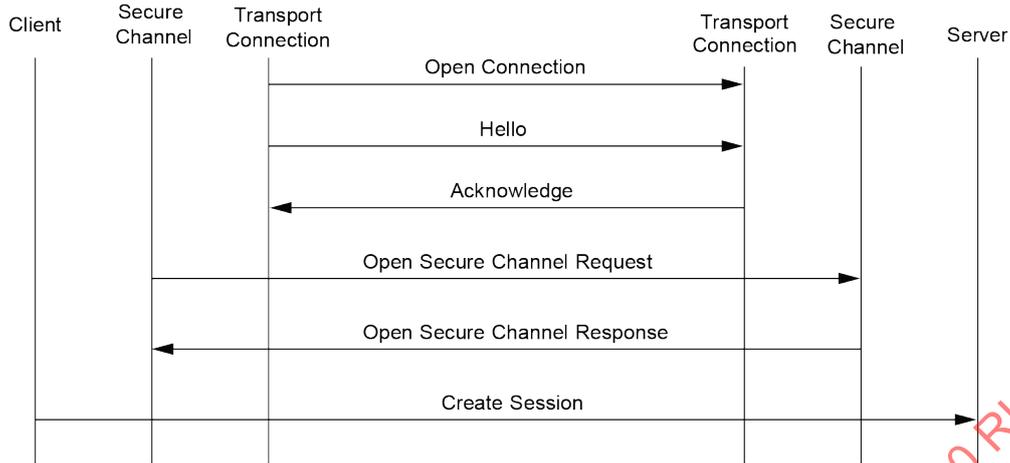


Figure 13 – Client initiated OPC UA Connection Protocol connection

The sequence of *Messages* when establishing a *Server* initiated OPC UA Connection Protocol connection is shown in Figure 14.

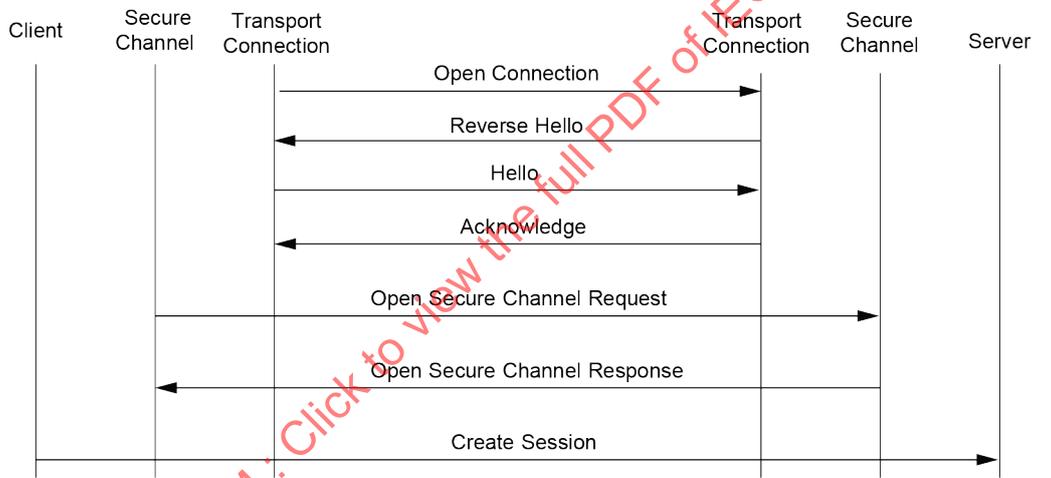


Figure 14 – Server initiated OPC UA Connection Protocol connection

The *Server* application does not do any processing while the *SecureChannel* is negotiated; however, the *Server* application shall to provide the *Stack* with the list of trusted *Certificates*. The *Stack* shall provide notifications to the *Server* application whenever it receives an *OpenSecureChannel* request. These notifications shall include the *OpenSecureChannel* or *Error* response returned to the *Client*.

The *Server* needs to be configured and enabled by an administrator to connect to one or more *Clients*. For each *Client*, the administrator shall provide an *ApplicationUri* and an *EndpointUrl* for the *Client*. If the *Client EndpointUrl* is not known, the administrator may provide the *EndpointUrl* for a GDS (see IEC 62541-12) which knows about the *Client*. The *Server* should expect that it will take some time for a *Client* to respond to a *ReverseHello*. Once a *Client* closes a *SecureChannel* or if the socket is closed without establishing a *SecureChannel*, the *Server* shall create a new socket and send a new *ReverseHello* message. Administrators may limit the number of simultaneous sockets that a *Server* will create.

7.1.4 Closing a connection

The *Client* closes the connection by sending a *CloseSecureChannel* request and closing the socket gracefully. When the *Server* receives this *Message*, it shall release all resources allocated for the channel. The body of the *CloseSecureChannel* request is empty. The *Server* does not send a *CloseSecureChannel* response.

If security verification fails for the *CloseSecureChannel Message*, then the *Server* shall report the error and close the socket. ~~The *Server* shall allow the *Client* to attempt to reconnect.~~

The sequence of *Messages* when closing an OPC UA TCP Connection Protocol connection is shown in Figure 15.

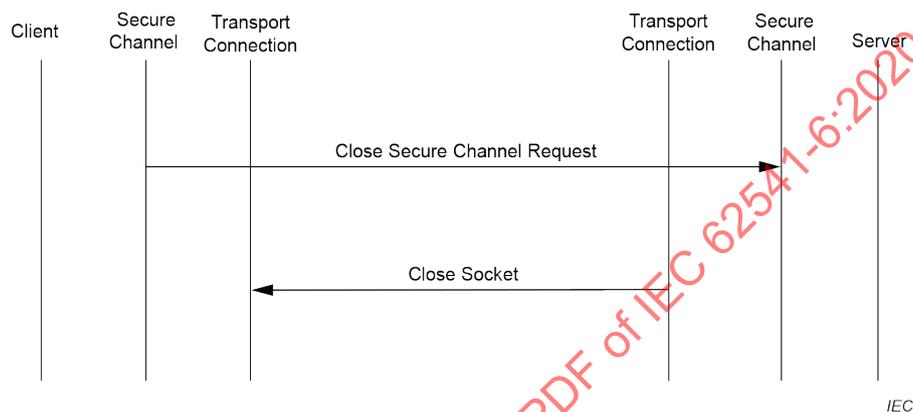


Figure 15 – Closing a OPC UA TCP Connection Protocol connection

The *Server* application does not do any processing when the *SecureChannel* is closed; however, the *Stack* shall provide notifications to the *Server* application whenever a *CloseSecureChannel* request is received or when the *Stack* cleans up an abandoned *SecureChannel*.

7.1.5 Error handling

~~When a fatal error occurs the *Server* shall send an *Error Message* to the *Client* and close the socket. When a *Client* encounters one of these errors, it shall also close the socket but does not send an *Error Message*. After the socket is closed a *Client* shall try to reconnect automatically using the mechanisms described in 7.1.6.~~

~~The possible OPC UA TCP errors are defined in Table 41.~~

Table 41 – OPC UA TCP error codes

Name	Description
TcpServerTooBusy	The <i>Server</i> cannot process the request because it is too busy. It is up to the <i>Server</i> to determine when it needs to return this <i>Message</i> . A <i>Server</i> can control the how frequently a <i>Client</i> reconnects by waiting to return this error.
TcpMessageTypeInvalid	The type of the <i>Message</i> specified in the header invalid. Each <i>Message</i> starts with a 4 byte sequence of ASCII values that identifies the <i>Message</i> type. The <i>Server</i> returns this error if the <i>Message</i> type is not accepted. Some of the <i>Message</i> types are defined by the <i>SecureChannel</i> layer.
TcpSecureChannelUnknown	The <i>SecureChannelId</i> and/or <i>TokenId</i> are not currently in use. This error is reported by the <i>SecureChannel</i> layer.
TcpMessageTooLarge	The size of the <i>Message</i> specified in the header is too large. The <i>Server</i> returns this error if the <i>Message</i> size exceeds its maximum buffer size or the receive buffer size negotiated during the Hello/Acknowledge exchange.
TcpTimeout	A timeout occurred while accessing a resource. It is up to the <i>Server</i> to determine when a timeout occurs.
TcpNotEnoughResources	There are not enough resources to process the request. The <i>Server</i> returns this error when it runs out of memory or encounters similar resource problems. A <i>Server</i> can control the how frequently a <i>Client</i> reconnects by waiting to return this error.
TcpInternalError	An internal error occurred. This should only be returned if an unexpected configuration or programming error occurs.
TcpUrlRejected	The <i>Server</i> does not recognize the <i>EndpointUrl</i> specified.
SecurityChecksFailed	The <i>Message</i> was rejected because it could not be verified.
RequestInterrupted	The request could not be sent because of a network interruption.
RequestTimeout	Timeout occurred while processing the request.
SecureChannelClosed	The secure channel has been closed.
SecureChannelTokenUnknown	The <i>SecurityToken</i> has expired or is not recognized.
CertificateUntrusted	The sender <i>Certificate</i> is not trusted by the receiver.
CertificateTimelInvalid	The sender <i>Certificate</i> has expired or is not yet valid.
CertificateIssuerTimelInvalid	The issuer for the sender <i>Certificate</i> has expired or is not yet valid.
CertificateUseNotAllowed	The sender's <i>Certificate</i> may not be used for establishing a secure channel.
CertificateIssuerUseNotAllowed	The issuer <i>Certificate</i> may not be used as a <i>Certificate Authority</i> .
CertificateRevocationUnknown	Could not verify the revocation status of the sender's <i>Certificate</i> .
CertificateIssuerRevocationUnknown	Could not verify the revocation status of the issuer <i>Certificate</i> .
CertificateRevoked	The sender <i>Certificate</i> has been revoked by the issuer.
IssuerCertificateRevoked	The issuer <i>Certificate</i> has been revoked by its issuer.
CertificateUnknown	The receiver <i>Certificate</i> thumbprint is not recognized by the receiver.

The numeric values for these error codes are defined in A.2.

When a fatal error occurs, the *Server* shall send an *Error Message* to the *Client* and closes the *TransportConnection* gracefully. When the *Client* receives an *Error Message* it reports the error to the application and closes the *TransportConnection* gracefully. If a *Client* encounters a fatal error, it shall report the error to the application and send a *CloseSecureChannel Message*. The *Server* shall close the *TransportConnection* gracefully when it receives the *CloseSecureChannel Message*.

The possible OPC UA Connection Protocol errors are defined in Table 57.

Table 57 – OPC UA Connection Protocol error codes

Name	Description
Bad_TcpServerTooBusy	The <i>Server</i> cannot process the request because it is too busy. It is up to the <i>Server</i> to determine when it needs to return this <i>Message</i> . A <i>Server</i> can control the how frequently a <i>Client</i> reconnects by waiting to return this error.
Bad_TcpMessageTypeInvalid	The type of the <i>Message</i> specified in the header invalid. Each <i>Message</i> starts with a 4-byte sequence of ASCII values that identifies the <i>Message</i> type. The <i>Server</i> returns this error if the <i>Message</i> type is not accepted. Some of the <i>Message</i> types are defined by the <i>SecureChannel</i> layer.
Bad_TcpSecureChannelUnknown	The <i>SecureChannelId</i> and/or <i>TokenId</i> are not currently in use. This error is reported by the <i>SecureChannel</i> layer.
Bad_TcpMessageTooLarge	The size of the <i>Message</i> specified in the header is too large. The <i>Server</i> returns this error if the <i>Message</i> size exceeds its maximum buffer size or the receive buffer size negotiated during the Hello/Acknowledge exchange.
Bad_Timeout	A timeout occurred while accessing a resource. It is up to the <i>Server</i> to determine when a timeout occurs.
Bad_TcpNotEnoughResources	There are not enough resources to process the request. The <i>Server</i> returns this error when it runs out of memory or encounters similar resource problems. A <i>Server</i> can control the how frequently a <i>Client</i> reconnects by waiting to return this error.
Bad_TcpInternalError	An internal error occurred. This should only be returned if an unexpected configuration or programming error occurs.
Bad_TcpEndpointUrlInvalid	The <i>Server</i> does not recognize the <i>EndpointUrl</i> specified.
Bad_SecurityChecksFailed	The <i>Message</i> was rejected because it could not be verified.
Bad_RequestInterrupted	The request could not be sent because of a network interruption.
Bad_RequestTimeout	Timeout occurred while processing the request.
Bad_SecureChannelClosed	The secure channel has been closed.
Bad_SecureChannelTokenUnknown	The <i>SecurityToken</i> has expired or is not recognized.
Bad_CertificateUntrusted	The sender <i>Certificate</i> is not trusted by the receiver.
Bad_CertificateTimeInvalid	The sender <i>Certificate</i> has expired or is not yet valid.
Bad_CertificateIssuerTimeInvalid	The issuer for the sender <i>Certificate</i> has expired or is not yet valid.
Bad_CertificateUseNotAllowed	The sender's <i>Certificate</i> may not be used for establishing a secure channel.
Bad_CertificateIssuerUseNotAllowed	The issuer <i>Certificate</i> may not be used as a <i>Certificate Authority</i> .
Bad_CertificateRevocationUnknown	Could not verify the revocation status of the sender's <i>Certificate</i> .
Bad_CertificateIssuerRevocationUnknown	Could not verify the revocation status of the issuer <i>Certificate</i> .
Bad_CertificateRevoked	The sender <i>Certificate</i> has been revoked by the issuer.
Bad_IssuerCertificateRevoked	The issuer <i>Certificate</i> has been revoked by its issuer.
Bad_SequenceNumberInvalid	This sequence number on the message was not valid.

The numeric values for these error codes are defined in A.2.

NOTE The 'Tcp' prefix for some of the error codes in Table 57 was chosen when TCP/IP was the only implementation of the OPC UA Connection Protocol. These codes are used with any implementation of the OPC UA Connection Protocol.

7.1.6 Error recovery

Once the *SecureChannel* has been established, the *Client* shall go into an error recovery state whenever the socket breaks or if the *Server* returns an OPC UA TCP Error Message as defined in Table 40. While in this state the *Client* periodically attempts to reconnect to the *Server*. If the reconnect succeeds the *Client* sends a *Hello* followed by an *OpenSecureChannel* request (see 6.7.4) that re-authenticates the *Client* and associates the new socket with the existing *SecureChannel*.

The *Client* shall wait between reconnect attempts. The first reconnect shall happen immediately. After that, the wait period should start as 1 second and increase gradually to a maximum of 2 minutes. One sequence would double the period each attempt until reaching the maximum. In other words, the *Client* would use the following wait periods: { 0, 1, 2, 4, 8, 16, 32, 64, 120, 120, ...}. The *Client* shall keep attempting to reconnect until the *SecureChannel* is closed or after the period equal to the *RevisedLifetime* of the last *SecurityToken* elapses.

The *Stack* in the *Server* should not discard responses if there is no connection immediately available. It should wait and see if the *Client* creates a new socket. It is up to the *Server* stack implementation to decide how long it will wait and how many responses it is willing to hold onto.

The *Stack* in the *Client* shall not fail requests that have already been sent and are waiting for a response when the socket is closed. However, these requests may timeout and report a *Bad_TcpRequestTimeout* error to the *Application*. If the *Client* sends a new request the stack shall either buffer the request or return a *Bad_TcpRequestInterrupted* error. The *Client* can stop the reconnect process by closing the *SecureChannel*.

The *Server* may abandon the *SecureChannel* before a *Client* is able to reconnect. If this happens the *Client* will get a *Bad_TcpSecureChannelUnknown* error in response to the *OpenSecureChannel* request. The *Stack* shall return this error to the *Application* that can attempt to create a new *SecureChannel*.

The negotiated buffer sizes should never change when a connection is recovered; however, the buffer sizes are negotiated before the *Server* knows whether the socket is being used for an existing *SecureChannel* or a new one. A *Client* shall treat this as a fatal error, close the *SecureChannel* and returns an *Bad_TcpSecureChannelClosed* error to the *Application*.

The sequence of Messages when recovering an OPC UA TCP connection is shown in Figure 17.

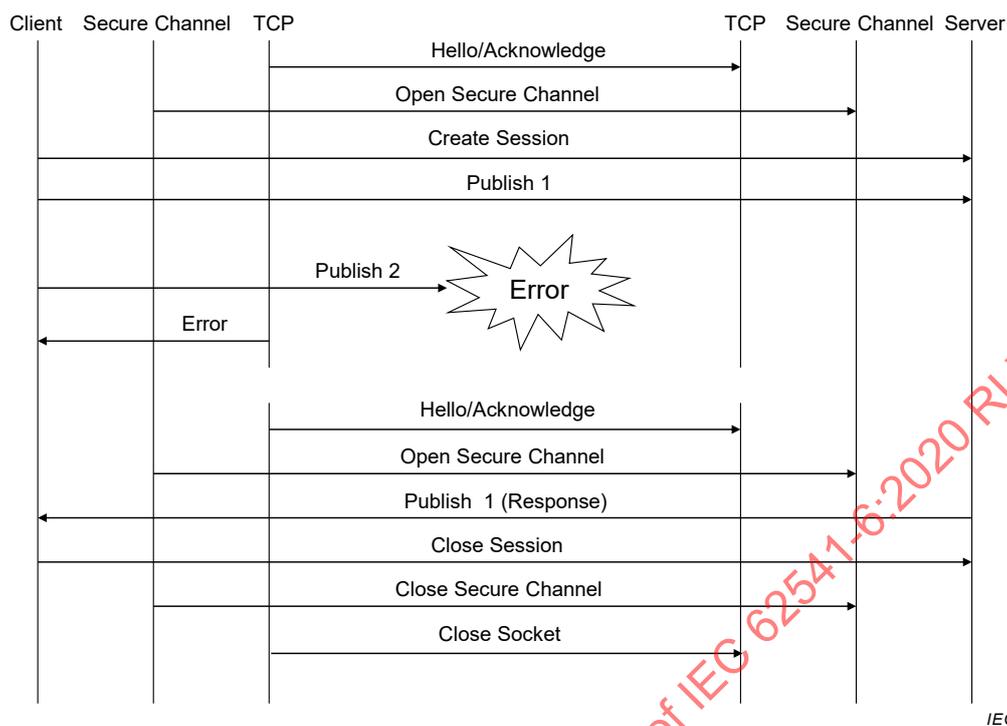


Figure 17 – Recovering an OPC UA TCP connection

7.2 OPC UA TCP

TCP/IP is a ubiquitous protocol that provides full-duplex communication between two applications. A socket is the *TransportConnection* in the TCP/IP implementation of the OPC UA Connection Protocol.

The URL scheme for endpoints using OPC UA TCP is 'opc.tcp'.

The *TransportProfileUri* shall be a URI for the TCP transport defined in IEC 62541-7.

7.3 SOAP/HTTP

7.2.1 Overview

SOAP describes an XML-based syntax for exchanging Messages between Applications. OPC UA Messages are exchanged using SOAP by serializing the OPC UA Messages using one of the supported encodings described in Clause 5 and inserting that encoded Message into the body of the SOAP Message.

All OPC UA Applications that support the SOAP/HTTP transport shall support SOAP 1.2 as described in SOAP Part 1.

All OPC UA Messages are exchanged using the request-response Message exchange pattern defined in SOAP Part 2 even if the OPC UA service does not specify any output parameters. In these cases, the Server shall return an empty response Message that tells the Client that no errors occurred.

WS-I Basic Profile 1.1 defines best practices when using SOAP Messages which will help ensure interoperability. All OPC UA implementations shall conform to this specification.

HTTP is the network communication protocol used to exchange SOAP Messages. An OPC UA service request Message is always sent by the Client in the body of an HTTP POST request.

~~The Server returns an OPC UA response Message in the body of the HTTP response. The HTTP binding for SOAP is described completely in SOAP Part 2.~~

~~OPC UA does not define any SOAP headers; however, SOAP Messages containing OPC UA Messages will include headers used by the other WS specifications in the stack.~~

~~SOAP faults are returned only for errors that occurred with in the SOAP stack. Errors that occur within in the Application are returned as OPC UA error response Messages as described in IEC 62541-4.~~

~~WS Addressing defines standard headers used to route SOAP Messages through intermediaries. Implementations shall support the WS Addressing headers listed Table 42.~~

Table 42 – WS Addressing headers

Header	Request	Response
wsa:To	Required	Optional
wsa:From	Optional	Optional
wsa:ReplyTo	Required	Not-Used
wsa:Action	Required	Required
wsa:MessageID	Required	Optional
wsa:RelatesTo	Not-Used	Required

~~Note that WS Addressing defines standard URIs to use in the ReplyTo and From headers when a Client does not have an externally accessible endpoint. In these cases, the SOAP response is always returned to the Client using the same communication channel that sent the request.~~

~~OPC UA Servers shall ignore the wsa:FaultTo header if it is specified in a request.~~

7.2.2 XML Encoding

~~The OPC UA XML Encoding specifies a way to represent an OPC UA Message as an XML element. This element is added to the SOAP Message as the only child of the SOAP body element.~~

~~If an error occurs in the Server while parsing the request body, the Server may return a SOAP fault or it may return an OPC UA error response.~~

~~The SOAP Action associated with an XML encoded request Message always has the form:~~

~~`http://opcfoundation.org/UA/2008/02/Services.wsdl/<service name>`~~

~~Where <service name> is the name of the OPC UA Service being invoked.~~

~~The SOAP Action associated with an XML encoded response Message always has the form:~~

~~`http://opcfoundation.org/UA/2008/02/Services.wsdl/<service name>Response`~~

7.2.3 OPC UA Binary Encoding

~~The OPC UA Binary Encoding specifies a way to represent an OPC UA Message as a sequence of bytes. These bytes sequences shall be encoded in the SOAP body using the Base64 data format.~~

~~The Base64 data format is a UTF-7 representation of binary data that is less efficient than raw binary data, however, many OPC UA Applications that exchange Messages using SOAP will~~

~~find that encoding OPC UA Messages in OPC UA Binary and then encoding the binary in Base64 is more efficient than encoding everything in XML.~~

~~The WSDL definition for a OPC UA Binary encoded request Message is:~~

```
<xs:element name="InvokeServiceRequest" type="xs:base64Binary" nillable="true" />

<wsdl:message name="InvokeServiceMessage">
  <wsdl:part name="input" element="s0:InvokeServiceRequest"/>
</wsdl:message>
```

~~The SOAP Action associated with an OPC UA Binary encoded request Message always has the form:~~

~~<http://opcfoundation.org/UA/2008/02/Services.wsdl/InvokeService>~~

~~The WSDL definition for an OPC UA Binary encoded response Message is:~~

```
<xs:element name="InvokeServiceResponse" type="xs:base64Binary" nillable="true" />

<wsdl:message name="InvokeServiceResponseMessage">
  <wsdl:part name="output" element="s0:InvokeServiceResponse"/>
</wsdl:message>
```

~~The SOAP Action associated with an OPC UA Binary encoded response Message always has the form:~~

~~<http://opcfoundation.org/UA/2008/02/Services.wsdl/InvokeServiceResponse>~~

NOTE Deprecated in Version 1.03 because WS-SecureConversation has not been widely adopted by industry.

7.4 OPC UA HTTPS

7.4.1 Overview

HTTPS refers HTTP Messages exchanged over a SSL/TLS connection (see HTTPS). The syntax of the HTTP Messages does not change and the only difference is a TLS connection is created instead of a TCP/IP connection. This implies that profiles which use this transport can also be used with HTTP when security is not a concern.

HTTPS is a protocol that provides transport security. This means all bytes are secured as they are sent without considering the Message boundaries. Transport security can only work for point to point communication and does not allow untrusted intermediaries or proxy servers to handle traffic.

The *SecurityPolicy* shall be specified, however, it only affects the algorithms used for signing the *Nonces* during the *CreateSession/ActivateSession* handshake. A *SecurityPolicy* of *None* indicates that the *Nonces* do not need to be signed. The *SecurityMode* is set to *Sign* unless the *SecurityPolicy* is *None*; in this case the *SecurityMode* shall be set to *None*. If a *UserIdentityToken* is to be encrypted, it shall be explicitly specified in the *UserTokenPolicy*.

An HTTP Header called 'OPCUA-SecurityPolicy' is used by the *Client* to tell the *Server* what *SecurityPolicy* it is using if there are multiple choices available. The value of the header is the URI for the *SecurityPolicy*. If the *Client* omits the header, then the *Server* shall assume a *SecurityPolicy* of *None*.

All HTTPS communications via a URL shall be treated as a single *SecureChannel* that is shared by multiple *Clients*. *Stacks* shall provide a unique identifier for the *SecureChannel* which allows applications correlate a request with a *SecureChannel*. This means that

Sessions can only be considered secure if the *AuthenticationToken* (see IEC 62541-4) is long (>20 bytes) and HTTPS encryption is enabled.

The cryptography algorithms used by HTTPS have no relationship to the *EndpointDescription SecurityPolicy* and are determined by the policies set for HTTPS and are outside the scope of OPC UA.

Figure 16 illustrates a few scenarios where the HTTPS transport could be used.

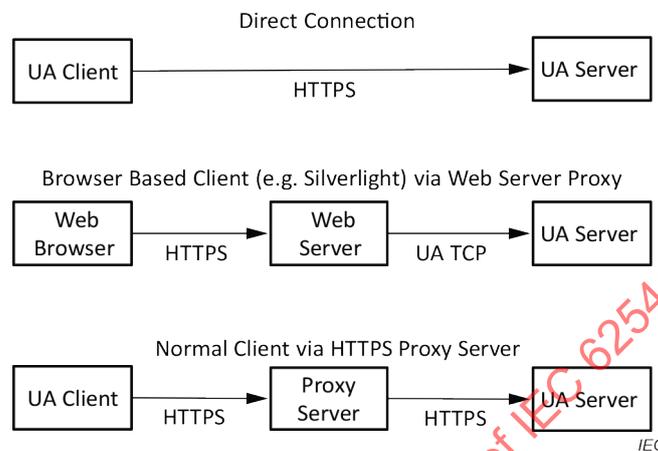


Figure 16 – Scenarios for the HTTPS Transport

In some scenarios, HTTPS communication will rely on an intermediary which is not trusted by the applications. If this is the case, then the HTTPS transport cannot be used to ensure security and the applications will ~~have~~ need to establish a secure tunnel like a VPN before attempting any OPC UA related communication.

Applications which support the HTTPS transport shall support HTTP 1.1 and SSL/TLS 1.0.

Some HTTPS implementations require that all *Servers* have a *Certificate* with a Common Name (CN) that matches the DNS name of the *Server* machine. This means that a *Server* with multiple DNS names will need multiple HTTPS certificates. If multiple *Servers* are on the same machine they may share HTTPS certificates. This means that *ApplicationCertificates* are not the same as HTTPS *Certificates*. Applications which use the HTTPS transport and require application authentication shall check application *Certificates* during the *CreateSession/ActivateSession* handshake.

HTTPS *Certificates* can be automatically generated; however, this will cause problems for *Clients* operating inside a restricted environment such as a web browser. Therefore, HTTPS certificates should be issued by an authority which is accepted by all web browsers which need to access the *Server*. The set of *Certificate* authorities accepted by the web browsers is determined by the organization that manages the *Client* machines. *Client* applications that are not running inside a web may use the trust list that is used for application *Certificates*.

HTTPS connections have an unpredictable lifetime. Therefore, *Servers* ~~must~~ need to rely on the *AuthenticationToken* passed in the *RequestHeader* to determine the identity of the *Client*. This means the *AuthenticationToken* shall be a randomly generated value with at least 32 bytes of data and HTTPS with signing and encryption shall always be used.

HTTPS allows *Clients* to have certificates; however, they are not required by the HTTPS transport. A *Server* shall allow *Clients* to connect without ~~an HTTPS~~ providing a *Certificate* during negotiation of the HTTPS connection.

HTTP 1.1 supports *Message* chunking where the Content-Length header in the request response is set to "chunked" and each chunk is prefixed by its size in bytes. All applications that support the HTTPS transport shall support HTTP chunking.

The URL scheme for endpoints using the HTTPS transport is 'opc.https'. Note that 'https' is the generic URL scheme for the underlying transport. The opc prefix specifies that the endpoint accepts OPC UA messages as defined in 7.4.

7.4.2 Session-less Services

Session-less Services (see IEC 62541-4) may be invoked via HTTPS POST. The HTTP *Authorization* header in the *Request* shall have a Bearer token which is an *Access Token* provided by the *Authorization Service*. The Content-type of the HTTP request shall specify the encoding of the body. If the Content-type is application/opcua+uabinary then the body is encoded using the OPC UA Binary encoding (see 7.4.4). If the Content-type is application/opcua+uajson then body is encoded using the reversible form of the JSON encoding (see 7.4.5).

Note that the Content-type for OPC UA Binary encoded bodies for Session-less Services is different from the Content-type for Session-based Services specified in 7.4.4.

7.4.3 XML Encoding

This ~~*TransportProfile*~~ *TransportProtocol* implements the OPC UA *Services* using a SOAP request-response message pattern over an HTTPS connection.

~~The body of the HTTP Messages shall be a SOAP 1.2 Message (see SOAP Part 1). WS-Addressing headers are optional. The contents of the SOAP body and the SOAP action are the same as specified in 7.2.2 and 7.2.3.~~

The body of the HTTP *Messages* shall be a SOAP 1.2 *Message* (see SOAP Part 1). WS-Addressing headers are optional (see WS Addressing).

The OPC UA XML Encoding specifies a way to represent an OPC UA *Message* as an XML element. This element is added to the SOAP *Message* as the only child of the SOAP body element. If an error occurs in the *Server* while parsing the request body, the *Server* may return a SOAP fault or it may return an OPC UA error response.

The SOAP Action associated with an XML encoded request *Message* always has the form:

```
http://opcfoundation.org/UA/2008/02/Services.wsdl/<service name>
```

where <service name> is the name of the OPC UA *Service* being invoked.

The SOAP Action associated with an XML encoded response *Message* always has the form:

```
http://opcfoundation.org/UA/2008/02/Services.wsdl/<service name>Response
```

All requests shall be HTTP POST requests. The Content-type shall be "application/soap+xml" and the charset and action parameters shall be specified. The charset parameter shall be "utf-8" and the action parameter shall be the URI for the SOAP action.

An example HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Content-Type: application/soap+xml; charset="utf-8";
    action="http://opcfoundation.org/UA/2008/02/Services.wsdl/Read"
Content-Length: nnnn
```

The action parameter appears on the same line as the Content-Type declaration.

An example request *Message* (see 7.2.3):

```
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope">
  <s:Body>
    <ReadRequest xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
      ...
    </ReadRequest>
  </s:Body>
</s:Envelope>
```

An example HTTP response header is:

```
HTTP/1.1 200 OK
Content-Type: application/soap+xml; charset="utf-8";
  action="http://opcfoundation.org/UA/2008/02/Services.wsdl/ReadResponse"
Content-Length: nnnn
```

The action parameter appears on the same line as the Content-Type declaration.

An example response *Message*:

```
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope">
  <s:Body>
    <ReadResponse xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
      ...
    </ReadResponse>
  </s:Body>
</s:Envelope>
```

7.4.4 OPC UA Binary Encoding

This *TransportProfile* *TransportProtocol* implements the OPC UA *Services* using an OPC UA Binary Encoded *Messages* exchanged over an HTTPS connection.

Applications which support the *HTTPS Profile* shall support HTTP 1.1.

The body of the HTTP *Messages* shall be OPC UA Binary encoded blob. The Content-type shall be "application/octet-stream".

An example HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Content-Type: application/octet-stream;
Content-Length: nnnn
```

An example HTTP response header is:

```
HTTP/1.1 200 OK
Content-Type: application/octet-stream;
Content-Length: nnnn
```

The *Message* body is the request or response structure encoded as an *ExtensionObject* in OPC UA Binary. The *Authorization* header is only used for *Session-less Service* calls (see 7.4.2).

If the OPC UA Binary Encoding is used for a *Session-less Service* the HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Authorization: Bearer <base64-encoded-token-data>
Content-Type: application/opcu+uabinary;
```

```
Content-Length: nnnn
```

7.4.5 JSON Encoding

This *TransportProtocol* implements the OPC UA *Services* using JSON encoded *Messages* exchanged over an HTTPS connection.

Applications which support the HTTPS *Profile* shall support HTTP 1.1.

The body of the HTTP *Messages* shall be OPC UA JSON Encoded. The Content-type shall be "application/opcua+uajson".

An example HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Authorization: Bearer <base64-encoded-token-data>
Content-Type: application/opcua+uajson;
Content-Length: nnnn
```

An example HTTP response header is:

```
HTTP/1.1 200 OK
Content-Type: application/opcua+uajson;
Content-Length: nnnn
```

7.5 WebSockets

7.5.1 Overview

This *TransportProtocol* sends OPC UA Connection Protocol messages over WebSockets.

WebSockets is a bi-directional protocol for communication via a web server which is commonly used by browser based applications to allow the web server to asynchronously send information to the client. WebSockets uses the same default port as HTTP or HTTPS and initiates communication with an HTTP request. This makes it very useful in environments where firewalls limit traffic to the ports used by HTTP or HTTPS.

WebSockets uses HTTP, however, in practice a WebSocket connection is only initiated with a HTTP GET request and the web server provides an HTTP response. After that exchange, all traffic uses the binary framing protocol defined by RFC 6455.

A *Server* that supports the WebSockets transport shall publish one or more *Endpoints* with the scheme 'opc-wss'. The *TransportProfileUri* shall be one of the URIs for WebSockets transports defined in IEC 62541-7. The *TransportProfileUri* specifies the encoding and security protocol used to construct the OPC UA messages sent via the *WebSocket*.

The *SecurityMode* and *SecurityPolicyUri* of the *Endpoint* control the security applied to the messages sent via the *WebSocket*. This allows the messages to be secure even if the *WebSocket* connection is established via untrusted HTTPS proxies.

Figure 17 summarizes the complete process for establishing communication over a *WebSocket*.

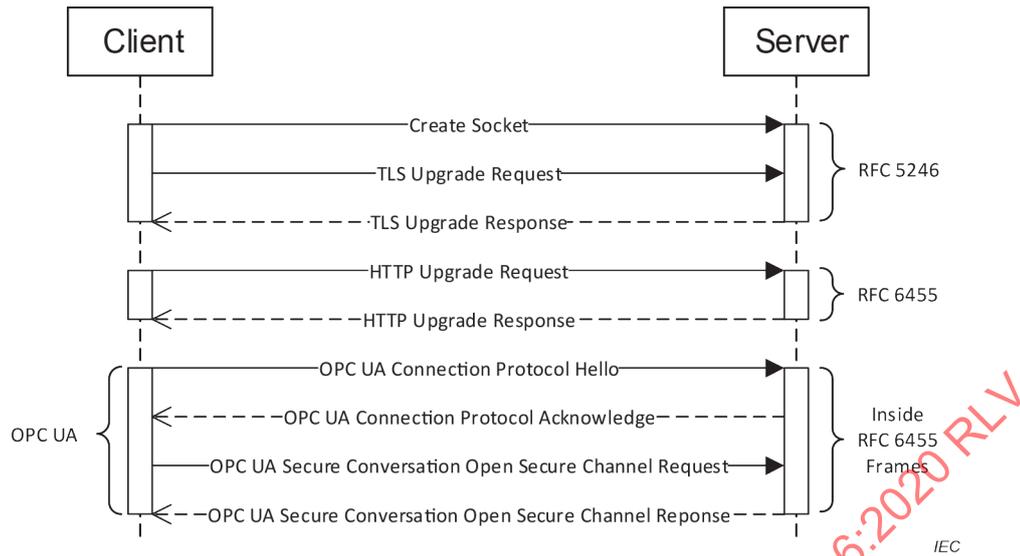


Figure 17 – Setting up Communication over a WebSocket

Figure 17 assumes the opcua+uacp protocol mapping (see 7.5.2).

7.5.2 Protocol Mapping

The WebSocket protocol allows clients to request that servers use specific sub-protocols with the "Sec-WebSocket-Protocol" header in the WebSocket handshake defined in RFC 6455. The protocols defined by this document are shown in Table 58.

Table 58 – WebSocket Protocols Mappings

Protocol	Description
opcua+uacp	Each WebSocket frame is a <i>MessageChunk</i> as defined in 6.7.2. After the WebSocket is created, the handshake described in 7.1.3 is used to negotiate the maximum size of the <i>MessageChunk</i> . The maximum size for a buffer needed to receive a WebSocket frame is the maximum length of a <i>MessageChunk</i> plus the maximum size for the WebSocket frame header. When using this protocol the payload in each frame is binary (OpCode 0x2 in RFC 6455).
opcua+uajson	Each WebSocket frame is a <i>Message</i> encoded using the JSON encoding described in 5.4.9. There is no mechanism to negotiate the maximum frame size. If the receiver encounters a frame that exceeds its internal limits, it shall close the WebSocket connection and provide a 1009 status code as described in RFC 6455. When using this protocol the payload in each frame is text (OpCode 0x1 in RFC 6455).

Each WebSocket protocol mapping defined has a *TransportProfileUri* defined in IEC 62541-7.

The *Client* shall request a protocol. If the *Server* does not support the protocol requested by the *Client*, the *Client* shall close the connection and report an error.

7.5.3 Security

WebSockets requires that the *Server* have a *Certificate*, however, the *Client* may have a *Certificate*. The *Server Certificate* should have the domain name as the common name component of the subject name; however, *Clients* that are able to override the *Certificate* validation procedure can choose to accept *Certificates* with a domain mismatch.

When using the WebSockets transport from a web browser, the browser environment may impose additional restrictions. For example, the web browser may require the *Server* have a valid TLS *Certificate* that is issued by CA that is installed in the *Trust List* for the web browser.

To support these *Clients*, a *Server* may use a *TLS Certificate* that does not conform to the requirements for an *ApplicationInstance Certificate*. In these cases, the *TLS Certificate* is only used for TLS negotiation and the *Server* shall use a valid *ApplicationInstance Certificate* for other interactions that require one. *Servers* shall allow administrators to specify a *Certificate* for use with TLS that is different from the *ApplicationInstance Certificate*.

Clients running in a browser environment specify the 'Origin' HTTP header during the *WebSocket* upgrade handshake. *Servers* should return the 'Access-Control-Allow-Origin' to indicate that the connection is allowed.

Any *Client* that does not run in a web browser environment and supports the *WebSockets* transport shall accept OPC UA *Application Instance Certificate* as the *TLS Certificate* provided the correct domain is specified in the *subjectAltName* field.

A *Client* may use its *Application Instance Certificate* as the *TLS Certificate* and *Servers* shall accept those *Certificates* if they are valid according to the OPC UA *Certificate* validation rules.

Some operating systems will not give the application any control over the set of algorithms that TLS will negotiate. In some cases, this set will be based on the needs of web browsers and will not be appropriate for the needs of an *OPC UA Application*. If this is a concern, applications should use OPC UA Secure Conversation in addition to TLS.

Clients that support the *WebSocket* transport shall support explicit configuration of an HTTPS proxy. When using an HTTPS proxy the *Client* shall first send an HTTP CONNECT message (see HTTP) before starting the *WebSocket* protocol handshake. Note that explicit HTTPS proxies allow for man-in-the-middle attacks. This threat may be mitigated by using OPC UA Secure Conversation in addition to TLS.

7.6 Well known addresses

The *Local Discovery Server* (LDS) is an OPC UA *Server* that implements the *Discovery Service Set* defined in IEC 62541-4. If an LDS is installed on a machine it shall use one or more of the well-known addresses defined in Table 59.

Table 59 – Well known addresses for Local Discovery Servers

Transport Mapping	URL	Notes
SOAP/HTTP	http://localhost/UADiscovery	May require integration with a web Server like IIS.
SOAP/HTTP	http://localhost:52601/UADiscovery	Alternate if it Port 80 cannot be used by the LDS.
OPC UA TCP	opc.tcp://localhost:4840/UADiscovery	
OPC UA WebSockets	opc.wss://localhost:443/UADiscovery	
OPC UA HTTPS	https://localhost:4843443/UADiscovery	

OPC UA applications that make use of the LDS shall allow administrators to change the well-known addresses used within a system.

The *Endpoint* used by *Servers* to register with the LDS shall be the base address with the path "/registration" appended to it (e.g. http://localhost/UADiscovery/registration). OPC UA *Servers* shall allow administrators to configure the address to use for registration.

Each OPC UA *Server* application implements the *Discovery Service Set*. If the OPC UA *Server* requires a different address for this *Endpoint*, it shall create the address by appending the path "/discovery" to its base address.

8 Normative Contracts

8.1 OPC Binary Schema

The normative contract for the OPC UA Binary Encoded *Messages* is an OPC Binary Schema. This file defines the structure of all types and *Messages*. The syntax for an OPC Binary Type Schema is described in IEC 62541-3. This schema captures normative names for types and their fields as well the order the fields appear when encoded. The data type of each field is also captured.

8.2 XML Schema and WSDL

The normative contract for the OPC UA XML encoded *Messages* is an XML Schema. This file defines the structure of all types and *Messages*. This schema captures normative names for types and their fields as well the order the fields appear when encoded. The data type of each field is also captured.

The normative contract for *Message* sent via the SOAP/HTTP *TransportProtocol* is a WSDL that includes XML Schema for the OPC UA XML encoded *Messages*. It also defines the port types for OPC UA *Servers* and *DiscoveryServers*.

Links to the WSDL and XML Schema files can be found in Annex D.

8.3 Information Model Schema

Annex F defines the schema to be used for Information Models.

8.4 Formal definition of UA Information Model

Annex B defines the OPC UA NodeSet.

8.5 Constants

Annex A defines constants for Attribute Ids, Status Codes and numeric NodeIds.

8.6 DataType encoding

Annex C defines the binary encoding for all DataTypes and Messages.

8.7 Security configuration

Annex E defines a schema for security settings.

Annex A (normative)

Constants

A.1 Attribute Ids

Table A.1 shows Identifiers assigned to Attributes.

Table A.1 – Identifiers assigned to Attributes

Attribute	Identifier
NodId	1
NodeClass	2
BrowseName	3
DisplayName	4
Description	5
WriteMask	6
UserWriteMask	7
IsAbstract	8
Symmetric	9
InverseName	10
ContainsNoLoops	11
EventNotifier	12
Value	13
DataType	14
ValueRank	15
ArrayDimensions	16
AccessLevel	17
UserAccessLevel	18
MinimumSamplingInterval	19
Historizing	20
Executable	21
UserExecutable	22
DataTypeDefinition	23
RolePermissions	24
UserRolePermissions	25
AccessRestrictions	26
AccessLevelEx	27

A.2 Status Codes

Clause A.2 defines the numeric identifiers for all of the StatusCodes defined by the OPC UA Specification. The identifiers are specified in a CSV file with the following syntax:

<SymbolName>, <Code>, <Description>

Where the *SymbolName* is the literal name for the error code that appears in the specification and the *Code* is the hexadecimal value for the *StatusCode* (see IEC 62541-4). The severity associated with a particular code is specified by the prefix (*Good*, *Uncertain* or *Bad*).

The CSV released with this version of the standards can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/StatusCode.csv>

NOTE The latest CSV that is compatible with this version of the standard can be found here:

<http://www.opcfoundation.org/UA/schemas/StatusCode.csv>

A.3 Numeric Node Ids

Clause A.3 defines the numeric identifiers for all of the numeric *NodeIds* defined by the OPC UA Specification. The identifiers are specified in a CSV file with the following syntax:

<SymbolName>, <Identifier>, <NodeClass>

Where the *SymbolName* is either the *BrowseName* of a *Type Node* or the *BrowsePath* for an *Instance Node* that appears in the specification and the *Identifier* is numeric value for the *NodeId*.

The *BrowsePath* for an instance *Node* is constructed by appending the *BrowseName* of the instance *Node* to *BrowseName* for the containing instance or type. A '_' character is used to separate each *BrowseName* in the path. For example, IEC 62541-5 defines the *ServerType ObjectType Node* which has the *NamespaceArray Property*. The *SymbolName* for the *NamespaceArray InstanceDeclaration* within the *ServerType* declaration is: *ServerType_NamespaceArray*. IEC 62541-5 also defines a standard instance of the *ServerType ObjectType* with the *BrowseName* 'Server'. The *BrowseName* for the *NamespaceArray Property* of the standard *Server Object* is: *Server_NamespaceArray*.

The *NamespaceUri* for all *NodeIds* defined here is <http://opcfoundation.org/UA/>

The CSV released with this version of the standards can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/NodeIds.csv>

NOTE The latest CSV that is compatible with this version of the standard can be found here:

<http://www.opcfoundation.org/UA/schemas/NodeIds.csv>

Annex B (normative)

OPC UA Nodeset

The OPC UA NodeSet includes the complete Information Model defined in this document. It follows the XML Information Model schema syntax defined in Annex F and can thus be read and processed by a computer program.

The Information Model Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/Opc.Ua.NodeSet2.xml>

NOTE The latest Information Model schema that is compatible with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/Opc.Ua.NodeSet2.xml>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Annex C (normative)

Type declarations for the OPC UA native Mapping

Annex C defines the OPC UA Binary encoding for all *DataTypes* and *Messages* defined in this document. The schema used to describe the type is defined in IEC 62541-3.

The OPC UA Binary Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/Opc.Ua.Types.bsd.xml>

NOTE The latest file that is compatible with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/Opc.Ua.Types.bsd.xml>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Annex D (normative)

WSDL for the XML Mapping

D.1 XML Schema

Clause D.1 defines the XML Schema for all DataTypes and Messages defined in this series of OPC UA standards.

The XML Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/Opc.Ua.Types.xsd>

NOTE The latest file that is compatible with this document can be found here:

<http://www.opcfoundation.org/UA/2008/02/Types.xsd>

D.2 WSDL Port Types

Clause D.2 defines the WSDL Operations and Port Types for all Services defined in IEC 62541-4.

The WSDL released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/Opc.Ua.Services.wsdl>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2008/02/Services.wsdl>

This WSDL imports the XML Schema defined in D.1.

D.3 WSDL Bindings

Clause D.3 defines the WSDL Bindings for all Services defined in IEC 62541-4.

The WSDL released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/Opc.Ua.Endpoints.wsdl>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2008/02/Endpoints.wsdl>

This WSDL imports the WSDL defined in D.2.

Annex E (normative)

Security settings management

E.1 Overview

All OPC UA applications shall support security; however, this requirement means that Administrators need to configure the security settings for the OPC UA application. Annex E describes an XML Schema which can be used to read and update the security settings for an OPC UA application. All OPC UA applications may support configuration by importing/exporting documents that conform to the schema (called the *SecuredApplication* schema) defined in Annex E.

The XML Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/SecuredApplication.xsd>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2011/03/SecuredApplication.xsd>

The *SecuredApplication* schema can be supported in two ways:

- 1) Providing an XML configuration file that can be edited directly;
- 2) Providing an import/export utility that can be run as required;

If the application supports direct editing of an XML configuration file, then that file shall have exactly one element with the local name 'SecuredApplication' and URI equal to the *SecuredApplication* schema URI. A third party configuration utility shall be able to parse the XML file, read and update the 'SecuredApplication' element. The administrator shall ensure that only authorized administrators can update this file. The following is an example of a configuration that can be directly edited:

```
<s1:SampleConfiguration xmlns:s1="http://acme.com/UA/Sample/Configuration.xsd">
  <ApplicationName>ACME UA Server</ApplicationName>
  <ApplicationUri>urn:myfactory.com:Machine54:ACME UA Server</ApplicationUri>

  <!-- any number of application specific elements -->

  <SecuredApplication xmlns="http://opcfoundation.org/UA/2011/03/SecuredApplication.xsd">
    <ApplicationName>ACME UA Server</ApplicationName>
    <ApplicationUri>urn:myfactory.com:Machine54:ACME UA Server</ApplicationUri>
    <ApplicationType>Server_0</ApplicationType>
    <ApplicationCertificate>
      <StoreType>Windows</StoreType>
      <StorePath>LocalMachine\My</StorePath>
      <SubjectName>ACME UA Server</SubjectName>
    </ApplicationCertificate>
  </SecuredApplication>

  <!-- any number of application specific elements -->

  <DisableHiResClock>true</DisableHiResClock>
</s1:SampleConfiguration>
```

If an application provides an import/export utility, then the import/export file shall be a document that conforms to the *SecuredApplication* schema. The administrator shall ensure that only authorized administrators can run the utility. The following is an example of a file used by an import/export utility:

```

<?xml version="1.0" encoding="utf-8" ?>
<SecuredApplication xmlns="http://opcfoundation.org/UA/2011/03/SecuredApplication.xsd">
  <ApplicationName>ACME UA Server</ApplicationName>
  <ApplicationUri>urn:myfactory.com:Machine54:ACME UA Server</ApplicationUri>
  <ApplicationType>Server_0</ApplicationType>
  <ConfigurationMode>urn:acme.com:ACME Configuration Tool</ConfigurationMode>
  <LastExportTime>2011-03-04T13:34:12Z</LastExportTime>
  <ExecutableFile>%ProgramFiles%\ACME\Bin\ACME UA Server.exe</ExecutableFile>
  <ApplicationCertificate>
    <StoreType>Windows</StoreType>
    <StorePath>LocalMachine\My</StorePath>
    <SubjectName>ACME UA Server</SubjectName>
  </ApplicationCertificate>
  <TrustedCertificateStore>
    <StoreType>Windows</StoreType>
    <StorePath>LocalMachine\UA applications</StorePath>
    <!-- Offline CRL Checks by Default -->
    <ValidationOptions>16</ValidationOptions>
  </TrustedCertificateStore>
  <TrustedCertificates>
    <Certificates>
      <CertificateIdentifier>
        <SubjectName>CN=MyFactory CA</SubjectName>
        <!-- Online CRL Check for this CA -->
        <ValidationOptions>32</ValidationOptions>
      </CertificateIdentifier>
    </Certificates>
  </TrustedCertificates>
  <RejectedCertificatesStore>
    <StoreType>Directory</StoreType>
    <StorePath>%CommonApplicationData%\OPC Foundation\RejectedCertificates</StorePath>
  </RejectedCertificatesStore>
</SecuredApplication>

```

E.2 SecuredApplication

The *SecuredApplication* element specifies the security settings for an application. The elements contained in a *SecuredApplication* are described in Table E.1.

When an instance of a *SecuredApplication* is imported into an application the application updates its configuration based on the information contained within it. If unrecoverable errors occur during import an application shall not make any changes to its configuration and report the reason for the error.

The mechanism used to import or export the configuration depends on the application. Applications shall ensure that only authorized users are able to access this feature.

The *SecuredApplication* element may reference X.509 v3 Certificates which are contained in physical stores. Each application needs to decide whether it uses shared physical stores which the administrator can control directly by changing the location or private stores that can only be accessed via the import/export utility. If the application uses private stores, then the contents of these private stores shall be copied to the export file during export. If the import file references shared physical stores, then the import/export utility shall copy the contents of those stores to the private stores.

The import/export utility shall not export private keys. If the administrator wishes to assign a new public-private key to the application the administrator shall place the private in a store where it can be accessed by the import/export utility. The import/export utility is then responsible for ensuring it is securely moved to a location where the application can access it.

Table E.1 – SecuredApplication

Element	Type	Description
ApplicationName	String	A human readable name for the application. Applications shall allow this value to be read or changed.
ApplicationUri	String	A globally unique identifier for the instance of the application. Applications shall allow this value to be read or changed.
ApplicationType	ApplicationType	The type of application. May be one of <ul style="list-style-type: none"> • Server_0; • Client_1; • ClientAndServer_2; • DiscoveryServer_3; Application shall provide this value. Applications do not allow this value to be changed.
ProductName	String	A name for the product. Application shall provide this value. Applications do not allow this value to be changed.
ConfigurationMode	String	Indicates how the application should be configured. An empty or missing value indicates that the configuration file can be edited directly. The location of the configuration file shall be provided in this case. Any other value is a URI that identifies the configuration utility. The vendor documentation shall explain how to use this utility. Application shall provide this value. Applications do not allow this value to be changed.
LastExportTime	UtcTime	When the configuration was exported by the import/export utility. It may be omitted if applications allow direct editing of the security configuration.
ConfigurationFile	String	The full path to a configuration file used by the application. Applications do not provide this value if an import/export utility is used. Applications do not allow this value to be changed. Permissions set on this file shall control who has rights to change the configuration of the application.
ExecutableFile	String	The full path to an executable file for the application. Applications may not provide this value. Applications do not allow this value to be changed. Permissions set on this file shall control who has rights to launch the application.
ApplicationCertificate	CertificateIdentifier	The identifier for the <i>Application Instance Certificate</i> . Applications shall allow this value to be read or changed. This identifier may reference a <i>Certificate</i> store that contains the private key. If the private key is not accessible to outside applications this value shall contain the X.509 v3 <i>Certificate</i> for the application. If the configuration utility assigns a new private key this value shall reference the store where the private key is placed. The import/export utility may delete this private key if it moves it to a secure location accessible to the application. Applications shall allow Administrators to enter the password required to access the private key during the import operation. The exact mechanism depends on the application. Applications shall report an error if the ApplicationCertificate is not valid.

Element	Type	Description
TrustedCertificateStore	CertificateStore Identifier	<p>The location of the CertificateStore containing the Certificates of applications or <i>Certificate</i> Authorities (CAs) which can be trusted. Applications shall allow this value to be read or changed.</p> <p>This value shall be a reference to a physical store which can be managed separately from the application. Applications that support shared physical stores shall check this store for changes whenever they validate a <i>Certificate</i>.</p> <p>The Administrator is responsible for verifying the signature on all Certificates placed in this store. This means the application may trust Certificates in this store even if they cannot be verified back to a trusted root.</p> <p>Administrators shall place any CA certificates used to verify the signature in the UntrustedIssuerStore IssuerStore or the UntrustedIssuerList IssuerList. This will allow applications to properly verify the signatures.</p> <p>The application shall check the revocation status of the Certificates in this store if the <i>Certificate</i> was issued by a CA. The application shall look for the offline <i>Certificate</i> Revocation List (CRL) for a CA in the store where it found the CA <i>Certificate</i>.</p> <p>The location of an online CRL for CA shall be specified with the CRLDistributionPoints (OID= 2.5.29.31) X.509 v3 <i>Certificate</i> extension.</p> <p>The ValidationOptions parameter is used to specify which revocation list should be used for CAs in this store.</p>
TrustedCertificates	CertificateList	<p>A list of Certificates for applications for CAs that can be trusted. Applications shall allow this value to be read or changed.</p> <p>The value is an explicit list of Certificates which is private to the application. It is used when the application does not support shared physical <i>Certificate</i> stores or when Administrators need to specify ValidationOptions for individual Certificates.</p> <p>If the TrustedCertificateStore and the TrustedCertificates parameters are both specified, then the application shall use the TrustedCertificateStore for checking trust relationships. The TrustedCertificates parameter is only used to lookup ValidationOptions for individual Certificates. It may also be used to provide CRLs for CA certificates.</p> <p>If the TrustedCertificateStore is not specified, then TrustedCertificates parameter shall contain the complete X.509 v3 <i>Certificate</i> for each entry.</p>
IssuerStore	CertificateStore Identifier	<p>The location of the CertificateStore containing CA Certificates which are not trusted but are needed to check signatures on Certificates. Applications shall allow this value to be read or changed.</p> <p>This value shall be a reference to a physical store which can be managed separately from the application. Applications that support shared physical stores shall check this store for changes whenever they validate a <i>Certificate</i>.</p> <p>This store may also contain CRLs for the CAs.</p>
IssuerCertificates	CertificateList	<p>A list of Certificates for CAs which are not trusted but are needed to check signatures on Certificates. Applications shall allow this value to be read or changed.</p> <p>The value is an explicit list of Certificates which is private to the application. It is used when the application does not support shared physical <i>Certificate</i> stores or when Administrators need to specify ValidationOptions for individual Certificates.</p> <p>If the IssuerStore and the IssuerCertificates parameters are both specified, then the application shall use the IssuerStore for checking signatures. The IssuerCertificates parameter is only used to lookup ValidationOptions for individual Certificates. It may also be used to provide CRLs for CA certificates.</p>

Element	Type	Description
RejectedCertificatesStore	CertificateStore Identifier	<p>The location of the shared CertificateStore containing the Certificates of applications which were rejected.</p> <p>Applications shall allow this value to be read or changed.</p> <p>Applications shall add the DER encoded <i>Certificate</i> into this store whenever it rejects a <i>Certificate</i> because it is untrusted or if it failed one of the validation rules which can be suppressed (see Clause E.6).</p> <p>Applications shall not add a <i>Certificate</i> to this store if it was rejected for a reason that cannot be suppressed (e.g. <i>Certificate</i> revoked).</p>
BaseAddresses	String []	<p>A list of URLs for the <i>Endpoints</i> supported by a <i>Server</i>.</p> <p>Applications shall allow these values to be read or changed.</p> <p>If a <i>Server</i> does not support the scheme for a URL it shall ignore it.</p> <p>This list can have multiple entries for the same URL scheme. The first entry for a scheme is the base URL. The rest are assumed to be DNS aliases that point to the first URL.</p> <p>It is the responsibility of the Administrator to configure the network to route these aliases correctly.</p>
SecurityProfileUris	<p>SecurityProfile</p> <p>ProfileUri String</p> <p>Enabled Boolean</p> <p>[]</p>	<p>A list of security profiles <i>SecurityPolicyUris</i> supported by a <i>Server</i>. The URIs are defined as security <i>Profiles</i> in IEC 62541-7.</p> <p>Applications shall allow these values to be read or changed.</p> <p>Applications shall allow the <i>Enabled</i> flag to be changed for each <i>SecurityProfile</i> that it supports.</p> <p>If the <i>Enabled</i> flag is false, the <i>Server</i> shall not allow connections using the <i>SecurityProfile</i>.</p> <p>If a <i>Server</i> does not support a <i>SecurityProfile</i> it shall ignore it.</p>
Extensions	xs:any	<p>A list of vendor defined Extensions attached to the security settings.</p> <p>Applications shall ignore Extensions that they do not recognize.</p> <p>Applications that update a file containing Extensions shall not delete or modify extensions that they do not recognize.</p>

E.3 CertificateIdentifier

The *CertificateIdentifier* element describes an X.509 v3 *Certificate*. The *Certificate* can be provided explicitly within the element or the element can specify the location of the *CertificateStore* that contains the *Certificate*. The elements contained in a *CertificateIdentifier* are described in Table E.2.

Table E.2 – CertificateIdentifier

Element	Type	Description
StoreType	String	The type of CertificateStore that contains the <i>Certificate</i> . Predefined values are "Windows" and "Directory". If not specified, the RawData element shall be specified.
StorePath	String	The path to the CertificateStore. The syntax depends on the StoreType. If not specified, the RawData element shall be specified.
SubjectName	String	The SubjectName for the <i>Certificate</i> . The Common Name (CN) component of the SubjectName. The SubjectName represented as a string that complies with Section 3 of RFC 4514. Values that do not contain '=' characters are presumed to be the Common Name component.
Thumbprint	String	The SHA1 thumbprint <i>CertificateDigest</i> for the <i>Certificate</i> formatted as a hexadecimal string. Case is not significant.
RawData	ByteString	The DER encoded <i>Certificate</i> . The CertificateIdentifier is invalid if the information in the DER <i>Certificate</i> conflicts with the information specified in other fields. Import utilities shall reject configurations containing invalid <i>Certificates</i> . This field shall not be specified if the StoreType and StorePath are specified.
ValidationOptions	Int32	The options to use when validating the <i>Certificate</i> . The possible options are described in E.6.
OfflineRevocationList	ByteString	A <i>Certificate</i> Revocation List (CRL) associated with an Issuer <i>Certificate</i> . The format of a CRL is defined by RFC 3280. This field is only meaningful for Issuer <i>Certificates</i> .
OnlineRevocationList	String	A URL for an Online Revocation List associated with an Issuer <i>Certificate</i> . This field is only meaningful for Issuer <i>Certificates</i> .

A "Windows" StoreType specifies a Windows *Certificate* store.

The syntax of the StorePath has the form:

[\\HostName]StoreLocation[(ServiceName | UserSid)]\StoreName

where:

HostName – the name of the machine where the store resides.

StoreLocation – one of LocalMachine, CurrentUser, User or Service

ServiceName – the name of a Windows Service.

UserSid – the SID for a Windows user account.

StoreName – the name of the store (e.g. My, Root, Trust, CA, etc.).

Examples of Windows StorePaths are:

\\MYPC\LocalMachine\My

\CurrentUser\Trust

\\MYPC\Service\My UA Server\UA applications

\User\S-1-5-25\Root

A "Directory" StoreType specifies a directory on disk which contains files with DER encoded Certificates. The name of the file is the ~~SHA1 thumbprint~~ *CertificateDigest* for the *Certificate*. Only public keys may be placed in a "Directory" Store. The StorePath is an absolute file system path with a syntax that depends on the operating system.

If a "Directory" store contains a 'certs' subdirectory, then it is presumed to be a structured store with the subdirectories described in Table E.3.

Table E.3 – Structured directory store

Subdirectory	Description
certs	Contains the DER encoded X.509 v3 Certificates. The files shall have a .der file extension.
private	Contains the private keys. The format of the file may be application specific. PEM encoded files should have a .pem extension. PKCS#12 encoded files should have a .pfx extension. The root file name shall be the same as the corresponding public key file in the certs directory.
crl	Contains the DER encoded CRL for any CA Certificates found in the certs or ca directories. The files shall have a .crl file extension.

Each *Certificate* is uniquely identified by its Thumbprint. The SubjectName or the distinguished SubjectName may be used to identify a *Certificate* to a human; however, they are not unique. The SubjectName may be specified in conjunction with the Thumbprint or the RawData. If there is an inconsistency between the information provided, then the *CertificateIdentifier* is invalid. Invalid *CertificateIdentifiers* are handled differently depending on where they are used.

It is recommended that the SubjectName always be specified.

A *Certificate* revocation list (CRL) contains a list of certificates issued by a CA that are no longer trusted. These lists should be checked before an application can trust a *Certificate* issued by a trusted CA. The format of a CRL is defined by RFC 3280.

Offline CRLs are placed in a local *Certificate* store with the Issuer *Certificate*. Online CRLs may exist but the protocol depends on the system. An online CRL is identified by a URL.

E.4 CertificateStoreIdentifier

The *CertificateStoreIdentifier* element describes a physical store containing X.509 v3 Certificates. The elements contained in a *CertificateStoreIdentifier* are described in Table E.4.

Table E.4 – CertificateStoreIdentifier

Element	Type	Description
StoreType	String	The type of CertificateStore that contains the <i>Certificate</i> . Predefined values are "Windows" and "Directory".
StorePath	String	The path to the CertificateStore. The syntax depends on the StoreType. See E.3 for a description of the syntax for different StoreTypes.
ValidationOptions	Int32	The options to use when validating the Certificates contained in the store. The possible options are described in E.6.

All *Certificates* are placed in a physical store which can be protected from unauthorized access. The implementation of a store can vary and will depend on the application, development tool or operating system. A *Certificate* store may be shared by many applications on the same machine.

Each *Certificate* store is identified by a *StoreType* and a *StorePath*. The same path on different machines identifies a different store.

E.5 CertificateList

The *CertificateList* element is a list of *Certificates*. The elements contained in a *CertificateList* are described in Table E.5.

Table E.5 – CertificateList

Element	Type	Description
Certificates	CertificateIdentifier []	The list of Certificates contained in the Trust List
ValidationOptions	Int32	The options to use when validating the Certificates contained in the store. These options only apply to <i>Certificates</i> that have <i>ValidationOptions</i> with the <i>UseDefaultOptions</i> bit set. The possible options are described in E.6.

E.6 CertificateValidationOptions

The *CertificateValidationOptions* control the process used to validate a *Certificate*. Any *Certificate* can have validation options associated. If none are specified, the *ValidationOptions* for the store or list containing the *Certificate* are used. The possible options are shown in Table E.6. Note that suppressing any validation step can create security risks which are discussed in more detail in IEC TR 62541-2. An audit log entry shall be created if any error is ignored because a validation option is suppressed.

Table E.6 – CertificateValidationOptions

Field	Bit	Description
SuppressCertificateExpired	0	Ignore errors related to the validity time of the <i>Certificate</i> or its issuers.
SuppressHostNameInvalid	1	Ignore mismatches between the host name or <i>ApplicationUri</i> .
SuppressRevocationStatusUnknown	2	Ignore errors if the issuer's revocation list cannot be found.
CheckRevocationStatusOnline	3	<p>Check the revocation status online.</p> <p>If set the validator will look for the URL of the CRL Distribution Point in the <i>Certificate</i> and use the OCSP (Online Certificate Status Protocol RFC 6960) to determine if the <i>Certificate</i> has been revoked.</p> <p>If the CRL Distribution Point is not reachable then the validator will look for offline CRLs if the <i>CheckRevocationStatusOffline</i> bit is set. Otherwise, validation fails.</p> <p>This option is specified for Issuer <i>Certificates</i> and used when validating <i>Certificates</i> issued by that Issuer.</p>
CheckRevocationStatusOffline	4	<p>Check the revocation status offline.</p> <p>If set the validator will look a CRL in the <i>Certificate Store</i> where the CA <i>Certificate</i> was found.</p> <p>Validation fails if a CRL is not found.</p> <p>This option is specified for Issuer <i>Certificates</i> and used when validating <i>Certificates</i> issued by that <i>Issuer</i>.</p>
UseDefaultOptions	5	<p>If set the <i>CertificateValidationOptions</i> from the <i>CertificateList</i> shall be used.</p> <p>If a <i>Certificate</i> does not belong to a <i>CertificateList</i> then the default is 0 for all bits.</p>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Annex F (normative)

Information Model XML Schema

F.1 Overview

Information Model developers define standard *AddressSpaces* which are implemented by many *Servers*. There is a need for a standard syntax that Information Model developers can use to formally define their models in a form that can be read by a computer program. Annex F defines an XML-based schema for this purpose.

The XML Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.0204/UANodeSet.xsd>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2011/03/UANodeSet.xsd>

The schema document is the formal definition. The description in Annex F only discusses details of the semantics that cannot be captured in the schema document. Types which are self-describing are not discussed.

This schema can also be used to serialize (i.e. import or export) an arbitrary set of *Nodes* in the *Server Address Space*. This serialized form can be used to save *Server* state for use by the *Server* later or to exchange with other applications (e.g. to support offline configuration by a *Client*).

This schema only defines a way to represent the structure of *Nodes*. It is not intended to represent the numerous semantic rules which are defined in other parts of IEC 62541. Consumers of data serialized with this schema need to handle inputs that conform to the schema, however, do not conform to the OPC UA specification because of one or more semantic rule violations.

The tables defining the *DataTypes* in the specification have field names starting with a lowercase letter. The first letter shall be converted to upper case when the field names are formally defined in a *UANodeSet*.

F.2 UANodeSet

The *UANodeSet* is the root of the document. It defines a set of *Nodes*, their *Attributes* and *References*. *References* to *Nodes* outside of the document are allowed.

The structure of a *UANodeSet* is shown in Table F.1.

Table F.1 – UANodeSet

Element	Type	Description
NamespaceUris	UriTable	A list of <i>NamespaceUris</i> used in the <i>UANodeSet</i> .
ServerUris	UriTable	A list of <i>ServerUris</i> used in the <i>UANodeSet</i> .
Models	ModelTableEntry []	A list of Models that are defined in the <i>UANodeSet</i> along with any dependencies these models have.
ModelUri	String	The URI for the model. This URI should be one of the entries in the <i>NamespaceUris</i> table.
Version	String	The version of the model defined in the <i>UANodeSet</i> . This is a human readable string and not intended for programmatic comparisons.
PublicationDate	DateTime	When the model was published. This value is used for comparisons if the Model is defined in multiple <i>UANodeSet</i> files.
RolePermissions	RolePermissions []	The list of default <i>RolePermissions</i> for all <i>Nodes</i> in the model.
AccessRestrictions	AccessRestrictions	The default <i>AccessRestrictions</i> that apply to all <i>Nodes</i> in the model.
RequiredModels	ModelTableEntry []	A list of dependencies for the model. If the model requires a minimum version the <i>PublicationDate</i> shall be specified. Tools which attempt to resolve these dependencies may accept any <i>PublicationDate</i> after this date.
Aliases	AliasTable	A list of <i>Aliases</i> used in the <i>UANodeSet</i> .
Extensions	xs:any	An element containing any vendor defined extensions to the <i>UANodeSet</i> .
LastModified	DateTime	The last time a document was modified.
<choice>	UObject UVariable UAMethod UAView UObjectType UVariableType UADatatype UReferenceType	The <i>Nodes</i> in the <i>UANodeSet</i> .

The *NamespaceUri* *NamespaceUris* is a list of URIs for namespaces used in the *UANodeSet*. The *NamespaceIndexes* used in *NodeId*, *ExpandedNodeIds* and *QualifiedNames* identify an element in this list. The first index is always 1 (0 is always the OPC UA namespace).

The *ServerUris* is a list of URIs for *Servers* referenced in the *UANodeSet*. The *ServerIndex* in *ExpandedNodeIds* identifies an element in this list. The first index is always 1 (0 is always the current *Server*).

The *Models* element specifies the *Models* which are formally defined by the *UANodeSet*. It includes version information as well as information about any dependencies which the model may have. If a *Model* is defined in the *UANodeSet* then the file shall also define an instance of the *NamespaceMetadataType ObjectType*. See IEC 62541-5 for more information.

The *Aliases* are a list of string substitutions for *NodeIds*. *Aliases* can be used to make the file more readable by allowing a string like 'HasProperty' in place of a numeric *NodeId* (i=46). *Aliases* are optional.

The *Extensions* are free form XML data that can be used to attach vendor defined data to the *UANodeSet*.

F.3 UANode

A *UANode* is an abstract base type for all *Nodes*. It defines the base set of *Attributes* and the *References*. There are subtypes for each *NodeClass* defined in IEC 62541-4. Each of these subtypes defines XML elements and attributes for the OPC UA *Attributes* specific to the *NodeClass*. The fields in the *UANode* type are defined in Table F.2.

Table F.2 – UANode

Element	Type	Description
Nodeld	Nodeld	A <i>Nodeld</i> serialized as a <i>String</i> . The syntax of the serialized <i>String</i> is defined in 5.3.1.10.
BrowseName	QualifiedName	A <i>QualifiedName</i> serialized as a <i>String</i> with the form: <namespace index>:<name> Where the <i>NamespaceIndex</i> refers to the <i>NamespaceUris</i> table.
SymbolicName	String	A symbolic name for the <i>Node</i> that can be used as a class/field name in auto generated code. It should only be specified if the <i>BrowseName</i> cannot be used for this purpose. This field does not appear in the <i>AddressSpace</i> and is intended for use by design tools. Only letters, digits or the underscore ('_') are permitted.
WriteMask	WriteMask	The value of the <i>WriteMask</i> Attribute.
UserWriteMask	WriteMask	The value of the <i>UserWriteMask</i> Attribute. Still in schema but no longer used.
AccessRestrictions	AccessRestrictions	The <i>AccessRestrictions</i> that apply to the <i>Node</i> .
DisplayName	LocalizedText []	A list of <i>DisplayNames</i> for the <i>Node</i> in different locales. There shall be only one entry per locale.
Description	LocalizedText []	The list of the <i>Descriptions</i> for the <i>Node</i> in different locales. There shall be only one entry per locale.
Category	String []	A list of identifiers used to group related <i>UANodes</i> together for use by tools that create/edit <i>UANodeSet</i> files.
Documentation	String	Additional non-localized documentation for use by tools that create/edit <i>UANodeSet</i> files.
References	Reference []	The list of <i>References</i> for the <i>Node</i> .
RolePermissions	RolePermissions []	The list of <i>RolePermissions</i> for the <i>Node</i> .
Extensions	xs:any	An element containing any vendor defined extensions to the <i>UANode</i> .

The *Extensions* are free form XML data that can be used to attach vendor defined data to the *UANode*.

Array values are denoted with [], however, in the XML Schema arrays are mapped to a complex type starting with the 'ListOf' prefix.

A *UANodeSet* is expected to contain many *UANodes* which reference each other. Tools that create *UANodeSets* should not add *Reference* elements for both directions in order to minimize the size of the XML file. Tools that read the *UANodeSets* shall automatically add reverse references unless reverse references are not appropriate given the *ReferenceType* semantics. *HasTypeDefinition* and *HasModellingRule* are two examples where it is not appropriate to add reverse references.

Note that a *UANodeSet* represents a collection of *Nodes* in an address space. This implies that any instances shall include the fully inherited *InstanceDeclarationHierarchy* as defined in IEC 62541-3.

F.4 Reference

The *Reference* type specifies a *Reference* for a *Node*. The *Reference* can be forward or inverse. Only one direction for each *Reference* needs to be in a *UANodeSet*. The other direction shall be added automatically during any import operation. The fields in the *Reference* type are defined in Table F.3.

Table F.3 – Reference

Element	Type	Description
Nodeld	Nodeld	The <i>Nodeld</i> of the target of the <i>Reference</i> serialized as a <i>String</i> . The syntax of the serialized <i>String</i> is defined in 5.3.1.11 (<i>ExpandedNodeld</i>). This value can be replaced by an <i>Alias</i> .
ReferenceType	Nodeld	The <i>Nodeld</i> of the <i>ReferenceType</i> serialized as a <i>String</i> . The syntax of the serialized <i>String</i> is defined in 5.3.1.10 (<i>Nodeld</i>). This value can be replaced by an <i>Alias</i> .
IsForward	Boolean	If TRUE, the <i>Reference</i> is a forward reference.

F.5 RolePermission

The *RolePermission* type specifies the *Permissions* granted to *Role* for a *Node*. The fields in the *RolePermission* type are defined in Table F.4.

Table F.4 – RolePermission

Element	Type	Description
Nodeld	Nodeld	The <i>Nodeld</i> of the <i>Role</i> which has the <i>Permissions</i> .
Permissions	UInt32	A bitmask specifying the <i>Permissions</i> granted to the <i>Role</i> . The bitmask values the <i>Permissions</i> bits defined in IEC 62541-3.

F.6 UAType

A *UAType* is a subtype of the *UANode* defined in F.3. It is the base type for the types defined in Table F.5.

Table F.5 – UANodeSet Type Nodes

Subtype	Description
UAObjectType	Defines an <i>ObjectType Node</i> as described in IEC 62541-3.
UAVariableType	Defines a <i>VariableType Node</i> as described in IEC 62541-3.
UADataType	Defines a <i>DataType Node</i> as described in IEC 62541-3.
UATypeReferenceType	Defines a <i>ReferenceType Node</i> as described in IEC 62541-3.

F.7 UAIInstance

A *UAIInstance* is a subtype of the *UANode* defined in F.3. It is the base type for the types defined in Table F.6. The fields in the *UAIInstance* type are defined in Table F.7. Subtypes of *UAIInstance* which have fields in addition to those defined in IEC 62541-3 are described in detail below.

Table F.6 – UANodeSet Instance Nodes

Subtype	Description
UAObject	Defines an <i>Object Node</i> as described in IEC 62541-3.
UAVariable	Defines a <i>Variable Node</i> as described in IEC 62541-3.
UAMethod	Defines a <i>Method Node</i> as described in IEC 62541-3.
UAView	Defines a <i>View Node</i> as described in IEC 62541-3.

Table F.7 – UAIInstance

Element	Type	Description
All of the fields from the <i>UANode</i> type described in F.3.		
ParentNodeId	NodeId	The <i>NodeId</i> of the <i>Node</i> that is the parent of the <i>Node</i> within the information model. This field is used to indicate that a tight coupling exists between the <i>Node</i> and its parent (e.g. when the parent is deleted the child is deleted as well). This information does not appear in the <i>AddressSpace</i> and is intended for use by design tools.

F.8 UAVariable

A *UAVariable* is a subtype of the *UAIInstance* defined in. It represents a Variable Node. The fields in the *UAVariable* type are defined in Table F.8.

Table F.8 – UAVariable

Element	Type	Description
All of the fields from the <i>UAInstance</i> type described in F.7.		
Value	Variant	The Value of the Node encoding using the UA XML wire encoding.
Translation	TranslationType []	A list of translations for the Value if the Value is a LocalizedText or a structure containing LocalizedTexts. This field may be omitted. If the Value is an array the number of elements in this array shall match the number of elements in the Value. Extra elements are ignored. If the Value is a scalar, then there is one element in this array. If the Value is a structure, then each element contains translations for one or more fields identified by a name. See the TranslationType for more information.
DataType	Nodeld	The data type of the value.
ValueRank	ValueRank	The value rank. If not specified, the default value is -1 (Scalar).
ArrayDimensions	ArrayDimensions	The number of dimensions in an array value.
AccessLevel	AccessLevel	The access level.
UserAccessLevel	AccessLevel	The access level for the current user. Still in schema but no longer used.
MinimumSamplingInterval	Duration	The minimum sampling interval.
Historizing	Boolean	Whether history is being archived.

F.9 UAMethod

A *UAMethod* is a subtype of the *UAInstance* defined in F.7. It represents a Method Node. The fields in the *UAMethod* type are defined in Table F.9.

Table F.9 – UAMethod

Element	Type	Description
All of the fields from the <i>UAInstance</i> type described in F.7.		
MethodDeclarationId	Nodeld	May be specified for <i>Method Nodes</i> that are a target of a <i>HasComponent</i> reference from a single <i>Object Node</i> . It is the <i>Nodeld</i> of the <i>UAMethod</i> with the same <i>BrowseName</i> contained in the <i>TypeDefinition</i> associated with the <i>Object Node</i> . If the <i>TypeDefinition</i> overrides a <i>Method</i> inherited from a base <i>ObjectType</i> then this attribute shall reference the <i>Method Node</i> in the subtype.
UserExecutable	Boolean	Still in schema but no longer used.
ArgumentDescription	UAMethodArgument []	A list of <i>Descriptions</i> for the <i>Method Node Arguments</i> . Each entry has a Name which uniquely identifies the <i>Argument</i> that the <i>Descriptions</i> apply to. There shall only be one entry per Name. Each entry also has a list of <i>Descriptions</i> for the <i>Argument</i> in different locales. There shall be only one entry per locale per <i>Argument</i> .

F.10 TranslationType

A *TranslationType* contains additional translations for *LocalizedTexts* used in the *Value* of a *Variable*. The fields in the *TranslationType* are defined in Table F.10. If multiple *Arguments* existed there would be a Translation element for each *Argument*.

The type can have two forms depending on whether the *Value* is a *LocalizedText* or a *Structure* containing *LocalizedTexts*. If it is a *LocalizedText* it contains a simple list of translations. If it is a *Structure*, it contains a list of fields which each contain a list of translations. Each field is identified by a Name which is unique within the structure. The mapping between the Name and the *Structure* requires an understanding of the *Structure* encoding. If the *Structure* field is encoded as a *LocalizedText* with UA XML, then the name is the unqualified path to the XML element where names in the path are separated by '/'. For example, a structure with a nested structure containing a *LocalizedText* could have a path like "Server/ApplicationName".

The following example illustrates how translations for the Description field in the *Argument Structure* are represented in XML:

```
<Value>
  <ListOfExtensionObject xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
    <ExtensionObject>
      <TypeId>
        <Identifier>i=297</Identifier>
      </TypeId>
      <Body>
        <Argument>
          <Name>ConfigData</Name>
          <DataType>
            <Identifier>i=15</Identifier>
          </DataType>
          <ValueRank>-1</ValueRank>
          <ArrayDimensions />
          <Description>
            <Text>[English Translation for Description]</Text>
          </Description>
        </Argument>
      </Body>
    </ExtensionObject>
  </ListOfExtensionObject>
</Value>
<Translation>
  <Field Name="Description">
    <Text Locale="de-DE">[German Translation for Description]</Text>
    <Text Locale="fr-FR">[French Translation for Description]</Text>
  </Field>
</Translation>
```

If multiple *Arguments* existed there would be a Translation element for each *Argument*.

Table F.10 – TranslationType

Element	Type	Description
Text	LocalizedText []	An array of translations for the Value. It only appears if the <i>Value</i> is a <i>LocalizedText</i> or an array of <i>LocalizedText</i> .
Field	StructureTranslationType []	An array of structure fields which have translations. It only appears if the <i>Value</i> is a <i>Structure</i> or an array of <i>Structures</i> .
Name	String	The name of the field. This uniquely identifies the field within the structure. The exact mapping depends on the encoding of the structure.
Text	LocalizedText []	An array of translations for the structure field.

F.11 UADatatype

A *UADatatype* is a subtype of the *UAType* defined in F.6. It defines a *DataType Node*. The fields in the *UADatatype* type are defined in Table F.11.

Table F.11 – UADatatype

Element	Type	Description
All of the fields from the <i>UANode</i> type described in F.3.		
Definition	DataTypeDefinition	An abstract definition of the data type that can be used by design tools to create code that can serialize the data type in XML and/or Binary forms. It does not appear in the <i>AddressSpace</i> . This is only used to define subtypes of the <i>Structure</i> or <i>Enumeration DataTypes</i> .

F.12 DataTypeDefinition

A *DataTypeDefinition* defines an abstract representation of a *UADatatype* that can be used by design tools to automatically create serialization code. The fields in the *DataTypeDefinition* type are defined in Table F.12.

Table F.12 – DataTypeDefinition

Element	Type	Description
Name	QualifiedName	A unique name for the data type. This field is only specified for nested <i>DataTypeDefinitions</i> . The <i>BrowseName</i> of the <i>DataType Node</i> is used otherwise.
SymbolicName	String	A symbolic name for the data type that can be used as a class/structure name in autogenerated code. It should only be specified if the <i>Name</i> cannot be used for this purpose. Only letters, digits or the underscore ('_') are permitted. This field is only specified for nested <i>DataTypeDefinitions</i> . The <i>SymbolicName</i> of the <i>DataType Node</i> is used otherwise.
BaseType	QualifiedName	The name of any base type. Note that the <i>BaseType</i> can refer to types defined in other files. The <i>NamespaceUri</i> associated with the <i>Name</i> should indicate where to look for the <i>BaseType</i> definition. This field is only specified for nested <i>DataTypeDefinitions</i> . The <i>HasSubtype Reference</i> of the <i>DataType Node</i> is used otherwise.
Fields	DataTypeField[]	The list of fields that make up the data type. This definition assumes the structure has a sequential layout. For enumerations the fields are simply a list of values. Unions are not supported.

Element	Type	Description
Name	QualifiedName	A unique name for the data type. This name should be the same as the <i>BrowseName</i> or the containing <i>Data Type</i> .
SymbolicName	String	A symbolic name for the data type that can be used as a class/structure name in autogenerated code. It should only be specified if the <i>Name</i> cannot be used for this purpose. Only letters, digits or the underscore ('_') are permitted and the first character shall be a letter. This field is only specified for nested <i>Data Type Definitions</i> . The <i>SymbolicName</i> of the <i>Data Type Node</i> is used otherwise.
BaseType	QualifiedName	Not used. Kept in schema for backward compatibility.
IsUnion	Boolean	This flag indicates if the data type represents a union. Only one of the Fields defined for the data type is encoded into a value. This field is optional. The default value is false. If this value is true, the first field is the switch value.
IsOptionSet	Boolean	This flag indicates that the data type defines the <i>OptionSetValues Property</i> . This field is optional. The default value is false.
Fields	Data TypeField []	The list of fields that make up the data type. This definition assumes the structure has a sequential layout. For enumerations, the fields are simply a list of values. This list does not include fields inherited from a base data type.

F.13 Data TypeField

A *Data TypeField* defines an abstract representation of a field within a *UAData Type* that can be used by design tools to automatically create serialization code. The fields in the *Data TypeField* type are defined in Table F.13.

Table F.13 – Data TypeField

Element	Type	Description
Name	String	A name for the field that is unique within the <i>Data Type Definition</i> .
SymbolicName	String	A symbolic name for the field that can be used in autogenerated code. It should only be specified if the <i>Name</i> cannot be used for this purpose. Only letters, digits or the underscore ('_') are permitted.
Data Type	NodeId	The <i>NodeId</i> of the <i>Data Type</i> for the field. This <i>NodeId</i> can refer to another <i>Node</i> with its own <i>Data Type Definition</i> . This field is not specified for subtypes of <i>Enumeration</i> .
ValueRank	Int32	The value rank for the field. It shall be <i>Scalar</i> (-1) or a fixed rank <i>Array</i> (>=1). This field is not specified for subtypes of <i>Enumeration</i> .
Description	LocalizedText[]	A description for the field in multiple locales.
Definition	Data TypeDefinition	The field is a structure with a layout specified by the definition. This field is optional. This field allows designers to create nested structures without defining a new <i>Data Type Node</i> for each structure. This field is not specified for subtypes of <i>Enumeration</i> .
Value	Int32	The value associated with the field. This field is only specified for subtypes of <i>Enumeration</i> .

Element	Type	Description
Name	String	A name for the field that is unique within the <i>DataTypeDefinition</i> .
SymbolicName	String	A symbolic name for the field that can be used in autogenerated code. It should only be specified if the <i>Name</i> cannot be used for this purpose. Only letters, digits or the underscore ('_') are permitted.
DisplayName	LocalizedText []	A display name for the field in multiple locales.
DataType	NodeId	The <i>NodeId</i> of the <i>DataType</i> for the field. This <i>NodeId</i> can refer to another <i>Node</i> with its own <i>DataTypeDefinition</i> . This field is not specified for subtypes of <i>Enumeration</i> .
ValueRank	Int32	The value rank for the field. It shall be <i>Scalar</i> (-1) or a fixed rank <i>Array</i> (≥ 1). This field is not specified for subtypes of <i>Enumeration</i> .
ArrayDimensions	String	The maximum length of an array. This field is a comma separated list of unsigned integer values. The list has a number of elements equal to the <i>ValueRank</i> . The value is 0 if the maximum is not known for a dimension. This field is not specified if the <i>ValueRank</i> ≤ 0 . This field is not specified for subtypes of <i>Enumeration</i> or for <i>DataTypes</i> with the <i>OptionSetValues Property</i> .
MaxStringLength	UInt32	The maximum length of a <i>String</i> or <i>ByteString</i> value. If not known the value is 0. The value is 0 if the <i>DataType</i> is not <i>String</i> or <i>ByteString</i> . If the <i>ValueRank</i> > 0 the maximum applies to each element in the array. This field is not specified for subtypes of <i>Enumeration</i> or for <i>DataTypes</i> with the <i>OptionSetValues Property</i> .
Description	LocalizedText []	A description for the field in multiple locales.
Value	Int32	The value associated with the field. This field is only specified for subtypes of <i>Enumeration</i> and <i>OptionSet DataTypes</i> . For <i>OptionSets</i> the value is the number of the bit associated with the field.
IsOptional	Boolean	The field indicates if a data type field in a structure is optional. This field is optional. The default value is false. This field is not specified for subtypes of <i>Enumeration</i> and <i>Union</i> .

F.14 Variant

The *Variant* type specifies the value for a *Variable* or *VariableType Node*. This type is the same as the type defined in 5.3.1.17. As a result, the functions used to serialize *Variants* during *Service* calls can be used to serialize *Variant* in this file syntax.

Variants can contain *NodeIds*, *ExpandedNodeIds* and *QualifiedNames* which ~~must~~ shall be modified so the *NamespaceIndexes* and *ServerIndexes* reference the *NamespaceUri* and *ServerUri* tables in the *UANodeSet*.

Variants can also contain *ExtensionObjects* which contain an *EncodingId* and a *Structure* with fields which could be are *NodeIds*, *ExpandedNodeIds* or *QualifiedNames*. The *NamespaceIndexes* and *ServerIndexes* in these fields shall also reference the tables in the *UANodeSet*.

F.15 Example ~~(Informative)~~

An example of the *UANodeSet* can be found below.

This example defines the *Nodes* for an *InformationModel* with the URI of "http://sample.com/Instances". This example references *Nodes* defined in the base OPC UA *InformationModel* and an *InformationModel* with the URI "http://sample.com/Types".

The XML namespaces declared at the top include the URIs for the *Namespaces* referenced in the document because the document includes *Complex Data*. Documents without *Complex Data* would not have these declarations.

```
<UANodeSet
xmlns:s1="http://sample.com/Instances"
xmlns:s0="http://sample.com/Types"
xmlns:uax="http://opcfoundation.org/UA/2008/02/Types.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns="http://opcfoundation.org/UA/2011/03/UANodeSet.xsd">
```

The *NamespaceUris* table includes all *Namespaces* referenced in the document except for the base OPC UA *InformationModel*. A *NamespaceIndex* of 1 refers to the URI "http://sample.com/Instances".

```
<NamespaceUris>
  <Uri>http://sample.com/Instances</Uri>
  <Uri>http://sample.com/Types</Uri>
</NamespaceUris>
```

The *Aliases* table is provided to enhance readability. There are no rules for what is included. A useful guideline would include standard *ReferenceTypes* and *DataTypes* if they are referenced in the document.

```
<Aliases>
  <Alias Alias="HasComponent">i=47</Alias>
  <Alias Alias="HasProperty">i=46</Alias>
  <Alias Alias="HasSubtype">i=45</Alias>
  <Alias Alias="HasTypeDefinition">i=40</Alias>
</Aliases>
```

The *BicycleType* is a *DataType Node* that inherits from a *DataType* defined in another *InformationModel* (ns=2;i=314). It is assumed that any application importing this file will already know about the referenced *InformationModel*. A *Server* could map the references onto another OPC UA *Server* by adding a *ServerIndex* to *TargetNode NodeIds*. The structure of the *DataType* is defined by the *Definition* element. This information can be used by code generators to automatically create serializers for the *DataType*.

```
<UADatatype NodeId="ns=1;i=365" BrowseName="1:BicycleType">
  <DisplayName>BicycleType</DisplayName>
  <References>
    <Reference ReferenceType="HasSubtype" IsForward="false">ns=2;i=314</Reference>
  </References>
  <Definition Name="BicycleType" BaseType="0:1:BicycleType">
    <Field Name="NoOfGears" DataType="UInt32" />
    <Field Name="ManufacturerName" DataType="QualifiedName" />
  </Definition>
</UADatatype>
```

This *Node* is an instance of an *Object TypeDefinition Node* defined in another *InformationModel* (ns=2;i=341). It has a single *Property* which is declared later in the document.

```
<UAObject NodeId="ns=1;i=375" BrowseName="1:DriverOfTheMonth" ParentNodeId="ns=1;i=281">
```

```

<DisplayName>DriverOfTheMonth</DisplayName>
<References>
  <Reference ReferenceType="HasProperty">ns=1;i=376</Reference>
  <Reference ReferenceType="HasTypeDefinition">ns=2;i=341</Reference>
  <Reference ReferenceType="HasComponent" IsForward="false">ns=1;i=281</Reference>
</References>
</UAObject>

```

This Node is an instance of a *Variable TypeDefinition Node* defined in base OPC UA *InformationModel* (i=68). The *DataType* is the base type for the *BicycleType DataType*. The *AccessLevels* declare the *Variable* as *Readable* and *Writeable*. The *ParentNodeId* indicates that this Node is tightly coupled with the Parent (*DriverOfTheMonth*) and will be deleted if the Parent is deleted.

```

<UAVariable NodeId="ns=1;i=376" BrowseName="2:PrimaryVehicle"
  ParentNodeId="ns=1;i=375" DataType="ns=2;i=314" AccessLevel="3" UserAccessLevel="3">
  <DisplayName>PrimaryVehicle</DisplayName>
  <References>
    <Reference ReferenceType="HasTypeDefinition">i=68</Reference>
    <Reference ReferenceType="HasProperty" IsForward="false">ns=1;i=375</Reference>
  </References>

```

This *Value* is an instance of a *BicycleType DataType*. It is wrapped in an *ExtensionObject* which declares that the value is serialized using the *Default XML DataTypeEncoding* for the *DataType*. The *Value* could be serialized using the *Default Binary DataTypeEncoding* but that would result in a document that cannot be edited by hand. No matter which *DataTypeEncoding* is used, the *NamespaceIndex* used in the *ManufacturerName* field refers to the *NamespaceUris* table in this document. The application is responsible for changing whatever value it needs to be when the document is loaded by an application.

```

<Value>
  <ExtensionObject xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
    <TypeId>
      <Identifier>ns=1;i=366</Identifier>
    </TypeId>
    <Body>
      <s1:BicycleType>
        <s0:Make>Trek</s0:Make>
        <s0:Model>Compact</s0:Model>
        <s1:NoOfGears>10</s1:NoOfGears>
        <s1:ManufacturerNameManufacturerName>
          <uax:NamespaceIndex>1</uax:NamespaceIndex>
          <uax:Name>Hello</uax:Name>
        </s1:ManufacturerNameManufacturerName>
      </s1:BicycleType>
    </Body>
  </ExtensionObject>
</Value>
</UAVariable>

```

These are the *DataTypeEncoding Nodes* for the *BicycleType DataType*.

```

<UAObject NodeId="ns=1;i=366" BrowseName="Default XML">
  <DisplayName>Default XML</DisplayName>
  <References>
    <Reference ReferenceType="HasEncoding" IsForward="false">ns=1;i=365</Reference>
    <Reference ReferenceType="HasDescription">ns=1;i=367</Reference>
    <Reference ReferenceType="HasTypeDefinition">i=76</Reference>
  </References>
</UAObject>
<UAObject NodeId="ns=1;i=370" BrowseName="Default Binary">
  <DisplayName>Default Binary</DisplayName>
  <References>
    <Reference ReferenceType="HasEncoding" IsForward="false">ns=1;i=365</Reference>
    <Reference ReferenceType="HasDescription">ns=1;i=371</Reference>
    <Reference ReferenceType="HasTypeDefinition">i=76</Reference>
  </References>
</UAObject>

```

This is the *DataTypeDescription Node* for the *Default XML DataTypeEncoding* of the *BicycleType DataType*. The *Value* is one of the built-in types.

```
<UAVariable NodeId="ns=1;i=367" BrowseName="1:BicycleType" DataType="String">
  <DisplayName>BicycleType</DisplayName>
  <References>
    <Reference ReferenceType="HasTypeDefinition">i=69</Reference>
    <Reference ReferenceType="HasComponent" IsForward="false">ns=1;i=341</Reference>
  </References>
  <Value>
    <uax:String>//xs:element[@name='BicycleType']</uax:String>
  </Value>
</UAVariable>
```

This is the ~~*DataTypeDescription*~~ *DataTypeDictionary Node* for the *DataTypeDescription* declared above. The XML Schema document is a UTF-8 document stored as a ~~base64 string~~ *xs:base64Binary* value (see Base64). This allows *Clients* to read the schema for all *DataTypes* which belong to the *DataTypeDictionary*. The value of *DataTypeDescription Node* for each *DataType* contains a XPath query that will find the correct definition inside the schema document.

```
<UAVariable NodeId="ns=1;i=341" BrowseName="1:Quickstarts.DataTypes.Instances"
DataType="ByteString">
  <DisplayName>Quickstarts.DataTypes.Instances</DisplayName>
  <References>
    <Reference ReferenceType="HasProperty">ns=1;i=343</Reference>
    <Reference ReferenceType="HasComponent">ns=1;i=367</Reference>
    <Reference ReferenceType="HasComponent" IsForward="false">i=92</Reference>
    <Reference ReferenceType="HasTypeDefinition">i=72</Reference>
  </References>
  <Value>
    <uax:ByteString>PHhz...WlhPg==</uax:ByteString>
  </Value>
</UAVariable>
```

F.16 UANodeSetChanges

The *UANodeSetChanges* is the root of a document that contains a set of changes to an *AddressSpace*. It is expected that a single file will contain either a *UANodeSet* or a *UANodeSetChanges* element at the root. It provides a list of *Nodes/References* to add and/or a list *Nodes/References* to delete. The *UANodeSetChangesStatus* structure defined in F.22 is produced when a *UANodeSetChanges* document is applied to an *AddressSpace*.

The elements of the type are defined in Table F.14.

Table F.14 – UANodeSetChanges

Element	Type	Description
NamespaceUris	UriTable	Same as described in Table F.1.
ServerUris	UriTable	Same as described in Table F.1.
Models	ModelTableEntry []	Same as described in Table F.1.
Aliases	AliasTable	Same as described in Table F.1.
Extensions	xs:any	Same as described in Table F.1.
LastModified	DateTime	Same as described in Table F.1.
NodesToAdd	NodesToAdd	A list of new <i>Nodes</i> to add to the <i>AddressSpace</i> .
ReferencesToAdd	ReferencesToChange	A list of new <i>References</i> to add to the <i>AddressSpace</i> .
NodesToDelete	NodesToDelete	A list of <i>Nodes</i> to delete from the <i>AddressSpace</i> .
ReferencesToDelete	ReferencesToChange	A list of <i>References</i> to delete from the <i>AddressSpace</i> .

The *Models* element specifies the version of one or more *Models* which the *UANodeSetChanges* file will create when it is applied to an existing Address Space. The *UANodeSetChanges* cannot be applied if the current version of the *Model* in the Address Space is higher. The *RequiredModels* sub-element (see Table F.1) specifies the versions of *Models* which shall already exist before the *UANodeSetChanges* file can be applied. When checking dependencies, the version of the *Model* in the existing Address Space shall exactly match the required version.

If a *UANodeSetChanges* file modifies types and there are existing instances of the types in the AddressSpace, then the *Server* shall automatically modify the instances to conform to the new type or generate an error.

A *UANodeSetChanges* file is processed as a single operation. This allows mandatory *Nodes* or *References* to be replaced by specifying a *Node/Reference* to delete and a *Node/Reference* to add.

F.17 NodesToAdd

The *NodesToAdd* type specifies a list of *Nodes* to add to an *AddressSpace*. The structure of these *Nodes* is defined by the *UANodeSet* type in Table F.1.

The elements of the type are defined in Table F.15.

Table F.15 – NodesToAdd

Element	Type	Description
<choice>	UAObject UAVariable UAMethod UAView UAObjectType UAVariableType UADatatype UAReferenceType	The <i>Nodes</i> to add to the <i>AddressSpace</i> .

When adding *Nodes*, *References* can be specified as part of the *Node* definition or as a separate *ReferencesToAdd*.

Note that *References* to *Nodes* that could exist are always allowed. In other words, a *Node* is never rejected simply because it has a reference to an unknown *Node*.

Reverse *References* are added automatically when deemed practical by the processor.

F.18 ReferencesToChange

The *ReferencesToChange* type specifies a list of *References* to add to or remove from an *AddressSpace*.

The elements of the type are defined in Table F.16.

Table F.16 – ReferenceToChange

Element	Type	Description
Reference	ReferenceToChange	A <i>Reference</i> to add to the <i>AddressSpace</i> .

F.19 ReferenceToChange

The *ReferenceToChange* type specifies a single *Reference* to add to or remove from an *AddressSpace*.

The elements of the type are defined in Table F.17.

Table F.17 – ReferenceToChange

Element	Type	Description
Source	Nodeld	The identifier for the source <i>Node</i> of the <i>Reference</i> .
ReferenceType	Nodeld	The identifier for the type of the <i>Reference</i> .
IsForward	Boolean	TRUE if the <i>Reference</i> is a forward reference.
Target	Nodeld	The identifier for the target <i>Node</i> of the <i>Reference</i> .

References to *Nodes* that could exist are always allowed. In other words, a *Reference* is never rejected simply because the target is unknown *Node*.

The source of the *Reference* shall exist in the *AddressSpace* or in the *UANodeSetChanges* document being processed.

Reverse *References* are added when deemed practical by the processor.

F.20 NodesToDelete

The *NodesToDelete* type specifies a list of *Nodes* to remove from an *AddressSpace*.

The elements of the type are defined in Table F.18.

Table F.18 – NodesToDelete

Element	Type	Description
Node	NodeToDelete	A <i>Node</i> to delete from the <i>AddressSpace</i> .

F.21 NodeToDelete

The *NodeToDelete* type specifies a *Node* to remove from an *AddressSpace*.

The elements of the type are defined in Table F.19.

Table F.19 – ReferencesToChange

Element	Type	Description
Node	NodeId	The identifier for the <i>Node</i> to delete.
DeleteReverseReferences	Boolean	If TRUE, then <i>References</i> to the <i>Node</i> are deleted as well.

F.22 UANodeSetChangesStatus

The *UANodeSetChangesStatus* is the root of a document that is produced when a *UANodeSetChanges* document is processed.

The elements of the type are defined in Table F.20.

Table F.20 – UANodeSetChangesStatus

Element	Type	Description
NamespaceUri	UriTable	Same as described in Table F.1.
ServerUri	UriTable	Same as described in Table F.1.
Aliases	AliasTable	Same as described in Table F.1.
Extensions	xs:any	Same as described in Table F.1.
Version	String	Same as described in Table F.1.
LastModified	DateTime	Same as described in Table F.1.
TransactionId	String	A globally unique identifier from the original <i>UANodeSetChanges</i> document.
NodesToAdd	NodeSetStatusList	A list of results for the <i>NodesToAdd</i> specified in the original document. The list is empty if all elements were processed successfully.
ReferencesToAdd	NodeSetStatusList	A list of results for the <i>ReferencesToAdd</i> specified in the original document. The list is empty if all elements were processed successfully.
NodesToDelete	NodeSetStatusList	A list of results for the <i>NodesToDelete</i> specified in the original document. The list is empty if all elements were processed successfully.
ReferencesToDelete	NodeSetStatusList	A list of results for the <i>ReferencesToDelete</i> specified in the original document. The list is empty if all elements were processed successfully.

F.23 NodeSetStatusList

The *NodeSetStatusList* type specifies a list of results produced when applying a *UANodeSetChanges* document to an *AddressSpace*.

If no errors occurred this list is empty.

If one or more errors occur, then this list contains one element for each operation specified in the original document.

The elements of the type are defined in Table F.21.

Table F.21 – NodeSetStatusList

Element	Type	Description
Result	NodeSetStatus	The result of a single operation.

F.24 NodeSetStatus

The *NodeSetStatus* type specifies a single result produced when applying an operation specified in a *UANodeSetChanges* document to an *AddressSpace*.

The elements of the type are defined in Table F.22.

Table F.22 – NodeSetStatus

Element	Type	Description
Code	StatusCode	The result of the operation. The possible StatusCodes are defined in IEC 62541-4.
Details	String	A string providing information that is not conveyed by the StatusCode. This is not a human readable string for the StatusCode.

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Bibliography

X200: ISO/IEC 7498-1 (ITU-T Rec. X.200), *Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model*

WS Addressing: Web Services Addressing (WS-Addressing)
<http://www.w3.org/Submission/ws-addressing/>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

INTERNATIONAL STANDARD

NORME INTERNATIONALE



**OPC unified architecture –
Part 6: Mappings**

**Architecture unifiée OPC –
Partie 6: Mappings**

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

CONTENTS

FOREWORD	8
1 Scope	11
2 Normative references	11
3 Terms, definitions, abbreviated terms and symbols.....	13
3.1 Terms and definitions.....	13
3.2 Abbreviated terms and symbols	14
4 Overview	14
5 Data encoding	16
5.1 General.....	16
5.1.1 Overview	16
5.1.2 Built-in Types	16
5.1.3 Guid	17
5.1.4 ByteString.....	17
5.1.5 ExtensionObject	17
5.1.6 Variant.....	18
5.1.7 Decimal	18
5.2 OPC UA Binary	19
5.2.1 General	19
5.2.2 Built-in Types	19
5.2.3 Decimal	30
5.2.4 Enumerations	30
5.2.5 Arrays.....	30
5.2.6 Structures.....	31
5.2.7 Structures with optional fields	33
5.2.8 Unions	35
5.2.9 Messages	36
5.3 OPC UA XML.....	37
5.3.1 Built-in Types	37
5.3.2 Decimal	43
5.3.3 Enumerations	43
5.3.4 Arrays.....	44
5.3.5 Structures.....	44
5.3.6 Structures with optional fields	45
5.3.7 Unions	45
5.3.8 Messages	46
5.4 OPC UA JSON.....	46
5.4.1 General	46
5.4.2 Built-in Types	46
5.4.3 Decimal	52
5.4.4 Enumerations	52
5.4.5 Arrays.....	52
5.4.6 Structures.....	53
5.4.7 Structures with optional fields	53
5.4.8 Unions	54
5.4.9 Messages	54
6 Message SecurityProtocols	55

6.1	Security handshake	55
6.2	Certificates	56
6.2.1	General	56
6.2.2	Application Instance Certificate.....	57
6.2.3	Certificate Chains	58
6.3	Time synchronization	58
6.4	UTC and International Atomic Time (TAI).....	58
6.5	Issued User Identity Tokens.....	58
6.5.1	Kerberos.....	58
6.5.2	JSON Web Token (JWT).....	59
6.5.3	OAuth2	60
6.6	WS Secure Conversation	62
6.7	OPC UA Secure Conversation	62
6.7.1	Overview	62
6.7.2	MessageChunk structure	62
6.7.3	MessageChunks and error handling.....	67
6.7.4	Establishing a SecureChannel.....	67
6.7.5	Deriving keys.....	69
6.7.6	Verifying Message security	70
7	TransportProtocols	71
7.1	OPC UA Connection Protocol.....	71
7.1.1	Overview	71
7.1.2	Message structure	72
7.1.3	Establishing a connection	75
7.1.4	Closing a connection	77
7.1.5	Error handling.....	77
7.2	OPC UA TCP	79
7.3	SOAP/HTTP.....	79
7.4	OPC UA HTTPS.....	79
7.4.1	Overview	79
7.4.2	Session-less Services.....	81
7.4.3	XML Encoding	81
7.4.4	OPC UA Binary Encoding	82
7.4.5	JSON Encoding	82
7.5	WebSockets.....	83
7.5.1	Overview	83
7.5.2	Protocol Mapping.....	84
7.5.3	Security	84
7.6	Well known addresses	85
8	Normative Contracts	86
8.1	OPC Binary Schema	86
8.2	XML Schema and WSDL.....	86
8.3	Information Model Schema.....	86
8.4	Formal definition of UA Information Model.....	86
8.5	Constants	86
8.6	DataType encoding.....	86
8.7	Security configuration	86
Annex A (normative)	Constants.....	87
A.1	Attribute Ids	87

A.2	Status Codes	87
A.3	Numeric Node Ids	88
Annex B (normative)	OPC UA Nodeset	89
Annex C (normative)	Type declarations for the OPC UA native Mapping	90
Annex D (normative)	WSDL for the XML Mapping	91
D.1	XML Schema	91
D.2	WDSL Port Types	91
D.3	WSDL Bindings	91
Annex E (normative)	Security settings management	92
E.1	Overview	92
E.2	SecuredApplication	93
E.3	CertificateIdentifier	96
E.4	CertificateStoreIdentifier	98
E.5	CertificateList	99
E.6	CertificateValidationOptions	99
Annex F (normative)	Information Model XML Schema	101
F.1	Overview	101
F.2	UANodeSet	101
F.3	UANode	103
F.4	Reference	104
F.5	RolePermission	104
F.6	UAType	104
F.7	UAInstance	105
F.8	UAVariable	105
F.9	UAMethod	106
F.10	TranslationType	106
F.11	UADatatype	107
F.12	DataTypeDefinition	108
F.13	DataTypeField	108
F.14	Variant	109
F.15	Example	110
F.16	UANodeSetChanges	112
F.17	NodesToAdd	113
F.18	ReferencesToChange	113
F.19	ReferenceToChange	114
F.20	NodesToDelete	114
F.21	NodeToDelete	114
F.22	UANodeSetChangesStatus	115
F.23	NodeSetStatusList	115
F.24	NodeSetStatus	115
Bibliography	117
Figure 1	– The OPC UA Stack Overview	15
Figure 2	– Encoding Integers in a binary stream	20
Figure 3	– Encoding Floating Points in a binary stream	20
Figure 4	– Encoding Strings in a binary stream	21
Figure 5	– Encoding Guids in a binary stream	22

Figure 6 – Encoding XmlElement in a binary stream	22
Figure 7 – A String NodeId.....	23
Figure 8 – A Two Byte NodeId	24
Figure 9 – A Four Byte NodeId.....	24
Figure 10 – Security handshake.....	55
Figure 11 – OPC UA Secure Conversation MessageChunk.....	63
Figure 12 – OPC UA Connection Protocol Message structure	72
Figure 13 – Client initiated OPC UA Connection Protocol connection.....	76
Figure 14 – Server initiated OPC UA Connection Protocol connection.....	76
Figure 15 – Closing a OPC UA Connection Protocol connection	77
Figure 16 – Scenarios for the HTTPS Transport.....	80
Figure 17 – Setting up Communication over a WebSocket	84
Table 1 – Built-in Data Types.....	16
Table 2 – Guid structure	17
Table 3 – Layout of Decimal	19
Table 4 – Supported Floating Point Types.....	20
Table 5 – NodeId components	22
Table 6 – NodeId DataEncoding values	23
Table 7 – Standard NodeId Binary DataEncoding.....	23
Table 8 – Two Byte NodeId Binary DataEncoding.....	24
Table 9 – Four Byte NodeId Binary DataEncoding.....	24
Table 10 – ExpandedNodeId Binary DataEncoding	25
Table 11 – DiagnosticInfo Binary DataEncoding.....	26
Table 12 – QualifiedName Binary DataEncoding	26
Table 13 – LocalizedText Binary DataEncoding	27
Table 14 – Extension Object Binary DataEncoding.....	28
Table 15 – Variant Binary DataEncoding.....	29
Table 16 – Data Value Binary DataEncoding.....	30
Table 17 – Sample OPC UA Binary Encoded structure.....	32
Table 18 – Sample OPC UA Binary Encoded Structure with optional fields	34
Table 19 – Sample OPC UA Binary Encoded Structure	35
Table 20 – XML Data Type Mappings for Integers.....	37
Table 21 – XML Data Type Mappings for Floating Points	37
Table 22 – Components of NodeId	39
Table 23 – Components of ExpandedNodeId	40
Table 24 – Components of Enumeration	44
Table 25 – JSON Object Definition for a NodeId	48
Table 26 – JSON Object Definition for an ExpandedNodeId	49
Table 27 – JSON Object Definition for a StatusCode	49
Table 28 – JSON Object Definition for a DiagnosticInfo	50
Table 29 – JSON Object Definition for a QualifiedName.....	50
Table 30 – JSON Object Definition for a LocalizedText.....	50

Table 31 – JSON Object Definition for an ExtensionObject	51
Table 32 – JSON Object Definition for a Variant	51
Table 33 – JSON Object Definition for a DataValue	52
Table 34 – JSON Object Definition for a Decimal	52
Table 35 – JSON Object Definition for a <i>Structure</i> with Optional Fields	53
Table 36 – JSON Object Definition for a Union	54
Table 37 – SecurityPolicy	56
Table 38 – Application Instance Certificate	57
Table 39 – Kerberos UserTokenPolicy	59
Table 40 – JWT UserTokenPolicy	59
Table 41 – JWT IssuerEndpointUrl Definition	60
Table 42 – Access Token Claims	61
Table 43 – OPC UA Secure Conversation Message header	63
Table 44 – Asymmetric algorithm Security header	64
Table 45 – Symmetric algorithm Security header	65
Table 46 – Sequence header	65
Table 47 – OPC UA Secure Conversation Message footer	66
Table 48 – OPC UA Secure Conversation Message abort body	67
Table 49 – OPC UA Secure Conversation OpenSecureChannel Service	68
Table 50 – PRF inputs for RSA based SecurityPolicies	70
Table 51 – Cryptography key generation parameters	70
Table 52 – OPC UA Connection Protocol Message header	72
Table 53 – OPC UA Connection Protocol Hello Message	73
Table 54 – OPC UA Connection Protocol Acknowledge Message	74
Table 55 – OPC UA Connection Protocol Error Message	74
Table 56 – OPC UA Connection Protocol ReverseHello Message	75
Table 57 – OPC UA Connection Protocol error codes	78
Table 58 – WebSocket Protocols Mappings	84
Table 59 – Well known addresses for Local Discovery Servers	85
Table A.1 – Identifiers assigned to Attributes	87
Table E.1 – SecuredApplication	94
Table E.2 – CertificateIdentifier	97
Table E.3 – Structured directory store	98
Table E.4 – CertificateStoreIdentifier	99
Table E.5 – CertificateList	99
Table E.6 – CertificateValidationOptions	100
Table F.1 – UANodeSet	102
Table F.2 – UANode	103
Table F.3 – Reference	104
Table F.4 – RolePermission	104
Table F.5 – UANodeSet Type Nodes	104
Table F.6 – UANodeSet Instance Nodes	105
Table F.7 – UAInstance	105

Table F.8 – UVariable..... 106

Table F.9 – UAMethod..... 106

Table F.10 – TranslationType..... 107

Table F.11 – UADatatype..... 108

Table F.12 – DataTypeDefinition..... 108

Table F.13 – DataTypeField..... 109

Table F.14 – UANodeSetChanges..... 112

Table F.15 – NodesToAdd..... 113

Table F.16 – ReferencesToChange..... 113

Table F.17 – ReferencesToChange..... 114

Table F.18 – NodesToDelete..... 114

Table F.19 – ReferencesToChange..... 114

Table F.20 – UANodeSetChangesStatus..... 115

Table F.21 – NodeSetStatusList..... 115

Table F.22 – NodeSetStatus..... 116

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

INTERNATIONAL ELECTROTECHNICAL COMMISSION

OPC UNIFIED ARCHITECTURE –

Part 6: Mappings

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 62541-6 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) Encodings:

- added JSON encoding for PubSub (non-reversible);
- added JSON encoding for Client/Server (reversible);
- added support for optional fields in structures;
- added support for Unions.

- b) Transport mappings:
- added WebSocket secure connection – WSS;
 - added support for reverse connectivity;
 - added support for session-less service invocation in HTTPS.
- c) Deprecated Transport (missing support on most platforms):
- SOAP/HTTP with WS-SecureConversation (all encodings).
- d) Added mapping for JSON Web Token.
- e) Added support for Unions to NodeSet Schema.
- f) Added batch operations to add/delete nodes to/from NodeSet Schema.
- g) Added support for multi-dimensional arrays outside of Variants.
- h) Added binary representation for Decimal data types.
- i) Added mapping for an OAuth2 Authorization Framework.

The text of this International Standard is based on the following documents:

FDIS	Report on voting
65E/718/FDIS	65E/734/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 IEC 62541, 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 62541-6:2020

OPC UNIFIED ARCHITECTURE –

Part 6: Mappings

1 Scope

This part of IEC 62541 specifies the OPC Unified Architecture (OPC UA) mapping between the security model described in IEC TR 62541-2, the abstract service definitions specified in IEC 62541-4, the data structures defined in IEC 62541-5 and the physical network protocols that can be used to implement the OPC UA specification.

2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. 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 62541-4, *OPC Unified Architecture – Part 4: Services*

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

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

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

ISO 8601-1:2019, *Date and time – Representations for information interchange – Part 1: Basic rules*

XML Schema Part 2: XML Schema Part 2: Datatypes
<http://www.w3.org/TR/xmlschema-2/>

SOAP Part 1: SOAP Version 1.2 Part 1: Messaging Framework
<http://www.w3.org/TR/soap12-part1/>

SSL/TLS: RFC 5246 – The TLS Protocol Version 1.2
<http://tools.ietf.org/html/rfc5246.txt>

X.509 v3: ISO/IEC 9594-8 (ITU-T Rec. X.509), *Information technology – Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks*

HTTP: RFC 2616 – Hypertext Transfer Protocol – HTTP/1.1
<http://www.ietf.org/rfc/rfc2616.txt>

HTTPS: RFC 2818 – HTTP Over TLS
<http://www.ietf.org/rfc/rfc2818.txt>

Base64: RFC 3548 – The Base16, Base32, and Base64 Data Encodings
<http://www.ietf.org/rfc/rfc3548.txt>

X690: ISO/IEC 8825-1 (ITU-T Rec. X.690), *Information technology – ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)*

IEEE-754: Standard for Floating-Point Arithmetic

HMAC: HMAC – Keyed-Hashing for Message Authentication
<http://www.ietf.org/rfc/rfc2104.txt>

PKCS #1: PKCS #1 – RSA Cryptography Specifications Version 2.0
<http://www.ietf.org/rfc/rfc2437.txt>

PKCS #12: PKCS #12 – Personal Information Exchange Syntax v1.1
<http://www.ietf.org/rfc/rfc7292.txt>

FIPS 180-4: Secure Hash Standard (SHS)
<https://csrc.nist.gov/publications/detail/fips/180/4/final>

FIPS 197: Advanced Encryption Standard (AES)
<https://csrc.nist.gov/publications/detail/fips/197/final>

UTF-8: UTF-8, a transformation format of ISO 10646
<http://www.ietf.org/rfc/rfc3629.txt>

RFC 3280: RFC 3280 – X.509 Public Key Infrastructure Certificate and CRL Profile
<http://www.ietf.org/rfc/rfc3280.txt>

RFC 4514: RFC 4514 – LDAP: String Representation of Distinguished Names
<http://www.ietf.org/rfc/rfc4514.txt>

NTP: RFC 1305 – Network Time Protocol (Version 3) Specification, Implementation and Analysis
<http://www.ietf.org/rfc/rfc1305.txt>

Kerberos: Web Services Security – Kerberos Token Profile 1.1
<http://docs.oasis-open.org/wss/v1.1/wss-v1.1-spec-os-KerberosTokenProfile.pdf>

RFC 1738: RFC 1738 – Uniform Resource Locators (URL)
<http://www.ietf.org/rfc/rfc1738.txt>

RFC 2141: RFC 2141 – URN Syntax
<http://www.ietf.org/rfc/rfc2141.txt>

RFC 6455: RFC 6455 – The WebSocket Protocol
<http://www.ietf.org/rfc/rfc6455.txt>

RFC 7159: The JavaScript Object Notation (JSON) Data Interchange Format
<http://www.ietf.org/rfc/rfc7159.txt>

RFC 7523: JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants
<https://tools.ietf.org/rfc/rfc7523.txt>

RFC 6749: The OAuth 2.0 Authorization Framework
<http://www.ietf.org/rfc/rfc6749.txt>

OpenID-Core: OpenID Connect Core 1.0
http://openid.net/specs/openid-connect-core-1_0.html

OpenID-Discovery: OpenID Connect Discovery 1.0
https://openid.net/specs/openid-connect-discovery-1_0.html

RFC 6960: RFC 6960 – X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP
<https://www.ietf.org/rfc/rfc6960.txt>

3 Terms, definitions, abbreviated terms and symbols

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

CertificateDigest

short identifier used to uniquely identify an X.509 v3 *Certificate*

Note 1 to entry: This is the SHA1 hash of DER encoded form of the *Certificate*.

3.1.2

DataEncoding

way to serialize OPC UA *Messages* and data structures

3.1.3

DevelopmentPlatform

suite of tools and/or programming languages used to create software

3.1.4

Mapping

specification on how to implement an OPC UA feature with a specific technology

Note 1 to entry: For example, the OPC UA Binary Encoding is a *Mapping* that specifies how to serialize OPC UA data structures as sequences of bytes.

3.1.5

SecurityProtocol

protocol which ensures the integrity and privacy of UA *Messages* that are exchanged between OPC UA applications

3.1.6**StackProfile**

combination of *DataEncodings*, *SecurityProtocol* and *TransportProtocol Mappings*

Note 1 to entry: OPC UA applications implement one or more *StackProfiles* and can only communicate with OPC UA applications that support a *StackProfile* that they support.

3.1.7**TransportConnection**

full-duplex communication link established between OPC UA applications

Note 1 to entry: A TCP/IP socket is an example of a *TransportConnection*.

3.1.8**TransportProtocol**

way to exchange serialized OPC UA *Messages* between OPC UA applications

3.2 Abbreviated terms and symbols

API	application programming interface
ASN.1	Abstract Syntax Notation #1 (used in X690)
CSV	comma separated value (file format)
ECC	elliptic curve cryptography
HTTP	Hypertext Transfer Protocol
HTTPS	Secure Hypertext Transfer Protocol
IPSec	Internet Protocol Security
OID	object identifier (used with ASN.1)
PRF	pseudo random function
RSA	Rivest, Shamir and Adleman [public key encryption system]
SHA1	secure hash algorithm
SOAP	Simple Object Access Protocol
SSL	Secure Sockets Layer (defined in SSL/TLS)
TCP	Transmission Control Protocol
TLS	Transport Layer Security (defined in SSL/TLS)
UA	Unified Architecture
UACP	OPC UA Connection Protocol
UASC	OPC UA Secure Conversation
WS-*	XML Web Services specifications
XML	eXtensible Markup Language

4 Overview

Other parts of the IEC 62541 series are written to be independent of the technology used for implementation. This approach means OPC UA is a flexible specification that will continue to be applicable as technology evolves. On the other hand, this approach means that it is not possible to build an OPC UA application with the information contained in IEC TR 62541-1 through to IEC 62541-5 because important implementation details have been left out.

This document defines *Mappings* between the abstract specifications and technologies that can be used to implement them. The *Mappings* are organized into three groups: *DataEncodings*, *SecurityProtocols* and *TransportProtocols*. Different *Mappings* are combined together to create *StackProfiles*. All OPC UA applications shall implement at least one

StackProfile and can only communicate with other OPC UA applications that implement the same *StackProfile*.

This document defines the *DataEncodings* in Clause 5, the *SecurityProtocols* in 5.4 and the *TransportProtocols* in 6.7.6. The *StackProfiles* are defined in IEC 62541-7.

All communication between OPC UA applications is based on the exchange of *Messages*. The parameters contained in the *Messages* are defined in IEC 62541-4; however, their format is specified by the *DataEncoding* and *TransportProtocol*. For this reason, each *Message* defined in IEC 62541-4 shall have a normative description which specifies exactly what shall be put on the wire. The normative descriptions are defined in the annexes.

A *Stack* is a collection of software libraries that implement one or more *StackProfiles*. The interface between an OPC UA application and the *Stack* is a non-normative API which hides the details of the *Stack* implementation. An API depends on a specific *DevelopmentPlatform*. Note that the datatypes exposed in the API for a *DevelopmentPlatform* may not match the datatypes defined by the specification because of limitations of the *DevelopmentPlatform*. For example, the Java programming language does not support an unsigned integer which means that any Java API will need to map unsigned integers onto a signed integer type.

Figure 1 illustrates the relationships between the different concepts defined in this document.

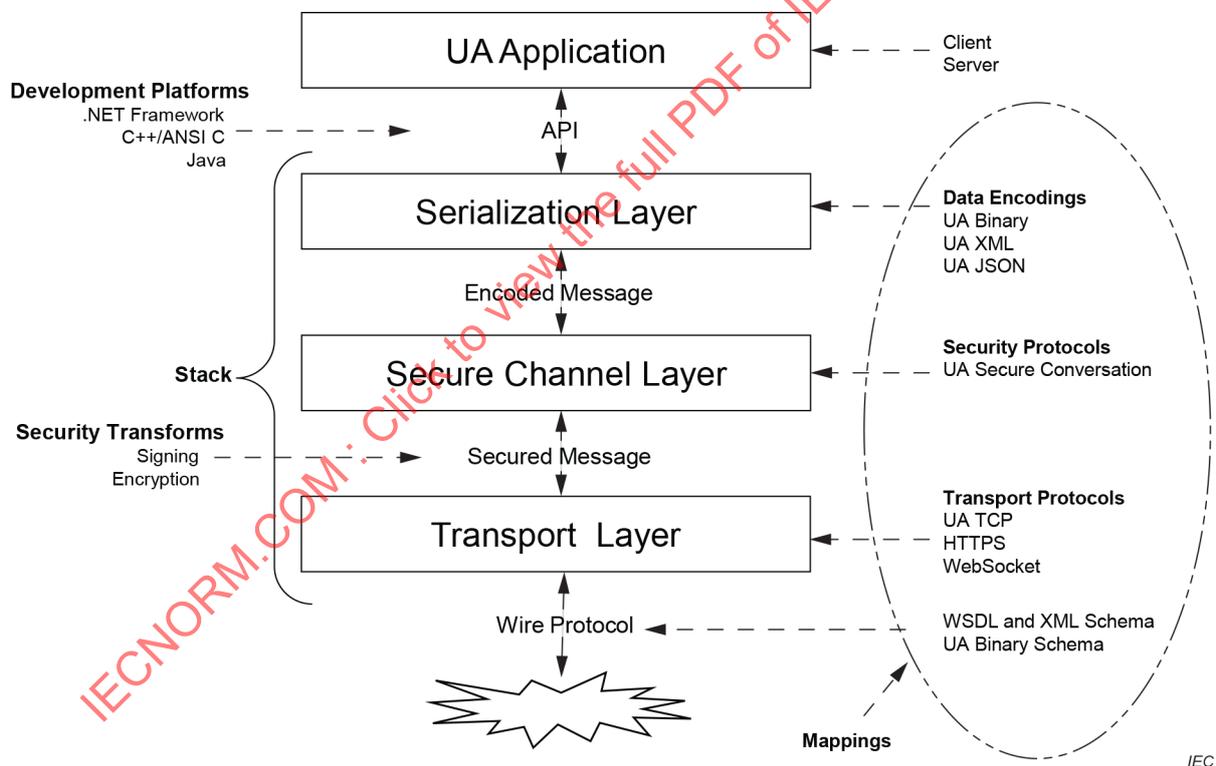


Figure 1 – The OPC UA Stack Overview

The layers described in this document do not correspond to layers in the OSI 7-layer model [X200]. Each OPC UA *StackProfile* should be treated as a single Layer 7 (application) protocol that is built on an existing Layer 5, 6 or 7 protocol such as TCP/IP, TLS or HTTP. The *SecureChannel* layer is always present even if the *SecurityMode* is *None*. In this situation, no security is applied but the *SecurityProtocol* implementation shall maintain a logical channel with a unique identifier. Users and administrators are expected to understand that a *SecureChannel* with *SecurityMode* set to *None* cannot be trusted unless the application is operating on a physically secure network or a low-level protocol such as IPSec is being used.

5 Data encoding

5.1 General

5.1.1 Overview

This document defines three data encodings: OPC UA Binary, OPC UA XML and OPC UA JSON. It describes how to construct *Messages* using each of these encodings.

5.1.2 Built-in Types

All OPC UA *DataEncodings* are based on rules that are defined for a standard set of built-in types. These built-in types are then used to construct structures, arrays and *Messages*. The built-in types are described in Table 1.

Table 1 – Built-in Data Types

ID	Name	Description
1	Boolean	A two-state logical value (true or false).
2	SByte	An integer value between –128 and 127 inclusive.
3	Byte	An integer value between 0 and 255 inclusive.
4	Int16	An integer value between –32 768 and 32 767 inclusive.
5	UInt16	An integer value between 0 and 65 535 inclusive.
6	Int32	An integer value between –2 147 483 648 and 2 147 483 647 inclusive.
7	UInt32	An integer value between 0 and 4 294 967 295 inclusive.
8	Int64	An integer value between –9 223 372 036 854 775 808 and 9 223 372 036 854 775 807 inclusive.
9	UInt64	An integer value between 0 and 18 446 744 073 709 551 615 inclusive.
10	Float	An IEEE single precision (32 bit) floating point value.
11	Double	An IEEE double precision (64 bit) floating point value.
12	String	A sequence of Unicode characters.
13	DateTime	An instance in time.
14	Guid	A 16-byte value that can be used as a globally unique identifier.
15	ByteString	A sequence of octets.
16	XmlElement	An XML element.
17	NodeId	An identifier for a node in the address space of an OPC UA <i>Server</i> .
18	ExpandedNodeId	A <i>NodeId</i> that allows the namespace URI to be specified instead of an index.
19	StatusCode	A numeric identifier for an error or condition that is associated with a value or an operation.
20	QualifiedName	A name qualified by a namespace.
21	LocalizedText	Human readable text with an optional locale identifier.
22	ExtensionObject	A structure that contains an application specific data type that may not be recognized by the receiver.
23	DataValue	A data value with an associated status code and timestamps.
24	Variant	A union of all of the types specified above.
25	DiagnosticInfo	A structure that contains detailed error and diagnostic information associated with a <i>StatusCode</i> .

Most of these data types are the same as the abstract types defined in IEC 62541-3 and IEC 62541-4. However, the *ExtensionObject* and *Variant* types are defined in this document. In addition, this document defines a representation for the *Guid* type defined in IEC 62541-3.

5.1.3 Guid

A *Guid* is a 16-byte globally unique identifier with the layout shown in Table 2.

Table 2 – Guid structure

Component	Data Type
Data1	UInt32
Data2	UInt16
Data3	UInt16
Data4	Byte [8]

Guid values may be represented as a string in this form:

```
<Data1>--<Data2>--<Data3>--<Data4[0:1]>--<Data4[2:7]>
```

where Data1 is 8 characters wide, Data2 and Data3 are 4 characters wide and each *Byte* in Data4 is 2 characters wide. Each value is formatted as a hexadecimal number with padded zeros. A typical *Guid* value would look like this when formatted as a string:

```
C496578A-0DFE-4B8F-870A-745238C6AEAE
```

5.1.4 ByteString

A *ByteString* is structurally the same as a one-dimensional array of *Byte*. It is represented as a distinct built-in data type because it allows encoders to optimize the transmission of the value. However, some *DevelopmentPlatforms* will not be able to preserve the distinction between a *ByteString* and a one-dimensional array of *Byte*.

If a decoder for *DevelopmentPlatform* cannot preserve the distinction it shall convert all one-dimensional arrays of *Byte* to *ByteStrings*.

Each element in a one-dimensional array of *ByteString* can have a different length which means is structurally different from a two-dimensional array of *Byte* where the length of each dimension is the same. This means decoders shall preserve the distinction between two or more dimension arrays of *Byte* and one or more dimension arrays of *ByteString*.

If a *DevelopmentPlatform* does not support unsigned integers, then it will need to represent *ByteStrings* as arrays of *SByte*. In this case, the requirements for *Byte* would then apply to *SByte*.

5.1.5 ExtensionObject

An *ExtensionObject* is a container for any *Structured DataTypes* which cannot be encoded as one of the other built-in data types. The *ExtensionObject* contains a complex value serialized as a sequence of bytes or as an XML element. It also contains an identifier which indicates what data it contains and how it is encoded.

Structured DataTypes are represented in a *Server* address space as sub-types of the *Structure DataType*. The *DataEncodings* available for any given *Structured DataTypes* are represented as a *DataTypeEncoding Object* in the *Server AddressSpace*. The *NodeId* for the *DataTypeEncoding Object* is the identifier stored in the *ExtensionObject*. IEC 62541-3 describes how *DataTypeEncoding Nodes* are related to other *Nodes* of the *AddressSpace*.

Server implementers should use namespace qualified numeric *NodeIds* for any *DataTypeEncoding Objects* they define. This will minimize the overhead introduced by packing *Structured DataType* values into an *ExtensionObject*.

ExtensionObjects and *Variants* allow unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

5.1.6 Variant

A *Variant* is a union of all built-in data types including an *ExtensionObject*. *Variants* can also contain arrays of any of these built-in types. *Variants* are used to store any value or parameter with a data type of *BaseDataType* or one of its subtypes.

Variants can be empty. An empty *Variant* is described as having a null value and should be treated like a null column in a SQL database. A null value in a *Variant* may not be the same as a null value for data types that support nulls such as *Strings*. Some *DevelopmentPlatforms* may not be able to preserve the distinction between a null for a *DataType* and a null for a *Variant*; therefore, applications shall not rely on this distinction. This requirement also means that if an *Attribute* supports the writing of a null value it shall also support writing of an empty *Variant* and vice versa.

Variants can contain arrays of *Variants* but they cannot directly contain another *Variant*.

DiagnosticInfo types only have meaning when returned in a response message with an associated *StatusCode* and table of strings. As a result, *Variants* cannot contain instances of *DiagnosticInfo*.

Values of *Attributes* are always returned in instances of *DataValues*. Therefore, the *DataType* of an *Attribute* cannot be a *DataValue*. *Variants* can contain *DataValue* when used in other contexts such as *Method Arguments* or *PubSub Messages*. The *Variant* in a *DataValue* cannot, directly or indirectly, contain another *DataValue*.

Variables with a *DataType* of *BaseDataType* are mapped to a *Variant*, however, the *ValueRank* and *ArrayDimensions Attributes* place restrictions on what is allowed in the *Variant*. For example, if the *ValueRank* is *Scalar* then the *Variant* may only contain scalar values.

ExtensionObjects and *Variants* allow unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

5.1.7 Decimal

A *Decimal* is a high-precision signed decimal number. It consists of an arbitrary precision integer unscaled value and an integer scale. The scale is the power of ten that is applied to the unscaled value.

A *Decimal* has the fields described in Table 3.

Table 3 – Layout of Decimal

Field	Type	Description
Typeld	Nodeld	The identifier for the <i>Decimal DataType</i> .
Encoding	Byte	This value is always 1.
Length	Int32	The length of the <i>Decimal</i> . If the length is less than or equal to 0 then the <i>Decimal</i> value is 0.
Scale	Int16	A signed integer representing the power of ten used to scale the value. i.e. the decimal number of the value multiplied by $10^{-\text{scale}}$ The integer is encoded starting with the least significant bit.
Value	Byte [*]	A 2-complement signed integer representing the unscaled value. The number of bits is inferred from the length of the <i>length</i> field. If the number of bits is 0 then the value is 0. The integer is encoded with the least significant byte first.

When a *Decimal* is encoded in a *Variant*, the built-in type is set to *ExtensionObject*. Decoders that do not understand the *Decimal* type shall treat it like any other unknown *Structure* and pass it on to the application. Decoders that do understand the *Decimal* can parse the value and use any construct that is suitable for the *DevelopmentPlatform*.

If a *Decimal* is embedded in another *Structure* then the *DataTypeDefinition* for the field shall specify the *Nodeld* of the *Decimal Node* as the *DataType*. If a *Server* publishes an OPC Binary type description for the *Structure* then the type description shall set the *DataType* for the field to *ExtensionObject*.

5.2 OPC UA Binary

5.2.1 General

The OPC UA *Binary DataEncoding* is a data format developed to meet the performance needs of OPC UA applications. This format is designed primarily for fast encoding and decoding, however, the size of the encoded data on the wire was also a consideration.

The OPC UA *Binary DataEncoding* relies on several primitive data types with clearly defined encoding rules that can be sequentially written to or read from a binary stream. A structure is encoded by sequentially writing the encoded form of each field. If a given field is also a structure, then the values of its fields are written sequentially before writing the next field in the containing structure. All fields shall be written to the stream even if they contain null values. The encodings for each primitive type specify how to encode either a null or a default value for the type.

The OPC UA *Binary DataEncoding* does not include any type or field name information because all OPC UA applications are expected to have advance knowledge of the services and structures that they support. An exception is an *ExtensionObject* which provides an identifier and a size for the *Structured DataType* structure it represents. This allows a decoder to skip over types that it does not recognize.

5.2.2 Built-in Types

5.2.2.1 Boolean

A *Boolean* value shall be encoded as a single byte where a value of 0 (zero) is false and any non-zero value is true.

Encoders shall use the value of 1 to indicate a true value; however, decoders shall treat any non-zero value as true.

5.2.2.2 Integer

All integer types shall be encoded as little endian values where the least significant byte appears first in the stream.

Figure 2 illustrates how value 1 000 000 000 (Hex: 3B9ACA00) is encoded as a 32-bit integer in the stream.

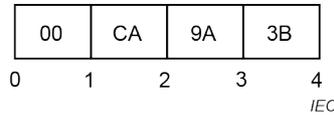


Figure 2 – Encoding Integers in a binary stream

5.2.2.3 Floating Point

All floating-point values shall be encoded with the appropriate IEEE-754 binary representation which has three basic components: the sign, the exponent, and the fraction. The bit ranges assigned to each component depend on the width of the type. Table 4 lists the bit ranges for the supported floating point types.

Table 4 – Supported Floating Point Types

Name	Width (bits)	Fraction	Exponent	Sign
Float	32	0 to 22	23 to 30	31
Double	64	0 to 51	52 to 62	63

In addition, the order of bytes in the stream is significant. All floating point values shall be encoded with the least significant byte appearing first (i.e. little endian).

Figure 3 illustrates how the value -6,5 (Hex: C0D00000) is encoded as a *Float*.

The floating-point type supports positive and negative infinity and not-a-number (NaN). The IEEE specification allows for multiple NaN variants; however, the encoders/decoders may not preserve the distinction. Encoders shall encode a NaN value as an IEEE quiet-NAN (000000000000F8FF) or (0000C0FF). Any unsupported types such as denormalized numbers shall also be encoded as an IEEE quiet-NAN. Any test for equality between NaN values always fails.

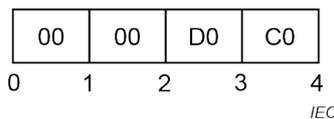


Figure 3 – Encoding Floating Points in a binary stream

5.2.2.4 String

All *String* values are encoded as a sequence of UTF-8 characters without a null terminator and preceded by the length in bytes.

The length in bytes is encoded as *Int32*. A value of -1 is used to indicate a 'null' string.

Figure 4 illustrates how the multilingual string "水Boy" is encoded in a byte stream.

Length				水			B	o	y	
06	00	00	00	E6	B0	B4	42	6F	79	
0	1	2	3	4	5	6	7	8	9	10

IEC

Figure 4 – Encoding Strings in a binary stream

5.2.2.5 DateTime

A *DateTime* value shall be encoded as a 64-bit signed integer (see 5.2.2.2) which represents the number of 100 nanosecond intervals since January 1, 1601 (UTC).

Not all *DevelopmentPlatforms* will be able to represent the full range of dates and times that can be represented with this *DataEncoding*. For example, the UNIX time_t structure only has a 1 second resolution and cannot represent dates prior to 1970. For this reason, a number of rules shall be applied when dealing with date/time values that exceed the dynamic range of a *DevelopmentPlatform*. These rules are:

- a) A date/time value is encoded as 0 if either
 - 1) The value is equal to or earlier than 1601-01-01 12:00AM UTC.
 - 2) The value is the earliest date that can be represented with the *DevelopmentPlatform's* encoding.
- b) A date/time is encoded as the maximum value for an *Int64* if either
 - 1) The value is equal to or greater than 9999-12-31 11:59:59PM UTC,
 - 2) The value is the latest date that can be represented with the *DevelopmentPlatform's* encoding.
- c) A date/time is decoded as the earliest time that can be represented on the platform if either
 - 1) The encoded value is 0,
 - 2) The encoded value represents a time earlier than the earliest time that can be represented with the *DevelopmentPlatform's* encoding.
- d) A date/time is decoded as the latest time that can be represented on the platform if either
 - 1) The encoded value is the maximum value for an *Int64*,
 - 2) The encoded value represents a time later than the latest time that can be represented with the *DevelopmentPlatform's* encoding.

These rules imply that the earliest and latest times that can be represented on a given platform are invalid date/time values and should be treated that way by applications.

A decoder shall truncate the value if a decoder encounters a *DateTime* value with a resolution that is greater than the resolution supported on the *DevelopmentPlatform*.

5.2.2.6 Guid

A *Guid* is encoded in a structure as shown in Table 2. Fields are encoded sequentially according to the data type for field.

Figure 5 illustrates how the *Guid* "72962B91-FA75-4AE6-8D28-B404DC7DAF63" is encoded in a byte stream.

Data1				Data2		Data3		Data4								
91	2B	96	72	75	FA	E6	4A	8D	28	B4	04	DC	7D	AF	63	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

IEC

Figure 5 – Encoding Guids in a binary stream

5.2.2.7 ByteString

A *ByteString* is encoded as sequence of bytes preceded by its length in bytes. The length is encoded as a 32-bit signed integer as described above.

If the length of the byte string is -1 then the byte string is 'null'.

5.2.2.8 XmlElement

An *XmlElement* is an XML fragment serialized as UTF-8 string and then encoded as *ByteString*.

Figure 6 illustrates how the *XmlElement* "<A>Hot水" is encoded in a byte stream.

Length				<A>			Hot			水							
0D	00	00	00	3C	41	3E	72	6F	74	E6	B0	B4	3C	3F	41	3E	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17

IEC

Figure 6 – Encoding XmlElement in a binary stream

A decoder may choose to parse the XML after decoding; if an unrecoverable parsing error occurs then the decoder should try to continue processing the stream. For example, if the *XmlElement* is the body of a *Variant* or an element in an array which is the body of a *Variant* then this error can be reported by setting the value of the *Variant* to the *StatusCode Bad_DecodingError*.

5.2.2.9 NodeId

The components of a *NodeId* are described the Table 5.

Table 5 – NodeId components

Name	Data Type	Description
Namespace	UInt16	The index for a namespace URI. An index of 0 is used for OPC UA defined <i>NodeIds</i> .
IdentifierType	Enumeration	The format and data type of the identifier. The value may be one of the following: NUMERIC - the value is an <i>UInteger</i> ; STRING - the value is <i>String</i> ; GUID - the value is a <i>Guid</i> ; OPAQUE - the value is a <i>ByteString</i> ;
Value	*	The identifier for a node in the address space of an OPC UA Server.

The *DataEncoding* of a *NodeId* varies according to the contents of the instance. For that reason, the first byte of the encoded form indicates the format of the rest of the encoded *NodeId*. The possible *DataEncoding* formats are shown in Table 6. Table 6 through Table 9 describe the structure of each possible format (they exclude the byte which indicates the format).

Table 6 – NodeId DataEncoding values

Name	Value	Description
Two Byte	0x00	A numeric value that fits into the two-byte representation.
Four Byte	0x01	A numeric value that fits into the four-byte representation.
Numeric	0x02	A numeric value that does not fit into the two or four byte representations.
String	0x03	A String value.
Guid	0x04	A Guid value.
ByteString	0x05	An opaque (ByteString) value.
NamespaceUri Flag	0x80	See discussion of <i>ExpandedNodeId</i> in 5.2.2.10.
ServerIndex Flag	0x40	See discussion of <i>ExpandedNodeId</i> in 5.2.2.10.

The standard *NodeId DataEncoding* has the structure shown in Table 7. The standard *DataEncoding* is used for all formats that do not have an explicit format defined.

Table 7 – Standard NodeId Binary DataEncoding

Name	Data Type	Description
Namespace	UInt16	The <i>NamespaceIndex</i> .
Identifier	*	The identifier which is encoded according to the following rules: NUMERIC UInt32 STRING String GUID Guid OPAQUE ByteString

An example of a String *NodeId* with Namespace = 1 and Identifier = "Hot水" is shown in Figure 7.

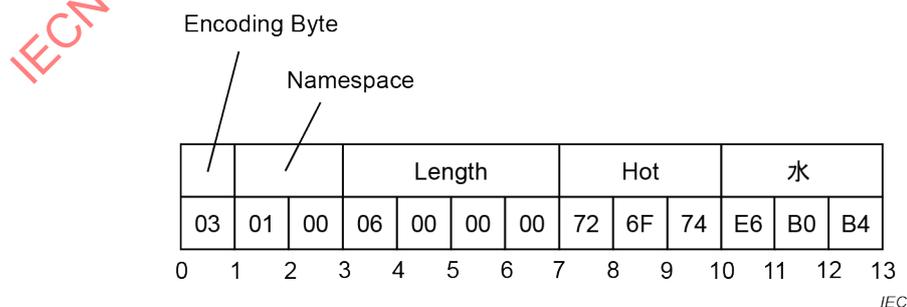


Figure 7 – A String NodeId

The Two Byte *NodeId DataEncoding* has the structure shown in Table 8.

Table 8 – Two Byte NodeId Binary DataEncoding

Name	Data Type	Description
Identifier	Byte	The <i>Namespace</i> is the default OPC UA namespace (i.e. 0). The <i>Identifier</i> Type is 'Numeric'. The <i>Identifier</i> shall be in the range 0 to 255.

An example of a Two Byte *NodeId* with Identifier = 72 is shown in Figure 8.



Figure 8 – A Two Byte NodeId

The Four Byte *NodeId DataEncoding* has the structure shown in Table 9.

Table 9 – Four Byte NodeId Binary DataEncoding

Name	Data Type	Description
Namespace	Byte	The <i>Namespace</i> shall be in the range 0 to 255.
Identifier	UInt16	The <i>Identifier</i> Type is 'Numeric'. The <i>Identifier</i> shall be an integer in the range 0 to 65 535.

An example of a Four Byte *NodeId* with Namespace = 5 and Identifier = 1 025 is shown in Figure 9.

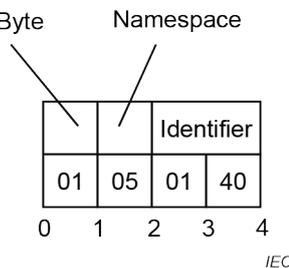


Figure 9 – A Four Byte NodeId

5.2.2.10 ExpandedNodeId

An *ExpandedNodeId* extends the *NodeId* structure by allowing the *NamespaceUri* to be explicitly specified instead of using the *NamespaceIndex*. The *NamespaceUri* is optional. If it is specified, then the *NamespaceIndex* inside the *NodeId* shall be ignored.

The *ExpandedNodeId* is encoded by first encoding a *NodeId* as described in 5.2.2.9 and then encoding *NamespaceUri* as a *String*.

An instance of an *ExpandedNodeId* may still use the *NamespaceIndex* instead of the *NamespaceUri*. In this case, the *NamespaceUri* is not encoded in the stream. The presence of the *NamespaceUri* in the stream is indicated by setting the *NamespaceUri* flag in the encoding format byte for the *NodeId*.

If the *NamespaceUri* is present, then the encoder shall encode the *NamespaceIndex* as 0 in the stream when the *NodeId* portion is encoded. The unused *NamespaceIndex* is included in the stream for consistency.

An *ExpandedNodeId* may also have a *ServerIndex* which is encoded as a *UInt32* after the *NamespaceUri*. The *ServerIndex* flag in the *NodeId* encoding byte indicates whether the *ServerIndex* is present in the stream. The *ServerIndex* is omitted if it is equal to zero.

The *ExpandedNodeId* encoding has the structure shown in Table 10.

Table 10 – ExpandedNodeId Binary DataEncoding

Name	Data Type	Description
NodeId	NodeId	The NamespaceUri and ServerIndex flags in the NodeId encoding indicate whether those fields are present in the stream.
NamespaceUri	String	Not present if null or Empty.
ServerIndex	UInt32	Not present if 0.

5.2.2.11 StatusCode

A *StatusCode* is encoded as a *UInt32*.

5.2.2.12 DiagnosticInfo

A *DiagnosticInfo* structure is described in IEC 62541-4. It specifies a number of fields that could be missing. For that reason, the encoding uses a bit mask to indicate which fields are actually present in the encoded form.

As described in IEC 62541-4, the *SymbolicId*, *NamespaceUri*, *LocalizedText* and *Locale* fields are indexes in a string table which is returned in the response header. Only the index of the corresponding string in the string table is encoded. An index of -1 indicates that there is no value for the string.

DiagnosticInfo allows unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

Table 11 – DiagnosticInfo Binary DataEncoding

Name	Data Type	Description
Encoding Mask	Byte	A bit mask that indicates which fields are present in the stream. The mask has the following bits: 0x01 Symbolic Id 0x02 Namespace 0x04 LocalizedText 0x08 Locale 0x10 Additional Info 0x20 InnerStatusCode 0x40 InnerDiagnosticInfo
SymbolicId	Int32	A symbolic name for the status code.
NamespaceUri	Int32	A namespace that qualifies the symbolic id.
Locale	Int32	The locale used for the localized text.
LocalizedText	Int32	A human readable summary of the status code.
Additional Info	String	Detailed application specific diagnostic information.
Inner StatusCode	StatusCode	A status code provided by an underlying system.
Inner DiagnosticInfo	DiagnosticInfo	Diagnostic info associated with the inner status code.

5.2.2.13 QualifiedName

A *QualifiedName* structure is encoded as shown in Table 12.

The abstract *QualifiedName* structure is defined in IEC 62541-3.

Table 12 – QualifiedName Binary DataEncoding

Name	Data Type	Description
NamespaceIndex	UInt16	The namespace index.
Name	String	The name.

5.2.2.14 LocalizedText

A *LocalizedText* structure contains two fields that could be missing. For that reason, the encoding uses a bit mask to indicate which fields are actually present in the encoded form.

The abstract *LocalizedText* structure is defined in IEC 62541-3.

Table 13 – LocalizedText Binary DataEncoding

Name	Data Type	Description
EncodingMask	Byte	A bit mask that indicates which fields are present in the stream. The mask has the following bits: 0x01 Locale 0x02 Text
Locale	String	The locale. Omitted is null or empty.
Text	String	The text in the specified locale. Omitted is null or empty.

5.2.2.15 ExtensionObject

An *ExtensionObject* is encoded as sequence of bytes prefixed by the *NodeId* of its *DataTypeEncoding* and the number of bytes encoded.

An *ExtensionObject* may be encoded by the application which means it is passed as a *ByteString* or an *XmlElement* to the encoder. In this case, the encoder will be able to write the number of bytes in the object before it encodes the bytes. However, an *ExtensionObject* may know how to encode/decode itself which means the encoder shall calculate the number of bytes before it encodes the object or it shall be able to seek backwards in the stream and update the length after encoding the body.

When a decoder encounters an *ExtensionObject* it shall check if it recognizes the *DataTypeEncoding* identifier. If it does, then it can call the appropriate function to decode the object body. If the decoder does not recognize the type it shall use the *Encoding* to determine if the body is a *ByteString* or an *XmlElement* and then decode the object body or treat it as opaque data and skip over it.

The serialized form of an *ExtensionObject* is shown in Table 14.

Table 14 – Extension Object Binary DataEncoding

Name	Data Type	Description
TypeId	NodeId	<p>The identifier for the <i>DataTypeEncoding</i> node in the <i>Server's AddressSpace</i>. <i>ExtensionObjects</i> defined by the OPC UA specification have a numeric node identifier assigned to them with a <i>NamespaceIndex</i> of 0. The numeric identifiers are defined in A.3.</p> <p>Decoders use this field to determine the syntax of the <i>Body</i>. For example, if this field is the <i>NodeId</i> of the <i>JSON Encoding Object</i> for a <i>DataType</i> then the <i>Body</i> is a <i>ByteString</i> containing a JSON document encoded as a UTF-8 string.</p>
Encoding	Byte	<p>An enumeration that indicates how the body is encoded.</p> <p>The parameter may have the following values:</p> <p>0x00 No body is encoded.</p> <p>0x01 The body is encoded as a <i>ByteString</i>.</p> <p>0x02 The body is encoded as a <i>XmlElement</i>.</p>
Length	Int32	<p>The length of the object body.</p> <p>The length shall be specified if the body is encoded.</p>
Body	Byte [*]	<p>The object body.</p> <p>This field contains the raw bytes for <i>ByteString</i> bodies.</p> <p>For <i>XmlElement</i> bodies this field contains the XML encoded as a UTF-8 string without any null terminator.</p> <p>Some binary encoded structures may have a serialized length that is not a multiple of 8 bits. Encoders shall append 0 bits to ensure the serialized length is a multiple of 8 bits. Decoders that understand the serialized format shall ignore the padding bits.</p>

ExtensionObjects are used in two contexts: as values contained in *Variant* structures or as parameters in OPC UA Messages.

A decoder may choose to parse an *XmlElement* body after decoding; if an unrecoverable parsing error occurs then the decoder should try to continue processing the stream. For example, if the *ExtensionObject* is the body of a *Variant* or an element in an array that is the body of *Variant* then this error can be reported by setting the value of the *Variant* to the *StatusCode Bad_DecodingError*.

5.2.2.16 Variant

A *Variant* is a union of the built-in types.

The structure of a *Variant* is shown in Table 15.

Table 15 – Variant Binary DataEncoding

Name	Data Type	Description
EncodingMask	Byte	<p>The type of data encoded in the stream. A value of 0 specifies a NULL and that no other fields are encoded. The mask has the following bits assigned:</p> <p>0:5 Built-in Type Id (see Table 1). 6 True if the Array Dimensions field is encoded. 7 True if an array of values is encoded.</p> <p>The Built-in Type Ids 26 through 31 are not currently assigned but may be used in the future. Decoders shall accept these IDs, assume the <i>Value</i> contains a <i>ByteString</i> and pass both onto the application. Encoders shall not use these IDs.</p>
ArrayLength	Int32	<p>The number of elements in the array. This field is only present if the array bit is set in the encoding mask.</p> <p>Multi-dimensional arrays are encoded as a one-dimensional array and this field specifies the total number of elements. The original array can be reconstructed from the dimensions that are encoded after the value field.</p> <p>Higher rank dimensions are serialized first. For example, an array with dimensions [2,2,2] is written in this order: [0,0,0], [0,0,1], [0,1,0], [0,1,1], [1,0,0], [1,0,1], [1,1,0], [1,1,1]</p>
Value	*	<p>The value encoded according to its built-in data type.</p> <p>If the array bit is set in the encoding mask, then each element in the array is encoded sequentially. Since many types have variable length encoding each element shall be decoded in order.</p> <p>The value shall not be a <i>Variant</i> but it could be an array of <i>Variants</i>.</p> <p>Many implementation platforms do not distinguish between one dimensional Arrays of <i>Bytes</i> and <i>ByteStrings</i>. For this reason, decoders are allowed to automatically convert an Array of <i>Bytes</i> to a <i>ByteString</i>.</p>
ArrayDimensions Length	Int32	<p>The number of dimensions. This field is only present if the ArrayDimensions flag is set in the encoding mask.</p>
ArrayDimensions	Int32[*]	<p>The length of each dimension encoded as a sequence of Int32 values This field is only present if the ArrayDimensions flag is set in the encoding mask. The lower rank dimensions appear first in the array.</p> <p>All dimensions shall be specified and shall be greater than zero.</p> <p>If <i>ArrayDimensions</i> are inconsistent with the <i>ArrayLength</i> then the decoder shall stop and raise a <i>Bad_DecodingError</i>.</p>

The types and their identifiers that can be encoded in a *Variant* are shown in Table 1.

5.2.2.17 DataValue

A *DataValue* is always preceded by a mask that indicates which fields are present in the stream.

The fields of a *DataValue* are described in Table 16.

Table 16 – Data Value Binary DataEncoding

Name	Data Type	Description
Encoding Mask	Byte	A bit mask that indicates which fields are present in the stream. The mask has the following bits: 0x01 False if the Value is <i>Null</i> . 0x02 False if the StatusCode is Good. 0x04 False if the Source Timestamp is <i>DateTime.MinValue</i> . 0x08 False if the Server Timestamp is <i>DateTime.MinValue</i> . 0x10 False if the Source PicoSeconds is 0. 0x20 False if the Server PicoSeconds is 0.
Value	VARIANT	The value. Not present if the Value bit in the EncodingMask is False.
Status	StatusCode	The status associated with the value. Not present if the StatusCode bit in the EncodingMask is False.
SourceTimestamp	DateTime	The source timestamp associated with the value. Not present if the SourceTimestamp bit in the EncodingMask is False.
SourcePicoSeconds	UInt16	The number of 10 picosecond intervals for the SourceTimestamp. Not present if the SourcePicoSeconds bit in the EncodingMask is False. If the source timestamp is missing the picoseconds are ignored.
ServerTimestamp	DateTime	The <i>Server</i> timestamp associated with the value. Not present if the ServerTimestamp bit in the EncodingMask is False.
ServerPicoSeconds	UInt16	The number of 10 picosecond intervals for the ServerTimestamp. Not present if the ServerPicoSeconds bit in the EncodingMask is False. If the <i>Server</i> timestamp is missing the picoseconds are ignored.

The *PicoSeconds* fields store the difference between a high-resolution timestamp with a resolution of 10 picoseconds and the *Timestamp* field value which only has a 100 ns resolution. The *PicoSeconds* fields shall contain values less than 10 000. The decoder shall treat values greater than or equal to 10 000 as the value '9999'.

5.2.3 Decimal

Decimals are encoded as described in 5.1.7.

A *Decimal* does not have a NULL value.

5.2.4 Enumerations

Enumerations are encoded as *Int32* values.

An *Enumeration* does not have a NULL value.

5.2.5 Arrays

One dimensional *Arrays* are encoded as a sequence of elements preceded by the number of elements encoded as an *Int32* value. If an *Array* is null, then its length is encoded as -1. An *Array* of zero length is different from an *Array* that is null so encoders and decoders shall preserve this distinction.

Multi-dimensional *Arrays* are encoded as an *Int32 Array* containing the dimensions followed by a list of all the values in the *Array*. The total number of values is equal to the product of the dimensions. The number of values is 0 if one or more dimensions are less than or equal to 0. The process for reconstructing the multi-dimensional array is described in 5.2.2.16.

5.2.6 Structures

Structures are encoded as a sequence of fields in the order that they appear in the definition. The encoding for each field is determined by the built-in type for the field.

All fields specified in the structure shall be encoded. If optional fields exist in the structure then see 5.2.7.

Structures do not have a null value. If an encoder is written in a programming language that allows structures to have null values, then the encoder shall create a new instance with default values for all fields and serialize that. Encoders shall not generate an encoding error in this situation.

The following is an example of a structure using C/C++ syntax:

```
struct Type2
{
    Int32 A;
    Int32 B;
};

struct Type1
{
    Int32 X;
    Byte NoOfY;
    Type2* Y;
    Int32 Z;
};
```

In the C/C++ example above, the Y field is a pointer to an array with a length stored in NoOfY. When encoding an array, the length is part of the array encoding so the NoOfY field is not encoded. That said, encoders and decoders use NoOfY during encoding.

An instance of *Type1*, which contains an array of two *Type2* instances would be encoded as 28-byte sequence. If the instance of *Type1* was encoded in an *ExtensionObject* it would have an additional prefix shown in Table 17 which would make the total length 37 bytes. The *TypeId*, *Encoding* and the *Length* are fields defined by the *ExtensionObject*. The encoding of the *Type2* instances do not include any type identifier because it is explicitly defined in *Type1*.

Table 17 – Sample OPC UA Binary Encoded structure

Field	Bytes	Value
Type Id	4	The identifier for the Type1 Binary Encoding Node
Encoding	1	0x1 for ByteString
Length	4	28
X	4	The value of field 'X'
Y.Length	4	2
Y.A	4	The value of field 'Y[0].A'
Y.B	4	The value of field 'Y[0].B'
Y.A	4	The value of field 'Y[1].A'
Y.B	4	The value of field 'Y[1].B'
Z	4	The value of field 'Z'

The Value of the *DataTypeDefinition Attribute* for a DataType Node describing Type1 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type1_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure without optional fields]
fields [0]	StructureField	
name	String	"X"
description	LocalizedText	Description of X
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"Y"
description	LocalizedText	Description of Y-Array
dataType	NodeId	NodeId of the Type2 DataType Node (e.g. "ns=3; s=MyType2")
valueRank	Int32	1 (OneDimension)
isOptional	Boolean	false
fields [2]	StructureField	
name	String	"Z"
description	LocalizedText	Description of Z
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false

The *Value* of the *DataTypeDefinition Attribute* for a *DataType Node* describing Type2 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type2_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure without optional fields]
fields [0]	StructureField	
name	String	"A"
description	LocalizedText	Description of A
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"B"
description	LocalizedText	Description of B
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false

5.2.7 Structures with optional fields

Structures with optional fields are encoded with an encoding mask preceding a sequence of fields in the order that they appear in the definition. The encoding for each field is determined by the data type for the field.

The *EncodingMask* is a 32-bit unsigned integer. Each optional field is assigned exactly one bit. The first optional field is assigned bit '0', the second optional field is assigned bit '1' and so on until all optional fields are assigned bits. A maximum of 32 optional fields can appear within a single *Structure*. Unassigned bits are set to 0 by encoders. Decoders shall report an error if unassigned bits are not 0.

The following is an example of a structure with optional fields using C++ syntax:

```
struct TypeA
{
    Int32 X;
    Int32* O1;
    SByte Y;
    Int32* O2;
};
```

O1 and O2 are optional fields which are NULL if not present.

An instance of *TypeA* which contains two mandatory (X and Y) and two optional (O1 and O2) fields would be encoded as a byte sequence. The length of the byte sequence depends on the available optional fields. An encoding mask field determines the available optional fields.

An instance of *TypeA* where field O2 is available and field O1 is not available would be encoded as a 13-byte sequence. If the instance of *TypeA* was encoded in an *ExtensionObject* it would have the encoded form shown in Table 18 and have a total length of 22 bytes. The length of the *TypeId*, *Encoding* and the *Length* are fields defined by the *ExtensionObject*.

Table 18 – Sample OPC UA Binary Encoded Structure with optional fields

Field	Bytes	Value
Type Id	4	The identifier for the TypeA Binary Encoding Node
Encoding	1	0x1 for ByteString
Length	4	13
EncodingMask	4	0x02 for O2
X	4	The value of X
Y	1	The value of Y
O2	4	The value of O2

If a *Structure* with optional fields is subtyped, the subtypes extend the *EncodingMask* defined for the parent.

The Value of the *DataTypeDefinition* Attribute for a *DataType* Node describing *TypeA* is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "TypeA_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	StructureWithOptionalFields_1 [Structure without optional fields]
fields [0]	StructureField	
name	String	"X"
description	LocalizedText	Description of X
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"O1"
description	LocalizedText	Description of O1
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true
fields [2]	StructureField	
name	String	"Y"
description	LocalizedText	Description of Z
dataType	NodeId	"i=2" [SByte]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [3]	StructureField	
name	String	"O2"
description	LocalizedText	Description of O2
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true

5.2.8 Unions

Unions are encoded as a switch field preceding one of the possible fields. The encoding for the selected field is determined by the data type for the field.

The switch field is encoded as a UInt32.

The switch field is the index of the available union fields starting with 1. If the switch field is 0 then no field is present. For any value greater than the number of defined union fields the encoders or decoders shall report an error.

A *Union* with no fields present has the same meaning as a NULL value. A *Union* with any field present is not a NULL value even if the value of the field itself is NULL.

The following is an example of a union using C/C++ syntax:

```
struct Type2
{
    Int32 A;
    Int32 B;
};

struct Type1
{
    Byte Selector;

    union
    {
        Int32 Field1;
        Type2 Field2;
    }
    Value;
};
```

In the C/C++ example above, the Selector and Value are semantically coupled to form a union. The order of the fields does not matter.

An instance of *Type1* would be encoded as a byte sequence. The length of the byte sequence depends on the selected field.

An instance of *Type1* where field *Field1* is available would be encoded as 8-byte sequence. If the instance of *Type1* was encoded in an *ExtensionObject* it would have the encoded form shown in Table 19 and it would have a total length of 17 bytes. The *TypeId*, *Encoding* and the *Length* are fields defined by the *ExtensionObject*.

Table 19 – Sample OPC UA Binary Encoded Structure

Field	Bytes	Value
Type Id	4	The identifier for Type1
Encoding	1	0x1 for ByteString
Length	4	8
SwitchValue	4	1 for Field1
Field1	4	The value of Field1

The *Value* of the *DataTypeDefinition Attribute* for a *DataType Node* describing Type1 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type1_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Union]
structureType	StructureType	Union_2 [Union]
fields [0]	StructureField	
name	String	"Field1"
description	LocalizedText	Description of Field1
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true
fields [1]	StructureField	
name	String	"Field2"
description	LocalizedText	Description of Field2
dataType	NodeId	NodeId of the Type2 DataType Node (e.g. "ns=3; s=MyType2")
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	true

The *Value* of the *DataTypeDefinition Attribute* for a *DataType Node* describing Type2 is:

Name	Type	Description
defaultEncodingId	NodeId	NodeId of the "Type2_Encoding_DefaultBinary" Node.
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure without optional fields]
fields [0]	StructureField	
name	String	"A"
description	LocalizedText	Description of A
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false
fields [1]	StructureField	
name	String	"B"
description	LocalizedText	Description of B
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalar)
isOptional	Boolean	false

5.2.9 Messages

Messages are *Structures* encoded as sequence of bytes prefixed by the *NodeId* of for the OPC UA Binary *DataTypeEncoding* defined for the Message.

Each OPC UA *Service* described in IEC 62541-4 has a request and response *Message*. The *DataTypeEncoding* IDs assigned to each *Service* are specified in Clause A.3.

5.3 OPC UA XML

5.3.1 Built-in Types

5.3.1.1 General

Most built-in types are encoded in XML using the formats defined in XML Schema Part 2 specification. Any special restrictions or usages are discussed below. Some of the built-in types have an XML Schema defined for them using the syntax defined in XML Schema Part 2.

The prefix *xs:* is used to denote a symbol defined by the XML Schema specification.

5.3.1.2 Boolean

A Boolean value is encoded as an *xs:boolean* value.

5.3.1.3 Integer

Integer values are encoded using one of the subtypes of the *xs:decimal* type. The mappings between the OPC UA integer types and XML schema data types are shown in Table 20.

Table 20 – XML Data Type Mappings for Integers

Name	XML Type
SByte	xs:byte
Byte	xs:unsignedByte
Int16	xs:short
UInt16	xs:unsignedShort
Int32	xs:int
UInt32	xs:unsignedInt
Int64	xs:long
UInt64	xs:unsignedLong

5.3.1.4 Floating Point

Floating point values are encoded using one of the XML floating point types. The mappings between the OPC UA floating point types and XML schema data types are shown in Table 21.

Table 21 – XML Data Type Mappings for Floating Points

Name	XML Type
Float	xs:float
Double	xs:double

The XML floating point type supports positive infinity (INF), negative infinity (-INF) and not-a-number (NaN).

5.3.1.5 String

A *String* value is encoded as an *xs:string* value.

5.3.1.6 DateTime

A *DateTime* value is encoded as an *xs:dateTime* value.

All *DateTime* values shall be encoded as UTC times or with the time zone explicitly specified.

Correct:

```
2002-10-10T00:00:00+05:00
2002-10-09T19:00:00Z
```

Incorrect:

```
2002-10-09T19:00:00
```

It is recommended that all *xs:dateTime* values be represented in UTC format.

The earliest and latest date/time values that can be represented on a *DevelopmentPlatform* have special meaning and shall not be literally encoded in XML.

The earliest date/time value on a *DevelopmentPlatform* shall be encoded in XML as '0001-01-01T00:00:00Z'.

The latest date/time value on a *DevelopmentPlatform* shall be encoded in XML as '9999-12-31T23:59:59Z'

If a decoder encounters a *xs:dateTime* value that cannot be represented on the *DevelopmentPlatform* it should convert the value to either the earliest or latest date/time that can be represented on the *DevelopmentPlatform*. The XML decoder should not generate an error if it encounters an out of range date value.

The earliest date/time value on a *DevelopmentPlatform* is equivalent to a null date/time value.

5.3.1.7 Guid

A *Guid* is encoded using the string representation defined in 5.1.3.

The XML schema for a *Guid* is:

```
<xs:complexType name="Guid">
  <xs:sequence>
    <xs:element name="String" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.8 ByteString

A *ByteString* value is encoded as an *xs:base64Binary* value (see Base64).

The XML schema for a *ByteString* is:

```
<xs:element name="ByteString" type="xs:base64Binary" nillable="true"/>
```

5.3.1.9 XmlElement

An *XmlElement* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="XmlElement">
  <xs:sequence>
    <xs:any minOccurs="0" maxOccurs="1" processContents="lax" />
  </xs:sequence>
</xs:complexType>
```

XmlElement may only be used inside *Variant* or *ExtensionObject* values.

5.3.1.10 NodeId

A *NodeId* value is encoded as an `xs:string` with the syntax:

```
ns=<namespaceindex>;<type>=<value>
```

The elements of the syntax are described in Table 22.

Table 22 – Components of NodeId

Field	Data Type	Description
<namespaceindex>	UInt16	The <i>NamespaceIndex</i> formatted as a base 10 number. If the index is 0 then the entire 'ns=0;' clause shall be omitted.
<type>	Enumeration	A flag that specifies the <i>IdentifierType</i> . The flag has the following values: i NUMERIC (UInt32) s STRING (String) g GUID (Guid) b OPAQUE (ByteString)
<value>	*	The <i>Identifier</i> encoded as string. The <i>Identifier</i> is formatted using the XML data type mapping for the <i>IdentifierType</i> . Note that the <i>Identifier</i> may contain any non-null UTF-8 character including whitespace.

Examples of *NodeIds*:

```
i=13
ns=10;i=-1
ns=10;s>Hello:World
g=09087e75-8e5e-499b-954f-f2a9603db28a
ns=1;b=M/RbKBsRVkePCePcx24oRA==
```

The XML schema for a *NodeId* is:

```
<xs:complexType name="NodeId">
  <xs:sequence>
    <xs:element name="Identifier" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.11 ExpandedNodeId

An *ExpandedNodeId* value is encoded as an `xs:string` with the syntax:

```
svr=<serverindex>;ns=<namespaceindex>;<type>=<value>
or
svr=<serverindex>;nsu=<uri>;<type>=<value>
```

The possible fields are shown in Table 23.

Table 23 – Components of ExpandedNodeId

Field	Data Type	Description
<serverindex>	UInt32	The <i>ServerIndex</i> formatted as a base 10 number. If the <i>ServerIndex</i> is 0 then the entire 'svr=0;' clause shall be omitted.
<namespaceindex>	UInt16	The <i>NamespaceIndex</i> formatted as a base 10 number. If the <i>NamespaceIndex</i> is 0 then the entire 'ns=0;' clause shall be omitted. The <i>NamespaceIndex</i> shall not be present if the URI is present.
<uri>	String	The <i>NamespaceUri</i> formatted as a string. Any reserved characters in the URI shall be replaced with a '%' followed by its 8 bit ANSI value encoded as two hexadecimal digits (case insensitive). For example, the character ';' would be replaced by '%3B'. The reserved characters are ';' and '%'. If the <i>NamespaceUri</i> is null or empty, then 'nsu=;' clause shall be omitted.
<type>	Enumeration	A flag that specifies the <i>IdentifierType</i> . This field is described in Table 22.
<value>	*	The <i>Identifier</i> encoded as string. This field is described in Table 22.

The XML schema for an *ExpandedNodeId* is:

```
<xs:complexType name="ExpandedNodeId">
  <xs:sequence>
    <xs:element name="Identifier" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.12 StatusCode

A *StatusCode* is encoded as an *xs:unsignedInt* with the following XML schema:

```
<xs:complexType name="StatusCode">
  <xs:sequence>
    <xs:element name="Code" type="xs:unsignedInt" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.13 DiagnosticInfo

An *DiagnosticInfo* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="DiagnosticInfo">
  <xs:sequence>
    <xs:element name="SymbolicId" type="xs:int" minOccurs="0" />
    <xs:element name="NamespaceUri" type="xs:int" minOccurs="0" />
    <xs:element name="Locale" type="xs:int" minOccurs="0"/>
    <xs:element name="LocalizedText" type="xs:int" minOccurs="0"/>
    <xs:element name="AdditionalInfo" type="xs:string" minOccurs="0"/>
    <xs:element name="InnerStatusCode" type="tns:StatusCode"
      minOccurs="0" />
    <xs:element name="InnerDiagnosticInfo" type="tns:DiagnosticInfo"
      minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

DiagnosticInfo allows unlimited nesting which could result in stack overflow errors even if the message size is less than the maximum allowed. Decoders shall support at least 100 nesting levels. Decoders shall report an error if the number of nesting levels exceeds what it supports.

5.3.1.14 QualifiedName

A *QualifiedName* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="QualifiedName">
  <xs:sequence>
    <xs:element name="NamespaceIndex" type="xs:int" minOccurs="0" />
    <xs:element name="Name" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.15 LocalizedText

A *LocalizedText* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="LocalizedText">
  <xs:sequence>
    <xs:element name="Locale" type="xs:string" minOccurs="0" />
    <xs:element name="Text" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.16 ExtensionObject

An *ExtensionObject* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="ExtensionObject">
  <xs:sequence>
    <xs:element name="TypeId" type="tns:NodeId" minOccurs="0" />
    <xs:element name="Body" minOccurs="0">
      <xs:complexType>
        <xs:sequence>
          <xs:any minOccurs="0" processContents="lax"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

The body of the *ExtensionObject* contains a single element which is either a *ByteString* or XML encoded *Structure*. A decoder can distinguish between the two by inspecting the top-level element. An element with the name *tns:ByteString* contains an OPC UA Binary encoded body. Any other name shall contain an OPC UA XML encoded body. The *TypeId* specifies the syntax of a *ByteString* body which could be UTF-8 encoded JSON, UA Binary or some other format.

The *TypeId* is the *NodeId* for the *Data Type Encoding Object*.

5.3.1.17 Variant

A *Variant* value is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="Variant">
  <xs:sequence>
    <xs:element name="Value" minOccurs="0" nillable="true">
      <xs:complexType>
        <xs:sequence>
          <xs:any minOccurs="0" processContents="lax"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

If the *Variant* represents a scalar value, then it shall contain a single child element with the name of the built-in type. For example, the single precision floating point value 3,141 5 would be encoded as:

```
<tns:Float>3.1415</tns:Float>
```

If the *Variant* represents a single dimensional array, then it shall contain a single child element with the prefix 'ListOf' and the name built-in type. For example, an *Array* of strings would be encoded as:

```
<tns:ListOfString>
  <tns:String>Hello</tns:String>
  <tns:String>World</tns:String>
</tns:ListOfString>
```

If the *Variant* represents a multidimensional *Array*, then it shall contain a child element with the name '*Matrix*' with the two sub-elements shown in this example:

```
<tns:Matrix>
  <tns:Dimensions>
    <tns:Int32>2</tns:Int32>
    <tns:Int32>2</tns:Int32>
  </tns:Dimensions>
  <tns:Elements>
    <tns:String>A</tns:String>
    <tns:String>B</tns:String>
    <tns:String>C</tns:String>
    <tns:String>D</tns:String>
  </tns:Elements>
</tns:Matrix>
```

In this example, the array has the following elements:

```
[0,0] = "A"; [0,1] = "B"; [1,0] = "C"; [1,1] = "D"
```

The elements of a multi-dimensional *Array* are always flattened into a single dimensional *Array* where the higher rank dimensions are serialized first. This single dimensional *Array* is encoded as a child of the 'Elements' element. The 'Dimensions' element is an *Array* of *Int32* values that specify the dimensions of the array starting with the lowest rank dimension. The multi-dimensional *Array* can be reconstructed by using the dimensions encoded. All dimensions shall be specified and shall be greater than zero. If the dimensions are inconsistent with the number of elements in the array, then the decoder shall stop and raise a *Bad_DecodingError*.

The complete set of built-in type names is found in Table 1.

5.3.1.18 DataValue

A *DataValue* value is encoded as a *xs:complexType* with the following XML schema:

```
<xs:complexType name="DataValue">
  <xs:sequence>
    <xs:element name="Value" type="tns:Variant" minOccurs="0"
      nillable="true" />
    <xs:element name="StatusCode" type="tns:StatusCode"
      minOccurs="0" />
    <xs:element name="SourceTimestamp" type="xs:dateTime"
      minOccurs="0" />
    <xs:element name="SourcePicoseconds" type="xs:unsignedShort"
      minOccurs="0"/>
    <xs:element name="ServerTimestamp" type="xs:dateTime"
      minOccurs="0" />
    <xs:element name="ServerPicoseconds" type="xs:unsignedShort"
      minOccurs="0"/>
  </xs:sequence>
</xs:complexType>
```

5.3.2 Decimal

A *Decimal Value* is encoded as an *xs:complexType* with the following XML schema:

```
<xs:complexType name="Decimal">
  <xs:sequence>
    <xs:element name="TypeId" type="tns:NodeId" minOccurs="0" />
    <xs:element name="Body" minOccurs="0">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="Scale" type="xs:unsignedShort" />
          <xs:element name="Value" type="xs:string" />
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

The *NodeId* is always the *NodeId* of the *Decimal DataType*. When encoded in a *Variant* the *Decimal* is encoded as an *ExtensionObject*. Arrays of *Decimals* are Arrays of *ExtensionObjects*.

The *Value* is a base-10 signed integer with no limit on size. See 5.1.7 for a description of the *Scale* and *Value* fields.

5.3.3 Enumerations

Enumerations that are used as parameters in the *Messages* defined in IEC 62541-4 are encoded as *xs:string* with the following syntax:

```
<symbol>_<value>
```

The elements of the syntax are described in Table 24.

Table 24 – Components of Enumeration

Field	Type	Description
<symbol>	String	The symbolic name for the enumerated value.
<value>	UInt32	The numeric value associated with enumerated value.

For example, the XML schema for the *NodeClass* enumeration is:

```
<xs:simpleType name="NodeClass">
  <xs:restriction base="xs:string">
    <xs:enumeration value="Unspecified_0" />
    <xs:enumeration value="Object_1" />
    <xs:enumeration value="Variable_2" />
    <xs:enumeration value="Method_4" />
    <xs:enumeration value="ObjectType_8" />
    <xs:enumeration value="VariableType_16" />
    <xs:enumeration value="ReferenceType_32" />
    <xs:enumeration value="DataType_64" />
    <xs:enumeration value="View_128" />
  </xs:restriction>
</xs:simpleType>
```

Enumerations that are stored in a *Variant* are encoded as an *Int32* value.

For example, any *Variable* could have a value with a *DataType* of *NodeClass*. In this case, the corresponding numeric value is placed in the *Variant* (e.g. *NodeClass Object* would be stored as a 1).

5.3.4 Arrays

One dimensional *Array* parameters are always encoded by wrapping the elements in a container element and inserting the container into the structure. The name of the container element should be the name of the parameter. The name of the element in the array shall be the type name.

For example, the *Read* service takes an array of *ReadValueIds*. The XML schema would look like:

```
<xs:complexType name="ListOfReadValueId">
  <xs:sequence>
    <xs:element name="ReadValueId" type="tns:ReadValueId"
      minOccurs="0" maxOccurs="unbounded" nillable="true" />
  </xs:sequence>
</xs:complexType>
```

The nillable attribute shall be specified because XML encoders will drop elements in arrays if those elements are empty.

Multi-dimensional *Array* parameters are encoded using the *Matrix* type defined in 5.3.1.17.

5.3.5 Structures

Structures are encoded as a *xs:complexType* with all of the fields appearing in a sequence. All fields are encoded as an *xs:element*. All elements have *minOccurs* set 0 to allow for compact XML representations. If an element is missing the default value for the field type is used. If the field type is a structure, the default value is an instance of the structure with default values for each contained field.

Types which have a NULL value defined shall have the *nillable="true"* flag set.

For example, the Read service has a *ReadValueId* structure in the request. The XML schema would look like:

```
<xs:complexType name="ReadValueId">
  <xs:sequence>
    <xs:element name="NodeId" type="tns:NodeId"
      minOccurs="0" nillable="true" />
    <xs:element name="AttributeId" type="xs:int" minOccurs="0" />
    <xs:element name="IndexRange" type="xs:string"
      minOccurs="0" nillable="true" />
    <xs:element name="DataEncoding" type="tns:NodeId"
      minOccurs="0" nillable="true" />
  </xs:sequence>
</xs:complexType>
```

5.3.6 Structures with optional fields

Structures with optional fields are encoded as an *xs:complexType* with all of the fields appearing in a sequence. The first element is a bit mask that specifies what fields are encoded. The bits in the mask are sequentially assigned to optional fields in the order they appear in the *Structure*.

To allow for compact XML, any field can be omitted from the XML so decoders shall assign default values based on the field type for any mandatory fields.

For example, the following *Structure* has one mandatory and two optional fields. The XML schema would look like:

```
<xs:complexType name="OptionalType">
  <xs:sequence>
    <xs:element name="EncodingMask" type="xs:unsignedLong" />
    <xs:element name="X" type="xs:int" minOccurs="0" />
    <xs:element name="O1" type="xs:int" minOccurs="0" />
    <xs:element name="Y" type="xs:byte" minOccurs="0" />
    <xs:element name="O2" type="xs:int" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

In the example above, the *EncodingMask* has a value of 3 if both O1 and O2 are encoded. Encoders shall set unused bits to 0 and decoders shall ignore unused bits.

If a *Structure* with optional fields is subtyped, the subtypes extend the *EncodingMask* defined for the parent.

5.3.7 Unions

Unions are encoded as an *xs:complexType* containing an *xs:sequence* with two entries.

The first entry in the sequence is the *SwitchField xs:element* and specifies a numeric value which identifies which element in the *xs:choice* is encoded. The name of the element may be any valid text.

The second entry in the sequence is an *xs:choice* which specifies the possible fields. The order in the *xs:choice* determines the value of the *SwitchField* when that choice is encoded. The first element has a *SwitchField* value of 1 and the last value has a *SwitchField* equal to the number of choices.

No additional elements in the sequence are permitted. If the *SwitchField* is missing or 0 then the union has a NULL value. Encoders or decoders shall report an error for any *SwitchField* value greater than the number of defined union fields.

For example, the following union has two fields. The XML schema would look like:

```
<xs:complexType name="Type1">
  <xs:sequence>
    <xs:element name="SwitchField"
      type="xs:unsignedInt" minOccurs="0"/>
    <xs:choice>
      <xs:element name="Field1" type="xs:int" minOccurs="0"/>
      <xs:element name="Field2" type="tns:Field2" minOccurs="0"/>
    </xs:choice>
  </xs:sequence>
</xs:complexType>
```

5.3.8 Messages

Messages are encoded as an `xs:complexType`. The parameters in each *Message* are serialized in the same way the fields of a *Structure* are serialized.

5.4 OPC UA JSON

5.4.1 General

The JSON *DataEncoding* was developed to allow OPC UA applications to interoperate with web and enterprise software that use this format. The OPC UA JSON *DataEncoding* defines standard JSON representations for all OPC UA Built-In types.

The JSON format is defined in RFC 7159. It is partially self-describing because each field has a name encoded in addition to the value; however, JSON has no mechanism to qualify names with namespaces.

The JSON format does not have a published standard for a schema that can be used to describe the contents of a JSON document. However, the schema mechanisms defined in this document can be used to describe JSON documents. Specifically, the *DataTypeDescription* structure defined in IEC 62541-3 can define any JSON document that conforms to the rules described below.

Servers that support the JSON *DataEncoding* shall add *DataTypeEncoding Nodes* called "Default JSON" to all *DataTypes* which can be serialized with the JSON encoding. The *NodeIds* of these *Nodes* are defined by the information model which defines the *DataType*. These *NodeIds* are used in *ExtensionObjects* as described in 5.4.2.16.

There are two important use cases for the JSON encoding: Cloud applications which consume *PubSub* messages and JavaScript *Clients* (JSON is the preferred serialization format for JavaScript). For the Cloud application use case, the *PubSub* message needs to be self-contained which implies it cannot contain numeric references to an externally defined namespace table. Cloud applications also often rely on scripting languages to process the incoming messages, so artefacts in the *DataEncoding* that exist to ensure fidelity during decoding are not necessary. For this reason, this *DataEncoding* defines a 'non-reversible' form which is designed to meet the needs of Cloud applications. Applications, such as JavaScript Clients, which use the *DataEncoding* for communication with other OPC UA applications use the normal or 'reversible' form. The differences, if any, between the reversible and non-reversible forms are described for each type.

5.4.2 Built-in Types

5.4.2.1 General

Any value for a Built-In type that is NULL shall be encoded as the JSON literal 'null' if the value is an element of an array. If the NULL value is a field within a *Structure* or *Union*, the field shall not be encoded.

5.4.2.2 Boolean

A *Boolean* value shall be encoded as the JSON literal 'true' or 'false'.

5.4.2.3 Integer

Integer values other than *Int64* and *UInt64* shall be encoded as a JSON number.

Int64 and *UInt64* values shall be formatted as a decimal number encoded as a JSON string

(See the XML encoding of 64-bit values described in 5.3.1.3).

5.4.2.4 Floating point

Normal *Float* and *Double* values shall be encoded as a JSON number.

Special floating-point numbers such as positive infinity (INF), negative infinity (-INF) and not-a-number (NaN) shall be represented by the values "Infinity", "-Infinity" and "NaN" encoded as a JSON string. See 5.2.2.3 for more information on the different types of special floating-point numbers.

5.4.2.5 String

String values shall be encoded as JSON strings.

Any characters which are not allowed in JSON strings are escaped using the rules defined in RFC 7159.

5.4.2.6 DateTime

DateTime values shall be formatted as specified by ISO 8601-1:2019 and encoded as a JSON string.

DateTime values which exceed the minimum or maximum values supported on a platform shall be encoded as "0001-01-01T00:00:00Z" or "9999-12-31T23:59:59Z", respectively. During decoding, these values shall be converted to the minimum or maximum values supported on the platform.

DateTime values equal to "0001-01-01T00:00:00Z" are considered to be NULL values.

5.4.2.7 Guid

Guid values shall be formatted as described in 5.1.3 and encoded as a JSON string.

5.4.2.8 ByteString

ByteString values shall be formatted as a Base64 text and encoded as a JSON string.

Any characters which are not allowed in JSON strings are escaped using the rules defined in RFC 7159.

5.4.2.9 XmlElement

XmlElement value shall be encoded as a String as described in 5.3.1.9.

5.4.2.10 NodeId

NodeId values shall be encoded as a JSON object with the fields defined in Table 25.

The abstract *NodeId* structure is defined in IEC 62541-3 and has three fields: *Identifier*, *IdentifierType* and *NamespaceIndex*. The representation of these abstract fields is described in Table 25.

Table 25 – JSON Object Definition for a NodeId

Name	Description
IdType	The <i>IdentifierType</i> encoded as a JSON number. Allowed values are: 0 – <i>UInt32 Identifier</i> encoded as a JSON number. 1 – A <i>String Identifier</i> encoded as a JSON string. 2 – A <i>Guid Identifier</i> encoded as described in 5.4.2.7. 3 – A <i>ByteString Identifier</i> encoded as described in 5.4.2.8. This field is omitted for <i>UInt32</i> identifiers.
Id	The <i>Identifier</i> . The value of the <i>id</i> field specifies the encoding of this field.
Namespace	The <i>NamespaceIndex</i> for the <i>NodeId</i> . The field is encoded as a JSON number for the reversible encoding. The field is omitted if the <i>NamespaceIndex</i> equals 0. For the non-reversible encoding, the field is the <i>NamespaceUri</i> associated with the <i>NamespaceIndex</i> , encoded as a JSON string. A <i>NamespaceIndex</i> of 1 is always encoded as a JSON number.

5.4.2.11 ExpandedNodeId

ExpandedNodeId values shall be encoded as a JSON object with the fields defined in Table 26.

The abstract *ExpandedNodeId* structure is defined in IEC 62541-3 and has five fields: *Identifier*, *IdentifierType*, *NamespaceIndex*, *NamespaceUri* and *ServerIndex*. The representation of these abstract fields is described in Table 26.

IECNORM.COM . Click to view the full PDF of IEC 62541-6:2020 RLV

Table 26 – JSON Object Definition for an ExpandedNodeId

Name	Description
IdType	The <i>IdentifierType</i> encoded as a JSON number. Allowed values are: 0 – <i>UInt32 Identifier</i> encoded as a JSON number. 1 – A <i>String Identifier</i> encoded as a JSON string. 2 – A <i>Guid Identifier</i> encoded as described in 5.4.2.7. 3 – A <i>ByteString Identifier</i> encoded as described in 5.4.2.8. This field is omitted for UInt32 identifiers.
Id	The <i>Identifier</i> . The value of the 't' field specifies the encoding of this field.
Namespace	The <i>NamespaceIndex</i> or the <i>NamespaceUri</i> for the <i>ExpandedNodeId</i> . If the <i>NamespaceUri</i> is not specified, the <i>NamespaceIndex</i> is encoded with these rules: The field is encoded as a JSON number for the reversible encoding. The field is omitted if the <i>NamespaceIndex</i> equals 0. For the non-reversible encoding the field is the <i>NamespaceUri</i> associated with the <i>NamespaceIndex</i> encoded as a JSON string. A <i>NamespaceIndex</i> of 1 is always encoded as a JSON number. If the <i>NamespaceUri</i> is specified it is encoded as a JSON string in this field.
ServerUri	The <i>ServerIndex</i> for the <i>ExpandedNodeId</i> . This field is encoded as a JSON number for the reversible encoding. This field is omitted if the <i>ServerIndex</i> equals 0. For the non-reversible encoding, this field is the <i>ServerUri</i> associated with the <i>ServerIndex</i> portion of the <i>ExpandedNodeId</i> , encoded as a JSON string.

5.4.2.12 StatusCode

StatusCode values shall be encoded as a JSON number for the reversible encoding.

For the non-reversible form, *StatusCode* values shall be encoded as a JSON object with the fields defined in Table 27.

Table 27 – JSON Object Definition for a StatusCode

Name	Description
Code	The numeric code encoded as a JSON number. The <i>Code</i> is omitted if the numeric code is 0 (Good).
Symbol	The string literal associated with the numeric code encoded as JSON string. e.g. 0x80AB0000 has the associated literal "BadInvalidArgument". The <i>Symbol</i> is omitted if the numeric code is 0 (Good).

A *StatusCode* of Good (0) is treated like a NULL and not encoded. If it is an element of an JSON array it is encoded as the JSON literal 'null'.

5.4.2.13 DiagnosticInfo

DiagnosticInfo values shall be encoded as a JSON object with the fields shown in Table 28.

Table 28 – JSON Object Definition for a DiagnosticInfo

Name	Data Type	Description
SymbolicId	Int32	A symbolic name for the status code.
NamespaceUri	Int32	A namespace that qualifies the symbolic id.
Locale	Int32	The locale used for the localized text.
LocalizedText	Int32	A human readable summary of the status code.
AdditionalInfo	String	Detailed application-specific diagnostic information.
InnerStatusCode	StatusCode	A status code provided by an underlying system.
InnerDiagnosticInfo	DiagnosticInfo	Diagnostic info associated with the inner status code.

Each field is encoded using the rules defined for the built-in type specified in the Data Type column.

The *SymbolicId*, *NamespaceUri*, *Locale* and *LocalizedText* fields are encoded as JSON numbers which reference the *StringTable* contained in the *ResponseHeader*.

5.4.2.14 QualifiedName

QualifiedName values shall be encoded as a JSON object with the fields shown in Table 29.

The abstract *QualifiedName* structure is defined in IEC 62541-3 and has two fields *Name* and *NamespaceIndex*. The *NamespaceIndex* is represented by the *Uri* field in the JSON object.

Table 29 – JSON Object Definition for a QualifiedName

Name	Description
Name	The Name component of the <i>QualifiedName</i> .
Uri	The <i>NamespaceIndex</i> component of the <i>QualifiedName</i> encoded as a JSON number. The <i>Uri</i> field is omitted if the <i>NamespaceIndex</i> equals 0. For the non-reversible form, the <i>NamespaceUri</i> associated with the <i>NamespaceIndex</i> portion of the <i>QualifiedName</i> is encoded as JSON string unless the <i>NamespaceIndex</i> is 1 or if <i>NamespaceUri</i> is unknown. In these cases, the <i>NamespaceIndex</i> is encoded as a JSON number.

5.4.2.15 LocalizedText

LocalizedText values shall be encoded as a JSON object with the fields shown in Table 30.

The abstract *LocalizedText* structure is defined in IEC 62541-3 and has two fields *Text* and *Locale*.

Table 30 – JSON Object Definition for a LocalizedText

Name	Description
Locale	The <i>Locale</i> portion of <i>LocalizedText</i> values shall be encoded as a JSON string
Text	The <i>Text</i> portion of <i>LocalizedText</i> values shall be encoded as a JSON string.

For the non-reversible form, *LocalizedText* value shall be encoded as a JSON string containing the *Text* component.

5.4.2.16 ExtensionObject

ExtensionObject values shall be encoded as a JSON object with the fields shown in Table 31.

Table 31 – JSON Object Definition for an ExtensionObject

Name	Description
TypeId	The <i>NodeId</i> of a <i>DataTypeEncoding Node</i> formatted using the rules in 5.4.2.10.
Encoding	The format of the <i>Body</i> field encoded as a JSON number. This value is 0 if the body is <i>Structure</i> encoded as a JSON object (see 5.4.6). This value is 1 if the body is a <i>ByteString</i> value encoded as a JSON string (see 5.4.2.8). This value is 2 if the body is an <i>XmlElement</i> value encoded as a JSON string (see 5.4.2.9). This field is omitted if the value is 0.
Body	<i>Body</i> of the <i>ExtensionObject</i> . The type of this field is specified by the <i>Encoding</i> field. If the <i>Body</i> is empty, the <i>ExtensionObject</i> is NULL and is omitted or encoded as a JSON null.

For the non-reversible form, *ExtensionObject* values shall be encoded as a JSON object containing only the value of the *Body* field. The *TypeId* and *Encoding* fields are dropped.

5.4.2.17 Variant

Variant values shall be encoded as a JSON object with the fields shown in Table 32.

Table 32 – JSON Object Definition for a Variant

Name	Description
Type	The Built-in type for the value contained in the <i>Body</i> (see Table 1) encoded as a JSON number. If type is 0 (NULL) the <i>Variant</i> contains a NULL value and the containing JSON object shall be omitted or replaced by the JSON literal 'null' (when an element of a JSON array).
Body	If the value is a scalar it is encoded using the rules for type specified for the <i>Type</i> . If the value is a one-dimensional array it is encoded as a JSON array (see 5.4.5). Multi-dimensional arrays are encoded as a one dimensional JSON array which is reconstructed using the value of the <i>Dimensions</i> field (see 5.2.2.16).
Dimensions	The dimensions of the array encoded as a JSON array of JSON numbers. The <i>Dimensions</i> are omitted for scalar and one-dimensional array values.

For the non-reversible form, *Variant* values shall be encoded as a JSON object containing only the value of the *Body* field. The *Type* and *Dimensions* fields are dropped. Multi-dimensional arrays are encoded as a multi-dimensional JSON array as described in 5.4.5.

5.4.2.18 DataValue

DataValue values shall be encoded as a JSON object with the fields shown in Table 33.

Table 33 – JSON Object Definition for a DataValue

Name	Data Type	Description
Value	Variant	The value.
Status	StatusCode	The status associated with the value.
SourceTimestamp	DateTime	The source timestamp associated with the value.
SourcePicoSeconds	UInt16	The number of 10 picosecond intervals for the SourceTimestamp.
ServerTimestamp	DateTime	The <i>Server</i> timestamp associated with the value.
ServerPicoSeconds	UInt16	The number of 10 picosecond intervals for the ServerTimestamp.

If a field has a null or default value it is omitted. Each field is encoded using the rules defined for the built-in type specified in the Data Type column.

5.4.3 Decimal

Decimal values shall be encoded as a JSON object with the fields in Table 34.

Table 34 – JSON Object Definition for a Decimal

Name	Description
Scale	A JSON number with the scale applied to the Value.
Value	A JSON string with the Value encoded as a base-10 signed integer. (See the XML encoding of Integer values described in 5.3.1.3).

See 5.1.7 for a description of the *Scale* and *Value* fields.

5.4.4 Enumerations

Enumeration values shall be encoded as a JSON number for the reversible encoding.

For the non-reversible form, *Enumeration* values are encoded as literals with the value appended as a JSON string.

The format of the string literal is:

<name>_<value>

where the name is the enumeration literal and the value is the numeric value.

If the literal is not known to the encoder, the numeric value is encoded as a JSON string.

5.4.5 Arrays

One dimensional *Arrays* shall be encoded as JSON arrays.

If an element is NULL, the element shall be encoded as the JSON literal 'null'.

Otherwise, the element is encoded according to the rules defined for the type.

Multi-dimensional *Arrays* are encoded as nested JSON arrays. The outer array is the first dimension and the innermost array is the last dimension. For example, the following matrix

0 2 3
1 3 4

is encoded in JSON as

```
[ [0, 2, 3], [1, 3, 4] ]
```

5.4.6 Structures

Structures shall be encoded as JSON objects.

If the value of a field is NULL it shall be omitted from the encoding.

For example, instances of the structures:

```
struct Type2
{
  Int32 A;
  Int32 B;
  Char* C;
};

struct Type1
{
  Int32 X;
  Int32 NoOfY;
  Type2* Y;
  Int32 Z;
};
```

are represented in JSON as:

```
{
  "X":1234,
  "Y":[ { "A":1, "B":2, "C":"Hello" }, { "A":3, "B":4 } ],
  "Z":5678
}
```

where "C" is omitted from the second Type2 instance because it has a NULL value.

5.4.7 Structures with optional fields

Structures with optional fields shall be encoded as JSON objects as shown in Table 35.

Table 35 – JSON Object Definition for a *Structure* with Optional Fields

Name	Description
EncodingMask	A bit mask indicating what fields are encoded in the structure (see 5.2.7) This mask is encoded as a JSON number. The bits are sequentially assigned to optional fields in the order that they are defined.
<FieldName>	The fields' structure is encoded according to the rules defined for their <i>DataType</i> .

For the non-reversible form, *Structures* with optional fields are encoded like *Structures*.

If a *Structure* with optional fields is subtyped, the subtypes extend the *EncodingMask* defined for the parent.

The following is an example of a structure with optional fields using C++ syntax:

```

struct TypeA
{
  Int32 X;
  Int32* O1;
  SByte Y;
  Int32* O2;
};

```

O1 and O2 are optional fields where a NULL indicates that the field is not present.

Assume that O1 is not specified and the value of O2 is 0.

The reversible encoding would be:

```

{ "EncodingMask": 2, "X": 1, "Y": 2 }

```

where decoders would assign the default value of 0 to O2 since the mask bit is set, even though the field was omitted (this is the behaviour defined for the *Int32 DataType*). Decoders would mark O1 as 'not specified'.

5.4.8 Unions

Unions shall be encoded as JSON objects as shown in Table 36 for the reversible encoding.

Table 36 – JSON Object Definition for a Union

Name	Description
SwitchField	The identifier for the field in the Union which is encoded as a JSON number.
Value	The value of the field encoded using the rules that apply to the data type.

For the non-reversible form, *Union* values are encoded using the rule for the current value.

For example, instances of the union:

```

struct Union1
{
  Byte Selector;
  {
    Int32 A;
    Double B;
    Char* C;
  }
  Value;
};

```

would be represented in reversible form as:

```

{ "SwitchField":2, "Value":3.1415 }

```

In non-reversible form, it is represented as:

```

3.1415

```

5.4.9 Messages

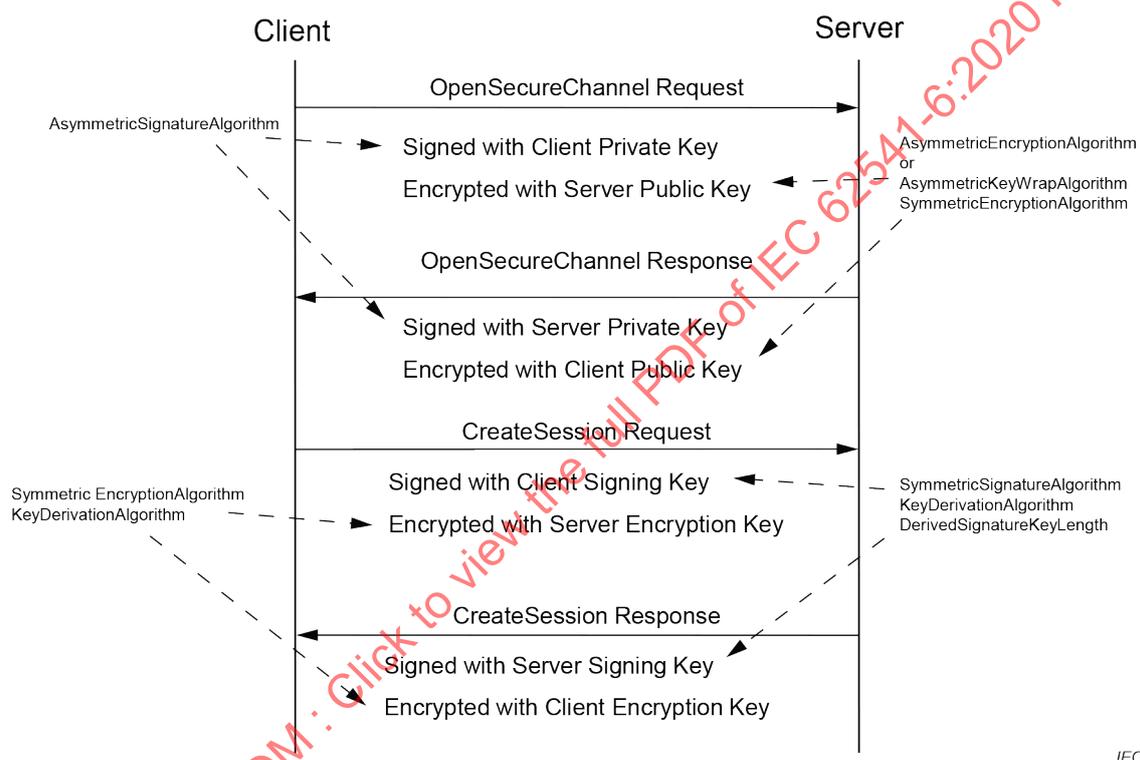
Messages are encoded *ExtensionObjects* (see 5.4.2.16).

6 Message SecurityProtocols

6.1 Security handshake

All *SecurityProtocols* shall implement the *OpenSecureChannel* and *CloseSecureChannel* services defined in IEC 62541-4. These *Services* specify how to establish a *SecureChannel* and how to apply security to *Messages* exchanged over that *SecureChannel*. The *Messages* exchanged and the security algorithms applied to them are shown in Figure 10.

SecurityProtocols shall support three *SecurityModes*: *None*, *Sign* and *SignAndEncrypt*. If the *SecurityMode* is *None* then no security is used and the security handshake shown in Figure 10 is not required. However, a *SecurityProtocol* implementation shall still maintain a logical channel and provide a unique identifier for the *SecureChannel*.



IEC

Figure 10 – Security handshake

Each *SecurityProtocol* mapping specifies exactly how to apply the security algorithms to the *Message*. A set of security algorithms that shall be used together during a security handshake is called a *SecurityPolicy*. IEC 62541-7 defines standard *SecurityPolicies* as parts of the standard *Profiles* which OPC UA applications are expected to support. IEC 62541-7 also defines a URI for each standard *SecurityPolicy*.

A *Stack* is expected to have built in knowledge of the *SecurityPolicies* that it supports. Applications specify the *SecurityPolicy* they wish to use by passing the URI to the *Stack*.

Table 37 defines the contents of a *SecurityPolicy*. Each *SecurityProtocol* mapping specifies how to use each of the parameters in the *SecurityPolicy*. A *SecurityProtocol* mapping may not make use of all of the parameters.

Table 37 – SecurityPolicy

Name	Description
PolicyUri	The URI assigned to the <i>SecurityPolicy</i> .
SymmetricSignatureAlgorithm	The symmetric signature algorithm to use.
SymmetricEncryptionAlgorithm	The symmetric encryption algorithm to use.
AsymmetricSignatureAlgorithm	The asymmetric signature algorithm to use.
AsymmetricEncryptionAlgorithm	The asymmetric encryption algorithm to use.
MinAsymmetricKeyLength	The minimum length, in bits, for an asymmetric key.
MaxAsymmetricKeyLength	The maximum length, in bits, for an asymmetric key.
KeyDerivationAlgorithm	The key derivation algorithm to use.
DerivedSignatureKeyLength	The length in bits of the derived key used for <i>Message</i> authentication.
CertificateSignatureAlgorithm	The asymmetric signature algorithm used to sign certificates.
SecureChannelNonceLength	The length, in bytes, of the <i>Nonces</i> exchanged when creating a <i>SecureChannel</i> .

The *KeyDerivationAlgorithm* is used to create the keys used to secure *Messages* sent over the *SecureChannel*. The length of the keys used for encryption is implied by the *SymmetricEncryptionAlgorithm*. The length of the keys used for creating *Signatures* is specified by the *DerivedSignatureKeyLength*.

The *MinAsymmetricKeyLength* and *MaxAsymmetricKeyLength* are constraints that apply to all *Certificates* (including *Issuers* in the chain). In addition, the key length of issued *Certificates* shall be less than or equal to the key length of the issuer *Certificate*. See 6.2.3 for information on *Certificate* chains.

The *CertificateKeyAlgorithm* and *EphemeralKeyAlgorithm* are used to generate new asymmetric key pairs used with *Certificates* and during the *SecureChannel* handshake. IEC 62541-7 specifies the bit lengths that need to be supported for each *SecurityPolicy*.

The *CertificateSignatureAlgorithm* applies to the *Certificate* and all *Issuer Certificates*. If a *CertificateSignatureAlgorithm* allows for more than one algorithm then the algorithms are listed in order of increasing priority. Each *Issuer* in a chain shall have an algorithm that is the same or higher priority than any *Certificate* it issues.

The *SecureChannelNonceLength* specifies the length of the *Nonces* exchanged when establishing a *SecureChannel* (see 6.7.4).

6.2 Certificates

6.2.1 General

OPC UA applications use *Certificates* to store the *Public Keys* needed for *Asymmetric Cryptography* operations. All *SecurityProtocols* use X.509 v3 *Certificates* (see X.509 v3) encoded using the DER format (see X690). *Certificates* used by OPC UA applications shall also conform to RFC 3280 which defines a profile for X.509 v3 *Certificates* when they are used as part of an Internet based application.

The *ServerCertificate* and *ClientCertificate* parameters used in the abstract *OpenSecureChannel* service are instances of the *Application Instance Certificate Data Type*. 6.2.2 describes how to create an X.509 v3 *Certificate* that can be used as an *Application Instance Certificate*.

6.2.2 Application Instance Certificate

An *Application Instance Certificate* is a *ByteString* containing the DER encoded form (see X690) of an X.509 v3 *Certificate*. This *Certificate* is issued by certifying authority and identifies an instance of an application running on a single host. The X.509 v3 fields contained in an *Application Instance Certificate* are described in Table 38. The fields are defined completely in RFC 3280.

Table 38 also provides a mapping from the RFC 3280 terms to the terms used in the abstract definition of an *Application Instance Certificate* defined in IEC 62541-4.

Table 38 – Application Instance Certificate

Name	IEC 62541-4 Parameter Name	Description
Application Instance Certificate		An X.509 v3 <i>Certificate</i> .
version	version	shall be "V3"
serialNumber	serialNumber	The serial number assigned by the issuer.
signatureAlgorithm	signatureAlgorithm	The algorithm used to sign the <i>Certificate</i> .
signature	signature	The signature created by the issuer.
issuer	issuer	The distinguished name of the <i>Certificate</i> used to create the signature. The <i>issuer</i> field is completely described in RFC 3280.
validity	validTo, validFrom	When the <i>Certificate</i> becomes valid and when it expires.
subject	subject	The distinguished name of the application <i>Instance</i> . The Common Name attribute shall be specified and should be the <i>productName</i> or a suitable equivalent. The Organization Name attribute shall be the name of the Organization that executes the application instance. This organization is usually not the vendor of the application. Other attributes may be specified. The <i>subject</i> field is completely described in RFC 3280.
subjectAltName	applicationUri, hostnames	The alternate names for the application <i>Instance</i> . Shall include a uniformResourceIdentifier which is equal to the <i>applicationUri</i> . The URI shall be a valid URL (see RFC 1738) or a valid URN (see RFC 2141). <i>Servers</i> shall specify a partial or a fully qualified <i>dnsName</i> or a static <i>IPAddress</i> which identifies the machine where the application <i>Instance</i> runs. Additional <i>dnsNames</i> may be specified if the machine has multiple names. The <i>subjectAltName</i> field is completely described in RFC 3280.
publicKey	publicKey	The public key associated with the <i>Certificate</i> .
keyUsage	keyUsage	Specifies how the <i>Certificate</i> key may be used. Shall include digitalSignature, nonRepudiation, keyEncipherment and dataEncipherment. Other key uses are allowed.
extendedKeyUsage	keyUsage	Specifies additional key uses for the <i>Certificate</i> . Shall specify 'serverAuth and/or clientAuth'. Other key uses are allowed.
authorityKeyIdentifier	(no mapping)	Provides more information about the key used to sign the <i>Certificate</i> . It shall be specified for <i>Certificates</i> signed by a CA. It should be specified for self-signed <i>Certificates</i> .

6.2.3 Certificate Chains

Any X.509 v3 *Certificate* may be signed by CA which means that validating the signature requires access to the X.509 v3 *Certificate* belonging to the signing CA. Whenever an application validates a signature it shall recursively build a chain of *Certificates* by finding the issuer *Certificate*, validating the *Certificate* and then repeating the process for the issuer *Certificate*. The chain ends with a self-signed *Certificate*.

The number of CAs used in a system should be small so it is common to install the necessary CAs on each machine with an OPC UA application. However, applications have the option of including a partial or complete chain whenever they pass a *Certificate* to a peer during the *SecureChannel* negotiation and during the *CreateSession/ActivateSession* handshake. All OPC UA applications shall accept partial or complete chains in any field that contains a DER encoded *Certificate*.

Chains are stored in a *ByteString* by simply appending the DER encoded form of the *Certificates*. The first *Certificate* shall be the end *Certificate* followed by its issuer. If the root CA is sent as part of the chain, the last *Certificate* is appended to the *ByteString*.

Chains are parsed by extracting the length of each Certificate from the DER encoding. For Certificates with lengths less than 65 535 bytes, it is an MSB encoded UInt16 starting at the third byte.

6.3 Time synchronization

All *SecurityProtocols* require that system clocks on communicating machines be reasonably synchronized in order to check the expiry times for *Certificates* or *Messages*. The amount of clock skew that can be tolerated depends on the system security requirements and applications shall allow administrators to configure the acceptable clock skew when verifying times. A suitable default value is 5 minutes.

The Network Time Protocol (NTP) provides a standard way to synchronize a machine clock with a time server on the network. Systems running on a machine with a full featured operating system like Windows or Linux will already support NTP or an equivalent. Devices running embedded operating systems should support NTP.

If a device operating system cannot practically support NTP then an OPC UA application can use the *Timestamps* in the *ResponseHeader* (see IEC 62541-4) to synchronize its clock. In this scenario, the OPC UA application will need to know the URL for a *Discovery Server* on a machine known to have the correct time. The OPC UA application or a separate background utility would call the *FindServers Service* and set its clock to the time specified in the *ResponseHeader*. This process will need to be repeated periodically because clocks can drift over time.

6.4 UTC and International Atomic Time (TAI)

All times in OPC UA are in UTC, however, UTC can include discontinuities due to leap seconds or repeating seconds added to deal with variations in the earth's orbit and rotation. *Servers* that have access to source for International Atomic Time (TAI) may choose to use this instead of UTC. That said, *Clients* always need to be prepared to deal with discontinuities due to the UTC or simply because the system clock is adjusted on the *Server* machine.

6.5 Issued User Identity Tokens

6.5.1 Kerberos

Kerberos *UserIdentityTokens* can be passed to the *Server* using the *IssuedIdentityToken*. The body of the token is an XML element that contains the WS-Security token as defined in the Kerberos Token Profile (Kerberos) specification.

Servers that support Kerberos authentication shall provide a *UserTokenPolicy* which specifies what version of the Kerberos Token Profile is being used, the Kerberos Realm and the Kerberos Principal Name for the *Server*. The Realm and Principal name are combined together with a simple syntax and placed in the *issuerEndpointUri* as shown in Table 39.

Table 39 – Kerberos UserTokenPolicy

Name	Description
tokenType	ISSUEDTOKEN_3
issuedTokenType	http://docs.oasis-open.org/wss/oasis-wss-kerberos-token-profile-1.1
issuerEndpointUri	A string with the form \\<realm>\<server principal name> where <realm> is the Kerberos realm name (e.g. Windows Domain); <server principal name> is the Kerberos principal name for the OPC UA Server.

The interface between the *Client* and *Server* applications and the Kerberos Authentication Service is application specific. The realm is the DomainName when using a Windows Domain controller as the Kerberos provider.

6.5.2 JSON Web Token (JWT)

JSON Web Token (JWT) *UserIdentityTokens* can be passed to the *Server* using the *IssuedIdentityToken*. The body of the token is a string that contains the JWT as defined in RFC 7159.

Servers that support JWT authentication shall provide a *UserTokenPolicy* which specifies the *Authorization Service* which provides the token and the parameters needed to access that service. The parameters are specified by a JSON object specified as the *issuerEndpointUri*. The contents of this JSON object are described in Table 41. The general *UserTokenPolicy* settings for JWT are defined in Table 40.

Table 40 – JWT UserTokenPolicy

Name	Description
tokenType	ISSUEDTOKEN_3
issuedTokenType	http://opcfoundation.org/UA/UserToken#JWT
issuerEndpointUri	For JWTs this is a JSON object with fields defined in Table 41.

Table 41 – JWT IssuerEndpointUrl Definition

Name	Type	Description
IssuerEndpointUrl	JSON object	Specifies the parameters for a JWT UserIdentityToken.
ua:resourceId	String	The URI identifying the <i>Server</i> to the <i>Authorization Service</i> . If not specified, the <i>Server's ApplicationUri</i> is used.
ua:authorityUrl	String	The base URL for the <i>Authorization Service</i> . This URL may be used to discover additional information about the authority. This field is equivalent to the "issuer" defined in OpenID-Discovery.
ua:authorityProfileUri	String	The profile that defines the interactions with the authority. If not specified, the URI is "http://opcfoundation.org/UA/Authorization#OAuth2".
ua:tokenEndpoint	String	A path relative to the base URL used to request <i>Access Tokens</i> . This field is equivalent to the "token_endpoint" defined in OpenID-Discovery.
ua:authorizationEndpoint	String	A path relative to the base URL used to validate user credentials. This field is equivalent to the "authorization_endpoint" defined in OpenID-Discovery.
ua:requestTypes	JSON array String	The list of request types supported by the authority. The possible values depend on the authorityProfileUri. IEC 62541-7 specifies the default for each authority profile defined.
ua:scopes	JSON array String	A list of Scopes that are understood by the <i>Server</i> . If not specified, the <i>Client</i> may be able to access any <i>Scope</i> supported by the <i>Authorization Service</i> . This field is equivalent to the "scopes_supported" defined in OpenID-Discovery.

6.5.3 OAuth2

6.5.3.1 General

The OAuth2 Authorization Framework (see RFC 6749) provides a web based mechanism to request claims based *Access Tokens* from an *Authorization Service* (AS) that is supported by many major companies providing cloud infrastructure. These *Access Tokens* are passed to a *Server* by a *Client* in a *UserIdentityToken* as described in IEC 62541-4.

The OpenID Connect specification (see OpenID) builds on the OAuth2 specification by defining the contents of the *Access Tokens* more strictly.

The OAuth2 specification supports a number of use cases (called 'flows') to handle different application requirements. The use cases that are relevant to OPC UA are discussed below.

6.5.3.2 Access Tokens

The JSON Web Token is the *Access Token* format which this document requires when using OAuth2. The JWT supports signatures using asymmetric cryptography which implies that *Servers* which accept the *Access Token* need to have access to the *Certificate* used by the *Authorization Service* (AS). The OpenID Connect Discovery specification is implemented by many AS products and provides a mechanism to fetch the AS *Certificate* via an HTTP request. If the AS does not support the discovery specification, then the signing *Certificate* will need to be provided to the *Server* when the location of the AS is added to the *Server* configuration.

Access Tokens expire and all *Servers* should revoke any privileges granted to the *Session* when the *Access Token* expires. If the *Server* allows for anonymous users, the *Server* should

allow the *Session* to stay open but treat it as an anonymous user. If the *Server* does not allow anonymous users, it should close the *Session* immediately.

Clients know when the *Access Token* will expire and should request a new *Access Token* and call *ActivateSession* before the old *Access Token* expires.

The JWT format allows the *Authorization Service* to insert any number of fields. The mandatory fields are defined in RFC 7159. Some additional fields are defined in Table 42 (see RFC 7523).

Table 42 – Access Token Claims

Field	Description
sub	The subject for the token. Usually the <i>client_id</i> which identifies the <i>Client</i> . If returned from an <i>Identity Provider</i> it may be a unique identifier for the user.
aud	The audience for the token. Usually the <i>resource_id</i> which identifies the <i>Server</i> or the <i>Server Application Uri</i> .
name	A human readable name for the <i>Client</i> application or user.
scp	A list of <i>Scopes</i> granted to the subject. Scopes apply to the <i>Access Token</i> and restrict how it may be used. Usually permissions or other restriction which limit access rights.
nonce	A nonce used to mitigate replay attacks. Shall be the value provided by the <i>Client</i> in the request.
groups	A list of groups which are assigned to the subject. Usually a list of unique identifiers for platform-specific security groups. For example, Azure AD user account groups may be returned in this claim.
roles	A list of roles which are assigned to the subject. Roles apply to the requestor and describe what the requestor can do with the resource. Usually a list of unique identifiers for roles known to the <i>Authorization Service</i> . These values are typically mapped to the <i>Roles</i> defined in IEC 62541-3.

6.5.3.3 Authorization Code

The authorization code flow is available to *Clients* which allow interaction with a human user. The *Client* application displays a window with a web browser which sends an HTTP GET to the *Identity Provider*. When the human user enters credentials that the *Identity Provider* validates, the *Identity Provider* returns an authorization code which is passed to the *Authorization Service*. The *Authorization Service* validates the code and returns an *Access Token* to the *Client*.

The complete flow is described in RFC 6749, 4.1.

A *requestType* of "authorization_code" in the *UserTokenPolicy* (see 6.5.2) means the *Authorization Service* supports the authorization code flow.

6.5.3.4 Refresh Token

The refresh token flow applies when a *Client* application has access to a refresh token returned in a previous response to an authorization code request. The refresh token allows applications to skip the step that requires human interaction with the *Identity Provider*. This flow is initiated when the *Client* sends the refresh token to *Authorization Service* which validates it and returns an *Access Token*. A *Client* that saves the refresh token for later use

shall use encryption or other means to ensure the refresh token cannot be accessed by unauthorized parties.

The complete flow is described in RFC 6749, Clause 6.

A *requestType* is not defined since support for refresh token is determined by checking the response to an authorization code request.

6.5.3.5 Client Credentials

The client credentials flow applies when a *Client* application cannot prompt a human user for input. This flow requires a secret known to the *Authorization Service* which the *Client* application can protect. This flow is initiated when the *Client* sends the *client_secret* to *Authorization Service* which validates it and returns an *Access Token*.

The complete flow is described in RFC 6749, 4.4.

A *requestType* of "client_credentials" in the *UserTokenPolicy* (see 6.5.2) means the *Authorization Service* supports the client credentials flow.

6.6 WS Secure Conversation

NOTE Deprecated in Version 1.03 because WS-SecureConversation has not been widely adopted by industry.

6.7 OPC UA Secure Conversation

6.7.1 Overview

OPC UA Secure Conversation (UASC) allows secure communication using binary encoded *Messages*.

UASC is designed to operate with different *TransportProtocols* that may have limited buffer sizes. For this reason, OPC UA Secure Conversation will break OPC UA *Messages* into several pieces (called '*MessageChunks*') that are smaller than the buffer size allowed by the *TransportProtocol*. UASC requires a *TransportProtocol* buffer size that is at least 8 192 bytes.

All security is applied to individual *MessageChunks* and not the entire OPC UA *Message*. A *Stack* that implements UASC is responsible for verifying the security on each *MessageChunk* received and reconstructing the original OPC UA *Message*.

All *MessageChunks* will have a 4-byte sequence assigned to them. These sequence numbers are used to detect and prevent replay attacks.

UASC requires a *TransportProtocol* that will preserve the order of *MessageChunks*, however, a UASC implementation does not necessarily process the *Messages* in the order that they were received.

6.7.2 MessageChunk structure

6.7.2.1 Overview

Figure 11 shows the structure of a *MessageChunk* and how security is applied to the *Message*.

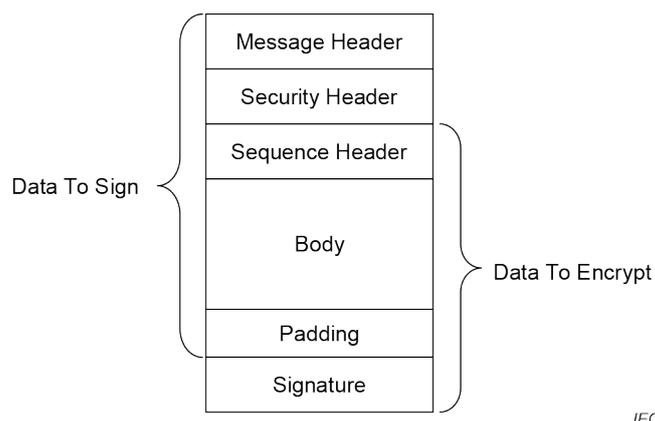


Figure 11 – OPC UA Secure Conversation MessageChunk

6.7.2.2 Message Header

Every *MessageChunk* has a *Message* header with the fields defined in Table 43.

Table 43 – OPC UA Secure Conversation Message header

Name	Data Type	Description
MessageType	Byte [3]	A three byte ASCII code that identifies the <i>Message</i> type. The following values are defined at this time: MSG A <i>Message</i> secured with the keys associated with a channel. OPN OpenSecureChannel <i>Message</i> . CLO CloseSecureChannel <i>Message</i> .
IsFinal	Byte	A one byte ASCII code that indicates whether the <i>MessageChunk</i> is the final chunk in a <i>Message</i> . The following values are defined at this time: C An intermediate chunk. F The final chunk. A The final chunk (used when an error occurred and the <i>Message</i> is aborted). This field is only meaningful for MessageType of 'MSG' This field is always 'F' for other MessageTypes.
MessageSize	UInt32	The length of the <i>MessageChunk</i> , in bytes. The length starts from the beginning of the MessageType field.
SecureChannelId	UInt32	A unique identifier for the <i>SecureChannel</i> assigned by the <i>Server</i> . If a <i>Server</i> receives a SecureChannelId which it does not recognize it shall return an appropriate transport layer error. When a <i>Server</i> starts the first <i>SecureChannelId</i> used should be a value that is likely to be unique after each restart. This ensures that a <i>Server</i> restart does not cause previously connected <i>Clients</i> to accidentally 'reuse' <i>SecureChannels</i> that did not belong to them.

6.7.2.3 Security Header

The *Message* header is followed by a security header which specifies what cryptography operations have been applied to the *Message*. There are two versions of the security header which depend on the type of security applied to the *Message*. The security header used for asymmetric algorithms is defined in Table 44. Asymmetric algorithms are used to secure the *OpenSecureChannel Messages*. PKCS #1 defines a set of asymmetric algorithms that may be used by UASC implementations. The *AsymmetricKeyWrapAlgorithm* element of the *SecurityPolicy* structure defined in Table 37 is not used by UASC implementations.

Table 44 – Asymmetric algorithm Security header

Name	Data Type	Description
SecurityPolicyUriLength	Int32	The length of the <i>SecurityPolicyUri</i> in bytes. This value shall not exceed 255 bytes. If a URI is not specified this value may be 0 or -1. Other negative values are invalid.
SecurityPolicyUri	Byte [*]	The URI of the <i>Security Policy</i> used to secure the <i>Message</i> . This field is encoded as a UTF-8 string without a null terminator.
SenderCertificateLength	Int32	The length of the <i>SenderCertificate</i> in bytes. This value shall not exceed <i>MaxSenderCertificateSize</i> bytes. If a certificate is not specified this value may be 0 or -1. Other negative values are invalid.
SenderCertificate	Byte [*]	The X.509 v3 <i>Certificate</i> assigned to the sending application <i>Instance</i> . This is a DER encoded blob. The structure of an X.509 v3 <i>Certificate</i> is defined in X.509 v3. The DER format for a <i>Certificate</i> is defined in X690 This indicates what <i>Private Key</i> was used to sign the <i>MessageChunk</i> . The <i>Stack</i> shall close the channel and report an error to the application if the <i>SenderCertificate</i> is too large for the buffer size supported by the transport layer. This field shall be null if the <i>Message</i> is not signed. If the <i>Certificate</i> is signed by a CA, the DER encoded CA <i>Certificate</i> may be appended after the <i>Certificate</i> in the byte array. If the CA <i>Certificate</i> is also signed by another CA this process is repeated until the entire <i>Certificate</i> chain is in the buffer or if <i>MaxSenderCertificateSize</i> limit is reached (the process stops after the last whole <i>Certificate</i> that can be added without exceeding the <i>MaxSenderCertificateSize</i> limit). Receivers can extract the <i>Certificates</i> from the byte array by using the <i>Certificate</i> size contained in DER header (see X.509 v3). Receivers that do not handle <i>Certificate</i> chains shall ignore the extra bytes.
ReceiverCertificateThumbprintLength	Int32	The length of the <i>ReceiverCertificateThumbprint</i> in bytes. If encrypted the length of this field is 20 bytes. If not encrypted the value may be 0 or -1. Other negative values are invalid.
ReceiverCertificateThumbprint	Byte [*]	The thumbprint of the X.509 v3 <i>Certificate</i> assigned to the receiving application <i>Instance</i> . The thumbprint is the <i>CertificateDigest</i> of the DER encoded form of the <i>Certificate</i> . This indicates what public key was used to encrypt the <i>MessageChunk</i> . This field shall be null if the <i>Message</i> is not encrypted.

The receiver shall close the communication channel if any of the fields in the security header have invalid lengths.

The *SenderCertificate*, including any chains, shall be small enough to fit into a single *MessageChunk* and leave room for at least one byte of body information. The maximum size for the *SenderCertificate* can be calculated with this formula:

```

MaxSenderCertificateSize =
  MessageChunkSize -
  12 - // Header size
  4 - // SecurityPolicyUriLength
  SecurityPolicyUri - // UTF-8 encoded string
  4 - // SenderCertificateLength
  4 - // ReceiverCertificateThumbprintLength
  20 - // ReceiverCertificateThumbprint
  8 - // SequenceHeader size
  1 - // Minimum body size
  1 - // PaddingSize if present
  Padding - // Padding if present
  ExtraPadding - // ExtraPadding if present
  AsymmetricSignatureSize // If present

```

The *MessageChunkSize* depends on the transport protocol but shall be at least 8 192 bytes. The *AsymmetricSignatureSize* depends on the number of bits in the public key for the *SenderCertificate*. The *Int32FieldLength* is the length of an encoded Int32 value and it is always 4 bytes.

The security header used for symmetric algorithms defined in Table 45. Symmetric algorithms are used to secure all *Messages* other than the *OpenSecureChannel Messages*. FIPS 197 define symmetric encryption algorithms that UASC implementations may use. FIPS 180-4 and HMAC define some symmetric signature algorithms.

Table 45 – Symmetric algorithm Security header

Name	Data Type	Description
TokenId	UInt32	A unique identifier for the <i>SecureChannel SecurityToken</i> used to secure the <i>Message</i> . This identifier is returned by the <i>Server</i> in an <i>OpenSecureChannel</i> response <i>Message</i> . If a <i>Server</i> receives a <i>TokenId</i> which it does not recognize it shall return an appropriate transport layer error.

6.7.2.4 Sequence header

The security header is always followed by the sequence header which is defined in Table 46. The sequence header ensures that the first encrypted block of every *Message* sent over a channel will start with different data.

Table 46 – Sequence header

Name	Data Type	Description
SequenceNumber	UInt32	A monotonically increasing sequence number assigned by the sender to each <i>MessageChunk</i> sent over the <i>SecureChannel</i> .
RequestId	UInt32	An identifier assigned by the <i>Client</i> to OPC UA request <i>Message</i> . All <i>MessageChunks</i> for the request and the associated response use the same identifier.

A *SequenceNumber* may not be reused for any *TokenId*. The *SecurityToken* lifetime should be short enough to ensure that this never happens; however, if it does the receiver shall treat it as a transport error and force a reconnect.

The *SequenceNumber* shall start at 1 023 and monotonically increase for all *Messages* and shall wrap around when it is equal to 4 294 966 271 (UInt32.MaxValue – 1 024). The first number after the wrap around shall be less than 1 024. Note that this requirement means that a *SequenceNumber* does not reset when a new *TokenId* is issued. The *SequenceNumber* shall be incremented by exactly one for each *MessageChunk* sent. For backward compatibility, receivers shall accept *SequenceNumbers* less than 1 023 and greater than 4 294 966 271 provided they are in sequence. Administrators shall be able to disable this

backward compatibility. Receivers shall log a warning when a rollover does not conform to the current specification.

The sequence header is followed by the *Message* body which is encoded with the OPC UA Binary encoding as described in 5.2.9. The body may be split across multiple *MessageChunks*.

6.7.2.5 Message footer

Each *MessageChunk* also has a footer with the fields defined in Table 47.

Table 47 – OPC UA Secure Conversation Message footer

Name	Data Type	Description
PaddingSize	Byte	The number of padding bytes (not including the byte for the PaddingSize).
Padding	Byte [*]	Padding added to the end of the <i>Message</i> to ensure length of the data to encrypt is an integer multiple of the encryption block size. The value of each byte of the padding is equal to PaddingSize.
ExtraPaddingSize	Byte	The most significant byte of a two-byte integer used to specify the padding size when the key used to encrypt the message chunk is larger than 2 048 bits. This field is omitted if the key length is less than or equal to 2 048 bits.
Signature	Byte [*]	The signature for the <i>MessageChunk</i> . The signature includes the all headers, all <i>Message</i> data, the PaddingSize and the Padding.

The formula to calculate the amount of padding depends on the amount of data that needs to be sent (called *BytesToWrite*). The sender shall first calculate the maximum amount of space available in the *MessageChunk* (called *MaxBodySize*) using the following formula:

$$\text{MaxBodySize} = \text{PlainTextBlockSize} * \text{Floor} ((\text{MessageChunkSize} - \text{HeaderSize} - 1) / \text{CipherTextBlockSize}) - \text{SequenceHeaderSize} - \text{SignatureSize};$$

The *HeaderSize* includes the *MessageHeader* and the *SecurityHeader*. The *SequenceHeaderSize* is always 8 bytes.

During encryption a block with a size equal to *PlainTextBlockSize* is processed to produce a block with size equal to *CipherTextBlockSize*. These values depend on the encryption algorithm and may be the same.

The OPC UA *Message* can fit into a single chunk if *BytesToWrite* is less than or equal to the *MaxBodySize*. In this case the *PaddingSize* is calculated with this formula:

$$\text{PaddingSize} = \text{PlainTextBlockSize} - ((\text{BytesToWrite} + \text{SignatureSize} + 1) \% \text{PlainTextBlockSize});$$

If the *BytesToWrite* is greater than *MaxBodySize* the sender shall write *MaxBodySize* bytes with a *PaddingSize* of 0. The remaining *BytesToWrite* – *MaxBodySize* bytes shall be sent in subsequent *MessageChunks*.

The *PaddingSize* and *Padding* fields are not present if the *MessageChunk* is not encrypted.

The *Signature* field is not present if the *MessageChunk* is not signed.

6.7.3 MessageChunks and error handling

MessageChunks are sent as they are encoded. *MessageChunks* belonging to the same *Message* shall be sent sequentially. If an error occurs creating a *MessageChunk* then the sender shall send a final *MessageChunk* to the receiver that tells the receiver that an error occurred and that it should discard the previous chunks. The sender indicates that the *MessageChunk* contains an error by setting the *IsFinal* flag to 'A' (for Abort). Table 48 specifies the contents of the *Message* abort *MessageChunk*.

Table 48 – OPC UA Secure Conversation Message abort body

Name	Data Type	Description
Error	UInt32	The numeric code for the error. This shall be one of the values listed in Table 57.
Reason	String	A more verbose description of the error. This string shall not be more than 4 096 bytes. A <i>Client</i> shall ignore strings that are longer than this.

The receiver shall check the security on the abort *MessageChunk* before processing it. If everything is ok, then the receiver shall ignore the *Message* but shall not close the *SecureChannel*. The *Client* shall report the error back to the application as *StatusCode* for the request. If the *Client* is the sender, then it shall report the error without waiting for a response from the *Server*.

6.7.4 Establishing a SecureChannel

Most *Messages* require a *SecureChannel* to be established. A *Client* does this by sending an *OpenSecureChannel* request to the *Server*. The *Server* shall validate the *Message* and the *ClientCertificate* and return an *OpenSecureChannel* response. Some of the parameters defined for the *OpenSecureChannel* service are specified in the security header (see 6.7.2) instead of the body of the *Message*. Table 49 lists the parameters that appear in the body of the *Message*.

Note that IEC 62541-4 is an abstract specification which defines interfaces that can work with any protocol. This document provides a concrete implementation for specific protocols. This document is the normative reference for all protocols and takes precedence if there are differences with IEC 62541-4.

Table 49 – OPC UA Secure Conversation OpenSecureChannel Service

Name	Data Type
Request	
RequestHeader	RequestHeader
ClientProtocolVersion	UInt32
RequestType	SecurityTokenRequestType
SecurityMode	MessageSecurityMode
ClientNonce	ByteString
RequestedLifetime	UInt32
Response	
ResponseHeader	ResponseHeader
ServerProtocolVersion	UInt32
SecurityToken	ChannelSecurityToken
SecureChannelId	UInt32
TokenId	UInt32
CreatedAt	UtcTime
RevisedLifetime	UInt32
ServerNonce	ByteString

The *ClientProtocolVersion* and *ServerProtocolVersion* parameters are not defined in IEC 62541-4 and are added to the *Message* to allow backward compatibility if OPC UA-SecureConversation needs to be updated in the future. Receivers always accept numbers greater than the latest version that they support. The receiver with the higher version number is expected to ensure backward compatibility.

If OPC UA-SecureConversation is used with the OPC UA-TCP protocol (see 7.1) then the version numbers specified in the *OpenSecureChannel Messages* shall be the same as the version numbers specified in the OPC UA-TCP protocol *Hello/Acknowledge Messages*. The receiver shall close the channel and report a *Bad_ProtocolVersionUnsupported* error if there is a mismatch.

The *Server* shall return an error response as described in IEC 62541-4 if there are any errors with the parameters specified by the *Client*.

The *RevisedLifetime* tells the *Client* when it shall renew the *SecurityToken* by sending another *OpenSecureChannel* request. The *Client* shall continue to accept the old *SecurityToken* until it receives the *OpenSecureChannel* response. The *Server* shall accept requests secured with the old *SecurityToken* until that *SecurityToken* expires or until it receives a *Message* from the *Client* secured with the new *SecurityToken*. The *Server* shall reject renew requests if the *SenderCertificate* is not the same as the one used to create the *SecureChannel* or if there is a problem decrypting or verifying the signature. The *Client* shall abandon the *SecureChannel* if the *Certificate* used to sign the response is not the same as the *Certificate* used to encrypt the request. Note that datatype is a *UInt32* value representing the number of milliseconds instead of the *Double (Duration)* defined in IEC 62541-4. This optimization is possible because sub-millisecond timeouts are not supported.

The *OpenSecureChannel Messages* are signed and encrypted if the *SecurityMode* is not *None* (even if the *SecurityMode* is *Sign*).

The *Nonces* shall be cryptographic random numbers with a length specified by the *SecureChannelNonceLength* of the *SecurityPolicy*.

See IEC TR 62541-2 for more information on the requirements for random number generators. The *OpenSecureChannel Messages* are not signed or encrypted if the *SecurityMode* is *None*. The *Nonces* are ignored and should be set to null. The *SecureChannelId* and the *TokenId* are still assigned but no security is applied to *Messages* exchanged via the channel. The *SecurityToken* shall still be renewed before the *RevisedLifetime* expires. Receivers shall still ignore invalid or expired *TokenIds*.

If the communication channel breaks the *Server* shall maintain the *SecureChannel* long enough to allow the *Client* to reconnect. The *RevisedLifetime* parameter also tells the *Client* how long the *Server* will wait. If the *Client* cannot reconnect within that period it shall assume the *SecureChannel* has been closed.

The *AuthenticationToken* in the *RequestHeader* shall be set to null.

If an error occurs after the *Server* has verified *Message* security it shall return a *ServiceFault* instead of a *OpenSecureChannel* response. The *ServiceFault Message* is described in IEC 62541-4.

If the *SecurityMode* is not *None* then the *Server* shall verify that a *SenderCertificate* and a *ReceiverCertificateThumbprint* were specified in the *SecurityHeader*.

6.7.5 Deriving keys

Once the *SecureChannel* is established the *Messages* are signed and encrypted with keys derived from the *Nonces* exchanged in the *OpenSecureChannel* call. These keys are derived by passing the *Nonces* to a pseudo-random function which produces a sequence of bytes from a set of inputs. A pseudo-random function is represented by the following function declaration:

```
Byte[] PRF(  
    Byte[] secret,  
    Byte[] seed,  
    Int32 length,  
    Int32 offset)
```

Where *length* is the number of bytes to return and *offset* is a number of bytes from the beginning of the sequence.

The lengths of the keys that need to be generated depend on the *SecurityPolicy* used for the channel. The following information is specified by the *SecurityPolicy*:

- a) *SigningKeyLength* (from the *DerivedSignatureKeyLength*);
- b) *EncryptingKeyLength* (implied by the *SymmetricEncryptionAlgorithm*);
- c) *EncryptingBlockSize* (implied by the *SymmetricEncryptionAlgorithm*).

The pseudo random function requires a secret and a seed. These values are derived from the *Nonces* exchanged in the *OpenSecureChannel* request and response. Table 50 specifies how to derive the secrets and seeds when using RSA based *SecurityPolicies*.

Table 50 – PRF inputs for RSA based SecurityPolicies

Name	Derivation
ClientSecret	The value of the <i>ClientNonce</i> provided in the <i>OpenSecureChannel</i> request.
ClientSeed	The value of the <i>ClientNonce</i> provided in the <i>OpenSecureChannel</i> request.
ServerSecret	The value of the <i>ServerNonce</i> provided in the <i>OpenSecureChannel</i> response.
ServerSeed	The value of the <i>ServerNonce</i> provided in the <i>OpenSecureChannel</i> response.

The parameters passed to the pseudo random function are specified in Table 51.

Table 51 – Cryptography key generation parameters

Key	Secret	Seed	Length	Offset
ClientSigningKey	ServerSecret	ClientSeed	SigningKeyLength	0
ClientEncryptingKey	ServerSecret	ClientSeed	EncryptingKeyLength	SigningKeyLength
ClientInitializationVector	ServerSecret	ClientSeed	EncryptingBlockSize	SigningKeyLength+ EncryptingKeyLength
ServerSigningKey	ClientSecret	ServerSeed	SigningKeyLength	0
ServerEncryptingKey	ClientSecret	ServerSeed	EncryptingKeyLength	SigningKeyLength
ServerInitializationVector	ClientSecret	ServerSeed	EncryptingBlockSize	SigningKeyLength+ EncryptingKeyLength

The *Client* keys are used to secure *Messages* sent by the *Client*. The *Server* keys are used to secure *Messages* sent by the *Server*.

The SSL/TLS specification defines a pseudo random function called P_HASH which is used for this purpose. The function is iterated until it produces enough data for all of the required keys. The Offset in Table 51 references to the offset from the start of the generated data.

The P_ hash algorithm is defined as follows:

$$P_HASH(secret, seed) = HMAC_HASH(secret, A(1) + seed) + \\ HMAC_HASH(secret, A(2) + seed) + \\ HMAC_HASH(secret, A(3) + seed) + \dots$$

Where A(n) is defined as:

$$A(0) = seed \\ A(n) = HMAC_HASH(secret, A(n-1))$$

+ indicates that the results are appended to previous results.

where 'HASH' is a hash function such as SHA256. The hash function to use depends on the *SecurityPolicyUri*.

6.7.6 Verifying Message security

The contents of the *MessageChunk* shall not be interpreted until the *Message* is decrypted and the signature and sequence number verified.

If an error occurs during *Message* verification the receiver shall close the communication channel. If the receiver is the *Server*, it shall also send a transport error *Message* before closing the channel. Once the channel is closed the *Client* shall attempt to re-open the channel and request a new *SecurityToken* by sending an *OpenSecureChannel* request. The mechanism for sending transport errors to the *Client* depends on the communication channel.

The receiver shall first check the *SecureChannelId*. This value may be 0 if the *Message* is an *OpenSecureChannel* request. For other *Messages*, it shall report a *Bad_SecureChannelUnknown* error if the *SecureChannelId* is not recognized. If the *Message*

is an *OpenSecureChannel* request and the *SecureChannelId* is not 0 then the *SenderCertificate* shall be the same as the *SenderCertificate* used to create the channel.

If the *Message* is secured with asymmetric algorithms, then the receiver shall verify that it supports the requested *SecurityPolicy*. If the *Message* is the response sent to the *Client*, then the *SecurityPolicy* shall be the same as the one specified in the request. In the *Server*, the *SecurityPolicy* shall be the same as the one used to originally create the *SecureChannel*. The receiver shall check that the Certificate is trusted first and return *Bad_CertificateUntrusted* on error. The receiver shall then verify the *SenderCertificate* using the rules defined in IEC 62541-4. The receiver shall report the appropriate error if *Certificate* validation fails. The receiver shall verify the *ReceiverCertificateThumbprint* and report a *Bad_CertificateInvalid* error if it does not recognize it.

If the *Message* is secured with symmetric algorithms, then a *Bad_SecureChannelTokenUnknown* error shall be reported if the *TokenId* refers to a *SecurityToken* that has expired or is not recognized.

If decryption or signature validation fails, then a *Bad_SecurityChecksFailed* error is reported. If an implementation allows multiple *SecurityModes* to be used the receiver shall also verify that the *Message* was secured properly as required by the *SecurityMode* specified in the *OpenSecureChannel* request.

After the security validation is complete the receiver shall verify the *RequestId* and the *SequenceNumber*. If these checks fail a *Bad_SecurityChecksFailed* error is reported. The *RequestId* only needs to be verified by the *Client* since only the *Client* knows if it is valid or not. If the *SequenceNumber* is not valid, the receiver shall log a *Bad_SequenceNumberInvalid* error.

At this point the *SecureChannel* knows it is dealing with an authenticated *Message* that was not tampered with or resent. This means the *SecureChannel* can return secured error responses if any further problems are encountered.

Stacks that implement UASC shall have a mechanism to log errors when invalid *Messages* are discarded. This mechanism is intended for developers, systems integrators and administrators to debug network system configuration issues and to detect attacks on the network.

7 TransportProtocols

7.1 OPC UA Connection Protocol

7.1.1 Overview

OPC UA Connection Protocol (UACP) is an abstract protocol that establishes a full duplex channel between a *Client* and *Server*. Concrete implementations of the UACP can be built with any middleware that supports full-duplex exchange of messages including TCP/IP and WebSockets. The term "*TransportConnection*" describes the middleware-specific connection used to exchange messages. For example, a socket is the *TransportConnection* for TCP/IP. *TransportConnections* allow responses to be returned in any order and allow responses to be returned on a different physical *TransportConnection* if communication failures cause temporary interruptions.

The OPC UA Connection Protocol is designed to work with the *SecureChannel* implemented by a layer higher in the stack. For this reason, the OPC UA Connection Protocol defines its interactions with the *SecureChannel* in addition to the wire protocol.

7.1.2 Message structure

7.1.2.1 Overview

Figure 12 illustrates the structure of a *Message* placed on the wire. This also illustrates how the *Message* elements defined by the OPC UA Binary Encoding mapping (see 5.2) and the OPC UA Secure Conversation mapping (see 6.7) relate to the OPC UA Connection Protocol *Messages*.

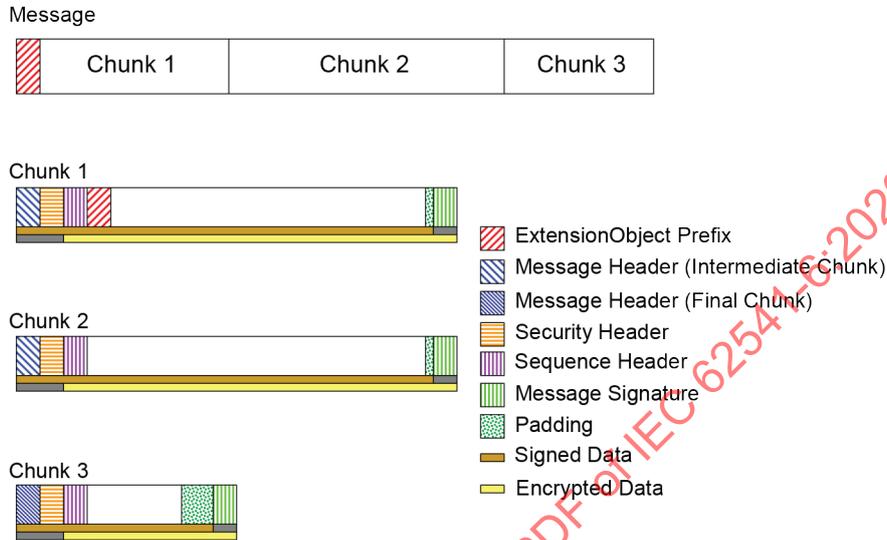


Figure 12 – OPC UA Connection Protocol Message structure

7.1.2.2 Message Header

Every OPC UA Connection Protocol *Message* has a header with the fields defined in Table 52.

Table 52 – OPC UA Connection Protocol Message header

Name	Type	Description
MessageType	Byte [3]	A three byte ASCII code that identifies the <i>Message</i> type. The following values are defined at this time: HEL a <i>Hello Message</i> . ACK an <i>Acknowledge Message</i> . ERR an <i>Error Message</i> . RHE a <i>ReverseHello Message</i> . The <i>SecureChannel</i> layer defines additional values which the OPC UA Connection Protocol layer shall accept.
Reserved	Byte [1]	Ignored. shall be set to the ASCII codes for 'F' if the <i>MessageType</i> is one of the values supported by the OPC UA Connection Protocol.
MessageSize	UInt32	The length of the <i>Message</i> , in bytes. This value includes the 8 bytes for the <i>Message</i> header.

The layout of the OPC UA Connection Protocol *Message* header is intentionally identical to the first 8 bytes of the OPC UA Secure Conversation *Message* header defined in Table 43. This allows the OPC UA Connection Protocol layer to extract the *SecureChannel Messages* from the incoming stream even if it does not understand their contents.

The OPC UA Connection Protocol layer shall verify the *MessageType* and make sure the *MessageSize* is less than the negotiated *ReceiveBufferSize* before passing any *Message* onto the *SecureChannel* layer.

7.1.2.3 Hello Message

The Hello *Message* has the additional fields shown in Table 53.

Table 53 – OPC UA Connection Protocol Hello Message

Name	Data Type	Description
ProtocolVersion	UInt32	The latest version of the UACP protocol supported by the <i>Client</i> . The <i>Server</i> may reject the <i>Client</i> by returning <i>Bad_ProtocolVersionUnsupported</i> . If the <i>Server</i> accepts the connection, it is responsible for ensuring that it returns <i>Messages</i> that conform to this version of the protocol. The <i>Server</i> shall always accept versions greater than what it supports.
ReceiveBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender can receive.
SendBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender will send.
MaxMessageSize	UInt32	The maximum size for any response <i>Message</i> . The <i>Server</i> shall abort the <i>Message</i> with a <i>Bad_ResponseTooLarge Error Message</i> if a response <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. The <i>Message</i> size is calculated using the unencrypted <i>Message</i> body. A value of zero indicates that the <i>Client</i> has no limit.
MaxChunkCount	UInt32	The maximum number of chunks in any response <i>Message</i> . The <i>Server</i> shall abort the <i>Message</i> with a <i>Bad_ResponseTooLarge Error Message</i> if a response <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. A value of zero indicates that the <i>Client</i> has no limit.
EndpointUrl	String	The URL of the <i>Endpoint</i> which the <i>Client</i> wished to connect to. The encoded value shall be less than 4 096 bytes. <i>Servers</i> shall return a <i>Bad_TcpEndpointUrlInvalid Error Message</i> and close the connection if the length exceeds 4 096 bytes or if it does not recognize the resource identified by the URL.

The *EndpointUrl* parameter is used to allow multiple *Servers* to share the same endpoint on a machine. The process listening (also known as the proxy) on the endpoint would connect to the *Server* identified by the *EndpointUrl* and would forward all *Messages* to the *Server* via this socket. If one socket closes, then the proxy shall close the other socket.

If the *Server* does not have sufficient resources to allow the establishment of a new *SecureChannel*, it shall immediately return a *Bad_TcpNotEnoughResources Error Message* and gracefully close the socket. *Clients* should not overload *Servers* that return this error by immediately trying to create a new *SecureChannel*.

7.1.2.4 Acknowledge Message

The *Acknowledge Message* has the additional fields shown in Table 54.

Table 54 – OPC UA Connection Protocol Acknowledge Message

Name	Type	Description
ProtocolVersion	UInt32	The latest version of the UACP protocol supported by the <i>Server</i> . If the <i>Client</i> accepts the connection, it is responsible for ensuring that it sends <i>Messages</i> that conform to this version of the protocol. The <i>Client</i> shall always accept versions greater than what it supports.
ReceiveBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender can receive. This value shall not be larger than what the <i>Client</i> requested in the Hello <i>Message</i> .
SendBufferSize	UInt32	The largest <i>MessageChunk</i> that the sender will send. This value shall not be larger than what the <i>Client</i> requested in the Hello <i>Message</i> .
MaxMessageSize	UInt32	The maximum size for any request <i>Message</i> . The <i>Client</i> shall abort the <i>Message</i> with a <i>Bad_RequestTooLarge StatusCode</i> if a request <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. The <i>Message</i> size is calculated using the unencrypted <i>Message</i> body. A value of zero indicates that the <i>Server</i> has no limit.
MaxChunkCount	UInt32	The maximum number of chunks in any request <i>Message</i> . The <i>Client</i> shall abort the <i>Message</i> with a <i>Bad_RequestTooLarge StatusCode</i> if a request <i>Message</i> exceeds this value. The mechanism for aborting <i>Messages</i> is described fully in 6.7.3. A value of zero indicates that the <i>Server</i> has no limit.

7.1.2.5 Error Message

The *Error Message* has the additional fields shown in Table 55.

Table 55 – OPC UA Connection Protocol Error Message

Name	Type	Description
Error	UInt32	The numeric code for the error. This shall be one of the values listed in Table 57.
Reason	String	A more verbose description of the error. This string shall not be more than 4 096 bytes. A <i>Client</i> shall ignore strings that are longer than this.

The socket is always closed gracefully by the *Client* after it receives an *Error Message*.

7.1.2.6 ReverseHello Message

The *ReverseHello Message* has the additional fields shown in Table 56.

Table 56 – OPC UA Connection Protocol ReverseHello Message

Name	Data Type	Description
ServerUri	String	The <i>ApplicationUri</i> of the <i>Server</i> which sent the <i>Message</i> . The encoded value shall be less than 4 096 bytes. <i>Client</i> shall return a <i>Bad_TcpEndpointUrlInvalid</i> error and close the connection if the length exceeds 4 096 bytes or if it does not recognize the <i>Server</i> identified by the URI.
EndpointUrl	String	The URL of the <i>Endpoint</i> which the <i>Client</i> uses when establishing the <i>SecureChannel</i> . This value shall be passed back to the <i>Server</i> in the <i>Hello Message</i> . The encoded value shall be less than 4 096 bytes. <i>Clients</i> shall return a <i>Bad_TcpEndpointUrlInvalid</i> error and close the connection if the length exceeds 4 096 bytes or if it does not recognize the resource identified by the URL. This value is a unique identifier for the <i>Server</i> which the <i>Client</i> may use to look up configuration information. It should be one of the URLs returned by the <i>GetEndpoints Service</i> .

For connection based protocols, such as TCP, the *ReverseHello Message* allows *Servers* behind firewalls with no open ports to connect to a *Client* and request that the *Client* establish a *SecureChannel* using the socket created by the *Server*.

For message based protocols the *ReverseHello Message* allows *Servers* to announce their presence to a *Client*. In this scenario, the *EndpointUrl* specifies the *Server's* protocol-specific address and any tokens required to access it.

7.1.3 Establishing a connection

Connections may be initiated by the *Client* or by the *Server* when they create a *TransportConnection* and establish a communication with their peer. If the *Client* creates the *TransportConnection*, the first *Message* sent shall be a *Hello* which specifies the buffer sizes that the *Client* supports. The *Server* shall respond with an *Acknowledge Message* which completes the buffer negotiation. The negotiated buffer size shall be reported to the *SecureChannel* layer. The negotiated *SendBufferSize* specifies the size of the *MessageChunks* to use for *Messages* sent over the connection.

If the *Server* creates the *TransportConnection* the first *Message* shall be a *ReverseHello* sent to the *Client*. If the *Client* accepts the connection, it sends a *Hello* message back to the *Server* which starts the buffer negotiation described for the *Client* initiated connection.

The *Hello/Acknowledge Messages* may only be sent once. If they are received again the receiver shall report an error and close the *TransportConnection*. Applications accepting incoming connections shall close any *TransportConnection* after a period of time if it does not receive a *Hello* or *ReverseHello Message*. This period of time shall be configurable and have a default value which does not exceed two minutes.

The *Client* sends the *OpenSecureChannel* request once it receives the *Acknowledge* back from the *Server*. If the *Server* accepts the new channel it shall associate the *TransportConnection* with the *SecureChannelId*. The *Server* uses this association to determine which *TransportConnection* to use when it needs to send a response to the *Client*. The *Client* does the same when it receives the *OpenSecureChannel* response.

The sequence of *Messages* when establishing a *Client* initiated OPC UA Connection Protocol connection is shown in Figure 13.

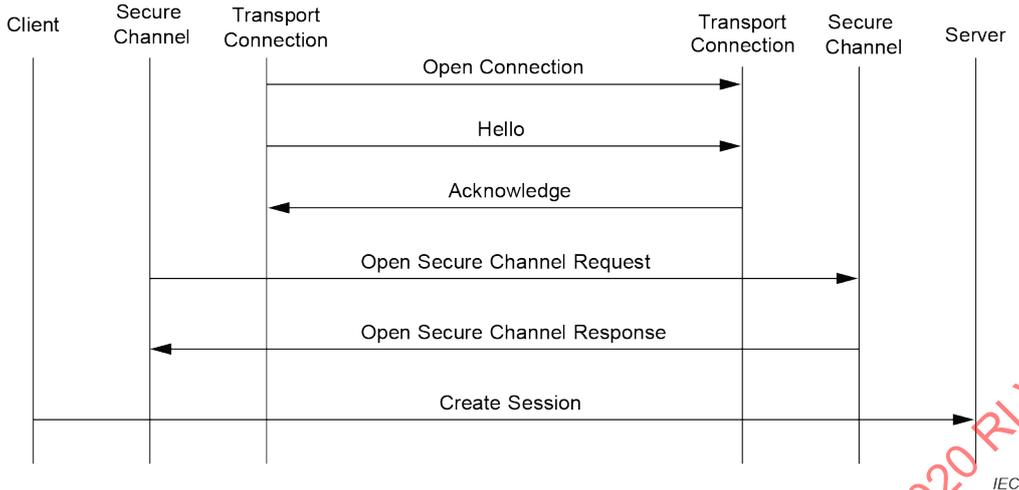


Figure 13 – Client initiated OPC UA Connection Protocol connection

The sequence of *Messages* when establishing a *Server* initiated OPC UA Connection Protocol connection is shown in Figure 14.

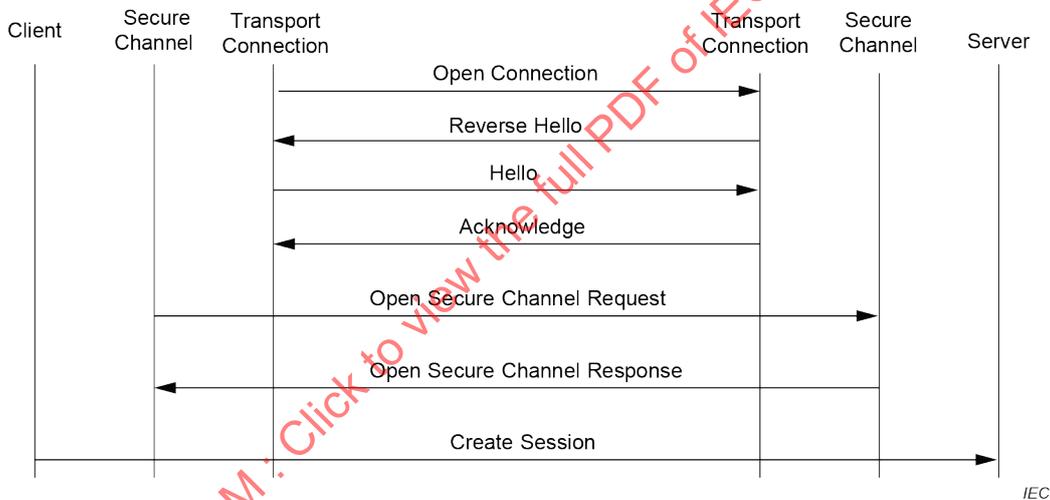


Figure 14 – Server initiated OPC UA Connection Protocol connection

The *Server* application does not do any processing while the *SecureChannel* is negotiated; however, the *Server* application shall to provide the *Stack* with the list of trusted *Certificates*. The *Stack* shall provide notifications to the *Server* application whenever it receives an *OpenSecureChannel* request. These notifications shall include the *OpenSecureChannel* or *Error* response returned to the *Client*.

The *Server* needs to be configured and enabled by an administrator to connect to one or more *Clients*. For each *Client*, the administrator shall provide an *ApplicationUri* and an *EndpointUrl* for the *Client*. If the *Client EndpointUrl* is not known, the administrator may provide the *EndpointUrl* for a GDS (see IEC 62541-12) which knows about the *Client*. The *Server* should expect that it will take some time for a *Client* to respond to a *ReverseHello*. Once a *Client* closes a *SecureChannel* or if the socket is closed without establishing a *SecureChannel*, the *Server* shall create a new socket and send a new *ReverseHello* message. Administrators may limit the number of simultaneous sockets that a *Server* will create.

7.1.4 Closing a connection

The *Client* closes the connection by sending a *CloseSecureChannel* request and closing the socket gracefully. When the *Server* receives this *Message*, it shall release all resources allocated for the channel. The body of the *CloseSecureChannel* request is empty. The *Server* does not send a *CloseSecureChannel* response.

If security verification fails for the *CloseSecureChannel Message*, then the *Server* shall report the error and close the socket.

The sequence of *Messages* when closing an OPC UA Connection Protocol connection is shown in Figure 15.

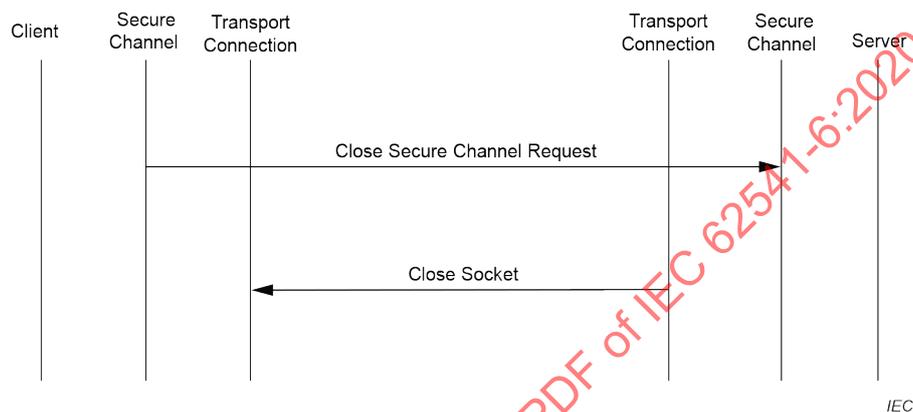


Figure 15 – Closing a OPC UA Connection Protocol connection

The *Server* application does not do any processing when the *SecureChannel* is closed; however, the *Stack* shall provide notifications to the *Server* application whenever a *CloseSecureChannel* request is received or when the *Stack* cleans up an abandoned *SecureChannel*.

7.1.5 Error handling

When a fatal error occurs, the *Server* shall send an *Error Message* to the *Client* and closes the *TransportConnection* gracefully. When the *Client* receives an *Error Message* it reports the error to the application and closes the *TransportConnection* gracefully. If a *Client* encounters a fatal error, it shall report the error to the application and send a *CloseSecureChannel Message*. The *Server* shall close the *TransportConnection* gracefully when it receives the *CloseSecureChannel Message*.

The possible OPC UA Connection Protocol errors are defined in Table 57.

Table 57 – OPC UA Connection Protocol error codes

Name	Description
Bad_TcpServerTooBusy	The <i>Server</i> cannot process the request because it is too busy. It is up to the <i>Server</i> to determine when it needs to return this <i>Message</i> . A <i>Server</i> can control the how frequently a <i>Client</i> reconnects by waiting to return this error.
Bad_TcpMessageTypeInvalid	The type of the <i>Message</i> specified in the header invalid. Each <i>Message</i> starts with a 4-byte sequence of ASCII values that identifies the <i>Message</i> type. The <i>Server</i> returns this error if the <i>Message</i> type is not accepted. Some of the <i>Message</i> types are defined by the <i>SecureChannel</i> layer.
Bad_TcpSecureChannelUnknown	The <i>SecureChannelId</i> and/or <i>TokenId</i> are not currently in use. This error is reported by the <i>SecureChannel</i> layer.
Bad_TcpMessageTooLarge	The size of the <i>Message</i> specified in the header is too large. The <i>Server</i> returns this error if the <i>Message</i> size exceeds its maximum buffer size or the receive buffer size negotiated during the Hello/Acknowledge exchange.
Bad_Timeout	A timeout occurred while accessing a resource. It is up to the <i>Server</i> to determine when a timeout occurs.
Bad_TcpNotEnoughResources	There are not enough resources to process the request. The <i>Server</i> returns this error when it runs out of memory or encounters similar resource problems. A <i>Server</i> can control the how frequently a <i>Client</i> reconnects by waiting to return this error.
Bad_TcpInternalError	An internal error occurred. This should only be returned if an unexpected configuration or programming error occurs.
Bad_TcpEndpointUrlInvalid	The <i>Server</i> does not recognize the <i>EndpointUrl</i> specified.
Bad_SecurityChecksFailed	The <i>Message</i> was rejected because it could not be verified.
Bad_RequestInterrupted	The request could not be sent because of a network interruption.
Bad_RequestTimeout	Timeout occurred while processing the request.
Bad_SecureChannelClosed	The secure channel has been closed.
Bad_SecureChannelTokenUnknown	The <i>SecurityToken</i> has expired or is not recognized.
Bad_CertificateUntrusted	The sender <i>Certificate</i> is not trusted by the receiver.
Bad_CertificateTimeInvalid	The sender <i>Certificate</i> has expired or is not yet valid.
Bad_CertificateIssuerTimeInvalid	The issuer for the sender <i>Certificate</i> has expired or is not yet valid.
Bad_CertificateUseNotAllowed	The sender's <i>Certificate</i> may not be used for establishing a secure channel.
Bad_CertificateIssuerUseNotAllowed	The issuer <i>Certificate</i> may not be used as a <i>Certificate Authority</i> .
Bad_CertificateRevocationUnknown	Could not verify the revocation status of the sender's <i>Certificate</i> .
Bad_CertificateIssuerRevocationUnknown	Could not verify the revocation status of the issuer <i>Certificate</i> .
Bad_CertificateRevoked	The sender <i>Certificate</i> has been revoked by the issuer.
Bad_IssuerCertificateRevoked	The issuer <i>Certificate</i> has been revoked by its issuer.
Bad_SequenceNumberInvalid	This sequence number on the message was not valid.

The numeric values for these error codes are defined in A.2.

NOTE The 'Tcp' prefix for some of the error codes in Table 57 was chosen when TCP/IP was the only implementation of the OPC UA Connection Protocol. These codes are used with any implementation of the OPC UA Connection Protocol.

7.2 OPC UA TCP

TCP/IP is a ubiquitous protocol that provides full-duplex communication between two applications. A socket is the *TransportConnection* in the TCP/IP implementation of the OPC UA Connection Protocol.

The URL scheme for endpoints using OPC UA TCP is 'opc.tcp'.

The *TransportProfileUri* shall be a URI for the TCP transport defined in IEC 62541-7.

7.3 SOAP/HTTP

NOTE Deprecated in Version 1.03 because WS-SecureConversation has not been widely adopted by industry.

7.4 OPC UA HTTPS

7.4.1 Overview

HTTPS refers HTTP *Messages* exchanged over a SSL/TLS connection (see HTTPS). The syntax of the HTTP *Messages* does not change and the only difference is a TLS connection is created instead of a TCP/IP connection. This implies that profiles which use this transport can also be used with HTTP when security is not a concern.

HTTPS is a protocol that provides transport security. This means all bytes are secured as they are sent without considering the *Message* boundaries. Transport security can only work for point to point communication and does not allow untrusted intermediaries or proxy servers to handle traffic.

The *SecurityPolicy* shall be specified, however, it only affects the algorithms used for signing the *Nonces* during the *CreateSession/ActivateSession* handshake. A *SecurityPolicy* of *None* indicates that the *Nonces* do not need to be signed. The *SecurityMode* is set to *Sign* unless the *SecurityPolicy* is *None*; in this case the *SecurityMode* shall be set to *None*. If a *UserIdentityToken* is to be encrypted, it shall be explicitly specified in the *UserTokenPolicy*.

An HTTP Header called 'OPCUA-SecurityPolicy' is used by the *Client* to tell the *Server* what *SecurityPolicy* it is using if there are multiple choices available. The value of the header is the URI for the *SecurityPolicy*. If the *Client* omits the header, then the *Server* shall assume a *SecurityPolicy* of *None*.

All HTTPS communications via a URL shall be treated as a single *SecureChannel* that is shared by multiple *Clients*. *Stacks* shall provide a unique identifier for the *SecureChannel* which allows applications correlate a request with a *SecureChannel*. This means that *Sessions* can only be considered secure if the *AuthenticationToken* (see IEC 62541-4) is long (>20 bytes) and HTTPS encryption is enabled.

The cryptography algorithms used by HTTPS have no relationship to the *EndpointDescription SecurityPolicy* and are determined by the policies set for HTTPS and are outside the scope of OPC UA.

Figure 16 illustrates a few scenarios where the HTTPS transport could be used.

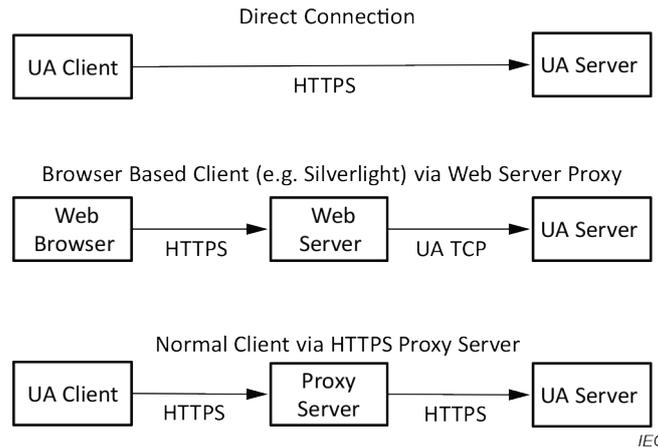


Figure 16 – Scenarios for the HTTPS Transport

In some scenarios, HTTPS communication will rely on an intermediary which is not trusted by the applications. If this is the case, then the HTTPS transport cannot be used to ensure security and the applications will need to establish a secure tunnel like a VPN before attempting any OPC UA related communication.

Applications which support the HTTPS transport shall support HTTP 1.1 and SSL/TLS 1.0.

Some HTTPS implementations require that all *Servers* have a *Certificate* with a Common Name (CN) that matches the DNS name of the *Server* machine. This means that a *Server* with multiple DNS names will need multiple HTTPS certificates. If multiple *Servers* are on the same machine they may share HTTPS certificates. This means that *ApplicationCertificates* are not the same as HTTPS *Certificates*. Applications which use the HTTPS transport and require application authentication shall check application *Certificates* during the *CreateSession/ActivateSession* handshake.

HTTPS *Certificates* can be automatically generated; however, this will cause problems for *Clients* operating inside a restricted environment such as a web browser. Therefore, HTTPS certificates should be issued by an authority which is accepted by all web browsers which need to access the *Server*. The set of *Certificate* authorities accepted by the web browsers is determined by the organization that manages the *Client* machines. *Client* applications that are not running inside a web may use the trust list that is used for application *Certificates*.

HTTPS connections have an unpredictable lifetime. Therefore, *Servers* need to rely on the *AuthenticationToken* passed in the *RequestHeader* to determine the identity of the *Client*. This means the *AuthenticationToken* shall be a randomly generated value with at least 32 bytes of data and HTTPS with signing and encryption shall always be used.

HTTPS allows *Clients* to have certificates; however, they are not required by the HTTPS transport. A *Server* shall allow *Clients* to connect without providing a *Certificate* during negotiation of the HTTPS connection.

HTTP 1.1 supports *Message* chunking where the Content-Length header in the request response is set to "chunked" and each chunk is prefixed by its size in bytes. All applications that support the HTTPS transport shall support HTTP chunking.

The URL scheme for endpoints using the HTTPS transport is 'opc.https'. Note that 'https' is the generic URL scheme for the underlying transport. The opc prefix specifies that the endpoint accepts OPC UA messages as defined in 7.4.

7.4.2 Session-less Services

Session-less Services (see IEC 62541-4) may be invoked via HTTPS POST. The HTTP *Authorization* header in the *Request* shall have a Bearer token which is an *Access Token* provided by the *Authorization Service*. The Content-type of the HTTP request shall specify the encoding of the body. If the Content-type is application/opcua+uabinary then the body is encoded using the OPC UA Binary encoding (see 7.4.4). If the Content-type is application/opcua+uajson then body is encoded using the reversible form of the JSON encoding (see 7.4.5).

Note that the Content-type for OPC UA Binary encoded bodies for Session-less Services is different from the Content-type for Session-based Services specified in 7.4.4.

7.4.3 XML Encoding

This *TransportProtocol* implements the OPC UA *Services* using a SOAP request-response message pattern over an HTTPS connection.

The body of the HTTP *Messages* shall be a SOAP 1.2 *Message* (see SOAP Part 1). WS-Addressing headers are optional (see WS Addressing).

The OPC UA XML Encoding specifies a way to represent an OPC UA *Message* as an XML element. This element is added to the SOAP *Message* as the only child of the SOAP body element. If an error occurs in the *Server* while parsing the request body, the *Server* may return a SOAP fault or it may return an OPC UA error response.

The SOAP Action associated with an XML encoded request *Message* always has the form:

```
http://opcfoundation.org/UA/2008/02/Services.wsdl/<service name>
```

where <service name> is the name of the OPC UA *Service* being invoked.

The SOAP Action associated with an XML encoded response *Message* always has the form:

```
http://opcfoundation.org/UA/2008/02/Services.wsdl/<service name>Response
```

All requests shall be HTTP POST requests. The Content-type shall be "application/soap+xml" and the charset and action parameters shall be specified. The charset parameter shall be "utf-8" and the action parameter shall be the URI for the SOAP action.

An example HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Content-Type: application/soap+xml; charset="utf-8";
    action="http://opcfoundation.org/UA/2008/02/Services.wsdl/Read"
Content-Length: nnnn
```

The action parameter appears on the same line as the Content-Type declaration.

An example request *Message*:

```
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope">
  <s:Body>
    <ReadRequest xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
      ...
    </ReadRequest>
  </s:Body>
</s:Envelope>
```

An example HTTP response header is:

```
HTTP/1.1 200 OK
Content-Type: application/soap+xml; charset="utf-8";
    action="http://opcfoundation.org/UA/2008/02/Services.wsdl/ReadResponse"
Content-Length: nnnn
```

The action parameter appears on the same line as the Content-Type declaration.

An example response *Message*:

```
<s:Envelope xmlns:s="http://www.w3.org/2003/05/soap-envelope">
  <s:Body>
    <ReadResponse xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
      ...
    </ReadResponse>
  </s:Body>
</s:Envelope>
```

7.4.4 OPC UA Binary Encoding

This *TransportProtocol* implements the OPC UA *Services* using an OPC UA Binary Encoded *Messages* exchanged over an HTTPS connection.

Applications which support the HTTPS *Profile* shall support HTTP 1.1.

The body of the HTTP *Messages* shall be OPC UA Binary encoded blob. The Content-type shall be "application/octet-stream".

An example HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Content-Type: application/octet-stream;
Content-Length: nnnn
```

An example HTTP response header is:

```
HTTP/1.1 200 OK
Content-Type: application/octet-stream;
Content-Length: nnnn
```

The *Message* body is the request or response structure encoded as an *ExtensionObject* in OPC UA Binary. The Authorization header is only used for Session-less Service calls (see 7.4.2).

If the OPC UA Binary Encoding is used for a Session-less *Service* the HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Authorization: Bearer <base64-encoded-token-data>
Content-Type: application/opcu+uabinary;
Content-Length: nnnn
```

7.4.5 JSON Encoding

This *TransportProtocol* implements the OPC UA *Services* using JSON encoded *Messages* exchanged over an HTTPS connection.

Applications which support the HTTPS *Profile* shall support HTTP 1.1.

The body of the HTTP *Messages* shall be OPC UA JSON Encoded. The Content-type shall be "application/opcu+uajson".

An example HTTP request header is:

```
POST /UA/SampleServer HTTP/1.1
Authorization: Bearer <base64-encoded-token-data>
Content-Type: application/opcua+uajson;
Content-Length: nnnn
```

An example HTTP response header is:

```
HTTP/1.1 200 OK
Content-Type: application/opcua+uajson;
Content-Length: nnnn
```

7.5 WebSockets

7.5.1 Overview

This *TransportProtocol* sends OPC UA Connection Protocol messages over WebSockets.

WebSockets is a bi-directional protocol for communication via a web server which is commonly used by browser based applications to allow the web server to asynchronously send information to the client. WebSockets uses the same default port as HTTP or HTTPS and initiates communication with an HTTP request. This makes it very useful in environments where firewalls limit traffic to the ports used by HTTP or HTTPS.

WebSockets uses HTTP, however, in practice a WebSocket connection is only initiated with a HTTP GET request and the web server provides an HTTP response. After that exchange, all traffic uses the binary framing protocol defined by RFC 6455.

A *Server* that supports the WebSockets transport shall publish one or more *Endpoints* with the scheme 'opc.wss'. The *TransportProfileUri* shall be one of the URIs for WebSockets transports defined in IEC 62541-7. The *TransportProfileUri* specifies the encoding and security protocol used to construct the OPC UA messages sent via the *WebSocket*.

The *SecurityMode* and *SecurityPolicyUri* of the *Endpoint* control the security applied to the messages sent via the *WebSocket*. This allows the messages to be secure even if the *WebSocket* connection is established via untrusted HTTPS proxies.

Figure 17 summarizes the complete process for establishing communication over a *WebSocket*.

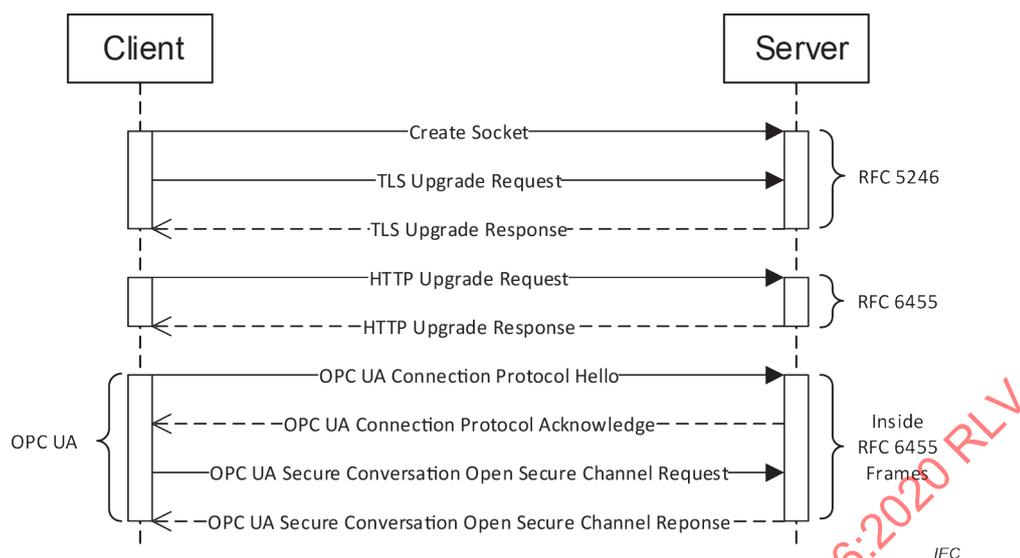


Figure 17 – Setting up Communication over a WebSocket

Figure 17 assumes the opcua+uacp protocol mapping (see 7.5.2).

7.5.2 Protocol Mapping

The WebSocket protocol allows clients to request that servers use specific sub-protocols with the "Sec-WebSocket-Protocol" header in the WebSocket handshake defined in RFC 6455. The protocols defined by this document are shown in Table 58.

Table 58 – WebSocket Protocols Mappings

Protocol	Description
opcua+uacp	Each WebSocket frame is a <i>MessageChunk</i> as defined in 6.7.2. After the WebSocket is created, the handshake described in 7.1.3 is used to negotiate the maximum size of the <i>MessageChunk</i> . The maximum size for a buffer needed to receive a WebSocket frame is the maximum length of a <i>MessageChunk</i> plus the maximum size for the WebSocket frame header. When using this protocol the payload in each frame is binary (OpCode 0x2 in RFC 6455).
opcua+uajson	Each WebSocket frame is a <i>Message</i> encoded using the JSON encoding described in 5.4.9. There is no mechanism to negotiate the maximum frame size. If the receiver encounters a frame that exceeds its internal limits, it shall close the WebSocket connection and provide a 1009 status code as described in RFC 6455. When using this protocol the payload in each frame is text (OpCode 0x1 in RFC 6455).

Each WebSocket protocol mapping defined has a *TransportProfileUri* defined in IEC 62541-7.

The *Client* shall request a protocol. If the *Server* does not support the protocol requested by the *Client*, the *Client* shall close the connection and report an error.

7.5.3 Security

WebSockets requires that the *Server* have a *Certificate*, however, the *Client* may have a *Certificate*. The *Server Certificate* should have the domain name as the common name component of the subject name; however, *Clients* that are able to override the *Certificate* validation procedure can choose to accept *Certificates* with a domain mismatch.

When using the WebSockets transport from a web browser, the browser environment may impose additional restrictions. For example, the web browser may require the *Server* have a valid TLS *Certificate* that is issued by CA that is installed in the *Trust List* for the web browser.

To support these *Clients*, a *Server* may use a *TLS Certificate* that does not conform to the requirements for an *ApplicationInstance Certificate*. In these cases, the *TLS Certificate* is only used for TLS negotiation and the *Server* shall use a valid *ApplicationInstance Certificate* for other interactions that require one. *Servers* shall allow administrators to specify a *Certificate* for use with TLS that is different from the *ApplicationInstance Certificate*.

Clients running in a browser environment specify the 'Origin' HTTP header during the *WebSocket* upgrade handshake. *Servers* should return the 'Access-Control-Allow-Origin' to indicate that the connection is allowed.

Any *Client* that does not run in a web browser environment and supports the *WebSockets* transport shall accept OPC UA *Application Instance Certificate* as the *TLS Certificate* provided the correct domain is specified in the *subjectAltName* field.

A *Client* may use its *Application Instance Certificate* as the *TLS Certificate* and *Servers* shall accept those *Certificates* if they are valid according to the OPC UA *Certificate* validation rules.

Some operating systems will not give the application any control over the set of algorithms that TLS will negotiate. In some cases, this set will be based on the needs of web browsers and will not be appropriate for the needs of an *OPC UA Application*. If this is a concern, applications should use OPC UA Secure Conversation in addition to TLS.

Clients that support the *WebSocket* transport shall support explicit configuration of an HTTPS proxy. When using an HTTPS proxy the *Client* shall first send an HTTP CONNECT message (see HTTP) before starting the *WebSocket* protocol handshake. Note that explicit HTTPS proxies allow for man-in-the-middle attacks. This threat may be mitigated by using OPC UA Secure Conversation in addition to TLS.

7.6 Well known addresses

The *Local Discovery Server* (LDS) is an OPC UA *Server* that implements the *Discovery Service Set* defined in IEC 62541-4. If an LDS is installed on a machine it shall use one or more of the well-known addresses defined in Table 59.

Table 59 – Well known addresses for Local Discovery Servers

Transport Mapping	URL	Notes
OPC UA TCP	opc.tcp://localhost:4840/UADiscovery	
OPC UA WebSockets	opc.wss://localhost:443/UADiscovery	
OPC UA HTTPS	https://localhost:443/UADiscovery	

OPC UA applications that make use of the LDS shall allow administrators to change the well-known addresses used within a system.

The *Endpoint* used by *Servers* to register with the LDS shall be the base address with the path "/registration" appended to it (e.g. http://localhost/UADiscovery/registration). OPC UA *Servers* shall allow administrators to configure the address to use for registration.

Each OPC UA *Server* application implements the *Discovery Service Set*. If the OPC UA *Server* requires a different address for this *Endpoint*, it shall create the address by appending the path "/discovery" to its base address.

8 Normative Contracts

8.1 OPC Binary Schema

The normative contract for the OPC UA Binary Encoded *Messages* is an OPC Binary Schema. This file defines the structure of all types and *Messages*. The syntax for an OPC Binary Type Schema is described in IEC 62541-3. This schema captures normative names for types and their fields as well the order the fields appear when encoded. The data type of each field is also captured.

8.2 XML Schema and WSDL

The normative contract for the OPC UA XML encoded *Messages* is an XML Schema. This file defines the structure of all types and *Messages*. This schema captures normative names for types and their fields as well the order the fields appear when encoded. The data type of each field is also captured.

The normative contract for *Message* sent via the SOAP/HTTP *TransportProtocol* is a WSDL that includes XML Schema for the OPC UA XML encoded *Messages*. It also defines the port types for OPC UA *Servers* and *DiscoveryServers*.

Links to the WSDL and XML Schema files can be found in Annex D.

8.3 Information Model Schema

Annex F defines the schema to be used for Information Models.

8.4 Formal definition of UA Information Model

Annex B defines the OPC UA NodeSet.

8.5 Constants

Annex A defines constants for Attribute Ids, Status Codes and numeric NodeIds.

8.6 DataType encoding

Annex C defines the binary encoding for all DataTypes and Messages.

8.7 Security configuration

Annex E defines a schema for security settings.

Annex A (normative)

Constants

A.1 Attribute Ids

Table A.1 shows Identifiers assigned to Attributes.

Table A.1 – Identifiers assigned to Attributes

Attribute	Identifier
NodId	1
NodeClass	2
BrowseName	3
DisplayName	4
Description	5
WriteMask	6
UserWriteMask	7
IsAbstract	8
Symmetric	9
InverseName	10
ContainsNoLoops	11
EventNotifier	12
Value	13
DataType	14
ValueRank	15
ArrayDimensions	16
AccessLevel	17
UserAccessLevel	18
MinimumSamplingInterval	19
Historizing	20
Executable	21
UserExecutable	22
DataTypeDefinition	23
RolePermissions	24
UserRolePermissions	25
AccessRestrictions	26
AccessLevelEx	27

A.2 Status Codes

Clause A.2 defines the numeric identifiers for all of the StatusCodes defined by the OPC UA Specification. The identifiers are specified in a CSV file with the following syntax:

<SymbolName>, <Code>, <Description>

Where the *SymbolName* is the literal name for the error code that appears in the specification and the *Code* is the hexadecimal value for the *StatusCode* (see IEC 62541-4). The severity associated with a particular code is specified by the prefix (*Good*, *Uncertain* or *Bad*).

The CSV released with this version of the standards can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/StatusCode.csv>

NOTE The latest CSV that is compatible with this version of the standard can be found here:

<http://www.opcfoundation.org/UA/schemas/StatusCode.csv>

A.3 Numeric Node Ids

Clause A.3 defines the numeric identifiers for all of the numeric *NodeIds* defined by the OPC UA Specification. The identifiers are specified in a CSV file with the following syntax:

<SymbolName>, <Identifier>, <NodeClass>

Where the *SymbolName* is either the *BrowseName* of a *Type Node* or the *BrowsePath* for an *Instance Node* that appears in the specification and the *Identifier* is numeric value for the *NodeId*.

The *BrowsePath* for an instance *Node* is constructed by appending the *BrowseName* of the instance *Node* to *BrowseName* for the containing instance or type. A '_' character is used to separate each *BrowseName* in the path. For example, IEC 62541-5 defines the *ServerType ObjectType Node* which has the *NamespaceArray Property*. The *SymbolName* for the *NamespaceArray InstanceDeclaration* within the *ServerType* declaration is: *ServerType_NamespaceArray*. IEC 62541-5 also defines a standard instance of the *ServerType ObjectType* with the *BrowseName* 'Server'. The *BrowseName* for the *NamespaceArray Property* of the standard *Server Object* is: *Server_NamespaceArray*.

The *NamespaceUri* for all *NodeIds* defined here is <http://opcfoundation.org/UA/>

The CSV released with this version of the standards can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/NodeIds.csv>

NOTE The latest CSV that is compatible with this version of the standard can be found here:

<http://www.opcfoundation.org/UA/schemas/NodeIds.csv>

Annex B (normative)

OPC UA Nodeset

The OPC UA NodeSet includes the complete Information Model defined in this document. It follows the XML Information Model schema syntax defined in Annex F and can thus be read and processed by a computer program.

The Information Model Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/Opc.Ua.NodeSet2.xml>

NOTE The latest Information Model schema that is compatible with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/Opc.Ua.NodeSet2.xml>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Annex C (normative)

Type declarations for the OPC UA native Mapping

Annex C defines the OPC UA Binary encoding for all *DataTypes* and *Messages* defined in this document. The schema used to describe the type is defined in IEC 62541-3.

The OPC UA Binary Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/Opc.Ua.Types.bsd.xml>

NOTE The latest file that is compatible with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/Opc.Ua.Types.bsd.xml>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Annex D (normative)

WSDL for the XML Mapping

D.1 XML Schema

Clause D.1 defines the XML Schema for all DataTypes and Messages defined in this series of OPC UA standards.

The XML Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/Opc.Ua.Types.xsd>

NOTE The latest file that is compatible with this document can be found here:

<http://www.opcfoundation.org/UA/2008/02/Types.xsd>

D.2 WSDL Port Types

Clause D.2 defines the WSDL Operations and Port Types for all Services defined in IEC 62541-4.

The WSDL released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/Opc.Ua.Services.wsdl>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2008/02/Services.wsdl>

This WSDL imports the XML Schema defined in D.1.

D.3 WSDL Bindings

Clause D.3 defines the WSDL Bindings for all Services defined in IEC 62541-4.

The WSDL released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/Opc.Ua.Endpoints.wsdl>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2008/02/Endpoints.wsdl>

This WSDL imports the WSDL defined in D.2.

Annex E (normative)

Security settings management

E.1 Overview

All OPC UA applications shall support security; however, this requirement means that Administrators need to configure the security settings for the OPC UA application. Annex E describes an XML Schema which can be used to read and update the security settings for an OPC UA application. All OPC UA applications may support configuration by importing/exporting documents that conform to the schema (called the *SecuredApplication* schema) defined in Annex E.

The XML Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/SecuredApplication.xsd>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2011/03/SecuredApplication.xsd>

The *SecuredApplication* schema can be supported in two ways:

- 1) Providing an XML configuration file that can be edited directly;
- 2) Providing an import/export utility that can be run as required;

If the application supports direct editing of an XML configuration file, then that file shall have exactly one element with the local name 'SecuredApplication' and URI equal to the *SecuredApplication* schema URI. A third party configuration utility shall be able to parse the XML file, read and update the 'SecuredApplication' element. The administrator shall ensure that only authorized administrators can update this file. The following is an example of a configuration that can be directly edited:

```
<s1:SampleConfiguration xmlns:s1="http://acme.com/UA/Sample/Configuration.xsd">
  <ApplicationName>ACME UA Server</ApplicationName>
  <ApplicationUri>urn:myfactory.com:Machine54:ACME UA Server</ApplicationUri>

  <!-- any number of application specific elements -->

  <SecuredApplication xmlns="http://opcfoundation.org/UA/2011/03/SecuredApplication.xsd">
    <ApplicationName>ACME UA Server</ApplicationName>
    <ApplicationUri>urn:myfactory.com:Machine54:ACME UA Server</ApplicationUri>
    <ApplicationType>Server_0</ApplicationType>
    <ApplicationCertificate>
      <StoreType>Windows</StoreType>
      <StorePath>LocalMachine\My</StorePath>
      <SubjectName>ACME UA Server</SubjectName>
    </ApplicationCertificate>
  </SecuredApplication>

  <!-- any number of application specific elements -->

  <DisableHiResClock>true</DisableHiResClock>
</s1:SampleConfiguration>
```

If an application provides an import/export utility, then the import/export file shall be a document that conforms to the *SecuredApplication* schema. The administrator shall ensure that only authorized administrators can run the utility. The following is an example of a file used by an import/export utility:

```

<?xml version="1.0" encoding="utf-8" ?>
<SecuredApplication xmlns="http://opcfoundation.org/UA/2011/03/SecuredApplication.xsd">
  <ApplicationName>ACME UA Server</ApplicationName>
  <ApplicationUri>urn:myfactory.com:Machine54:ACME UA Server</ApplicationUri>
  <ApplicationType>Server_0</ApplicationType>
  <ConfigurationMode>urn:acme.com:ACME Configuration Tool</ConfigurationMode>
  <LastExportTime>2011-03-04T13:34:12Z</LastExportTime>
  <ExecutableFile>%ProgramFiles%\ACME\Bin\ACME UA Server.exe</ExecutableFile>
  <ApplicationCertificate>
    <StoreType>Windows</StoreType>
    <StorePath>LocalMachine\My</StorePath>
    <SubjectName>ACME UA Server</SubjectName>
  </ApplicationCertificate>
  <TrustedCertificateStore>
    <StoreType>Windows</StoreType>
    <StorePath>LocalMachine\UA applications</StorePath>
    <!-- Offline CRL Checks by Default -->
    <ValidationOptions>16</ValidationOptions>
  </TrustedCertificateStore>
  <TrustedCertificates>
    <Certificates>
      <CertificateIdentifier>
        <SubjectName>CN=MyFactory CA</SubjectName>
        <!-- Online CRL Check for this CA -->
        <ValidationOptions>32</ValidationOptions>
      </CertificateIdentifier>
    </Certificates>
  </TrustedCertificates>
  <RejectedCertificatesStore>
    <StoreType>Directory</StoreType>
    <StorePath>%CommonApplicationData%\OPC Foundation\RejectedCertificates</StorePath>
  </RejectedCertificatesStore>
</SecuredApplication>

```

E.2 SecuredApplication

The *SecuredApplication* element specifies the security settings for an application. The elements contained in a *SecuredApplication* are described in Table E.1.

When an instance of a *SecuredApplication* is imported into an application the application updates its configuration based on the information contained within it. If unrecoverable errors occur during import an application shall not make any changes to its configuration and report the reason for the error.

The mechanism used to import or export the configuration depends on the application. Applications shall ensure that only authorized users are able to access this feature.

The *SecuredApplication* element may reference X.509 v3 Certificates which are contained in physical stores. Each application needs to decide whether it uses shared physical stores which the administrator can control directly by changing the location or private stores that can only be accessed via the import/export utility. If the application uses private stores, then the contents of these private stores shall be copied to the export file during export. If the import file references shared physical stores, then the import/export utility shall copy the contents of those stores to the private stores.

The import/export utility shall not export private keys. If the administrator wishes to assign a new public-private key to the application the administrator shall place the private in a store where it can be accessed by the import/export utility. The import/export utility is then responsible for ensuring it is securely moved to a location where the application can access it.

Table E.1 – SecuredApplication

Element	Type	Description
ApplicationName	String	A human readable name for the application. Applications shall allow this value to be read or changed.
ApplicationUri	String	A globally unique identifier for the instance of the application. Applications shall allow this value to be read or changed.
ApplicationType	ApplicationType	The type of application. May be one of <ul style="list-style-type: none"> • Server_0; • Client_1; • ClientAndServer_2; • DiscoveryServer_3; Application shall provide this value. Applications do not allow this value to be changed.
ProductName	String	A name for the product. Application shall provide this value. Applications do not allow this value to be changed.
ConfigurationMode	String	Indicates how the application should be configured. An empty or missing value indicates that the configuration file can be edited directly. The location of the configuration file shall be provided in this case. Any other value is a URI that identifies the configuration utility. The vendor documentation shall explain how to use this utility. Application shall provide this value. Applications do not allow this value to be changed.
LastExportTime	UtcTime	When the configuration was exported by the import/export utility. It may be omitted if applications allow direct editing of the security configuration.
ConfigurationFile	String	The full path to a configuration file used by the application. Applications do not provide this value if an import/export utility is used. Applications do not allow this value to be changed. Permissions set on this file shall control who has rights to change the configuration of the application.
ExecutableFile	String	The full path to an executable file for the application. Applications may not provide this value. Applications do not allow this value to be changed. Permissions set on this file shall control who has rights to launch the application.
ApplicationCertificate	CertificateIdentifier	The identifier for the <i>Application Instance Certificate</i> . Applications shall allow this value to be read or changed. This identifier may reference a <i>Certificate</i> store that contains the private key. If the private key is not accessible to outside applications this value shall contain the X.509 v3 <i>Certificate</i> for the application. If the configuration utility assigns a new private key this value shall reference the store where the private key is placed. The import/export utility may delete this private key if it moves it to a secure location accessible to the application. Applications shall allow Administrators to enter the password required to access the private key during the import operation. The exact mechanism depends on the application. Applications shall report an error if the ApplicationCertificate is not valid.

Element	Type	Description
TrustedCertificateStore	CertificateStore Identifier	<p>The location of the CertificateStore containing the Certificates of applications or <i>Certificate</i> Authorities (CAs) which can be trusted.</p> <p>Applications shall allow this value to be read or changed.</p> <p>This value shall be a reference to a physical store which can be managed separately from the application. Applications that support shared physical stores shall check this store for changes whenever they validate a <i>Certificate</i>.</p> <p>The Administrator is responsible for verifying the signature on all Certificates placed in this store. This means the application may trust Certificates in this store even if they cannot be verified back to a trusted root.</p> <p>Administrators shall place any CA certificates used to verify the signature in the IssuerStore or the IssuerList. This will allow applications to properly verify the signatures.</p> <p>The application shall check the revocation status of the Certificates in this store if the <i>Certificate</i> was issued by a CA. The application shall look for the offline <i>Certificate</i> Revocation List (CRL) for a CA in the store where it found the CA <i>Certificate</i>.</p> <p>The location of an online CRL for CA shall be specified with the CRLDistributionPoints (OID= 2.5.29.31) X.509 v3 <i>Certificate</i> extension.</p> <p>The ValidationOptions parameter is used to specify which revocation list should be used for CAs in this store.</p>
TrustedCertificates	CertificateList	<p>A list of Certificates for applications for CAs that can be trusted.</p> <p>Applications shall allow this value to be read or changed.</p> <p>The value is an explicit list of Certificates which is private to the application. It is used when the application does not support shared physical <i>Certificate</i> stores or when Administrators need to specify ValidationOptions for individual Certificates.</p> <p>If the TrustedCertificateStore and the TrustedCertificates parameters are both specified, then the application shall use the TrustedCertificateStore for checking trust relationships. The TrustedCertificates parameter is only used to lookup ValidationOptions for individual Certificates. It may also be used to provide CRLs for CA certificates.</p> <p>If the TrustedCertificateStore is not specified, then TrustedCertificates parameter shall contain the complete X.509 v3 <i>Certificate</i> for each entry.</p>
IssuerStore	CertificateStore Identifier	<p>The location of the CertificateStore containing CA Certificates which are not trusted but are needed to check signatures on Certificates.</p> <p>Applications shall allow this value to be read or changed.</p> <p>This value shall be a reference to a physical store which can be managed separately from the application. Applications that support shared physical stores shall check this store for changes whenever they validate a <i>Certificate</i>.</p> <p>This store may also contain CRLs for the CAs.</p>
IssuerCertificates	CertificateList	<p>A list of Certificates for CAs which are not trusted but are needed to check signatures on Certificates.</p> <p>Applications shall allow this value to be read or changed.</p> <p>The value is an explicit list of Certificates which is private to the application. It is used when the application does not support shared physical <i>Certificate</i> stores or when Administrators need to specify ValidationOptions for individual Certificates.</p> <p>If the IssuerStore and the IssuerCertificates parameters are both specified, then the application shall use the IssuerStore for checking signatures. The IssuerCertificates parameter is only used to lookup ValidationOptions for individual Certificates. It may also be used to provide CRLs for CA certificates.</p>

Element	Type	Description
RejectedCertificatesStore	CertificateStore Identifier	<p>The location of the shared CertificateStore containing the Certificates of applications which were rejected.</p> <p>Applications shall allow this value to be read or changed.</p> <p>Applications shall add the DER encoded <i>Certificate</i> into this store whenever it rejects a <i>Certificate</i> because it is untrusted or if it failed one of the validation rules which can be suppressed (see Clause E.6).</p> <p>Applications shall not add a <i>Certificate</i> to this store if it was rejected for a reason that cannot be suppressed (e.g. <i>Certificate</i> revoked).</p>
BaseAddresses	String []	<p>A list of URLs for the <i>Endpoints</i> supported by a <i>Server</i>.</p> <p>Applications shall allow these values to be read or changed.</p> <p>If a <i>Server</i> does not support the scheme for a URL it shall ignore it.</p> <p>This list can have multiple entries for the same URL scheme. The first entry for a scheme is the base URL. The rest are assumed to be DNS aliases that point to the first URL.</p> <p>It is the responsibility of the Administrator to configure the network to route these aliases correctly.</p>
SecurityProfileUris	SecurityProfile []	<p>A list of <i>SecurityPolicyUris</i> supported by a <i>Server</i>. The URIs are defined as security <i>Profiles</i> in IEC 62541-7.</p> <p>Applications shall allow these values to be read or changed.</p> <p>Applications shall allow the <i>Enabled</i> flag to be changed for each <i>SecurityProfile</i> that it supports.</p> <p>If the <i>Enabled</i> flag is false, the <i>Server</i> shall not allow connections using the <i>SecurityProfile</i>.</p> <p>If a <i>Server</i> does not support a <i>SecurityProfile</i> it shall ignore it.</p>
Extensions	xs:any	<p>A list of vendor defined Extensions attached to the security settings.</p> <p>Applications shall ignore Extensions that they do not recognize.</p> <p>Applications that update a file containing Extensions shall not delete or modify extensions that they do not recognize.</p>

E.3 CertificateIdentifier

The *CertificateIdentifier* element describes an X.509 v3 *Certificate*. The *Certificate* can be provided explicitly within the element or the element can specify the location of the *CertificateStore* that contains the *Certificate*. The elements contained in a *CertificateIdentifier* are described in Table E.2.

Table E.2 – CertificateIdentifier

Element	Type	Description
StoreType	String	The type of CertificateStore that contains the <i>Certificate</i> . Predefined values are "Windows" and "Directory". If not specified, the RawData element shall be specified.
StorePath	String	The path to the CertificateStore. The syntax depends on the StoreType. If not specified, the RawData element shall be specified.
SubjectName	String	The SubjectName for the <i>Certificate</i> . The Common Name (CN) component of the SubjectName. The SubjectName represented as a string that complies with Section 3 of RFC 4514. Values that do not contain '=' characters are presumed to be the Common Name component.
Thumbprint	String	The <i>CertificateDigest</i> for the <i>Certificate</i> formatted as a hexadecimal string. Case is not significant.
RawData	ByteString	The DER encoded <i>Certificate</i> . The CertificateIdentifier is invalid if the information in the DER <i>Certificate</i> conflicts with the information specified in other fields. Import utilities shall reject configurations containing invalid Certificates. This field shall not be specified if the StoreType and StorePath are specified.
ValidationOptions	Int32	The options to use when validating the <i>Certificate</i> . The possible options are described in E.6.
OfflineRevocationList	ByteString	A <i>Certificate</i> Revocation List (CRL) associated with an Issuer <i>Certificate</i> . The format of a CRL is defined by RFC 3280. This field is only meaningful for Issuer Certificates.
OnlineRevocationList	String	A URL for an Online Revocation List associated with an Issuer <i>Certificate</i> . This field is only meaningful for Issuer Certificates.

A "Windows" StoreType specifies a Windows *Certificate* store.

The syntax of the StorePath has the form:

[\\HostName\]StoreLocation[(ServiceName | UserSid)]\StoreName

where:

HostName – the name of the machine where the store resides.

StoreLocation – one of LocalMachine, CurrentUser, User or Service

ServiceName – the name of a Windows Service.

UserSid – the SID for a Windows user account.

StoreName – the name of the store (e.g. My, Root, Trust, CA, etc.).

Examples of Windows StorePaths are:

\\MYPC\LocalMachine\My

\CurrentUser\Trust

\\MYPC\Service\My UA Server\UA applications

\User\S-1-5-25\Root

A "Directory" StoreType specifies a directory on disk which contains files with DER encoded Certificates. The name of the file is the *CertificateDigest* for the *Certificate*. Only public keys may be placed in a "Directory" Store. The StorePath is an absolute file system path with a syntax that depends on the operating system.

If a "Directory" store contains a 'certs' subdirectory, then it is presumed to be a structured store with the subdirectories described in Table E.3.

Table E.3 – Structured directory store

Subdirectory	Description
certs	Contains the DER encoded X.509 v3 Certificates. The files shall have a .der file extension.
private	Contains the private keys. The format of the file may be application specific. PEM encoded files should have a .pem extension. PKCS#12 encoded files should have a .pfx extension. The root file name shall be the same as the corresponding public key file in the certs directory.
crl	Contains the DER encoded CRL for any CA Certificates found in the certs or ca directories. The files shall have a .crl file extension.

Each *Certificate* is uniquely identified by its Thumbprint. The SubjectName or the distinguished SubjectName may be used to identify a *Certificate* to a human; however, they are not unique. The SubjectName may be specified in conjunction with the Thumbprint or the RawData. If there is an inconsistency between the information provided, then the *CertificateIdentifier* is invalid. Invalid *CertificateIdentifiers* are handled differently depending on where they are used.

It is recommended that the SubjectName always be specified.

A *Certificate* revocation list (CRL) contains a list of certificates issued by a CA that are no longer trusted. These lists should be checked before an application can trust a *Certificate* issued by a trusted CA. The format of a CRL is defined by RFC 3280.

Offline CRLs are placed in a local *Certificate* store with the Issuer *Certificate*. Online CRLs may exist but the protocol depends on the system. An online CRL is identified by a URL.

E.4 CertificateStoreIdentifier

The *CertificateStoreIdentifier* element describes a physical store containing X.509 v3 Certificates. The elements contained in a *CertificateStoreIdentifier* are described in Table E.4.

Table E.4 – CertificateStoreIdentifier

Element	Type	Description
StoreType	String	The type of CertificateStore that contains the <i>Certificate</i> . Predefined values are "Windows" and "Directory".
StorePath	String	The path to the CertificateStore. The syntax depends on the StoreType. See E.3 for a description of the syntax for different StoreTypes.
ValidationOptions	Int32	The options to use when validating the Certificates contained in the store. The possible options are described in E.6.

All *Certificates* are placed in a physical store which can be protected from unauthorized access. The implementation of a store can vary and will depend on the application, development tool or operating system. A *Certificate* store may be shared by many applications on the same machine.

Each *Certificate* store is identified by a *StoreType* and a *StorePath*. The same path on different machines identifies a different store.

E.5 CertificateList

The *CertificateList* element is a list of *Certificates*. The elements contained in a *CertificateList* are described in Table E.5.

Table E.5 – CertificateList

Element	Type	Description
Certificates	CertificateIdentifier []	The list of Certificates contained in the Trust List
ValidationOptions	Int32	The options to use when validating the Certificates contained in the store. These options only apply to <i>Certificates</i> that have <i>ValidationOptions</i> with the <i>UseDefaultOptions</i> bit set. The possible options are described in E.6.

E.6 CertificateValidationOptions

The *CertificateValidationOptions* control the process used to validate a *Certificate*. Any *Certificate* can have validation options associated. If none are specified, the *ValidationOptions* for the store or list containing the *Certificate* are used. The possible options are shown in Table E.6. Note that suppressing any validation step can create security risks which are discussed in more detail in IEC TR 62541-2. An audit log entry shall be created if any error is ignored because a validation option is suppressed.

Table E.6 – CertificateValidationOptions

Field	Bit	Description
SuppressCertificateExpired	0	Ignore errors related to the validity time of the <i>Certificate</i> or its issuers.
SuppressHostNameInvalid	1	Ignore mismatches between the host name or <i>ApplicationUri</i> .
SuppressRevocationStatusUnknown	2	Ignore errors if the issuer's revocation list cannot be found.
CheckRevocationStatusOnline	3	<p>Check the revocation status online.</p> <p>If set the validator will look for the URL of the CRL Distribution Point in the <i>Certificate</i> and use the OCSP (RFC 6960) to determine if the <i>Certificate</i> has been revoked.</p> <p>If the CRL Distribution Point is not reachable then the validator will look for offline CRLs if the <i>CheckRevocationStatusOffline</i> bit is set. Otherwise, validation fails.</p> <p>This option is specified for Issuer <i>Certificates</i> and used when validating <i>Certificates</i> issued by that Issuer.</p>
CheckRevocationStatusOffline	4	<p>Check the revocation status offline.</p> <p>If set the validator will look a CRL in the <i>Certificate Store</i> where the CA <i>Certificate</i> was found.</p> <p>Validation fails if a CRL is not found.</p> <p>This option is specified for Issuer <i>Certificates</i> and used when validating <i>Certificates</i> issued by that <i>Issuer</i>.</p>
UseDefaultOptions	5	<p>If set the <i>CertificateValidationOptions</i> from the <i>CertificateList</i> shall be used.</p> <p>If a <i>Certificate</i> does not belong to a <i>CertificateList</i> then the default is 0 for all bits.</p>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 PLV

Annex F (normative)

Information Model XML Schema

F.1 Overview

Information Model developers define standard *AddressSpaces* which are implemented by many *Servers*. There is a need for a standard syntax that Information Model developers can use to formally define their models in a form that can be read by a computer program. Annex F defines an XML-based schema for this purpose.

The XML Schema released with this document can be found here:

<http://www.opcfoundation.org/UA/schemas/1.04/UANodeSet.xsd>

NOTE The latest file that is compatible with this document can be found here:

<http://opcfoundation.org/UA/2011/03/UANodeSet.xsd>

The schema document is the formal definition. The description in Annex F only discusses details of the semantics that cannot be captured in the schema document. Types which are self-describing are not discussed.

This schema can also be used to serialize (i.e. import or export) an arbitrary set of *Nodes* in the *Server Address Space*. This serialized form can be used to save *Server* state for use by the *Server* later or to exchange with other applications (e.g. to support offline configuration by a *Client*).

This schema only defines a way to represent the structure of *Nodes*. It is not intended to represent the numerous semantic rules which are defined in other parts of IEC 62541. Consumers of data serialized with this schema need to handle inputs that conform to the schema, however, do not conform to the OPC UA specification because of one or more semantic rule violations.

The tables defining the *DataTypes* in the specification have field names starting with a lowercase letter. The first letter shall be converted to upper case when the field names are formally defined in a *UANodeSet*.

F.2 UANodeSet

The *UANodeSet* is the root of the document. It defines a set of *Nodes*, their *Attributes* and *References*. *References* to *Nodes* outside of the document are allowed.

The structure of a *UANodeSet* is shown in Table F.1.

Table F.1 – UANodeSet

Element	Type	Description
NamespaceUris	UriTable	A list of <i>NamespaceUris</i> used in the <i>UANodeSet</i> .
ServerUris	UriTable	A list of <i>ServerUris</i> used in the <i>UANodeSet</i> .
Models	ModelTableEntry []	A list of Models that are defined in the <i>UANodeSet</i> along with any dependencies these models have.
ModelUri	String	The URI for the model. This URI should be one of the entries in the <i>NamespaceUris</i> table.
Version	String	The version of the model defined in the <i>UANodeSet</i> . This is a human readable string and not intended for programmatic comparisons.
PublicationDate	DateTime	When the model was published. This value is used for comparisons if the Model is defined in multiple <i>UANodeSet</i> files.
RolePermissions	RolePermissions []	The list of default <i>RolePermissions</i> for all <i>Nodes</i> in the model.
AccessRestrictions	AccessRestrictions	The default <i>AccessRestrictions</i> that apply to all <i>Nodes</i> in the model.
RequiredModels	ModelTableEntry []	A list of dependencies for the model. If the model requires a minimum version the <i>PublicationDate</i> shall be specified. Tools which attempt to resolve these dependencies may accept any <i>PublicationDate</i> after this date.
Aliases	AliasTable	A list of <i>Aliases</i> used in the <i>UANodeSet</i> .
Extensions	xs:any	An element containing any vendor defined extensions to the <i>UANodeSet</i> .
LastModified	DateTime	The last time a document was modified.
<choice>	UObject UVariable UAMethod UAView UObjectType UVariableType UADatatype UReferenceType	The <i>Nodes</i> in the <i>UANodeSet</i> .

The *NamespaceUris* is a list of URIs for namespaces used in the *UANodeSet*. The *NamespaceIndexes* used in *NodeId*, *ExpandedNodeIds* and *QualifiedNames* identify an element in this list. The first index is always 1 (0 is always the OPC UA namespace).

The *ServerUris* is a list of URIs for *Servers* referenced in the *UANodeSet*. The *ServerIndex* in *ExpandedNodeIds* identifies an element in this list. The first index is always 1 (0 is always the current *Server*).

The *Models* element specifies the Models which are formally defined by the *UANodeSet*. It includes version information as well as information about any dependencies which the model may have. If a Model is defined in the *UANodeSet* then the file shall also define an instance of the *NamespaceMetadataType ObjectType*. See IEC 62541-5 for more information.

The *Aliases* are a list of string substitutions for *NodeIds*. *Aliases* can be used to make the file more readable by allowing a string like 'HasProperty' in place of a numeric *NodeId* (i=46). *Aliases* are optional.

The *Extensions* are free form XML data that can be used to attach vendor defined data to the *UANodeSet*.

F.3 UANode

A *UANode* is an abstract base type for all *Nodes*. It defines the base set of *Attributes* and the *References*. There are subtypes for each *NodeClass* defined in IEC 62541-4. Each of these subtypes defines XML elements and attributes for the OPC UA *Attributes* specific to the *NodeClass*. The fields in the *UANode* type are defined in Table F.2.

Table F.2 – UANode

Element	Type	Description
Nodeld	Nodeld	A <i>Nodeld</i> serialized as a <i>String</i> . The syntax of the serialized <i>String</i> is defined in 5.3.1.10.
BrowseName	QualifiedName	A <i>QualifiedName</i> serialized as a <i>String</i> with the form: <namespace index>:<name> Where the <i>NamespaceIndex</i> refers to the <i>NamespaceUris</i> table.
SymbolicName	String	A symbolic name for the <i>Node</i> that can be used as a class/field name in auto generated code. It should only be specified if the <i>BrowseName</i> cannot be used for this purpose. This field does not appear in the <i>AddressSpace</i> and is intended for use by design tools. Only letters, digits or the underscore ('_') are permitted.
WriteMask	WriteMask	The value of the <i>WriteMask</i> Attribute.
UserWriteMask	WriteMask	Still in schema but no longer used.
AccessRestrictions	AccessRestrictions	The <i>AccessRestrictions</i> that apply to the <i>Node</i> .
DisplayName	LocalizedText []	A list of <i>DisplayNames</i> for the <i>Node</i> in different locales. There shall be only one entry per locale.
Description	LocalizedText []	The list of the <i>Descriptions</i> for the <i>Node</i> in different locales. There shall be only one entry per locale.
Category	String []	A list of identifiers used to group related <i>UANodes</i> together for use by tools that create/edit <i>UANodeSet</i> files.
Documentation	String	Additional non-localized documentation for use by tools that create/edit <i>UANodeSet</i> files.
References	Reference []	The list of <i>References</i> for the <i>Node</i> .
RolePermissions	RolePermissions []	The list of <i>RolePermissions</i> for the <i>Node</i> .
Extensions	xs:any	An element containing any vendor defined extensions to the <i>UANode</i> .

The *Extensions* are free form XML data that can be used to attach vendor defined data to the *UANode*.

Array values are denoted with [], however, in the XML Schema arrays are mapped to a complex type starting with the 'ListOf' prefix.

A *UANodeSet* is expected to contain many *UANodes* which reference each other. Tools that create *UANodeSets* should not add *Reference* elements for both directions in order to minimize the size of the XML file. Tools that read the *UANodeSets* shall automatically add reverse references unless reverse references are not appropriate given the *ReferenceType* semantics. *HasTypeDefinition* and *HasModellingRule* are two examples where it is not appropriate to add reverse references.

Note that a *UANodeSet* represents a collection of *Nodes* in an address space. This implies that any instances shall include the fully inherited *InstanceDeclarationHierarchy* as defined in IEC 62541-3.

F.4 Reference

The *Reference* type specifies a *Reference* for a *Node*. The *Reference* can be forward or inverse. Only one direction for each *Reference* needs to be in a *UANodeSet*. The other direction shall be added automatically during any import operation. The fields in the *Reference* type are defined in Table F.3.

Table F.3 – Reference

Element	Type	Description
Nodeld	Nodeld	The <i>Nodeld</i> of the target of the <i>Reference</i> serialized as a <i>String</i> . The syntax of the serialized <i>String</i> is defined in 5.3.1.11 (<i>ExpandedNodeld</i>). This value can be replaced by an <i>Alias</i> .
ReferenceType	Nodeld	The <i>Nodeld</i> of the <i>ReferenceType</i> serialized as a <i>String</i> . The syntax of the serialized <i>String</i> is defined in 5.3.1.10 (<i>Nodeld</i>). This value can be replaced by an <i>Alias</i> .
IsForward	Boolean	If TRUE, the <i>Reference</i> is a forward reference.

F.5 RolePermission

The *RolePermission* type specifies the *Permissions* granted to *Role* for a *Node*. The fields in the *RolePermission* type are defined in Table F.4.

Table F.4 – RolePermission

Element	Type	Description
Nodeld	Nodeld	The <i>Nodeld</i> of the <i>Role</i> which has the <i>Permissions</i> .
Permissions	UInt32	A bitmask specifying the <i>Permissions</i> granted to the <i>Role</i> . The bitmask values the <i>Permissions</i> bits defined in IEC 62541-3.

F.6 UAType

A *UAType* is a subtype of the *UANode* defined in F.3. It is the base type for the types defined in Table F.5.

Table F.5 – UANodeSet Type Nodes

Subtype	Description
UAObjectType	Defines an <i>ObjectType Node</i> as described in IEC 62541-3.
UAVariableType	Defines a <i>VariableType Node</i> as described in IEC 62541-3.
UADataType	Defines a <i>DataType Node</i> as described in IEC 62541-3.
UATypeReferenceType	Defines a <i>ReferenceType Node</i> as described in IEC 62541-3.

F.7 UAIInstance

A *UAIInstance* is a subtype of the *UANode* defined in F.3. It is the base type for the types defined in Table F.6. The fields in the *UAIInstance* type are defined in Table F.7. Subtypes of *UAIInstance* which have fields in addition to those defined in IEC 62541-3 are described in detail below.

Table F.6 – UANodeSet Instance Nodes

Subtype	Description
UAObject	Defines an <i>Object Node</i> as described in IEC 62541-3.
UAVariable	Defines a <i>Variable Node</i> as described in IEC 62541-3.
UAMethod	Defines a <i>Method Node</i> as described in IEC 62541-3.
UAView	Defines a <i>View Node</i> as described in IEC 62541-3.

Table F.7 – UAIInstance

Element	Type	Description
All of the fields from the <i>UANode</i> type described in F.3.		
ParentNodeId	NodeId	The <i>NodeId</i> of the <i>Node</i> that is the parent of the <i>Node</i> within the information model. This field is used to indicate that a tight coupling exists between the <i>Node</i> and its parent (e.g. when the parent is deleted the child is deleted as well). This information does not appear in the <i>AddressSpace</i> and is intended for use by design tools.

F.8 UAVariable

A *UAVariable* is a subtype of the *UAIInstance* defined in. It represents a Variable Node. The fields in the *UAVariable* type are defined in Table F.8.

Table F.8 – UVariable

Element	Type	Description
All of the fields from the <i>UInstance</i> type described in F.7.		
Value	Variant	The Value of the Node encoding using the UA XML wire encoding.
Translation	TranslationType []	A list of translations for the Value if the Value is a LocalizedText or a structure containing LocalizedTexts. This field may be omitted. If the Value is an array the number of elements in this array shall match the number of elements in the Value. Extra elements are ignored. If the Value is a scalar, then there is one element in this array. If the Value is a structure, then each element contains translations for one or more fields identified by a name. See the TranslationType for more information.
DataType	Nodeld	The data type of the value.
ValueRank	ValueRank	The value rank. If not specified, the default value is -1 (Scalar).
ArrayDimensions	ArrayDimensions	The number of dimensions in an array value.
AccessLevel	AccessLevel	The access level.
UserAccessLevel	AccessLevel	Still in schema but no longer used.
MinimumSamplingInterval	Duration	The minimum sampling interval.
Historizing	Boolean	Whether history is being archived.

F.9 UAMethod

A *UAMethod* is a subtype of the *UInstance* defined in F.7. It represents a Method Node. The fields in the *UAMethod* type are defined in Table F.9.

Table F.9 – UAMethod

Element	Type	Description
All of the fields from the <i>UInstance</i> type described in F.7.		
MethodDeclarationId	Nodeld	May be specified for <i>Method Nodes</i> that are a target of a <i>HasComponent</i> reference from a single <i>Object Node</i> . It is the <i>Nodeld</i> of the <i>UAMethod</i> with the same <i>BrowseName</i> contained in the <i>TypeDefinition</i> associated with the <i>Object Node</i> . If the <i>TypeDefinition</i> overrides a <i>Method</i> inherited from a base <i>ObjectType</i> then this attribute shall reference the <i>Method Node</i> in the subtype.
UserExecutable	Boolean	Still in schema but no longer used.
ArgumentDescription	UAMethodArgument []	A list of <i>Descriptions</i> for the <i>Method Node Arguments</i> . Each entry has a <i>Name</i> which uniquely identifies the <i>Argument</i> that the <i>Descriptions</i> apply to. There shall only be one entry per <i>Name</i> . Each entry also has a list of <i>Descriptions</i> for the <i>Argument</i> in different locales. There shall be only one entry per locale per <i>Argument</i> .

F.10 TranslationType

A *TranslationType* contains additional translations for *LocalizedTexts* used in the *Value* of a *Variable*. The fields in the *TranslationType* are defined in Table F.10. If multiple *Arguments* existed there would be a Translation element for each *Argument*.

The type can have two forms depending on whether the *Value* is a *LocalizedText* or a *Structure* containing *LocalizedTexts*. If it is a *LocalizedText* it contains a simple list of translations. If it is a *Structure*, it contains a list of fields which each contain a list of translations. Each field is identified by a Name which is unique within the structure. The mapping between the Name and the *Structure* requires an understanding of the *Structure* encoding. If the *Structure* field is encoded as a *LocalizedText* with UA XML, then the name is the unqualified path to the XML element where names in the path are separated by '/'. For example, a structure with a nested structure containing a *LocalizedText* could have a path like "Server/ApplicationName".

The following example illustrates how translations for the Description field in the *Argument Structure* are represented in XML:

```
<Value>
  <ListOfExtensionObject xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
    <ExtensionObject>
      <TypeId>
        <Identifier>i=297</Identifier>
      </TypeId>
      <Body>
        <Argument>
          <Name>ConfigData</Name>
          <DataType>
            <Identifier>i=15</Identifier>
          </DataType>
          <ValueRank>-1</ValueRank>
          <ArrayDimensions />
          <Description>
            <Text>[English Translation for Description]</Text>
          </Description>
        </Argument>
      </Body>
    </ExtensionObject>
  </ListOfExtensionObject>
</Value>
<Translation>
  <Field Name="Description">
    <Text Locale="de-DE">[German Translation for Description]</Text>
    <Text Locale="fr-FR">[French Translation for Description]</Text>
  </Field>
</Translation>
```

If multiple Arguments existed there would be a Translation element for each Argument.

Table F.10 – TranslationType

Element	Type	Description
Text	LocalizedText []	An array of translations for the Value. It only appears if the <i>Value</i> is a <i>LocalizedText</i> or an array of <i>LocalizedText</i> .
Field	StructureTranslationType []	An array of structure fields which have translations. It only appears if the <i>Value</i> is a <i>Structure</i> or an array of <i>Structures</i> .
Name	String	The name of the field. This uniquely identifies the field within the structure. The exact mapping depends on the encoding of the structure.
Text	LocalizedText []	An array of translations for the structure field.

F.11 UADatatype

A *UADatatype* is a subtype of the *UAType* defined in F.6. It defines a *Data Type Node*. The fields in the *UADatatype* type are defined in Table F.11.

Table F.11 – UADatatype

Element	Type	Description
All of the fields from the <i>UANode</i> type described in F.3.		
Definition	<i>DataTypeDefinition</i>	An abstract definition of the data type that can be used by design tools to create code that can serialize the data type in XML and/or Binary forms. It does not appear in the <i>AddressSpace</i> . This is only used to define subtypes of the <i>Structure</i> or <i>Enumeration DataTypes</i> .

F.12 DataTypeDefinition

A *DataTypeDefinition* defines an abstract representation of a *UADatatype* that can be used by design tools to automatically create serialization code. The fields in the *DataTypeDefinition* type are defined in Table F.12.

Table F.12 – DataTypeDefinition

Element	Type	Description
Name	<i>QualifiedName</i>	A unique name for the data type. This name should be the same as the <i>BrowseName</i> or the containing <i>DataType</i> .
SymbolicName	String	A symbolic name for the data type that can be used as a class/structure name in autogenerated code. It should only be specified if the <i>Name</i> cannot be used for this purpose. Only letters, digits or the underscore ('_') are permitted and the first character shall be a letter. This field is only specified for nested <i>DataTypeDefinitions</i> . The <i>SymbolicName</i> of the <i>DataType Node</i> is used otherwise.
BaseType	<i>QualifiedName</i>	Not used. Kept in schema for backward compatibility.
IsUnion	Boolean	This flag indicates if the data type represents a union. Only one of the Fields defined for the data type is encoded into a value. This field is optional. The default value is false. If this value is true, the first field is the switch value.
IsOptionSet	Boolean	This flag indicates that the data type defines the <i>OptionSetValues Property</i> . This field is optional. The default value is false.
Fields	<i>DataTypeField []</i>	The list of fields that make up the data type. This definition assumes the structure has a sequential layout. For enumerations, the fields are simply a list of values. This list does not include fields inherited from a base data type.

F.13 DataTypeField

A *DataTypeField* defines an abstract representation of a field within a *UADatatype* that can be used by design tools to automatically create serialization code. The fields in the *DataTypeField* type are defined in Table F.13.

Table F.13 – DataTypeField

Element	Type	Description
Name	String	A name for the field that is unique within the <i>DataTypeDefinition</i> .
SymbolicName	String	A symbolic name for the field that can be used in autogenerated code. It should only be specified if the <i>Name</i> cannot be used for this purpose. Only letters, digits or the underscore ('_') are permitted.
DisplayName	LocalizedText []	A display name for the field in multiple locales.
DataType	NodeId	The <i>NodeId</i> of the <i>DataType</i> for the field. This <i>NodeId</i> can refer to another <i>Node</i> with its own <i>DataTypeDefinition</i> . This field is not specified for subtypes of <i>Enumeration</i> .
ValueRank	Int32	The value rank for the field. It shall be <i>Scalar</i> (-1) or a fixed rank <i>Array</i> (≥ 1). This field is not specified for subtypes of <i>Enumeration</i> .
ArrayDimensions	String	The maximum length of an array. This field is a comma separated list of unsigned integer values. The list has a number of elements equal to the <i>ValueRank</i> . The value is 0 if the maximum is not known for a dimension. This field is not specified if the <i>ValueRank</i> ≤ 0 . This field is not specified for subtypes of <i>Enumeration</i> or for <i>DataTypes</i> with the <i>OptionSetValues Property</i> .
MaxStringLength	UInt32	The maximum length of a <i>String</i> or <i>ByteString</i> value. If not known the value is 0. The value is 0 if the <i>DataType</i> is not <i>String</i> or <i>ByteString</i> . If the <i>ValueRank</i> > 0 the maximum applies to each element in the array. This field is not specified for subtypes of <i>Enumeration</i> or for <i>DataTypes</i> with the <i>OptionSetValues Property</i> .
Description	LocalizedText []	A description for the field in multiple locales.
Value	Int32	The value associated with the field. This field is only specified for subtypes of <i>Enumeration</i> and <i>OptionSet DataTypes</i> . For <i>OptionSets</i> the value is the number of the bit associated with the field.
IsOptional	Boolean	The field indicates if a data type field in a structure is optional. This field is optional. The default value is false. This field is not specified for subtypes of <i>Enumeration</i> and <i>Union</i> .

F.14 Variant

The *Variant* type specifies the value for a *Variable* or *VariableType Node*. This type is the same as the type defined in 5.3.1.17. As a result, the functions used to serialize *Variants* during *Service* calls can be used to serialize *Variant* in this file syntax.

Variants can contain *NodeIds*, *ExpandedNodeIds* and *QualifiedNames* which shall be modified so the *NamespaceIndexes* and *ServerIndexes* reference the *NamespaceUri* and *ServerUri* tables in the *UANodeSet*.

Variants can also contain *ExtensionObjects* which contain an *EncodingId* and a *Structure* with fields which could be *NodeIds*, *ExpandedNodeIds* or *QualifiedNames*. The *NamespaceIndexes* and *ServerIndexes* in these fields shall also reference the tables in the *UANodeSet*.

F.15 Example

An example of the *UANodeSet* can be found below.

This example defines the *Nodes* for an *InformationModel* with the URI of "http://sample.com/Instances". This example references *Nodes* defined in the base OPC UA *InformationModel* and an *InformationModel* with the URI "http://sample.com/Types".

The XML namespaces declared at the top include the URIs for the *Namespaces* referenced in the document because the document includes *Complex Data*. Documents without *Complex Data* would not have these declarations.

```
<UANodeSet
xmlns:s1="http://sample.com/Instances"
xmlns:s0="http://sample.com/Types"
xmlns:uax="http://opcfoundation.org/UA/2008/02/Types.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns="http://opcfoundation.org/UA/2011/03/UANodeSet.xsd">
```

The *NamespaceUris* table includes all *Namespaces* referenced in the document except for the base OPC UA *InformationModel*. A *NamespaceIndex* of 1 refers to the URI "http://sample.com/Instances".

```
<NamespaceUris>
  <Uri>http://sample.com/Instances</Uri>
  <Uri>http://sample.com/Types</Uri>
</NamespaceUris>
```

The *Aliases* table is provided to enhance readability. There are no rules for what is included. A useful guideline would include standard *ReferenceTypes* and *DataTypes* if they are referenced in the document.

```
<Aliases>
  <Alias Alias="HasComponent">i=47</Alias>
  <Alias Alias="HasProperty">i=46</Alias>
  <Alias Alias="HasSubtype">i=45</Alias>
  <Alias Alias="HasTypeDefinition">i=40</Alias>
</Aliases>
```

The *BicycleType* is a *DataType Node* that inherits from a *DataType* defined in another *InformationModel* (ns=2;i=314). It is assumed that any application importing this file will already know about the referenced *InformationModel*. A *Server* could map the references onto another OPC UA *Server* by adding a *ServerIndex* to *TargetNode NodeIds*. The structure of the *DataType* is defined by the *Definition* element. This information can be used by code generators to automatically create serializers for the *DataType*.

```
<UADatatype NodeId="ns=1;i=365" BrowseName="1:BicycleType">
  <DisplayName>BicycleType</DisplayName>
  <References>
    <Reference ReferenceType="HasSubtype" IsForward="false">ns=2;i=314</Reference>
  </References>
  <Definition Name="BicycleType">
    <Field Name="NoOfGears" DataType="UInt32" />
    <Field Name="ManufacturerName" DataType="QualifiedName" />
  </Definition>
</UADatatype>
```

This *Node* is an instance of an *Object TypeDefinition Node* defined in another *InformationModel* (ns=2;i=341). It has a single *Property* which is declared later in the document.

```
<UAObject NodeId="ns=1;i=375" BrowseName="1:DriverOfTheMonth" ParentNodeId="ns=1;i=281">
```

```

<DisplayName>DriverOfTheMonth</DisplayName>
<References>
  <Reference ReferenceType="HasProperty">ns=1;i=376</Reference>
  <Reference ReferenceType="HasTypeDefinition">ns=2;i=341</Reference>
  <Reference ReferenceType="HasComponent" IsForward="false">ns=1;i=281</Reference>
</References>
</UAObject>

```

This *Node* is an instance of a *Variable TypeDefinition Node* defined in base OPC UA *InformationModel* (i=68). The *DataType* is the base type for the *BicycleType DataType*. The *AccessLevels* declare the *Variable* as *Readable* and *Writeable*. The *ParentNodeId* indicates that this *Node* is tightly coupled with the Parent (*DriverOfTheMonth*) and will be deleted if the Parent is deleted.

```

<UAVariable NodeId="ns=1;i=376" BrowseName="2:PrimaryVehicle"
  ParentNodeId="ns=1;i=375" DataType="ns=2;i=314" AccessLevel="3" UserAccessLevel="3">
  <DisplayName>PrimaryVehicle</DisplayName>
  <References>
    <Reference ReferenceType="HasTypeDefinition">i=68</Reference>
    <Reference ReferenceType="HasProperty" IsForward="false">ns=1;i=375</Reference>
  </References>

```

This *Value* is an instance of a *BicycleType DataType*. It is wrapped in an *ExtensionObject* which declares that the value is serialized using the *Default XML DataTypeEncoding* for the *DataType*. The *Value* could be serialized using the *Default Binary DataTypeEncoding* but that would result in a document that cannot be edited by hand. No matter which *DataTypeEncoding* is used, the *NamespaceIndex* used in the *ManufactureName* field refers to the *NamespaceUri* table in this document. The application is responsible for changing whatever value it needs to be when the document is loaded by an application.

```

<Value>
  <ExtensionObject xmlns="http://opcfoundation.org/UA/2008/02/Types.xsd">
    <TypeId>
      <Identifier>ns=1;i=366</Identifier>
    </TypeId>
    <Body>
      <s1:BicycleType>
        <s0:Make>Trek</s0:Make>
        <s0:Model>Compact</s0:Model>
        <s1:NoOfGears>10</s1:NoOfGears>
        <s1:ManufactureName>
          <uax:NamespaceIndex>1</uax:NamespaceIndex>
          <uax:Name>Hello</uax:Name>
        </s1:ManufactureName>
      </s1:BicycleType>
    </Body>
  </ExtensionObject>
</Value>
</UAVariable>

```

These are the *DataTypeEncoding Nodes* for the *BicycleType DataType*.

```

<UAObject NodeId="ns=1;i=366" BrowseName="Default XML">
  <DisplayName>Default XML</DisplayName>
  <References>
    <Reference ReferenceType="HasEncoding" IsForward="false">ns=1;i=365</Reference>
    <Reference ReferenceType="HasDescription">ns=1;i=367</Reference>
    <Reference ReferenceType="HasTypeDefinition">i=76</Reference>
  </References>
</UAObject>
<UAObject NodeId="ns=1;i=370" BrowseName="Default Binary">
  <DisplayName>Default Binary</DisplayName>
  <References>
    <Reference ReferenceType="HasEncoding" IsForward="false">ns=1;i=365</Reference>
    <Reference ReferenceType="HasDescription">ns=1;i=371</Reference>
    <Reference ReferenceType="HasTypeDefinition">i=76</Reference>
  </References>
</UAObject>

```

This is the *DataTypeDescription Node* for the *Default XML DataTypeEncoding* of the *BicycleType DataType*. The *Value* is one of the built-in types.

```
<UAVariable NodeId="ns=1;i=367" BrowseName="1:BicycleType" DataType="String">
  <DisplayName>BicycleType</DisplayName>
  <References>
    <Reference ReferenceType="HasTypeDefinition">i=69</Reference>
    <Reference ReferenceType="HasComponent" IsForward="false">ns=1;i=341</Reference>
  </References>
  <Value>
    <uax:String>//xs:element[@name='BicycleType']</uax:String>
  </Value>
</UAVariable>
```

This is the *DataTypeDictionary Node* for the *DataTypeDescription* declared above. The XML Schema document is a UTF-8 document stored as *xs:base64Binary* value (see Base64). This allows *Clients* to read the schema for all *DataTypes* which belong to the *DataTypeDictionary*. The value of *DataTypeDescription Node* for each *DataType* contains a XPath query that will find the correct definition inside the schema document.

```
<UAVariable NodeId="ns=1;i=341" BrowseName="1:Quickstarts.DataTypes.Instances"
DataType="ByteString">
  <DisplayName>Quickstarts.DataTypes.Instances</DisplayName>
  <References>
    <Reference ReferenceType="HasProperty">ns=1;i=343</Reference>
    <Reference ReferenceType="HasComponent">ns=1;i=367</Reference>
    <Reference ReferenceType="HasComponent" IsForward="false">i=92</Reference>
    <Reference ReferenceType="HasTypeDefinition">i=72</Reference>
  </References>
  <Value>
    <uax:ByteString>PHhz...WlhPg==</uax:ByteString>
  </Value>
</UAVariable>
```

F.16 UANodeSetChanges

The *UANodeSetChanges* is the root of a document that contains a set of changes to an *AddressSpace*. It is expected that a single file will contain either a *UANodeSet* or a *UANodeSetChanges* element at the root. It provides a list of *Nodes/References* to add and/or a list *Nodes/References* to delete. The *UANodeSetChangesStatus* structure defined in F.22 is produced when a *UANodeSetChanges* document is applied to an *AddressSpace*.

The elements of the type are defined in Table F.14.

Table F.14 – UANodeSetChanges

Element	Type	Description
NamespaceUri	UriTable	Same as described in Table F.1.
ServerUri	UriTable	Same as described in Table F.1.
Models	ModelTableEntry []	Same as described in Table F.1.
Aliases	AliasTable	Same as described in Table F.1.
Extensions	xs:any	Same as described in Table F.1.
LastModified	DateTime	Same as described in Table F.1.
NodesToAdd	NodesToAdd	A list of new <i>Nodes</i> to add to the <i>AddressSpace</i> .
ReferencesToAdd	ReferencesToChange	A list of new <i>References</i> to add to the <i>AddressSpace</i> .
NodesToDelete	NodesToDelete	A list of <i>Nodes</i> to delete from the <i>AddressSpace</i> .
ReferencesToDelete	ReferencesToChange	A list of <i>References</i> to delete from the <i>AddressSpace</i> .

The *Models* element specifies the version of one or more *Models* which the *UANodeSetChanges* file will create when it is applied to an existing *Address Space*. The *UANodeSetChanges* cannot be applied if the current version of the *Model* in the *Address Space* is higher. The *RequiredModels* sub-element (see Table F.1) specifies the versions of

Models which shall already exist before the *UANodeSetChanges* file can be applied. When checking dependencies, the version of the *Model* in the existing Address Space shall exactly match the required version.

If a *UANodeSetChanges* file modifies types and there are existing instances of the types in the AddressSpace, then the *Server* shall automatically modify the instances to conform to the new type or generate an error.

A *UANodeSetChanges* file is processed as a single operation. This allows mandatory *Nodes* or *References* to be replaced by specifying a *Node/Reference* to delete and a *Node/Reference* to add.

F.17 NodesToAdd

The *NodesToAdd* type specifies a list of *Nodes* to add to an *AddressSpace*. The structure of these *Nodes* is defined by the *UANodeSet* type in Table F.1.

The elements of the type are defined in Table F.15.

Table F.15 – NodesToAdd

Element	Type	Description
<choice>	UAObject UAVariable UAMethod UAView UAObjectType UAVariableType UADataType UAReferenceType	The <i>Nodes</i> to add to the <i>AddressSpace</i> .

When adding *Nodes*, *References* can be specified as part of the *Node* definition or as a separate *ReferencesToAdd*.

Note that *References* to *Nodes* that could exist are always allowed. In other words, a *Node* is never rejected simply because it has a reference to an unknown *Node*.

Reverse *References* are added automatically when deemed practical by the processor.

F.18 ReferencesToChange

The *ReferencesToChange* type specifies a list of *References* to add to or remove from an *AddressSpace*.

The elements of the type are defined in Table F.16.

Table F.16 – ReferencesToChange

Element	Type	Description
Reference	ReferenceToChange	A <i>Reference</i> to add to the <i>AddressSpace</i> .

F.19 ReferenceToChange

The *ReferenceToChange* type specifies a single *Reference* to add to or remove from an *AddressSpace*.

The elements of the type are defined in Table F.17.

Table F.17 – ReferencesToChange

Element	Type	Description
Source	Nodeld	The identifier for the source <i>Node</i> of the <i>Reference</i> .
ReferenceType	Nodeld	The identifier for the type of the <i>Reference</i> .
IsForward	Boolean	TRUE if the <i>Reference</i> is a forward reference.
Target	Nodeld	The identifier for the target <i>Node</i> of the <i>Reference</i> .

References to *Nodes* that could exist are always allowed. In other words, a *Reference* is never rejected simply because the target is unknown *Node*.

The source of the *Reference* shall exist in the *AddressSpace* or in the *UANodeSetChanges* document being processed.

Reverse *References* are added when deemed practical by the processor.

F.20 NodesToDelete

The *NodesToDelete* type specifies a list of *Nodes* to remove from an *AddressSpace*.

The elements of the type are defined in Table F.18.

Table F.18 – NodesToDelete

Element	Type	Description
Node	NodeToDelete	A <i>Node</i> to delete from the <i>AddressSpace</i> .

F.21 NodeToDelete

The *NodeToDelete* type specifies a *Node* to remove from an *AddressSpace*.

The elements of the type are defined in Table F.19.

Table F.19 – ReferencesToChange

Element	Type	Description
Node	Nodeld	The identifier for the <i>Node</i> to delete.
DeleteReverseReferences	Boolean	If TRUE, then <i>References</i> to the <i>Node</i> are deleted as well.

F.22 UANodeSetChangesStatus

The *UANodeSetChangesStatus* is the root of a document that is produced when a *UANodeSetChanges* document is processed.

The elements of the type are defined in Table F.20.

Table F.20 – UANodeSetChangesStatus

Element	Type	Description
NamespaceUris	UriTable	Same as described in Table F.1.
ServerUris	UriTable	Same as described in Table F.1.
Aliases	AliasTable	Same as described in Table F.1.
Extensions	xs:any	Same as described in Table F.1.
Version	String	Same as described in Table F.1.
LastModified	DateTime	Same as described in Table F.1.
TransactionId	String	A globally unique identifier from the original <i>UANodeSetChanges</i> document.
NodesToAdd	NodeSetStatusList	A list of results for the <i>NodesToAdd</i> specified in the original document. The list is empty if all elements were processed successfully.
ReferencesToAdd	NodeSetStatusList	A list of results for the <i>ReferencesToAdd</i> specified in the original document. The list is empty if all elements were processed successfully.
NodesToDelete	NodeSetStatusList	A list of results for the <i>NodesToDelete</i> specified in the original document. The list is empty if all elements were processed successfully.
ReferencesToDelete	NodeSetStatusList	A list of results for the <i>ReferencesToDelete</i> specified in the original document. The list is empty if all elements were processed successfully.

F.23 NodeSetStatusList

The *NodeSetStatusList* type specifies a list of results produced when applying a *UANodeSetChanges* document to an *AddressSpace*.

If no errors occurred this list is empty.

If one or more errors occur, then this list contains one element for each operation specified in the original document.

The elements of the type are defined in Table F.21.

Table F.21 – NodeSetStatusList

Element	Type	Description
Result	NodeSetStatus	The result of a single operation.

F.24 NodeSetStatus

The *NodeSetStatus* type specifies a single result produced when applying an operation specified in a *UANodeSetChanges* document to an *AddressSpace*.

The elements of the type are defined in Table F.22.

Table F.22 – NodeSetStatus

Element	Type	Description
Code	StatusCode	The result of the operation. The possible StatusCodes are defined in IEC 62541-4.
Details	String	A string providing information that is not conveyed by the StatusCode. This is not a human readable string for the StatusCode.

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Bibliography

X200: ISO/IEC 7498-1 (ITU-T Rec. X.200), *Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model*

WS Addressing: Web Services Addressing (WS-Addressing)
<http://www.w3.org/Submission/ws-addressing/>

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

SOMMAIRE

AVANT-PROPOS	124
1 Domaine d'application	127
2 Références normatives	127
3 Termes, définitions, termes abrégés et symboles	129
3.1 Termes et définitions	129
3.2 Termes abrégés et symboles	130
4 Vue d'ensemble	130
5 Codage de données	132
5.1 Généralités	132
5.1.1 Vue d'ensemble	132
5.1.2 Types intégrés	132
5.1.3 Guid	133
5.1.4 ByteString	134
5.1.5 ExtensionObject	134
5.1.6 Variant	135
5.1.7 Decimal	135
5.2 Codage OPC UA binaire	136
5.2.1 Généralités	136
5.2.2 Types intégrés	137
5.2.3 Décimaux	147
5.2.4 Enumérations	147
5.2.5 Matrices	147
5.2.6 Structures	148
5.2.7 Structures avec champs facultatifs	150
5.2.8 Unions	152
5.2.9 Messages	153
5.3 Codage OPC UA XML	154
5.3.1 Types intégrés	154
5.3.2 Décimaux	161
5.3.3 Enumérations	161
5.3.4 Matrices	162
5.3.5 Structures	162
5.3.6 Structures avec champs facultatifs	162
5.3.7 Unions	163
5.3.8 Messages	163
5.4 Codage OPC UA JSON	164
5.4.1 Généralités	164
5.4.2 Types intégrés	164
5.4.3 Décimaux	170
5.4.4 Enumérations	171
5.4.5 Matrices	171
5.4.6 Structures	171
5.4.7 Structures avec champs facultatifs	172
5.4.8 Unions	173
5.4.9 Messages	173
6 SecurityProtocols des messages	173

6.1	Protocole d'établissement de liaison de sécurité	173
6.2	Certificats	175
6.2.1	Généralités	175
6.2.2	Certificat d'instance d'application	176
6.2.3	Chaînes de Certificats	177
6.3	Synchronisation horaire	177
6.4	Temps universel coordonné (UTC) et Temps atomique international (TAI)	177
6.5	Jetons d'identité utilisateur émis	178
6.5.1	Kerberos	178
6.5.2	JWT (JSON Web Token)	178
6.5.3	OAuth2	179
6.6	Conversation sécurisée WS	181
6.7	Conversation OPC UA sécurisée	181
6.7.1	Vue d'ensemble	181
6.7.2	Structure de MessageChunk	182
6.7.3	MessageChunks et traitement d'erreurs	186
6.7.4	Etablissement d'un SecureChannel	186
6.7.5	Dérivation des clés	188
6.7.6	Vérification de la sécurité d'un message	189
7	TransportProtocols	190
7.1	Protocole de connexion OPC UA	190
7.1.1	Vue d'ensemble	190
7.1.2	Structure de message	191
7.1.3	Etablissement d'une connexion	194
7.1.4	Fermeture d'une connexion	196
7.1.5	Traitement d'erreurs	196
7.2	OPC UA TCP	198
7.3	SOAP/HTTP	198
7.4	OPC UA HTTPS	198
7.4.1	Vue d'ensemble	198
7.4.2	Services sans Session	200
7.4.3	Codage XML	200
7.4.4	Codage OPC UA binaire	201
7.4.5	Codage JSON	202
7.5	WebSockets	202
7.5.1	Vue d'ensemble	202
7.5.2	Mapping de protocole	203
7.5.3	Sécurité	203
7.6	Adresses notoires	204
8	Contrats normatifs	205
8.1	Schéma OPC binaire	205
8.2	Schéma XML et langage WSDL	205
8.3	Schéma du Modèle d'Information	205
8.4	Définition formelle du Modèle d'Information UA	205
8.5	Constantes	205
8.6	Encodage des Types de Données	205
8.7	Configuration de Sécurité	205
Annexe A (normative)	Constantes	206
A.1	Identificateurs d'attributs	206

A.2	Codes de statut.....	206
A.3	Identificateurs de nœud numériques	207
Annexe B (normative)	Nodeset OPC UA	208
Annexe C (normative)	Déclarations de type pour le mapping d'origine OPC UA	209
Annexe D (normative)	Langage WSDL pour le mapping XML	210
D.1	Schéma XML	210
D.2	Types d'accès WDSL	210
D.3	Liaisons WSDL	210
Annexe E (normative)	Gestion des paramètres de sécurité	211
E.1	Vue d'ensemble	211
E.2	SecuredApplication	212
E.3	CertificateIdentifier	216
E.4	CertificateStoreIdentifier	218
E.5	CertificateList.....	218
E.6	CertificateValidationOptions.....	218
Annexe F (normative)	Schéma XML du Modèle d'information.....	220
F.1	Vue d'ensemble	220
F.2	UANodeSet.....	220
F.3	UANode	222
F.4	Reference	223
F.5	RolePermission.....	223
F.6	UAType.....	223
F.7	UAIInstance	224
F.8	UAVariable	224
F.9	UAMethod.....	225
F.10	TranslationType	226
F.11	UADatatype	227
F.12	DataTypeDefinition	227
F.13	DataTypeField	228
F.14	Variant.....	229
F.15	Exemple.....	230
F.16	UANodeSetChanges	232
F.17	NodesToAdd	233
F.18	ReferencesToChange	234
F.19	ReferenceToChange	234
F.20	NodesToDelete	234
F.21	NodeToDelete.....	235
F.22	UANodeSetChangesStatus	235
F.23	NodeSetStatusList	236
F.24	NodeSetStatus.....	236
Bibliographie.....		237
Figure 1 – Vue d'ensemble des piles OPC UA.....		131
Figure 2 – Codage des entiers dans une séquence binaire		137
Figure 3 – Codage des virgules flottantes dans une séquence binaire		138
Figure 4 – Codage des chaînes dans une séquence binaire.....		138
Figure 5 – Codage des Guid dans une séquence binaire.....		139

Figure 6 – Codage d'un Elément Xml dans une séquence binaire	139
Figure 7 – Nodeld de chaîne	141
Figure 8 – Nodeld à deux octets	141
Figure 9 – Nodeld à quatre octets	142
Figure 10 – Protocole d'établissement de liaison de sécurité	174
Figure 11 – MessageChunk de conversation sécurisée OPC UA	182
Figure 12 – Structure d'un Message pour le Protocole de connexion OPC UA	191
Figure 13 – Connexion initiée par le Client via le Protocole de connexion OPC UA	195
Figure 14 – Connexion initiée par le Serveur via le Protocole de connexion OPC UA	195
Figure 15 – Fermeture d'une connexion via le Protocole de connexion OPC UA	196
Figure 16 – Scénarios pour le transport HTTPS	199
Figure 17 – Configuration de la communication via WebSocket	203
Tableau 1 – Types de données intégrés	133
Tableau 2 – Structure du Guid	133
Tableau 3 – Présentation d'un Décimal	136
Tableau 4 – Types à virgule flottante pris en charge	137
Tableau 5 – Composants de Nodeld	140
Tableau 6 – Valeurs de DataEncoding de Nodeld	140
Tableau 7 – DataEncoding binaire de Nodeld normalisé	140
Tableau 8 – DataEncoding binaire de Nodeld à deux octets	141
Tableau 9 – DataEncoding binaire de Nodeld à quatre octets	141
Tableau 10 – DataEncoding binaire d'ExpandedNodeld	142
Tableau 11 – DataEncoding binaire de DiagnosticInfo	143
Tableau 12 – DataEncoding binaire de QualifiedName	143
Tableau 13 – DataEncoding binaire de LocalizedText	144
Tableau 14 – DataEncoding binaire d'Objet d'extension	145
Tableau 15 – DataEncoding binaire de Variante	146
Tableau 16 – DataEncoding binaire de Valeur de données	147
Tableau 17 – Echantillon de structure à Codage OPC UA binaire	149
Tableau 18 – Echantillon de structure à Codage OPC UA binaire avec champs facultatifs	151
Tableau 19 – Echantillon de structure à Codage OPC UA binaire	152
Tableau 20 – Mappings des types de données XML pour les entiers	154
Tableau 21 – Mappings de types de données XML pour les virgules flottantes	154
Tableau 22 – Composants de Nodeld	156
Tableau 23 – Composants d'ExpandedNodeld	157
Tableau 24 – Composants d'énumération	161
Tableau 25 – Définition d'Objet JSON pour un Nodeld	166
Tableau 26 – Définition d'Objet JSON pour un ExpandedNodeld	167
Tableau 27 – Définition d'Objet JSON pour un StatusCode	167
Tableau 28 – Définition d'Objet JSON pour une DiagnosticInfo	168
Tableau 29 – Définition d'Objet JSON pour un QualifiedName	168

Tableau 30 – Définition d'objet JSON pour un LocalizedText	169
Tableau 31 – Définition d'Objet JSON pour un ExtensionObject.....	169
Tableau 32 – Définition d'Objet JSON pour une Variante	170
Tableau 33 – Définition d'Objet JSON pour une DataValue	170
Tableau 34 – Définition d'Objet JSON pour un Décimal.....	171
Tableau 35 – Définition d'Objet JSON pour une <i>Structure</i> comportant des champs facultatifs	172
Tableau 36 – Définition d'Objet JSON pour une Union	173
Tableau 37 – SecurityPolicy.....	175
Tableau 38 – Certificat d'instance d'application.....	176
Tableau 39 – UserTokenPolicy Kerberos	178
Tableau 40 – UserTokenPolicy JWT	178
Tableau 41 – Définition d'IssuerEndpointUrl JWT.....	179
Tableau 42 – Revendications de Jeton d'accès	180
Tableau 43 – En-tête de message de conversation OPC UA sécurisée.....	182
Tableau 44 – En-tête de sécurité d'algorithme asymétrique	183
Tableau 45 – En-tête de sécurité d'algorithme symétrique	184
Tableau 46 – En-tête de séquence.....	184
Tableau 47 – Cartouche de message de conversation OPC UA sécurisée	185
Tableau 48 – Corps de l'abandon de message de conversation OPC UA sécurisée	186
Tableau 49 – Service "OpenSecureChannel" pour une conversation OPC UA sécurisée	187
Tableau 50 – Saisies de PRF pour les SecurityPolicies reposant sur RSA	189
Tableau 51 – Paramètres de génération de clés de cryptographie	189
Tableau 52 – En-tête de Message pour le Protocole de connexion OPC UA	191
Tableau 53 – Message d'accueil pour le Protocole de connexion OPC UA.....	192
Tableau 54 – Message d'acquiescement pour le Protocole de connexion OPC UA.....	193
Tableau 55 – Message d'erreur pour le Protocole de connexion OPC UA.....	193
Tableau 56 – Message ReverseHello pour le Protocole de connexion OPC UA.....	194
Tableau 57 – Codes d'erreur pour le Protocole de connexion OPC UA.....	197
Tableau 58 – Mappings de protocoles WebSocket	203
Tableau 59 – Adresses notoires pour les serveurs "Découverte" locaux.....	204
Tableau A.1 – Identificateurs affectés aux attributs	206
Tableau E.1 – SecuredApplication	213
Tableau E.2 – Identificateur de certificat	216
Tableau E.3 – Mémoire de répertoire structurée	217
Tableau E.4 – CertificateStoreIdentifier	218
Tableau E.5 – CertificateList.....	218
Tableau E.6 – CertificateValidationOptions	219
Tableau F.1 – UANodeSet	221
Tableau F.2 – UANode	222
Tableau F.3 – Référence	223
Tableau F.4 – RolePermission	223
Tableau F.5 – Nœuds de type UANodeSet.....	223

Tableau F.6 – Nœuds d'instance UANodeSet	224
Tableau F.7 – UAInstance	224
Tableau F.8 – UAVariable	225
Tableau F.9 – UAMethod	225
Tableau F.10 – TranslationType.....	227
Tableau F.11 – UADatatype.....	227
Tableau F.12 – DataTypeDefinition	228
Tableau F.13 – DataTypeField	229
Tableau F.14 – UANodeSetChanges.....	233
Tableau F.15 – NodesToAdd	233
Tableau F.16 – ReferencesToChange.....	234
Tableau F.17 – ReferencesToChange.....	234
Tableau F.18 – NodesToDelete.....	235
Tableau F.19 – ReferencesToChange.....	235
Tableau F.20 – UANodeSetChangesStatus.....	235
Tableau F.21 – NodeSetStatusList.....	236
Tableau F.22 – NodeSetStatus	236

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

COMMISSION ÉLECTROTECHNIQUE INTERNATIONALE

ARCHITECTURE UNIFIÉE OPC –

Partie 6: Mappings

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 62541-6 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) codages:

- ajout du codage JSON pour PubSub (irréversible);
- ajout du codage JSON pour le Client/Serveur (réversible);
- ajout de la prise en charge des champs facultatifs dans les structures;

- ajout de la prise en charge des Unions;
- b) mappings de transport:
- ajout de la connexion sécurisée WebSocket (WSS);
 - ajout de la prise en charge de la connectivité inversée;
 - ajout de la prise en charge de l'invocation de service sans session dans HTTPS;
- c) transport déconseillé (absence de prise en charge sur la plupart des plateformes):
- SOAP/HTTP avec WS-SecureConversation (tous les codages);
- d) ajout du mapping pour JSON Web Token;
- e) ajout de la prise en charge des Unions pour le Schéma de NodeSet;
- f) ajout d'opérations par lots permettant d'ajouter/de supprimer des nœuds au niveau du Schéma de NodeSet;
- g) ajout de la prise en charge des matrices multidimensionnelles à l'extérieur des Variantes;
- h) ajout d'une représentation binaire pour les types de données Décimaux;
- i) ajout du mapping pour le Cadre d'autorisation OAuth2.

Le texte de cette Norme internationale est issu des documents suivants:

FDIS	Rapport de vote
65E/718/FDIS	65E/734/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 l'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 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, à quelques exceptions près, 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 *OPC Unified Architecture*, 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 62541-6:2020

ARCHITECTURE UNIFIÉE OPC –

Partie 6: Mappings

1 Domaine d'application

La présente partie de l'IEC 62541 spécifie les mappings de l'Architecture unifiée OPC (OPC UA) entre le modèle de sécurité décrit dans l'IEC TR 62541-2, les définitions de services abstraits spécifiées dans l'IEC 62541-4, les structures de données définies dans l'IEC 62541-5 et les protocoles de réseaux physiques qui peuvent être utilisés pour mettre en œuvre la spécification OPC UA.

2 Références normatives

Les documents ci-après, dans leur intégralité ou non, sont des références normatives indispensables à l'application 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 62541-4, *Architecture unifiée OPC – Partie 4: Services*

IEC 62541-5, *Architecture unifiée OPC – Partie 5: Modèle d'information*

IEC 62541-7, *Architecture unifiée OPC – Partie 7: Profils*

IEC 62541-12, *Architecture unifiée OPC – Partie 12: Services globaux et de découverte*

ISO 8601-1:2019, *Date et heure – Représentations pour l'échange d'information – Partie 1: Règles de base*

Schéma XML Partie 2: XML Schema Part 2: Datatypes
<http://www.w3.org/TR/xmlschema-2/>

SOAP Partie 1: SOAP Version 1.2 Part 1: Messaging Framework
<http://www.w3.org/TR/soap12-part1/>

SSL/TLS: RFC 5246 – The TLS Protocol Version 1.2
<http://tools.ietf.org/html/rfc5246.txt>

X.509 v3: ISO/IEC 9594-8 (ITU-T Rec. X.509), *Information technology – Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks* (disponible en anglais seulement)

HTTP: RFC 2616 – Hypertext Transfer Protocol – HTTP/1.1
<http://www.ietf.org/rfc/rfc2616.txt>

HTTPS: RFC 2818 – HTTP Over TLS
<http://www.ietf.org/rfc/rfc2818.txt>

Base64: RFC 3548 – The Base16, Base32, and Base64 Data Encodings
<http://www.ietf.org/rfc/rfc3548.txt>

X690: ISO/IEC 8825-1 (ITU-T Rec. X.690), *Information technology – ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)* (disponible en anglais seulement)

IEEE-754: Standard for Floating-Point Arithmetic
<http://grouper.ieee.org/groups/754/>

HMAC: HMAC – Keyed-Hashing for Message Authentication
<http://www.ietf.org/rfc/rfc2104.txt>

PKCS #1: PKCS #1 – RSA Cryptography Specifications Version 2.0
<http://www.ietf.org/rfc/rfc2437.txt>

PKCS #12: PKCS #12: Personal Information Exchange Syntax
<http://www.ietf.org/rfc/rfc7292.txt>

FIPS 180-4: Secure Hash Standard (SHS)
<https://csrc.nist.gov/publications/detail/fips/180/4/final>

FIPS 197: Advanced Encryption Standard (AES)
<https://csrc.nist.gov/publications/detail/fips/197/final>

UTF-8: UTF-8, a transformation format of ISO 10646
<http://www.ietf.org/rfc/rfc3629.txt>

RFC 3280: RFC 3280 X.509 Public Key Infrastructure Certificate and CRL Profile
<http://www.ietf.org/rfc/rfc3280.txt>

RFC 4514: RFC 4514 – LDAP: String Representation of Distinguished Names
<http://www.ietf.org/rfc/rfc4514.txt>

NTP: RFC 1305 – Network Time Protocol (Version 3) Specification, Implementation and Analysis
<http://www.ietf.org/rfc/rfc1305.txt>

Kerberos: Web Services Security – Kerberos Token Profile 1.1
<http://docs.oasis-open.org/wss/v1.1/wss-v1.1-spec-os-KerberosTokenProfile.pdf>

RFC 1738: RFC 1738 – Uniform Resource Locators (URL)
<http://www.ietf.org/rfc/rfc1738.txt>

RFC 2141: RFC 2141 – URN Syntax
<http://www.ietf.org/rfc/rfc2141.txt>

RFC 6455: RFC 6455 – The WebSocket Protocol
<http://www.ietf.org/rfc/rfc6455.txt>

RFC 7159: The JavaScript Object Notation (JSON) Data Interchange Format
<http://www.ietf.org/rfc/rfc7159.txt>

RFC 7523: JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants
<https://tools.ietf.org/rfc/rfc7523.txt>

RFC 6749: The OAuth 2.0 Authorization Framework
<http://www.ietf.org/rfc/rfc6749.txt>

OpenID-Core: OpenID Connect Core 1.0
http://openid.net/specs/openid-connect-core-1_0.html

OpenID-Discovery: OpenID Connect Discovery 1.0
https://openid.net/specs/openid-connect-discovery-1_0.html

RFC 6960: RFC 6960 – X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP
<https://www.ietf.org/rfc/rfc6960.txt>

3 Termes, définitions, termes abrégés et symboles

3.1 Termes et définitions

Pour les besoins du présent document, les termes et définitions donnés dans 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

CertificateDigest

identificateur court utilisé pour identifier de manière unique un *Certificat* X.509 v3

Note 1 à l'article: Il s'agit du hachage SHA1 de la forme à codage DER du *Certificat*.

3.1.2

DataEncoding

méthode de sérialisation des *Messages* et des structures de données OPC UA

3.1.3

DevelopmentPlatform

suite d'outils et/ou langages de programmation utilisés pour créer des logiciels

3.1.4

Mapping

spécification de la méthode de mise en œuvre d'une fonctionnalité OPC UA avec une technologie spécifique

Note 1 à l'article: Par exemple, le codage OPC UA binaire est un *Mapping* qui spécifie comment sérialiser les structures de données OPC UA en séquences d'octets.

3.1.5**SecurityProtocol**

protocole garantissant l'intégrité et la confidentialité des *Messages* UA échangés entre des applications OPC UA

3.1.6**StackProfile**

combinaison de *Mappings DataEncodings*, *SecurityProtocol* et *TransportProtocol*

Note 1 à l'article: Les applications OPC UA mettent en œuvre un ou plusieurs *StackProfiles* et peuvent communiquer uniquement avec des applications OPC UA qui prennent en charge le même *StackProfile* qu'elles.

3.1.7**TransportConnection**

liaison de communication en duplex intégral établie entre les applications OPC UA

Note 1 à l'article: Un connecteur logiciel TCP/IP est un exemple de *TransportConnection*.

3.1.8**TransportProtocol**

méthode d'échange des *Messages* OPC UA sérialisés entre des applications OPC UA

3.2 Termes abrégés et symboles

API	application programming interface (interface de programmation d'application)
ASN.1	Abstract Syntax Notation (Notation syntaxique abstraite #1) (utilisée dans la recommandation X690)
CSV	comma separated value (valeur séparée par des virgules) (format de fichier)
ECC	elliptic curve cryptography (cryptographie sur les courbes elliptiques)
HTTP	Hypertext Transfer Protocol (Protocole de Transport hypertexte)
HTTPS	Secure Hypertext Transfer Protocol (Protocole de Transport hypertexte sécurisé)
IPSec	Internet Protocol Security (Sécurité de protocole Internet)
OID	object identifier (identificateur d'objet) (utilisé avec ASN.1)
PRF	pseudo random function (fonction pseudoaléatoire)
RSA	Rivest, Shamir and Adleman (système de chiffrement à clé publique)
SHA1	secure hash algorithm (algorithme de hachage sécurisé)
SOAP	Single Object Access Protocol (Protocole d'accès d'objet simple – Protocole SOAP)
SSL	Secure Sockets Layer (Protocole de sécurisation – Protocole SSL) (défini en SSL/TLS)
TCP	Transmission Control Protocol (Protocole de contrôle de transmission)
TLS	Transport Layer Security (Sécurité de la couche transport) (défini en SSL/TLS)
UA	Unified Architecture (Architecture unifiée)
UACP	OPC UA Connection Protocol (Protocole de connexion OPC UA)
UASC	OPC UA Secure Conversation (Conversation sécurisée OPC UA)
WS-*	XML Web Services specifications (spécifications des services web XML)
XML	eXtensible Markup Language (Langage de balisage extensible)

4 Vue d'ensemble

Les autres parties de la série IEC 62541 sont rédigées de façon à être indépendantes de la technologie de mise en œuvre utilisée. Cette approche signifie qu'OPC UA est une spécification souple qui s'adapte aux évolutions technologiques. Par ailleurs, cette approche signifie qu'il n'est pas possible de constituer une application OPC UA à partir des informations

contenues dans les normes IEC TR 62541-1 à IEC 62541-5 du fait de l'omission d'informations de mise en œuvre importantes.

Le présent document définit des *Mappings* entre les spécifications abstraites et les technologies qui peuvent être utilisés pour la mise en œuvre. Les *Mappings* sont organisés en trois groupes: *DataEncodings*, *SecurityProtocols* et *TransportProtocols*. Différents *Mappings* sont associés pour créer des *StackProfiles*. Toutes les applications OPC UA doivent mettre en œuvre au moins un *StackProfile* et peuvent communiquer uniquement avec d'autres applications OPC UA qui mettent en œuvre le même *StackProfile*.

Le présent document définit les *DataEncodings* à l'Article 5, les *SecurityProtocols* au 5.4 et les *TransportProtocols* au 6.7.6. Les *StackProfiles* sont définis dans l'IEC 62541-7.

Toutes les communications entre les applications OPC UA reposent sur l'échange de *Messages*. Les paramètres contenus dans les *Messages* sont définis dans l'IEC 62541-4. Toutefois, leur format est spécifié par le *DataEncoding* et le *TransportProtocol*. Ainsi, chaque *Message* défini dans l'IEC 62541-4 doit avoir une description normative qui spécifie exactement ce qui doit être mis sur le réseau de communication. Les descriptions normatives sont définies dans les annexes.

Une *Pile* est un regroupement de bibliothèques logicielles qui mettent en œuvre un ou plusieurs *StackProfiles*. L'interface entre une application OPC UA et la *Pile* est une interface API non normative qui masque les informations de mise en œuvre de la *Pile*. Une interface API dépend d'une *DevelopmentPlatform* spécifique. Du fait des limites de la *DevelopmentPlatform*, les types de données présentés dans l'interface API pour une *DevelopmentPlatform* peuvent ne pas correspondre aux types de données définis par la spécification. Par exemple, le langage de programmation Java ne prend pas en charge un entier non signé, ce qui signifie qu'il est nécessaire que les interfaces API Java mettent en correspondance les entiers non signés avec un type d'entier signé.

La Figure 1 décrit la relation entre les différents concepts définis dans le présent document.

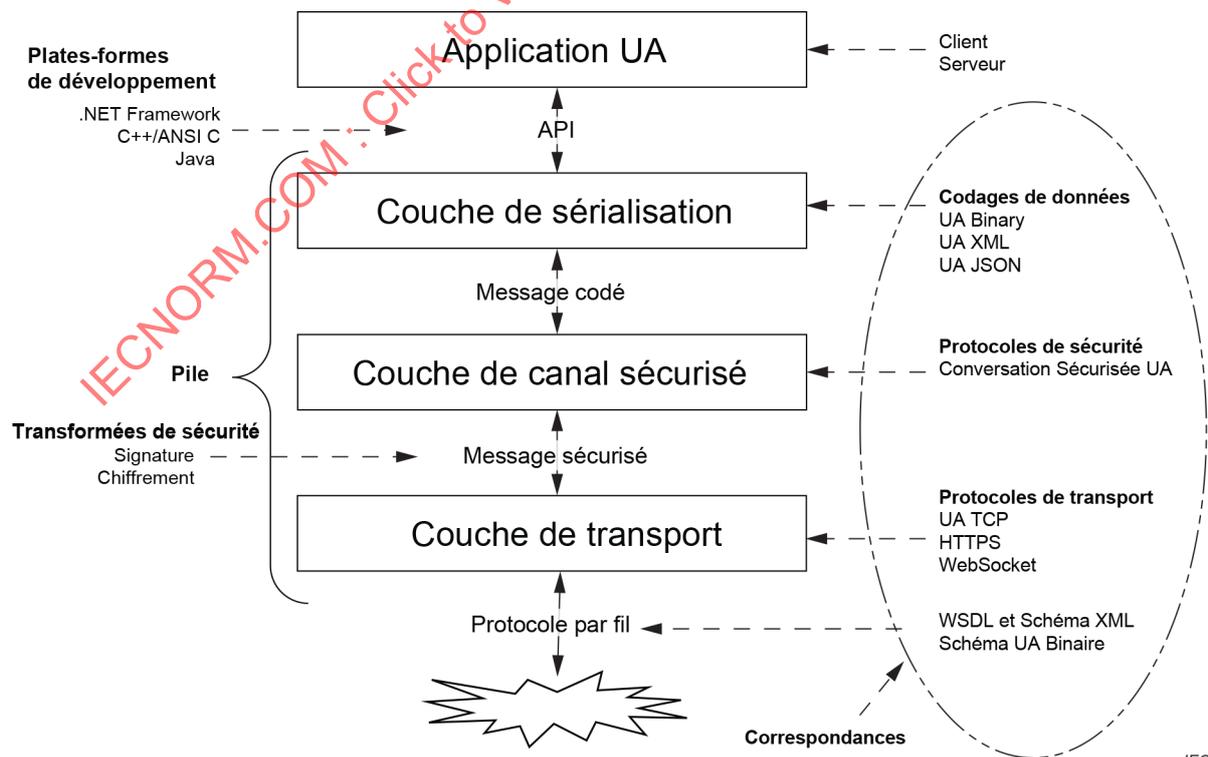


Figure 1 – Vue d'ensemble des piles OPC UA

Les couches décrites dans le présent document ne correspondent pas aux couches du modèle OSI 7 [X200]. Il convient de traiter chaque *StackProfile* OPC UA comme un protocole de Couche 7 (application) unique reposant sur un protocole existant (Couche 5, 6 ou 7), comme TCP/IP, TLS ou HTTP. La couche *SecureChannel* est toujours présente, même si le *SecurityMode* est *None*. Dans ce type de situation, aucune sécurité n'est appliquée, mais la mise en œuvre du *SecurityProtocol* doit maintenir un canal logique avec un identificateur unique. Les utilisateurs et les administrateurs sont supposés savoir qu'un *SecureChannel* dont le *SecurityMode* est réglé sur *None* ne peut pas être fiable, à moins que l'application fonctionne sur un réseau physiquement sécurisé ou qu'un protocole de bas niveau, comme IPSec, soit utilisé.

5 Codage de données

5.1 Généralités

5.1.1 Vue d'ensemble

Le présent document définit trois types de codages de données: le codage OPC UA binaire, le codage OPC UA XML et le codage OPC UA JSON. Elle décrit comment créer des *Messages* avec chacun de ces codages.

5.1.2 Types intégrés

Tous les *DataEncodings* OPC UA reposent sur des règles établies pour un ensemble normalisé de types intégrés. Ces types intégrés permettent alors de créer des structures, des matrices et des *Messages*. Les types intégrés sont décrits dans le Tableau 1.

IECNORM.COM : Click to view the full PDF of IEC 62541-6:2020 RLV

Tableau 1 – Types de données intégrés

ID	Name	Description
1	Boolean	Valeur logique binaire (vrai ou faux).
2	SByte	Valeur entière entre -128 et 127 inclus.
3	Byte	Valeur entière entre 0 et 255 inclus.
4	Int16	Valeur entière entre -32 768 et 32 767 inclus.
5	UInt16	Valeur entière entre 0 et 65 535 inclus.
6	Int32	Valeur entière entre -2 147 483 648 et 2 147 483 647 inclus.
7	UInt32	Valeur entière entre 0 et 4 294 967 295 inclus.
8	Int64	Valeur entière entre -9 223 372 036 854 775 808 et 9 223 372 036 854 775 807 inclus.
9	UInt64	Valeur entière entre 0 et 18 446 744 073 709 551 615 inclus.
10	Float	Valeur à virgule flottante (32 bits) en simple précision IEEE.
11	Double	Valeur à virgule flottante (64 bits) en double précision IEEE.
12	String	Séquence de caractères Unicode.
13	DateTime	Instance dans le temps.
14	Guid	Valeur à 16 octets qui peut être utilisée comme identificateur globalement unique.
15	ByteString	Séquence d'octets.
16	XmlElement	Élément XML.
17	NodeId	Identificateur d'un nœud dans l'espace d'adressage d'un <i>Serveur</i> OPC UA.
18	ExpandedNodeId	NodeId permettant de spécifier l'URI d'espace de noms en lieu et place d'un indice.
19	StatusCode	Identifiant numérique d'une erreur ou d'un état associé à une valeur ou une opération.
20	QualifiedName	Nom qualifié par un espace de noms.
21	LocalizedText	Texte lisible en clair avec un identificateur de lieu facultatif.
22	ExtensionObject	Structure contenant un type de données spécifique à l'application qui peut ne pas être reconnu par le récepteur.
23	DataValue	Valeur de données avec code de statut et horodatages associés.
24	Variant	Union de tous les types spécifiés ci-dessus.
25	DiagnosticInfo	Structure contenant des informations précises d'erreur et de diagnostic associées à un <i>StatusCode</i>

La plupart de ces types de données sont identiques aux types abstraits définis dans l'IEC 62541-3 et l'IEC 62541-4. Le présent document définit toutefois les types *ExtensionObject* et *Variant*. De plus, le présent document définit une représentation pour le type *Guid* défini dans l'IEC 62541-3.

5.1.3 Guid

Un *Guid* de 16 octets, dont la configuration est indiquée dans le Tableau 2.

Tableau 2 – Structure du Guid

Composant	Type de données
Data1	UInt32
Data2	UInt16
Data3	UInt16
Data4	Byte [8]

Les valeurs du *Guid* peuvent être représentées sous la forme de la chaîne suivante:

```
<Data1>-<Data2>-<Data3>-<Data4[0:1]>-<Data4[2:7]>
```

Lorsque *Data1* a une largeur de 8 caractères, *Data2* et *Data3* ont une largeur de 4 caractères et chaque *Octet* dans *Data4* a une largeur de 2 caractères. Chaque valeur est formatée sous la forme d'un nombre hexadécimal rempli de zéros. Une valeur de *Guid* type ressemble à ce qui suit lorsqu'il est au format chaîne:

```
C496578A-0DFE-4B8F-870A-745238C6AEAE
```

5.1.4 ByteString

La structure d'une *ByteString* est identique à celle d'une matrice unidimensionnelle d'*Octet*. Elle est représentée comme un type de données intégré distinct dans la mesure, où elle permet aux codeurs d'optimiser la transmission de la valeur. Cependant, certaines *DevelopmentPlatforms* ne sont pas capables de préserver la distinction entre une *ByteString* et une matrice unidimensionnelle d'*Octet*.

Si un décodeur applicable à la *DevelopmentPlatform* ne peut pas pérenniser cette distinction, il doit convertir toutes les matrices unidimensionnelles d'*Octet* en *ByteStrings*.

Chaque élément d'une matrice unidimensionnelle de *ByteString* peut avoir une longueur différente, ce qui signifie que sa structure est différente de celle d'une matrice bidimensionnelle d'*Octet* où la longueur de chaque dimension est identique. Cela signifie que les décodeurs doivent préserver la distinction entre deux matrices dimensionnelles d'*Octet* ou plus et une ou plusieurs matrices dimensionnelles de *ByteString*.

Si une *DevelopmentPlatform* ne prend pas en charge les entiers non signés, il est essentiel qu'elle représente les *ByteStrings* sous forme de matrices de type *SByte*. Dans ce cas, les exigences d'*Octet* s'appliqueraient alors à *SByte*.

5.1.5 ExtensionObject

Un *ExtensionObject* contient tous les *DataTypes Structurés* qui ne peuvent pas être codés comme l'un des autres types de données intégrés. L'*ExtensionObject* contient une valeur complexe sérialisée sous forme d'une séquence d'octets ou d'un élément XML. Il contient également un identificateur qui indique les données qu'il contient, ainsi que la méthode de codage utilisée.

Les *DataTypes structurés* sont représentés dans l'espace d'adressage du *Serveur* sous forme de sous-types du *DataType* de *Structure*. Les *DataEncodings* disponibles pour tous les *DataTypes Structurés* sont représentés sous la forme d'un *Objet "DataTypeEncoding"* dans l'*AddressSpace* du *Serveur*. Le *NodeId* de l'*Objet "DataTypeEncoding"* est l'identificateur archivé dans l'*ExtensionObject*. L'IEC 62541-3 décrit comment les *Nœuds "DataTypeEncoding"* sont liés aux autres *Nœuds* de l'*AddressSpace*.

Il convient que les ingénieurs qui implémentent des *Serveurs* utilisent des *NodeId* numériques qualifiés de l'espace de noms pour les *objets "DataTypeEncoding"* qu'ils définissent. Cela permet de réduire le plus possible la surcharge générée par le tassement des valeurs de *DataTypes structurés* dans un *ExtensionObject*.

Les *ExtensionObjects* et *Variantes* permettent une imbrication illimitée qui peut générer des erreurs de débordement de pile même si la taille du message est inférieure à la taille maximale admise. Les décodeurs doivent prendre en charge au moins 100 niveaux d'imbrication. Ils doivent signaler une erreur si le nombre de niveaux d'imbrication dépasse le nombre pris en charge.

5.1.6 Variant

Une *Variante* est l'union de tous les types de données intégrés, y compris un *ExtensionObject*. Les *Variantes* peuvent également contenir des matrices de ces types intégrés. Les *Variantes* servent à mémoriser toute valeur ou tout paramètre avec un type de données de *BaseDataType* ou l'un de ses sous-types.

Les *Variantes* peuvent être vides. Une *Variante* vide peut être décrite comme une variante ayant une valeur nulle, et il convient de la traiter comme une colonne nulle dans une base de données SQL. Une valeur nulle dans une *Variante* peut ne pas être identique à une valeur nulle pour les types de données qui prennent en charge les valeurs nulles, comme les *Chaînes*. Certaines *DevelopmentPlatforms* peuvent ne pas être capables de pérenniser la distinction entre une valeur nulle pour un *DataType* et une valeur nulle pour une *Variante*. Par conséquent, les applications ne doivent pas reposer sur cette distinction. Cette exigence signifie également que si un *Attribut* prend en charge l'écriture d'une valeur nulle, il doit également prendre en charge l'écriture d'une *Variante* vide et inversement.

Les *Variantes* peuvent contenir des matrices de *Variantes*, mais ne peuvent pas contenir directement une autre *Variante*.

Les types de *DiagnosticInfo* ne sont significatifs que lorsqu'ils sont renvoyés dans un message de réponse avec un *StatusCode* associé et un tableau de chaînes. De ce fait, les *Variantes* ne peuvent pas contenir d'instances de *DiagnosticInfo*.

Les *Valeurs* des *Attributs* sont toujours renvoyées dans des instances de *DataValues*. Par conséquent, le *DataType* d'un *Attribut* ne peut pas être une *DataValue*. Les *Variantes* peuvent contenir une *DataValue* lorsqu'elles sont utilisées dans d'autres contextes, comme les *Arguments de Méthode* ou les *Messages PubSub*. La *Variante* d'une *DataValue* ne peut pas, directement ou indirectement, contenir d'autres *DataValues*.

Les *Variables* avec un *DataType* de *BaseDataType* sont mises en correspondance avec une *Variante*; toutefois, les *Attributs ValueRank* et *ArrayDimensions* imposent des restrictions concernant le contenu admis de la *Variante*. Par exemple, si le *ValueRank* est *Scalaire*, alors la *Variante* ne peut contenir que des valeurs scalaires.

Les *ExtensionObjects* et *Variantes* permettent une imbrication illimitée qui peut générer des erreurs de débordement de pile même si la taille du message est inférieure à la taille maximale admise. Les décodeurs doivent prendre en charge au moins 100 niveaux d'imbrication. Ils doivent signaler une erreur si le nombre de niveaux d'imbrication dépasse le nombre pris en charge.

5.1.7 Decimal

Un *Décimal* est un nombre décimal signé de haute précision. Il se compose d'une valeur non échelonnée entière de précision arbitraire et d'une échelle d'entiers. L'échelle est égale à la puissance de dix appliquée à la valeur non échelonnée.

Un *Décimal* contient les champs décrits dans le Tableau 3.

Tableau 3 – Présentation d'un Décimal

Champ	Type	Description
Typeld	Nodeld	Identificateur du <i>Data Type Décimal</i> .
Encoding	Byte	Cette valeur est toujours égale à 1.
Length	Int32	Longueur du <i>Décimal</i> . Si la longueur est inférieure ou égale à 0, la valeur du <i>Décimal</i> est 0.
Scale	Int16	Entier signé représentant la puissance de dix utilisée pour échelonner la valeur. C'est-à-dire le nombre décimal de la valeur multiplié par 10^{-scale} L'entier est codé en commençant par le bit de poids faible.
Value	Byte [*]	Entier signé à 2 compléments représentant la valeur non échelonnée. Le nombre de bits est déduit de la longueur du champ <i>Length</i> . Si le nombre de bits est 0, la valeur est égale à 0. L'entier est codé en commençant par l'octet de poids faible.

Lorsqu'un *Décimal* est codé dans une *Variante*, le type intégré est défini sur *ExtensionObject*. Les décodeurs qui ne comprennent pas le type *Décimal* doivent le traiter comme n'importe quelle autre *Structure* inconnue et le transmettre à l'application. Les décodeurs qui comprennent le *Décimal* peuvent analyser la valeur et utiliser une construction adaptée à la *DevelopmentPlatform*.

Si un *Décimal* est imbriqué dans une autre *Structure*, la *DataTypeDefintion* pour le champ doit spécifier le *Nodeld* du *Nœud décimal* en tant que *Data Type*. Si un *Serveur* publie une description de type OPC binaire pour la *Structure*, la description du type doit définir le *Data Type* du champ sur *ExtensionObject*.

5.2 Codage OPC UA binaire

5.2.1 Généralités

Le *DataEncoding* OPC UA binaire est un format de données qui a été mis au point afin de satisfaire aux exigences de performance des applications OPC UA. Ce format est conçu principalement pour un codage et un décodage rapides. Toutefois, il a également été tenu compte de la taille des données codées mises sur le réseau.

Le *DataEncoding* OPC UA binaire repose sur plusieurs types de données primitives avec des règles de codage clairement définies, qui peuvent être écrites ou lues selon un ordre séquentiel à partir d'un flot de bits. Le codage d'une structure s'effectue par l'écriture de la forme codée de chaque champ selon un ordre séquentiel. Lorsqu'un champ donné est également une structure, les valeurs de ses champs sont écrites suivant un ordre séquentiel, avant d'écrire le champ suivant dans la structure qu'il contient. Tous les champs doivent être écrits dans la séquence même s'ils contiennent des valeurs nulles. Les codages applicables à chaque type de primitive spécifient comment coder soit une valeur nulle, soit une valeur par défaut pour le type concerné.

Le *DataEncoding* OPC UA binaire ne comprend aucune information de nom de type ou de champ, dans la mesure où toutes les applications OPC UA sont supposées avoir une connaissance avancée des services et structures qu'elles prennent en charge. Un *ExtensionObject* qui fournit un identificateur et une taille pour la structure des *DataTypes Structurés* qu'il représente constitue une exception. Cela permet à un décodeur d'ignorer les types qu'il ne reconnaît pas.

5.2.2 Types intégrés

5.2.2.1 Boolean

Une valeur de *Booléen* doit être codée comme un octet simple, où une valeur de 0 (zéro) est fausse et toute valeur non nulle est vraie.

Les codeurs doivent utiliser la valeur de 1 pour indiquer une valeur vraie; toutefois, les décodeurs doivent traiter toute valeur non nulle comme étant vraie.

5.2.2.2 Entier

Tous les types d'entiers doivent être codés comme des valeurs "little-endian" où l'octet le plus faible apparaît en premier dans la séquence des bits.

La Figure 2 décrit la méthode utilisée pour coder la valeur 1 000 000 000 (Hex: 3B9ACA00) comme un entier à 32 bits dans la séquence binaire.

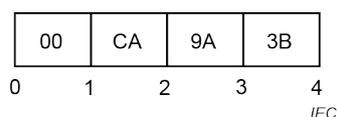


Figure 2 – Codage des entiers dans une séquence binaire

5.2.2.3 Virgule flottante

Toutes les valeurs à virgule flottante doivent être codées avec la représentation binaire IEEE-754 appropriée qui comporte trois composants de base: le signe, l'exposant et la fraction. Les plages de bits attribuées à chaque composant dépendent de la largeur du type. Le Tableau 4 donne la liste des plages de bits relatives aux types à virgule flottante pris en charge.

Tableau 4 – Types à virgule flottante pris en charge

Name	Largeur (bits)	Fraction	Exposant	Signe
Float	32	0 à 22	23 à 30	31
Double	64	0 à 51	52 à 62	63

De plus, l'ordre des octets dans la séquence est important. Toutes les valeurs à virgule flottante doivent être codées avec l'octet de poids faible apparaissant en premier (c'est-à-dire "little endian").

La Figure 3 décrit la méthode utilisée pour coder la valeur -6,5 (Hex: C0D00000) comme une *Float*.

Le type à virgule flottante prend en charge l'infinité positive et négative et "not-a-number" (NaN). La spécification IEEE autorise l'utilisation de plusieurs variantes NaN. Toutefois, les codeurs/décodeurs peuvent ne pas préserver la distinction. Les codeurs doivent coder une valeur NaN comme une NAN "quiet" IEEE (000000000000F8FF) ou (0000C0FF). Les types non pris en charge (tels que les nombres non normalisés) doivent également être codés sous la forme d'une NAN "quiet" IEEE. Tous les essais d'égalité entre les valeurs NaN échouent toujours.

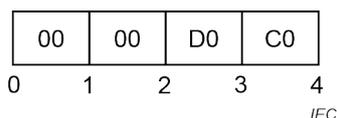


Figure 3 – Codage des virgules flottantes dans une séquence binaire

5.2.2.4 String

Toutes les valeurs de *Chaîne* sont codées sous forme de séquence de caractères UTF-8 sans terminateur nul et précédée par la longueur en octets.

La longueur en octets est codée sous la forme *Int32*. Une valeur de -1 permet d'indiquer une chaîne "null".

La Figure 4 décrit la méthode utilisée pour le codage de la chaîne multilingue "水Boy" dans une séquence d'octets.

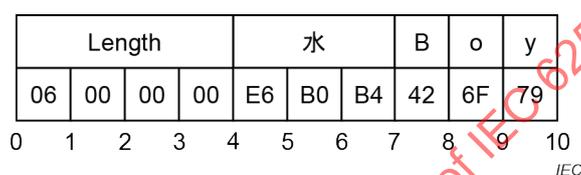


Figure 4 – Codage des chaînes dans une séquence binaire

5.2.2.5 DateTime

Une valeur *DateTime* doit être codée sous la forme d'un entier signé de 64 bits (voir 5.2.2.2) qui représente le nombre d'intervalles de 100 nanosecondes depuis le 1^{er} janvier 1601 (UTC).

Les *DevelopmentPlatforms* ne sont pas toutes capables de représenter la plage complète de dates et d'heures pouvant être représentée par ce *DataEncoding*. Par exemple, la structure *time_t* sous UNIX a uniquement une résolution de 1 seconde et ne peut représenter de dates antérieures à 1970. C'est pourquoi un certain nombre de règles doivent être appliquées lorsqu'il s'agit de valeurs date/heure au-delà de la plage dynamique d'une *DevelopmentPlatform*. Ces règles sont les suivantes:

- a) une valeur date/heure est codée 0 si:
 - 1) la valeur est inférieure ou égale à 1601-01-01 12:00AM UTC; ou
 - 2) la valeur constitue la date la plus proche qui peut être représentée avec le codage de la *DevelopmentPlatform*;
- b) une valeur date/heure est codée comme la valeur maximale d'un *Int64* si:
 - 1) la valeur est supérieure ou égale à 9999-12-31 11:59:59PM UTC; ou
 - 2) la valeur constitue la date la plus lointaine qui peut être représentée avec le codage de la *DevelopmentPlatform*;
- c) une valeur date/heure est décodée comme l'heure la plus proche qui peut être représentée sur la plateforme si:
 - 1) la valeur codée est 0; ou
 - 2) la valeur codée représente une heure antérieure à l'heure la plus proche qui peut être représentée avec le codage de la *DevelopmentPlatform*;
- d) une valeur date/heure est décodée comme l'heure la plus lointaine qui peut être représentée sur la plateforme si:
 - 1) la valeur codée est la valeur maximale d'un *Int64*; ou

- 2) la valeur codée représente une heure ultérieure à l'heure la plus lointaine qui peut être représentée avec le codage de la *DevelopmentPlatform*.

Ces règles impliquent que les heures les plus proches et les plus lointaines qui peuvent être représentées sur une plateforme donnée sont des valeurs de date/heure non valides, qu'il convient de traiter comme telles par les applications.

Un décodeur doit tronquer la valeur s'il rencontre une valeur de *DateTime* dont la résolution est supérieure à celle prise en charge sur la *DevelopmentPlatform*.

5.2.2.6 Guid

Un *Guid* est codé dans une structure comme indiqué dans le Tableau 2. Le codage des champs s'effectue selon un ordre séquentiel en fonction du type de données propre au champ.

La Figure 5 représente la méthode utilisée pour le codage du *Guid* "72962B91-FA75-4AE6-8D28-B404DC7DAF63" dans une séquence d'octets.

Data1				Data2		Data3		Data4								
91	2B	96	72	75	FA	E6	4A	8D	28	B4	04	DC	7D	AF	63	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

IEC

Figure 5 – Codage des Guid dans une séquence binaire

5.2.2.7 ByteString

Une *ByteString* est codée comme une séquence d'octets précédée de sa longueur en octets. La longueur est codée comme un entier signé de 32 bits comme décrit ci-dessus.

Si la longueur de la chaîne d'octets est –1, la chaîne d'octets est alors "null".

5.2.2.8 XmlElement

Un *XmlElement* est un fragment XML sérialisé sous forme d'une chaîne UTF-8, puis codé sous forme de *ByteString*.

La Figure 6 décrit la méthode utilisée pour le codage du *XmlElement* "<A>Hot水" dans une séquence d'octets.

Length				<A>			Hot			水							
0D	00	00	00	3C	41	3E	72	6F	74	E6	B0	B4	3C	3F	41	3E	
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17

IEC

Figure 6 – Codage d'un Élément Xml dans une séquence binaire

Un décodeur peut choisir d'analyser le code XML après le décodage; si une erreur d'analyse irrécupérable se produit, il convient que le décodeur tente de poursuivre le traitement de la séquence. Par exemple, si le *XmlElement* est le corps d'une *Variante* ou un élément dans une matrice qui est le corps d'une *Variante*, cette erreur peut être signalée en établissant la valeur de la *Variante* sur le *StatusCode Bad_DecodingError*.

5.2.2.9 Nodeld

Les composants d'un *Nodeld* sont décrits dans le Tableau 5.

Tableau 5 – Composants de Nodeld

Name	Type de données	Description
Namespace	UInt16	Indice d'un URI d'espace de noms. Un indice de 0 est utilisé pour les <i>Nodelds</i> OPC UA définis.
IdentifierType	Enumeration	Format et type de données de l'identificateur. La valeur peut être l'une des valeurs suivantes: NUMERIC - la valeur est un <i>UInteger</i> ; CHAÎNE - la valeur est une <i>String</i> ; GUID - la valeur est un <i>Guid</i> ; OPAQUE - la valeur est une <i>ByteString</i> .
Value	*	L'identificateur d'un nœud dans l'espace d'adressage d'un <i>Serveur</i> OPC UA.

Le *DataEncoding* d'un *Nodeld* varie selon le contenu de l'instance. Ainsi, le premier octet de la forme codée indique le format du reste du *Nodeld* codé. Les formats de *DataEncoding* possibles sont présentés dans le Tableau 6. Tableau 6 à Tableau 9 décrivent la structure de chaque format possible (ils excluent l'octet qui indique le format).

Tableau 6 – Valeurs de DataEncoding de Nodeld

Name	Value	Description
2 Octets	0x00	Valeur numérique ajustée à la représentation à deux octets.
4 Octets	0x01	Valeur numérique ajustée à la représentation à quatre octets.
Numeric	0x02	Valeur numérique non ajustée à la représentation à deux ou à quatre octets.
String	0x03	Valeur de chaîne.
Guid	0x04	Valeur de Guid.
ByteString	0x05	Valeur opaque (<i>ByteString</i>).
Fanion NamespaceUri	0x80	Voir présentation de <i>ExpandedNodeld</i> en 5.2.2.10.
Fanion ServerIndex	0x40	Voir présentation de <i>ExpandedNodeld</i> en 5.2.2.10.

La structure du *DataEncoding* du *Nodeld* normalisé est décrite dans le Tableau 7. Le *DataEncoding* normalisé est utilisé pour tous les formats non définis de manière explicite.

Tableau 7 – DataEncoding binaire de Nodeld normalisé

Name	Type de données	Description
Namespace	UInt16	Indice de l'espace de noms (<i>NamespaceIndex</i>).
Identifier	*	Identificateur codé selon les règles suivantes: NUMÉRIQUE UInt32 CHAÎNE String GUID Guid OPAQUE ByteString

Un exemple de *Nodeld* de Chaîne avec Namespace = 1 et présentation = "Hot水" est représenté à la Figure 7.

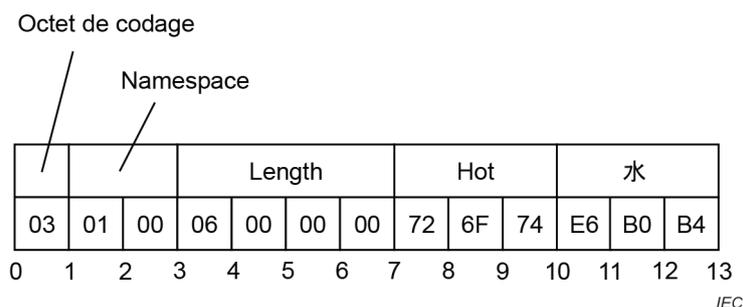


Figure 7 – Nodeld de chaîne

La structure du *DataEncoding* du *Nodeld* à deux octets est décrite dans le Tableau 8.

Tableau 8 – DataEncoding binaire de Nodeld à deux octets

Name	Type de données	Description
Identifiant	Byte	Le <i>Namespace</i> est l'espace de noms OPC UA par défaut (c'est-à-dire 0). Le type d' <i>identificateur</i> est "numérique". L' <i>identificateur</i> doit se situer dans la plage comprise entre 0 et 255.

Un exemple de *Nodeld* à deux octets avec Identifiant = 72 est représenté à la Figure 8.

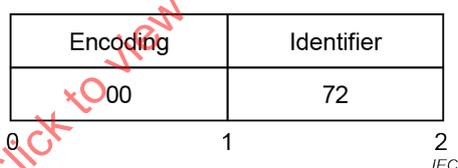


Figure 8 – Nodeld à deux octets

La structure du *DataEncoding* du *Nodeld* à quatre octets est décrite dans le Tableau 9.

Tableau 9 – DataEncoding binaire de Nodeld à quatre octets

Name	Type de données	Description
Namespace	Byte	Le <i>Namespace</i> doit se situer dans la plage comprise entre 0 et 255.
Identifiant	UInt16	Le type d' <i>identificateur</i> est "numérique". L' <i>identificateur</i> doit être un entier situé dans la plage comprise entre 0 et 65 535.

5.2.2.12 DiagnosticInfo

Une structure *DiagnosticInfo* est décrite dans l'IEC 62541-4. Elle spécifie un nombre de champs pouvant être manquants. Pour cette raison, le codage utilise un masque de bits pour indiquer quels champs sont effectivement présents dans la forme codée.

Comme décrit dans l'IEC 62541-4, les champs *SymbolicId*, *NamespaceUri*, *LocalizedText* et *Locale* constituent des indices dans un tableau de chaînes renvoyé dans l'en-tête de réponse. Seul l'indice de la chaîne correspondante dans le tableau de chaînes est codé. Un indice de -1 indique qu'il n'existe aucune valeur pour la chaîne.

Le type *DiagnosticInfo* permet une imbrication illimitée qui peut générer des erreurs de débordement de pile même si la taille du message est inférieure à la taille maximale admise. Les décodeurs doivent prendre en charge au moins 100 niveaux d'imbrication. Ils doivent signaler une erreur si le nombre de niveaux d'imbrication dépasse le nombre pris en charge.

Tableau 11 – DataEncoding binaire de DiagnosticInfo

Name	Type de données	Description
EncodingMask	Byte	Masque de bits indiquant les champs présents dans la séquence. Le masque comprend les bits suivants: 0x01 SymbolicId 0x02 Namespace 0x04 LocalizedText 0x08 Locale 0x10 AdditionalInfo 0x20 InnerStatusCode 0x40 InnerDiagnosticInfo
SymbolicId	Int32	Nom symbolique pour le code de statut.
NamespaceUri	Int32	Espace de noms qui qualifie l'identificateur symbolique.
Locale	Int32	Paramètre de lieu utilisé pour le texte localisé.
LocalizedText	Int32	Récapitulatif lisible en clair du code de statut.
AdditionalInfo	String	Informations de diagnostic précises, spécifiques à l'application.
InnerStatusCode	StatusCode	Code de statut fourni par un système sous-jacent.
InnerDiagnosticInfo	DiagnosticInfo	Information de diagnostic associée au code de statut interne.

5.2.2.13 QualifiedName

Une structure de *QualifiedName* est codée comme indiqué dans le Tableau 12.

La structure abstraite de *QualifiedName* est définie dans l'IEC 62541-3.

Tableau 12 – DataEncoding binaire de QualifiedName

Name	Type de données	Description
NamespaceIndex	UInt16	Indice de l'espace de noms.
Name	String	Nom.

5.2.2.14 LocalizedText

Une structure de *LocalizedText* contient deux champs pouvant être manquants. Pour cette raison, le codage utilise un masque de bits pour indiquer quels champs sont effectivement présents dans la forme codée.

La structure abstraite de *LocalizedText* est définie dans l'IEC 62541-3.

Tableau 13 – DataEncoding binaire de LocalizedText

Name	Type de données	Description
EncodingMask	Byte	Masque de bits indiquant les champs présents dans la séquence. Le masque comprend les bits suivants: 0x01 Locale 0x02 Text
Locale	String	Paramètre de lieu. S'il n'est pas pris en compte, cela signifie nul ou vide.
Text	String	Texte dans le paramètre de lieu spécifié. S'il n'est pas pris en compte, cela signifie nul ou vide.

5.2.2.15 ExtensionObject

Un *ExtensionObject* est codé comme une séquence d'octets dont le préfixe est le *NodeId* de son *DataTypeEncoding* et le nombre d'octets codés.

Un *ExtensionObject* peut être codé par l'application, ce qui signifie qu'il est transmis au codeur sous forme de *ByteString* ou *XmlElement*. Dans ce cas, le codeur est capable d'écrire le nombre d'octets dans l'objet avant de coder les octets. Toutefois, un *ExtensionObject* peut savoir se coder/décoder lui-même, ce qui signifie que le codeur doit calculer le nombre d'octets avant de coder l'objet, ou il doit être capable d'effectuer une rétrorecherche dans la séquence et d'actualiser la longueur après codage du corps de l'objet.

Lorsqu'un décodeur rencontre un *ExtensionObject*, il doit vérifier s'il reconnaît l'identificateur de *DataTypeEncoding*. Si c'est le cas, il peut demander à la fonction appropriée de décoder le corps de l'objet. Si le décodeur ne reconnaît pas le type, il doit utiliser le champ *Encoding* pour déterminer si le corps est une *ByteString* ou un *XmlElement*, puis décoder le corps de l'objet ou le traiter comme une donnée opaque et l'ignorer.

La forme sérialisée d'un *ExtensionObject* est spécifiée dans le Tableau 14.

Tableau 14 – DataEncoding binaire d'Objet d'extension

Name	Type de données	Description
TypeId	NodId	Identificateur du nœud de <i>DataTypeId</i> dans l' <i>AddressSpace</i> du <i>Serveur</i> . Les <i>ExtensionObjects</i> définis par la spécification OPC UA ont un identificateur de nœud numérique qui leur est attribué avec un <i>NamespaceIndex</i> de 0. Les identificateurs numériques sont définis en A.3. Les décodeurs utilisent ce champ pour déterminer la syntaxe du <i>Corps</i> . Par exemple, si ce champ est le <i>NodId</i> de l' <i>Objet de codage JSON</i> d'un <i>DataTypeId</i> , le <i>Corps</i> est une <i>ByteString</i> contenant un document JSON codé sous la forme d'une chaîne UTF-8.
Encoding	Byte	Énumération qui indique la méthode de codage du corps. Le paramètre peut avoir les valeurs suivantes: 0x00 Aucun corps n'est codé. 0x01 Le corps est codé comme une <i>ByteString</i> . 0x02 Le corps est codé comme un <i>XmlElement</i> .
Length	Int32	Longueur du corps de l'objet. La longueur doit être spécifiée si le corps est codé.
Body	Byte [*]	Corps de l'objet. Ce champ contient les octets bruts des corps de <i>ByteString</i> . Pour les corps <i>XmlElement</i> , ce champ contient l'élément XML codé sous forme de chaîne UTF-8 sans terminateur nul. Certaines structures à codage binaire peuvent avoir une longueur sérialisée qui n'est pas un multiple de 8 bits. Les codeurs doivent ajouter des bits 0 pour garantir que la longueur sérialisée est un multiple de 8 bits. Les décodeurs qui comprennent le format sérialisé doivent ignorer les bits de remplissage.

Les *ExtensionObjects* sont utilisés dans deux contextes: sous forme de valeurs contenues dans les structures de *Variante*s ou sous forme de paramètres dans les *Messages* OPC UA.

Un décodeur peut choisir d'analyser le corps d'un *XmlElement* après décodage; si une erreur d'analyse irrécupérable se produit, il convient que le décodeur tente de poursuivre le traitement de la séquence. Par exemple, si l'*ExtensionObject* est le corps d'une *Variante* ou un élément dans une matrice qui est le corps d'une *Variante*, cette erreur peut être signalée en définissant la valeur de la *Variante* sur le *StatusCode Bad_DecodingError*.

5.2.2.16 Variant

Une *Variante* est une union des types intégrés.

La structure d'une *Variante* est présentée dans le Tableau 15.

Tableau 15 – DataEncoding binaire de Variante

Name	Type de données	Description
EncodingMask	Byte	<p>Type de données codé dans la séquence. La valeur 0 indique une valeur nulle et qu'aucun autre champ n'est codé. Les bits suivants sont attribués au masque:</p> <p>0:5 Identificateur de type Intégré (voir Tableau 1).</p> <p>6 Vrai si le champ de Dimensions de Matrice est codé.</p> <p>7 Vrai si une matrice de valeurs est codée.</p> <p>Les identificateurs de type intégré 26 à 31 ne sont actuellement pas attribués, mais peuvent être utilisés ultérieurement. Les décodeurs doivent accepter ces identificateurs, prendre pour hypothèse que la <i>Valeur</i> contient une <i>ByteString</i> et les transmettre à l'application. Les codeurs ne doivent pas utiliser ces identificateurs.</p>
ArrayLength	Int32	<p>Nombre d'éléments dans la matrice.</p> <p>Ce champ est présent uniquement si le bit de matrice est établi dans le masque de codage.</p> <p>Les matrices multidimensionnelles sont codées sous forme d'une matrice unidimensionnelle et ce champ spécifie le nombre total d'éléments. La matrice d'origine peut être reconstituée à partir des dimensions codées d'après le champ de valeur.</p> <p>Les dimensions de rang supérieur sont sérialisées en premier lieu. Par exemple, une matrice de dimensions [2,2,2] s'écrit dans l'ordre suivant:</p> <p>[0,0,0], [0,0,1], [0,1,0], [0,1,1], [1,0,0], [1,0,1], [1,1,0], [1,1,1]</p>
Value	*	<p>Valeur codée selon son type de données intégré.</p> <p>Si le bit de matrice est établi dans le masque de codage, chaque élément dans la matrice est codé de manière séquentielle. Dans la mesure où de nombreux types ont un codage de longueur variable, chaque élément doit être décodé dans l'ordre.</p> <p>La valeur ne doit pas être une <i>Variante</i>, mais peut être une matrice de <i>Variantes</i>.</p> <p>De nombreuses plateformes de mise en œuvre ne font pas la différence entre les matrices unidimensionnelles <i>Bytes</i> et <i>ByteStrings</i>. Ainsi, les décodeurs peuvent convertir automatiquement une matrice <i>Bytes</i> en une <i>ByteString</i>.</p>
ArrayDimensions Length	Int32	<p>Nombre de dimensions.</p> <p>Ce champ est présent uniquement si le fanion <i>ArrayDimensions</i> est établi dans le masque de codage.</p>
ArrayDimensions	Int32[*]	<p>Longueur de chaque dimension codée sous forme de séquence de valeurs Int32</p> <p>Ce champ est présent uniquement si le fanion <i>ArrayDimensions</i> est établi dans le masque de codage. Les dimensions de rang inférieur apparaissent en premier dans la matrice.</p> <p>Toutes les dimensions doivent être spécifiées et doivent être supérieures à zéro.</p> <p>Si les <i>ArrayDimensionse</i> ne sont pas cohérentes avec l'<i>ArrayLength</i>, le décodeur doit s'arrêter et émettre le code de statut <i>Bad_DecodingError</i>.</p>

Les types et leurs identificateurs qui peuvent être codés dans une *Variante* sont présentés dans le Tableau 1.

5.2.2.17 DataValue

Une *DataValue* est toujours précédée d'un masque qui indique quels champs sont présents dans la séquence.

Les champs d'une *DataValue* sont décrits dans le Tableau 16.

Tableau 16 – DataEncoding binaire de Valeur de données

Name	Type de données	Description
EncodingMask	Byte	Masque de bits indiquant les champs présents dans la séquence. Le masque comprend les bits suivants: 0x01 Faux si la valeur est <i>Null</i> . 0x02 Faux si le <i>StatusCode</i> est <i>Good</i> . 0x04 Faux si l'horodatage source est <i>DateTime.MinValue</i> . 0x08 Faux si l'horodatage <i>Serveur</i> est <i>DateTime.MinValue</i> . 0x10 Faux si les <i>Picosecondes Source</i> sont égales à 0. 0x20 Faux si les <i>Picosecondes Serveur</i> sont égales à 0.
Value	Variant	Valeur. Absente si le bit de <i>Valeur</i> dans l' <i>EncodingMask</i> est à l'état <i>False</i> .
Status	StatusCode	Etat associé à la valeur. Absent si le bit de <i>StatusCode</i> dans l' <i>EncodingMask</i> est <i>False</i> .
SourceTimestamp	DateTime	Horodatage source associé à la valeur. Absent si le bit <i>SourceTimestamp</i> dans l' <i>EncodingMask</i> est <i>False</i> .
SourcePicoSeconds	UInt16	Nombre d'intervalles de 10 picosecondes pour le <i>SourceTimestamp</i> . Absent si le bit <i>SourcePicoSeconds</i> dans l' <i>EncodingMask</i> est <i>False</i> . En l'absence de l'horodatage source, les picosecondes ne sont pas prises en compte.
ServerTimestamp	DateTime	Horodatage <i>Serveur</i> associé à la valeur. Absent si le bit <i>ServerTimestamp</i> dans l' <i>EncodingMask</i> est <i>False</i> .
ServerPicoSeconds	UInt16	Nombre d'intervalles de 10 picosecondes pour le <i>ServerTimestamp</i> . Absent si le bit <i>ServerPicoSeconds</i> dans l' <i>EncodingMask</i> est <i>False</i> . En l'absence de l'horodatage <i>Serveur</i> , les picosecondes ne sont pas prises en compte.

Les champs *Picosecondes* mémorisent la différence entre un horodatage haute résolution avec une résolution de 10 picosecondes et la valeur du champ *Horodatage* ayant une résolution de 100 ns uniquement. Les champs *Picosecondes* doivent contenir des valeurs inférieures à 10 000. Le décodeur doit traiter les valeurs supérieures ou égales à 10 000 comme la valeur '9 999'.

5.2.3 Décimaux

Les *Décimaux* sont codés comme décrit au 5.1.7.

Un *Décimal* n'a pas de valeur nulle.

5.2.4 Enumérations

Les énumérations sont codées comme des valeurs *Int32*.

Une *énumération* n'a pas de valeur NULL.

5.2.5 Matrices

Les *Matrices* unidimensionnelles sont codées comme une séquence d'éléments précédée du nombre d'éléments codés comme une valeur *Int32*. Si une *Matrice* est nulle, sa longueur est

codée –1. Une *Matrice* de longueur nulle étant différente d'une *Matrice* nulle, les codeurs et les décodeurs doivent conserver cette distinction.

Les *Matrices* multidimensionnelles sont codées comme une *Matrice Int32* contenant les dimensions suivie d'une liste de toutes les valeurs contenues dans la *Matrice*. Le nombre total de valeurs est égal au produit des dimensions. Le nombre de valeurs est 0 si une ou plusieurs dimensions sont inférieures ou égales à 0. Le processus de reconstitution de la matrice multidimensionnelle est décrit en 5.2.2.16.

5.2.6 Structures

Les *Structures* sont codées sous forme d'une séquence de champs selon leur ordre d'apparition dans la définition. Le codage de chaque champ est déterminé par le type intégré propre au champ.

Tous les champs spécifiés dans la structure doivent être codés. Si la structure contient des champs facultatifs, voir 5.2.7.

Les *Structures* n'ont pas de valeur nulle. Si un codeur est écrit dans un langage de programmation qui permet aux structures d'avoir des valeurs nulles, le codeur doit créer une nouvelle instance comportant des valeurs par défaut pour tous les champs et procéder à une sérialisation. Les codeurs ne doivent pas générer une erreur de codage dans ce type de situation.

L'exemple suivant représente une structure utilisant la syntaxe C/C++:

```
struct Type2
{
    Int32 A;
    Int32 B;
};

struct Type1
{
    Int32 X;
    Byte NoOfY;
    Type2* Y;
    Int32 Z;
};
```

Dans l'exemple C/C++ ci-dessus, le champ Y est un pointeur dirigé vers une matrice dont la longueur est mémorisée dans NoOfY. Lors du codage d'une matrice, la longueur faisant partie du processus, le champ NoOfY n'est pas codé. Cela dit, les codeurs et les décodeurs utilisent NoOfY lors du codage.

Une instance de *Type1* qui contient une matrice de deux instances de *Type2* est codée comme une séquence de 28 octets. Si l'instance de *Type1* a été codée dans un *ExtensionObject*, elle possède un préfixe supplémentaire indiqué dans le Tableau 17, portant la longueur totale à 37 octets. Les champs *TypeId*, *Encoding* et *Length* sont définis par l'*ExtensionObject*. Le codage des instances de *Type2* n'inclut pas l'identificateur de type dans la mesure où il est défini de manière explicite dans le *Type1*.

Tableau 17 – Echantillon de structure à Codage OPC UA binaire

Champ	Octets	Value
TypeId	4	Identificateur du nœud de codage binaire Type1
Encoding	1	0x1 pour la ByteString
Length	4	28
X	4	Valeur du champ 'X'
Y.Length	4	2
Y.A	4	Valeur du champ 'Y[0].A'
Y.B	4	Valeur du champ 'Y[0].B'
Y.A	4	Valeur du champ 'Y[1].A'
Y.B	4	Valeur du champ 'Y[1].B'
Z	4	Valeur du champ 'Z'

La *Valeur* de l'Attribut *DataTypeDefinition* d'un Nœud de *DataType* de Type1 est:

Name	Type	Description
defaultEncodingId	NodeId	NodeId du nœud "Type1_Encoding_DefaultBinary".
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure sans champ facultatif]
champs [0]	StructureField	
name	String	"X"
description	LocalizedText	Description de X
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false
champs [1]	StructureField	
name	String	"Y"
description	LocalizedText	Description de Y-Array
dataType	NodeId	NodeId du nœud de DataType de Type2 (par ex. "ns=3; s=MyType2")
valueRank	Int32	1 (OneDimension)
isOptional	Boolean	false
champs [2]	StructureField	
name	String	"Z"
description	LocalizedText	Description de Z
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false

La Valeur de l'Attribut *DataTypeDefinition* d'un Nœud de *DataType* de Type2 est:

Name	Type	Description
defaultEncodingId	NodeId	NodeId du nœud "Type2_Encoding_DefaultBinary".
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure sans champ facultatif]
champs [0]	StructureField	
name	String	"A"
description	LocalizedText	Description de A
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false
champs [1]	StructureField	
name	String	"B"
description	LocalizedText	Description de B
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false

5.2.7 Structures avec champs facultatifs

Les *structures* comportant des champs facultatifs sont codées avec un masque de codage précédant une séquence de champs dans leur ordre d'apparition dans la définition. Le codage de chaque champ est déterminé par le type de données propre au champ.

L'*EncodingMask* est un entier non signé de 32 bits. Un bit exactement est attribué à chaque champ facultatif. Le bit '0' est attribué au premier champ facultatif, le bit '1' est attribué au deuxième champ facultatif, etc. jusqu'à ce que des bits soient attribués à tous les champs facultatifs. Un maximum de 32 champs facultatifs peut apparaître dans une seule *Structure*. Les bits non attribués sont établis sur 0 par les codeurs. Les décodeurs doivent signaler une erreur si les bits non attribués ne sont pas établis sur 0.

L'exemple suivant représente une structure comportant des champs facultatifs utilisant la syntaxe C++:

```
struct TypeA
{
  Int32 X;
  Int32* O1;
  SByte Y;
  Int32* O2;
};
```

O1 et O2 sont des champs facultatifs qui sont NULL s'ils sont absents.

Une instance de *TypeA* qui contient deux champs obligatoires (X et Y) et deux champs facultatifs (O1 et O2) est codée sous forme de séquence d'octets. La longueur de la séquence d'octets dépend des champs facultatifs disponibles. Un champ de masque de codage détermine les champs facultatifs disponibles.

Une instance de *TypeA* où le champ O2 est disponible et le champ O1 ne l'est pas est codée sous forme de séquence à 13 octets. Si l'instance de *TypeA* était codée dans un *ExtensionObject*, elle aurait la forme de codage présentée dans le Tableau 18 et aurait une

longueur totale de 22 octets. La longueur des champs *TypeId*, *Encoding* et *Length* est définie par l'*ExtensionObject*.

**Tableau 18 – Echantillon de structure
à Codage OPC UA binaire avec champs facultatifs**

Champ	Octets	Value
TypeId	4	Identificateur du nœud de codage binaire TypeA
Encoding	1	0x1 pour la ByteString
Length	4	13
EncodingMask	4	0x02 pour O2
X	4	Valeur de X
Y	1	Valeur de Y
O2	4	Valeur de O2

Si une *Structure* comportant des champs facultatifs est sous-typée, les sous-types étendent l'*EncodingMask* défini pour le parent.

La *Valeur* de l'*Attribut DataTypeDefinition* d'un *Nœud* de *DataType* de TypeA est:

Name	Type	Description
defaultEncodingId	NodeId	NodeId du nœud "TypeA_Encoding_DefaultBinary".
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	StructureWithOptionalFields_1 [Structure sans champ facultatif]
champs [0]	StructureField	
name	String	"X"
description	LocalizedText	Description de X
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false
champs [1]	StructureField	
name	String	"O1"
description	LocalizedText	Description de O1
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	true
champs [2]	StructureField	
name	String	"Y"
description	LocalizedText	Description de Z
dataType	NodeId	"i=2" [SByte]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false
champs [3]	StructureField	
name	String	"O2"
description	LocalizedText	Description de O2
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	true

5.2.8 Unions

Les *Unions* sont codées sous forme de champ de permutation précédant l'un des champs possibles. Le codage du champ choisi est déterminé par le type de données propre au champ.

Le champ de permutation est codé au format UInt32.

Le champ de permutation est l'indice des champs d'union disponibles commençant par 1. Si le champ de permutation est 0, aucun champ n'est présent. Pour toute autre valeur supérieure au nombre de champs d'union définis, les codeurs ou les décodeurs doivent signaler une erreur.

Une *Union* sans champ a la même signification qu'une valeur NULL. Une *Union* comportant des champs n'est pas une valeur NULL, même si la valeur du champ lui-même est NULL.

L'exemple suivant représente une union utilisant la syntaxe C/C++:

```

struct Type2
{
    Int32 A;
    Int32 B;
};

struct Type1
{
    Byte Selector;

    union
    {
        Int32 Field1;
        Type2 Field2;
    }
    Value;
};
    
```

Dans l'exemple C/C++ ci-dessus, le Sélecteur et la Valeur sont associés sur le plan sémantique pour former une union. L'ordre des champs n'a pas d'importance.

Une instance de *Type1* serait codée sous forme de séquence d'octets. La longueur de la séquence d'octets dépend du champ choisi.

Une instance de *Type1*, où le champ *Field1* est disponible, est codée sous forme de séquence à 8 octets. Si l'instance de *Type1* était codée dans un *ExtensionObject*, elle aurait la forme de codage décrite dans le Tableau 19 et aurait une longueur totale de 17 octets. Les champs *TypeId*, *Encoding* et *Length* sont des champs définis par l'*ExtensionObject*.

Tableau 19 – Echantillon de structure à Codage OPC UA binaire

Champ	Octets	Value
TypeId	4	Identificateur du Type1
Encoding	1	0x1 pour la ByteString
Length	4	8
SwitchValue	4	1 pour Field1
Field1	4	Valeur de Field1

La Valeur de l'Attribut *DataTypeDefinition* d'un Nœud de *DataType* de Type1 est:

Name	Type	Description
defaultEncodingId	NodeId	NodeId du nœud "Type1_Encoding_DefaultBinary".
baseDataType	NodeId	"i=22" [Union]
structureType	StructureType	Union_2 [Union]
champs [0]	StructureField	
name	String	"Field1"
description	LocalizedText	Description de Field1
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	true
champs [1]	StructureField	
name	String	"Field2"
description	LocalizedText	Description de Field2
dataType	NodeId	NodeId du nœud de <i>DataType</i> de Type2 (par ex. "ns=3; s=MyType2")
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	true

La Valeur de l'Attribut *DataTypeDefinition* d'un Nœud de *DataType* de Type2 est:

Name	Type	Description
defaultEncodingId	NodeId	NodeId du nœud "Type2_Encoding_DefaultBinary".
baseDataType	NodeId	"i=22" [Structure]
structureType	StructureType	Structure_0 [Structure sans champ facultatif]
champs [0]	StructureField	
name	String	"A"
description	LocalizedText	Description de A
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false
champs [1]	StructureField	
name	String	"B"
description	LocalizedText	Description de B
dataType	NodeId	"i=6" [Int32]
valueRank	Int32	-1 (Scalaire)
isOptional	Boolean	false

5.2.9 Messages

Les *Messages* sont des *Structures* codées sous forme de séquence d'octets préfixées par le *NodeId* du *DataTypeEncoding* OPC UA binaire défini pour le Message.

Chaque *Service* OPC UA décrit dans l'IEC 62541-4 a un *Message* de demande et de réponse. Les identificateurs du *DataTypeEncoding* attribués à chaque *Service* sont spécifiés à l'Article A.3.

5.3 Codage OPC UA XML

5.3.1 Types intégrés

5.3.1.1 Généralités

La plupart des types intégrés est codée en langage XML en utilisant les formats définis dans la spécification de Schéma XML Partie 2. Les restrictions ou utilisations spéciales sont traitées ci-dessous. Certains types intégrés comportent un Schéma XML défini pour eux au moyen de la syntaxe explicitée dans la Schéma XML Partie 2.

Le préfixe *xs:* sert à désigner un symbole défini par la spécification du Schéma XML.

5.3.1.2 Boolean

Une valeur booléenne est codée sous forme d'une valeur *xs:boolean*.

5.3.1.3 Entier

Les valeurs entières sont codées avec l'un des sous-types du type *xs:decimal*. Les mappings entre les types d'entier OPC UA et les types de données du schéma XML sont présentés dans le Tableau 20.

Tableau 20 – Mappings des types de données XML pour les entiers

Name	Type XML
SByte	xs:byte
Byte	xs:unsignedByte
Int16	xs:short
UInt16	xs:unsignedShort
Int32	xs:int
UInt32	xs:entier non signé
Int64	xs:long
UInt64	xs:unsignedLong

5.3.1.4 Virgule flottante

Les valeurs à virgule flottante sont codées avec l'un des types XML à virgule flottante. Les mappings entre les types à virgule flottante OPC UA et les types de données du schéma XML sont présentés dans le Tableau 21.

Tableau 21 – Mappings de types de données XML pour les virgules flottantes

Name	Type XML
Float	xs:float
Double	xs:double

Le type à virgule flottante XML prend en charge infinité positive (INF), infinité négative (-INF) et "not-a-number" (NaN).

5.3.1.5 String

Une valeur de *Chaîne* est codée comme une valeur *xs:string*.

5.3.1.6 DateTime

Une valeur de *DateTime* est codée comme une valeur *xs:DateTime*.

Toutes les valeurs de *DateTime* doivent être codées sous forme de temps UTC ou avec une spécification explicite du fuseau horaire.

Correct:

```
2002-10-10T00:00:00+05:00
2002-10-09T19:00:00Z
```

Incorrect:

```
2002-10-09T19:00:00
```

Il est recommandé de représenter toutes les valeurs de *xs:DateTime* sous format UTC.

La valeur date/heure la plus récente et la valeur date/heure la plus lointaine qui peuvent être représentées sur une *DevelopmentPlatform* ont une signification particulière et ne doivent pas être strictement codées au format XML.

La valeur date/heure la plus récente sur une *DevelopmentPlatform* doit être codée au format XML suivant: "0001-01-01T00:00:00Z".

La valeur de date/heure la plus ancienne sur une *DevelopmentPlatform* doit être codée au format XML suivant "9999-12-31T23:59:59Z"

Si un décodeur rencontre une valeur *xs:DateTime* qui ne peut pas être représentée sur la *DevelopmentPlatform*, il convient qu'il convertisse la valeur en date/heure la plus récente ou la plus ancienne qui peut être représentée sur la *DevelopmentPlatform*. Il convient que le décodeur XML ne génère aucune erreur s'il rencontre une valeur de date en dehors des valeurs définies.

La valeur de date/heure la plus proche sur une *DevelopmentPlatform* est équivalente à une valeur date/heure nulle.

5.3.1.7 Guid

Un *Guid* est codé au moyen de la représentation de chaîne définie en 5.1.3.

Le schéma XML applicable à un *Guid* est:

```
<xs:complexType name="Guid">
  <xs:sequence>
    <xs:element name="String" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.8 ByteString

Une valeur de *ByteString* est codée comme une valeur *xs:base64Binary* (voir Base64).

Le schéma XML applicable à une *ByteString* est: />

```
<xs:element name="ByteString" type="xs:base64Binary" nillable="true"/>
```

5.3.1.9 XmlElement

Une valeur *XmlElement* est codée comme une valeur *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="XmlElement">
  <xs:sequence>
    <xs:any minOccurs="0" maxOccurs="1" processContents="lax" />
  </xs:sequence>
</xs:complexType>
```

Les *XmlElements* peuvent être utilisés uniquement dans les valeurs de *Variante* ou d'*ExtensionObject*.

5.3.1.10 NodeId

Une valeur de *NodeId* est codée sous forme de valeur *xs:string* avec la syntaxe suivante:

```
ns=<indiceespace nom>;<type>=<valeur>
```

Les éléments de la syntaxe sont décrits dans le Tableau 22.

Tableau 22 – Composants de NodeId

Champ	Type de données	Description
<namespaceindex>	UInt16	<i>NamespaceIndex</i> formaté comme un nombre en base 10. Si l'indice est 0, l'entier 'ns=0;' ne doit pas être pris en compte.
<type>	Enumeration	Fanion de spécification de l' <i>IdentifieurType</i> . Le fanion a les valeurs suivantes: i NUMERIC (UInt32) s STRING (Chaîne) g GUID (Guid) b OPAQUE (ByteString)
<value>	*	<i>Identificateur</i> codé sous forme de chaîne. L' <i>identificateur</i> est formaté en utilisant le mapping de type de données XML propre à l' <i>IdentifieurType</i> . Noter que l' <i>identificateur</i> peut contenir tout caractère UTF-8 non nul, y compris un blanc.

Exemples de *NodeIds*:

```
i=13
ns=10;i=-1
ns=10;s>Hello:World
g=09087e75-8e5e-499b-954f-f2a9603db28a
ns=1;b=M/RbKBSRVkePCePcx24oRA==
```

Le schéma XML applicable à un *NodeId* est:

```
<xs:complexType name="NodeId">
  <xs:sequence>
    <xs:element name="Identifieur" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.11 ExpandedNodeId

Une valeur d'*ExpandedNodeId* est codée comme une valeur *xs:string* avec la syntaxe suivante:

```
svr=<indiceserveur>;ns=<indiceespacenom>;<type>=<valeur>
ou
svr=<indiceserveur>;nsu=<uri>;<type>=<valeur>
```

Les champs possibles sont indiqués dans le Tableau 23.

Tableau 23 – Composants d'ExpandedNodeId

Champ	Type de données	Description
<serverindex>	UInt32	<i>ServerIndex</i> formaté comme un nombre en base 10. Si le <i>ServerIndex</i> est 0, l'entier 'svr=0;' ne doit pas être pris en compte.
<namespaceindex>	UInt16	<i>NamespaceIndex</i> formaté comme un nombre en base 10. Si le <i>NamespaceIndex</i> est 0, alors l'article entier 'ns=0;' ne doit pas être pris en compte. Aucun <i>NamespaceIndex</i> ne doit être présent si l'URI est spécifié.
<uri>	String	Le <i>NamespaceUri</i> formaté comme une chaîne. Les caractères réservés éventuels de l'URI doivent être remplacés par un '%' suivi de sa valeur ANSI de 8 bits codée sous forme de deux chiffres hexadécimaux (insensibles à la casse). Par exemple, le caractère ';' serait remplacé par '%3B'. Les caractères réservés sont ';' et '%'. Si le <i>NamespaceUri</i> est nul ou vide, alors l'article 'nsu=;' ne doit pas être pris en compte.
<type>	Enumeration	Fanion de spécification de l' <i>IdentifierType</i> . Ce champ est décrit dans le Tableau 22.
<value>	*	<i>Identificateur</i> codé sous forme de chaîne. Ce champ est décrit dans le Tableau 22.

Le schéma XML applicable à un *ExpandedNodeId* est:

```
<xs:complexType name="ExpandedNodeId">
  <xs:sequence>
    <xs:element name="Identifieur" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.12 StatusCode

Un *StatusCode* est codé comme un *xs:unsignedInt* avec le schéma XML suivant:

```
<xs:complexType name="StatusCode">
  <xs:sequence>
    <xs:element name="Code" type="xs:unsignedInt" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.13 DiagnosticInfo

Une valeur *DiagnosticInfo* est codée comme un *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="DiagnosticInfo">
  <xs:sequence>
    <xs:element name="SymbolicId" type="xs:int" minOccurs="0" />
    <xs:element name="NamespaceUri" type="xs:int" minOccurs="0" />
    <xs:element name="Locale" type="xs:int" minOccurs="0"/>
    <xs:element name="LocalizedText" type="xs:int" minOccurs="0"/>
    <xs:element name="AdditionalInfo" type="xs:string" minOccurs="0"/>
    <xs:element name="InnerStatusCode" type="tns:StatusCode"
      minOccurs="0" />
    <xs:element name="InnerDiagnosticInfo" type="tns:DiagnosticInfo"
      minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

Le type *DiagnosticInfo* permet une imbrication illimitée qui peut générer des erreurs de débordement de pile même si la taille du message est inférieure à la taille maximale admise. Les décodeurs doivent prendre en charge au moins 100 niveaux d'imbrication. Ils doivent signaler une erreur si le nombre de niveaux d'imbrication dépasse le nombre pris en charge.

5.3.1.14 QualifiedName

Une valeur *QualifiedName* est codée comme un *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="QualifiedName">
  <xs:sequence>
    <xs:element name="NamespaceIndex" type="xs:int" minOccurs="0" />
    <xs:element name="Name" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.15 LocalizedText

Une valeur *LocalizedText* est codée comme un *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="LocalizedText">
  <xs:sequence>
    <xs:element name="Locale" type="xs:string" minOccurs="0" />
    <xs:element name="Text" type="xs:string" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

5.3.1.16 ExtensionObject

Une valeur *ExtensionObject* est codée comme un *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="ExtensionObject">
  <xs:sequence>
    <xs:element name="TypeId" type="tns:NodeId" minOccurs="0" />
    <xs:element name="Body" minOccurs="0">
      <xs:complexType>
        <xs:sequence>
          <xs:any minOccurs="0" processContents="lax"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

Le corps de l'*ExtensionObject* contient un seul élément qui est soit une *ByteString*, soit une *Structure* à codage XML. Un décodeur peut faire la différence entre ces deux éléments en examinant l'élément de niveau supérieur. Un élément portant le nom *tns:ByteString* contient un corps à codage OPC UA binaire. Tout autre nom doit contenir un corps à codage OPC UA XML. Le *TypeId* spécifie la syntaxe d'un corps de *ByteString* qui peut être au format JSON à codage UTF-8, UA binaire, ou autre.

Le *TypeId* est le *NodeId* pour l'*Objet DataTypeEncoding*.

5.3.1.17 Variant

Une valeur *Variant* est codée comme un *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="Variant">
  <xs:sequence>
    <xs:element name="Value" minOccurs="0" nillable="true">
      <xs:complexType>
        <xs:sequence>
          <xs:any minOccurs="0" processContents="lax"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

Si la *Variante* représente une valeur scalaire, elle doit alors contenir un seul élément enfant portant le nom du type intégré. Par exemple, la valeur à virgule flottante en simple précision 3,141 5 serait codée comme suit:

```
<tns:Float>3.1415</tns:Float>
```

Si la *Variante* représente une matrice unidimensionnelle, elle doit alors contenir un seul élément enfant portant le préfixe "ListOf" et le nom du type intégré. Par exemple, une *Matrice* de chaînes serait codée comme suit:

```
<tns:ListOfString>
  <tns:String>Hello</tns:String>
  <tns:String>World</tns:String>
</tns:ListOfString>
```

Si la *Variante* représente une *Matrice* multidimensionnelle, elle doit alors contenir un élément enfant portant le nom "*Matrix*" comportant les deux sous-éléments de cet exemple:

```
<tns:Matrix>
  <tns:Dimensions>
    <tns:Int32>2</tns:Int32>
    <tns:Int32>2</tns:Int32>
  </tns:Dimensions>
  <tns:Elements>
    <tns:String>A</tns:String>
    <tns:String>B</tns:String>
    <tns:String>C</tns:String>
    <tns:String>D</tns:String>
  </tns:Elements>
</tns:Matrix>
```

Dans l'exemple, la matrice comporte les éléments suivants:

```
[0,0] = "A"; [0,1] = "B"; [1,0] = "C"; [1,1] = "D"
```

Les éléments d'une *Matrice* multidimensionnelle sont toujours aplatis dans une *Matrice* unidimensionnelle, où les dimensions de rang supérieur sont sérialisées en premier. Cette *Matrice* unidimensionnelle est codée comme un enfant de l'élément "Eléments". L'élément "Dimensions" est une *Matrice* de valeurs *Int32* qui spécifient les dimensions de la matrice en commençant par la dimension de rang inférieur. La *Matrice* multidimensionnelle peut être reconstituée à partir des dimensions. Toutes les dimensions doivent être spécifiées et doivent être supérieures à zéro. Si les dimensions ne sont pas cohérentes avec le nombre d'éléments dans la matrice, le décodeur doit s'arrêter et émettre le code de statut *Bad_DecodingError*.

Le Tableau 1 donne l'ensemble complet des noms de types intégrés.

5.3.1.18 DataValue

Une valeur *DataValue* est codée comme un *xs:complexType* avec le schéma XML suivant:

```
<xs:complexType name="DataValue">
  <xs:sequence>
    <xs:element name="Value" type="tns:Variant" minOccurs="0"
      nillable="true" />
    <xs:element name="StatusCode" type="tns:StatusCode"
      minOccurs="0" />
    <xs:element name="SourceTimestamp" type="xs:dateTime"
      minOccurs="0" />
    <xs:element name="SourcePicoseconds" type="xs:unsignedShort"
      minOccurs="0"/>
    <xs:element name="ServerTimestamp" type="xs:dateTime"
      minOccurs="0" />
    <xs:element name="ServerPicoseconds" type="xs:unsignedShort"
      minOccurs="0"/>
  </xs:sequence>
</xs:complexType>
```

5.3.2 Décimaux

Une valeur *Decimal* est codée comme un `xs:complexType` avec le schéma XML suivant:

```
<xs:complexType name="Decimal">
  <xs:sequence>
    <xs:element name="TypeId" type="tns:NodeId" minOccurs="0" />
    <xs:element name="Body" minOccurs="0">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="Scale" type="xs:unsignedShort" />
          <xs:element name="Value" type="xs:string" />
        </xs:sequence>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

Le *NodeId* est toujours celui dont le *DataType* est *Décimal*. Lorsqu'il est codé dans une *Variante*, le *Décimal* est codé sous forme d'*ExtensionObject*. Les Matrices de *Décimaux* sont des *Matrices d'ExtensionObjects*.

Le champ de *Valeur* est un entier signé en base 10 sans limite de taille. Voir 5.1.7 pour une description des champs *Scale* et *Value*.

5.3.3 Enumérations

Les *énumérations* utilisées comme des paramètres dans les *Messages* définis dans l'IEC 62541-4 sont codées comme `xs:string` avec la syntaxe suivante:

```
<symbole>_<valeur>
```

Les éléments de la syntaxe sont décrits dans le Tableau 24.

Tableau 24 – Composants d'énumération

Champ	Type	Description
<symbol>	String	Nom symbolique de la valeur énumérée.
<value>	UInt32	Valeur numérique associée à la valeur énumérée.

Par exemple, le schéma XML applicable à l'énumération de *NodeClass* est:

```
<xs:simpleType name="NodeClass">
  <xs:restriction base="xs:string">
    <xs:enumeration value="Unspecified_0" />
    <xs:enumeration value="Object_1" />
    <xs:enumeration value="Variable_2" />
    <xs:enumeration value="Method_4" />
    <xs:enumeration value="ObjectType_8" />
    <xs:enumeration value="VariableType_16" />
    <xs:enumeration value="ReferenceType_32" />
    <xs:enumeration value="DataType_64" />
    <xs:enumeration value="View_128" />
  </xs:restriction>
</xs:simpleType>
```

Les *énumérations* stockées dans une *Variante* sont codées comme une valeur *Int32*.

Par exemple, toute *Variable* peut avoir une valeur avec un *DataType* de *NodeClass*. Dans ce cas, la valeur numérique correspondante est placée dans la *Variante* (par exemple, l'*Objet NodeClass* serait enregistré comme 1).

5.3.4 Matrices

Les paramètres d'une *Matrice* unidimensionnelle sont toujours codés en enveloppant des éléments dans un élément conteneur et en insérant ce dernier dans la structure. Il convient que le nom de l'élément conteneur soit le nom du paramètre. Le nom de l'élément dans la matrice doit être le nom de type.

Par exemple, le service *Lecture* utilise une matrice de *ReadValueIds*. Le schéma XML serait le suivant:

```
<xs:complexType name="ListOfReadValueId">
  <xs:sequence>
    <xs:element name="ReadValueId" type="tns:ReadValueId"
      minOccurs="0" maxOccurs="unbounded" nillable="true" />
  </xs:sequence>
</xs:complexType>
```

Les attributs nillables doivent être spécifiés, car les codeurs XML placent les éléments dans les matrices si ces éléments sont vides.

Les paramètres d'une *Matrice* multidimensionnelle sont codés en utilisant le type "*Matrix*" défini en 5.3.1.17.

5.3.5 Structures

Les structures sont codées sous la forme *xs:complexType*, tous les champs apparaissant de manière séquentielle. Tous les champs sont codés sous la forme *xs:element*. Tous les éléments comportent l'attribut *minOccurs* réglé sur 0 qui permet les représentations XML compactes. Si un élément est manquant, la valeur par défaut du type de champ est utilisée. Si le type de champ est une structure, la valeur par défaut est une instance de la structure dont chaque champ est réglé sur les valeurs par défaut.

Le fanion *nillable="true"* doit être défini pour les types ayant une valeur NULL.

Par exemple, la demande du service "Lecture" a une structure de *ReadValueId*. Le schéma XML serait le suivant:

```
<xs:complexType name="ReadValueId">
  <xs:sequence>
    <xs:element name="NodeId" type="tns:NodeId"
      minOccurs="0" nillable="true" />
    <xs:element name="AttributeId" type="xs:int" minOccurs="0" />
    <xs:element name="IndexRange" type="xs:string"
      minOccurs="0" nillable="true" />
    <xs:element name="DataEncoding" type="tns:NodeId"
      minOccurs="0" nillable="true" />
  </xs:sequence>
</xs:complexType>
```

5.3.6 Structures avec champs facultatifs

Les *Structures comportant des champs facultatifs* sont codées sous la forme *xs:complexType*, tous les champs apparaissant de manière séquentielle. Le premier élément est un masque de bits qui spécifie les champs codés. Les bits du masque sont attribués de manière séquentielle aux champs facultatifs dans l'ordre dans lequel ils apparaissent dans la *Structure*.

Pour permettre des représentations XML compactes, n'importe quel champ peut être exclu dans le code XML; de cette manière, les décodeurs doivent attribuer les valeurs par défaut en fonction du type de champ des champs obligatoires.

Par exemple, la *Structure* suivante comporte un champ obligatoire et deux champs facultatifs. Le schéma XML serait le suivant:

```
<xs:complexType name="OptionalType">
  <xs:sequence>
    <xs:element name="EncodingMask" type="xs:unsignedLong" />
    <xs:element name="X" type="xs:int" minOccurs="0" />
    <xs:element name="O1" type="xs:int" minOccurs="0" />
    <xs:element name="Y" type="xs:byte" minOccurs="0" />
    <xs:element name="O2" type="xs:int" minOccurs="0" />
  </xs:sequence>
</xs:complexType>
```

Dans l'exemple ci-dessus, l'EncodingMask a la valeur 3 si O1 et O2 sont codés. Les codeurs doivent régler les bits non utilisés sur 0 et les décodeurs doivent les ignorer.

Si une *Structure* comportant des champs facultatifs est sous-typée, les sous-types étendent l'EncodingMask défini pour le parent.

5.3.7 Unions

Les unions sont codées sous la forme *xs:complexType* contenant *xs:sequence* avec deux entrées.

La première entrée de la séquence est le *SwitchField xs:element* qui spécifie une valeur numérique identifiant l'élément codé dans *xs:choice*. Le nom de l'élément peut être n'importe quel texte valide.

La deuxième entrée de la séquence est *xs:choice* et spécifie les champs possibles. L'ordre dans *xs:choice* détermine la valeur du *SwitchField* lorsque ce choix est codé. Pour le premier élément, le *SwitchField* est réglé sur 1 et pour la dernière valeur, le *SwitchField* est égal au nombre de choix.

Aucun élément supplémentaire n'est admis dans la séquence. Si le *SwitchField* est manquant ou qu'il est réglé sur 0, l'union possède alors une valeur nulle. Les codeurs ou les décodeurs doivent signaler une erreur si une valeur *SwitchField* est supérieure au nombre de champs d'union définis.

Par exemple, l'union suivante comporte deux champs. Le schéma XML serait le suivant:

```
<xs:complexType name="Type1">
  <xs:sequence>
    <xs:element name="SwitchField"
      type="xs:unsignedInt" minOccurs="0"/>
    <xs:choice>
      <xs:element name="Field1" type="xs:int" minOccurs="0"/>
      <xs:element name="Field2" type="tns:Field2" minOccurs="0"/>
    </xs:choice>
  </xs:sequence>
</xs:complexType>
```

5.3.8 Messages

Les *Messages* sont codés comme un *xs:complexType*. Les paramètres de chaque *Message* sont sérialisés de la même manière que les champs d'une *Structure*.

5.4 Codage OPC UA JSON

5.4.1 Généralités

Le *DataEncoding* JSON a été développé pour permettre aux applications OPC UA d'interagir avec les logiciels web et les logiciels d'entreprise qui utilisent ce format. Le *DataEncoding* JSON OPC UA définit les représentations JSON normalisées pour tous les types intégrés OPC UA.

Le format JSON est défini dans la RFC 7159. Il s'agit d'un format partiellement autodéscripteur, car chaque champ possède un nom codé en plus de la valeur, mais JSON ne dispose d'aucun mécanisme pour qualifier les noms avec des espaces de noms.

Le format JSON ne possède aucune norme publiée relative à un schéma pouvant être utilisé pour décrire le contenu d'un document JSON. Les mécanismes schématiques définis dans le présent document peuvent toutefois être utilisés pour décrire les documents JSON. La structure *DataTypeIdDescription* définie dans l'IEC 62541-3 peut précisément définir n'importe quel document JSON qui satisfait aux règles décrites ci-dessous.

Les *Serveurs* qui prennent en charge le *DataEncoding* JSON doivent ajouter des *Nœuds* de *DataTypeIdEncoding* appelés "JSON par défaut" à tous les *DataTypes* pouvant être sérialisés avec le codage JSON. Les *NodeIds* de ces *Nœuds* sont définis par le modèle d'information qui spécifie le *DataTypeId*. Ces *NodeIds* sont utilisés dans les *ExtensionObjects* comme décrit au 5.4.2.16.

Deux cas d'utilisation sont importants pour le codage JSON: les applications Cloud qui utilisent des messages *PubSub* et les *Clients* JavaScript (JSON est le format de sérialisation privilégié pour JavaScript). En cas d'utilisation de l'application Cloud, il est nécessaire que le message *PubSub* soit autonome, ce qui signifie qu'il ne peut pas contenir de références numériques à un tableau d'espace de noms défini en externe. De même, les applications Cloud reposent souvent sur les langages de script pour traiter les messages entrants, de sorte que les artefacts du *DataEncoding* destinés à garantir la fidélité lors du décodage ne sont pas nécessaires. C'est pourquoi ce *DataEncoding* définit un format "irréversible" conçu pour répondre aux besoins des applications Cloud. Les applications, comme les *Clients* JavaScript, qui utilisent le *DataEncoding* pour communiquer avec d'autres applications OPC UA utilisent le format normal ou "réversible". Les éventuelles différences entre le format réversible et le format irréversible sont décrites pour chaque type.

5.4.2 Types intégrés

5.4.2.1 Généralités

Toute valeur NULL d'un type intégré doit être codée sous la forme du littéral JSON 'null' si la valeur est un élément d'une matrice. Si la valeur NULL est un champ contenu dans une *Structure* ou une *Union*, le champ ne doit pas être codé.

5.4.2.2 Boolean

Une valeur de type *Booléen* doit être codée sous la forme du littéral JSON "true" ou "false".

5.4.2.3 Entier

Les valeurs entières autres que *Int64* et *UInt64* doivent être codées sous la forme d'un nombre JSON.

Les valeurs *Int64* et *UInt64* doivent être formatées en nombre décimal codé sous la forme d'une chaîne JSON.

(Voir le codage XML des valeurs 64 bits décrit en 5.3.1.3).

5.4.2.4 Virgule flottante

Les valeurs normales *Float* et *Double* doivent être codées sous la forme d'un nombre JSON.

Les nombres spéciaux à virgule flottante de type infinité positive (INF), infinité négative (-INF) et "not-a-number" (NaN) doivent être représentés par les valeurs "Infinity", "-Infinity" et "NaN" codées sous la forme d'une chaîne JSON. Pour plus d'informations sur les différents types de nombres spéciaux à virgule flottante, voir 5.2.2.3.

5.4.2.5 String

Les valeurs de *Chaîne* doivent être codées sous la forme de chaînes JSON.

Un caractère d'échappement est ajouté à tout caractère non admis dans les chaînes JSON en utilisant les règles définies dans la RFC 7159.

5.4.2.6 DateTime

Les valeurs de *DateTime* doivent être formatées comme spécifié par l'ISO 8601-1:2019 et codées sous la forme d'une chaîne JSON.

Les valeurs *DateTime* qui dépassent les valeurs minimales ou maximales prises en charge sur une plateforme doivent respectivement être codées sous la forme "0001-01-01T00:00:00Z" ou "9999-12-31T23:59:59Z". Lors du décodage, ces valeurs doivent être converties en valeurs minimales ou maximales prises en charge sur la plateforme.

Les valeurs *DateTime* égales à "0001-01-01T00:00:00Z" sont prises comme valeurs nulles.

5.4.2.7 Guid

Les valeurs de *Guid* doivent être formatées comme décrit en 5.1.3 et codées sous la forme d'une chaîne JSON.

5.4.2.8 ByteString

Les valeurs de *ByteString* doivent être formatées en tant que texte Base64 et codées sous la forme d'une chaîne JSON.

Un caractère d'échappement est ajouté à tout caractère non admis dans les chaînes JSON en utilisant les règles définies dans la RFC 7159.

5.4.2.9 XmlElement

Une valeur de *XmlElement* doit être codée sous la forme d'une Chaîne, comme décrit au 5.3.1.9.

5.4.2.10 Nodeld

Les valeurs *Nodeld* doivent être codées sous la forme d'un objet JSON avec les champs définis dans le Tableau 25.

La structure abstraite de *Nodeld* est définie dans l'IEC 62541-3 et comporte trois champs: *Identificateur*, *IdentifieurType* et *INamespaceIndex*. La description de ces champs abstraits est fournie dans le Tableau 25.