

TECHNICAL REPORT

**Information technology – Fibre channel –
Part 314: Avionics environment – Remote direct memory access (FC-AE-RDMA)**

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-314:2013



THIS PUBLICATION IS COPYRIGHT PROTECTED
Copyright © 2013 ISO/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 ISO/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
Fax: +41 22 919 03 00
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 corrigenda or an amendment might have been published.

Useful links:

IEC publications search - www.iec.ch/searchpub

The advanced search enables you 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It also gives information on projects, replaced and withdrawn publications.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 on-line and also once a month by email.

Electropedia - www.electropedia.org

The world's leading online dictionary of electronic and electrical terms containing more than 30 000 terms and definitions in English and French, with equivalent terms in additional languages. Also known as the International Electrotechnical Vocabulary (IEV) on-line.

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: csc@iec.ch.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14655-3 14:2013

TECHNICAL REPORT

**Information technology – Fibre channel –
Part 314: Avionics environment – Remote direct memory access (FC-AE-RDMA)**

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

PRICE CODE



ICS 35.200

ISBN 978-2-83220-659-1

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

CONTENTS

FOREWORD.....	3
INTRODUCTION.....	5
1 Scope.....	6
2 Normative references.....	6
3 Terms, definitions, abbreviations and conventions.....	6
3.1 Terms and definitions.....	6
3.2 Abbreviations and acronyms.....	7
3.3 Editorial conventions.....	7
3.3.1 Overview.....	7
3.3.2 Binary notation.....	8
3.3.3 Hexadecimal notation.....	8
3.3.4 Applicability and use of this technical report.....	8
4 FC-AE-RDMA protocol.....	9
4.1 Introduction to FC-AE-RDMA.....	9
4.2 Remote direct memory access (RDMA) using FCP.....	10
4.2.1 FC-AE-RDMA modifications to FCP.....	10
4.2.2 Logical Unit.....	10
4.2.3 FC-AE-RDMA frame header.....	10
4.2.4 FC-AE-RDMA Features.....	11
Annex A (informative) Example FC-AE-RDMA avionics network profile.....	17
Bibliography.....	27
Figure 1 – FC-AE-RDMA parameter field usage – FCP_CMND IU only.....	11
Table 1 – Summary of implementation and use of features.....	9
Table 2 – FC-AE-RDMA Features.....	12
Table A.1 – FC-FS and FC-AL-2 Features for example avionics network.....	17

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-314:2013

INFORMATION TECHNOLOGY – FIBRE CHANNEL –

Part 314: Avionics environment – Remote direct memory access (FC-AE-RDMA)

FOREWORD

- 1) ISO (International Organization for Standardization) and IEC (International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards. Their preparation is entrusted to technical committees; any ISO and IEC member body interested in the subject dealt with may participate in this preparatory work. International governmental and non-governmental organizations liaising with ISO and IEC also participate in this preparation.
- 2) In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.
- 3) The formal decisions or agreements of IEC and ISO 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 and ISO member bodies.
- 4) IEC, ISO and ISO/IEC publications have the form of recommendations for international use and are accepted by IEC and ISO member bodies in that sense. While all reasonable efforts are made to ensure that the technical content of IEC, ISO and ISO/IEC publications is accurate, IEC or ISO cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 5) In order to promote international uniformity, IEC and ISO member bodies undertake to apply IEC, ISO and ISO/IEC publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any ISO/IEC publication and the corresponding national or regional publication should be clearly indicated in the latter.
- 6) ISO and IEC provide no marking procedure to indicate their approval and cannot be rendered responsible for any equipment declared to be in conformity with an ISO/IEC publication.
- 7) All users should ensure that they have the latest edition of this publication.
- 8) No liability shall attach to IEC or ISO or its directors, employees, servants or agents including individual experts and members of their technical committees and IEC or ISO member bodies 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 of, use of, or reliance upon, this ISO/IEC publication or any other IEC, ISO or ISO/IEC publications.
- 9) 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.
- 10) Attention is drawn to the possibility that some of the elements of this Technical Report may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

The main task of IEC technical committees is to prepare International Standards. However, a technical committee may propose the publication of a technical report when it has collected data of a different kind from that which is normally published as an International Standard, for example "state of the art".

ISO/IEC TR 14165-314, which is a technical report, has been prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

A list of all currently available parts of the ISO/IEC 14165 series, under the general title *Information technology – Fibre channel*, can be found on the IEC web site.

This Technical Report has been approved by vote of the member bodies, and the voting results may be obtained from the address given on the second title page.

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

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-314:2013

INTRODUCTION

The Fibre Channel SCSI-3 Remote Direct Memory Access (FC-AE-RDMA) Technical Report defines a set of features necessary to implement a real-time Fibre Channel network (switched fabric or arbitrated loop) supporting the FC-AE-RDMA Upper Level Protocol.

FC-AE-RDMA is intended to support bi-directional communication between two N_Ports in a constrained and carefully defined environment, typical of avionics applications. The intended usage is avionic command, control, instrumentation, simulation, signal processing and sensor/video data distribution. These application areas are characterized by a variety of requirements, among them a need for high reliability, fault tolerance and deterministic behaviour to support real-time control/response.

FC-AE-RDMA follows the SCSI-3 FCP standard in its definition of the services necessary to support low-latency, low overhead communication between elements of a mission-critical avionics system. The key feature of FC-AE-RDMA is that it allows an Initiator to read data from or write data to a remote Target memory in peer-to-peer mode (similar to SCSI-3 processor device type) with lower latency.

This technical report is divided into 4 clauses:

Clause 1 is the scope.

Clause 2 enumerates the normative references.

Clause 3 describes the terms, definitions, abbreviations, and conventions.

Clause 4 defines the FC-AE-RDMA Upper Level Protocol. This clause lists features defined in the SCSI-3 FCP standard and indicates whether the features are Required, Prohibited, Allowed, or Invocable in this Technical Report. This Technical Report places certain restrictions on SCSI-3 FCP in order to improve support for low latency, real-time applications. This clause also defines some new features for FC-AE-RDMA that are not defined in SCSI-3 FCP.

Annex A gives an example of a profile for the FC-FS and FC-AL-2 standards for an example avionics Fibre Channel network that uses FC-AE-RDMA.

INFORMATION TECHNOLOGY – FIBRE CHANNEL –

Part 314: Avionics environment – Remote direct memory access (FC-AE-RDMA)

1 Scope

This part of ISO/IEC 14165 defines the FC-AE-RDMA Upper Level Protocol. FC-AE-RDMA follows the SCSI-3 FCP standard in its definition of the services necessary to support low latency, low overhead communication between elements of a mission-critical avionics system.

This part of ISO/IEC 14165 is intended to serve as an implementation guide to maximize the likelihood of interoperability between conforming implementations. This technical report Prohibits or Requires features that are optional and Prohibits the use of some non-optional features that are referenced in some standards (see Clause 2).

In addition, this technical report simplifies implementations and their associated documentation, testing and support requirements.

This technical report does not define internal characteristics of conformant implementations. Nonetheless, it incorporates features from the standards listed in Clause 2.

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.

ISO/IEC 14165-122, *Information technology – Fibre channel – Part 122: Arbitrated loop-2 (FC-AL-2)*¹

ISO/IEC 14165-251, *Information technology – Fibre channel – Part 251: Framing and signalling (FC-FS)*²

ISO/IEC 14776-411, *Information technology – Small computer system interface-3 (SCSI-3) – Part 411: Architecture model (SCSI-3 SAM)*³

3 Terms, definitions, abbreviations and conventions

3.1 Terms and definitions

For the purposes of this document the following terms and definitions apply.

¹ INCITS.332 – 1999 – *Information technology – Fibre channel Arbitrated Loop-2 (FC-AL-2)*.

² ANSI INCITS 373 – 2003 *Information technology – Framing and Signalling (FC-FS)*.

³ ANSI INCITS 270 – 1996 *Information technology – SCSI-3 Architecture Model (SAM)*.

NOTE Some definitions from the glossary or body of other standards are included here for easy reference.

3.1.1

implicit loop initialization

method of defining and specifying the AL_PA and optional Position Map of L_Ports by means other than the explicit use of the Loop Initialization procedure defined in ISO/IEC 14165-122 (FC-AL-2)

Note 1 to entry: Specific methods of implicit Loop Initialization are not defined in this technical report.

3.1.2

initiator

device containing application clients that originate RDMA Messages to be processed by a Target

Note 1 to entry: In ANSI/INCITS X3.269-1996 (FCP) an initiator is defined as “a SCSI device containing application clients that originate device service requests and task management functions to be processed by a target SCSI device”.

3.1.3

logical unit

unit that selects among specific memory regions within a Target

Note 1 to entry: Defined in ANSI/INCITS X3.269-1996 (FCP) as “a Target resident entity that implements a device model and executes SCSI commands sent by an application client”.

3.1.4

RDMA message

equivalent to an FCP Exchange, except that it consists only of an FCP_CMND IU, 0 or more FCP_DATA IUs and an FCP_RSP IU

3.1.5

target

device that receives RDMA Messages and directs such messages to specified memory areas

Note 1 to entry: Defined in ANSI/INCITS X3.269-1996 (FCP) as “a SCSI device that receives SCSI commands and directs such commands to one or more logical units for execution”.

3.2 Abbreviations and acronyms

Abbreviations and acronyms applicable to this technical report are listed below. Abbreviations and acronyms for commonly used terms defined in referenced standards are not listed here.

FC-AE-RDMA	Fibre Channel – Avionics Environment – Remote Direct Memory Access, (this technical report)
FCP	Fibre Channel Protocol for SCSI
RDMA	Remote Direct Memory Access
IU	Information Unit
SCSI	Small Computer System Interface (as defined by the SCSI-3 family of standards)
ULP	Upper Level Protocol

3.3 Editorial conventions

3.3.1 Overview

In this technical report, a number of conditions, mechanisms, sequences, parameters, events, states, or similar terms that do not have their normal English meaning are printed with the following conventions.

- The first letter of each word in uppercase and the rest lowercase (e.g., Exchange, Class, etc.).
- A term consisting of multiple words, with the first letter of each word in uppercase and the rest lowercase and each word separated from the other by an underscore (_) character. A word may consist of an acronym or abbreviation, which would be printed in uppercase. (e.g., NL_Port, Transfer_Length, etc.).

All terms and words not conforming to the conventions noted above have the normal technical English meanings.

Numbered items in this technical report do not represent any priority. Any priority is explicitly indicated.

In all of the figures, tables and text of this technical report, the most significant bit of a binary quantity is shown on the left side. Exceptions to this convention are indicated in the appropriate clauses.

The term “shall” is used to indicate a mandatory rule. If such a rule is not followed, the results are unpredictable unless indicated otherwise.

The term “should” is used to indicate flexibility of choice with a strongly preferred alternative; equivalent to the phrase “it is strongly recommended”.

The term “may” is used to indicate flexibility of choice with no implied preference; equivalent to “may or may not”.

The fields or control bits that are not applicable shall be set as required by the defining standard or technical specification or technical report.

If a field or a control bit in a frame is specified as not meaningful, the entity that receives the frame shall not check that field or control bit.

In several tables within this report, there is a column on the right side of the table labelled “Notes”. These notes are NORMATIVE and shall be considered requirements of this report.

In the event of conflict between the text, tables and figures in this technical report, the following precedence shall be used: tables (highest), text and figures (lowest).

3.3.2 Binary notation

Binary notation may be used to represent some fields. Single bit fields are represented using the binary values 0 and 1. For multiple bit fields, the binary value is enclosed in single quotation marks followed by the letter b. For example, a four-byte field containing a binary value may be represented as '00000000 11111111 10011000 11111010'b.

3.3.3 Hexadecimal notation

Hexadecimal notation may be used to represent some fields. When this is done, the value is enclosed in single quotation marks and preceded by the word hex. For example, a four-byte field containing a binary value of '00000000 11111111 10011000 11111010'b is shown in hexadecimal format as hex '00 FF 98 FA'.

3.3.4 Applicability and use of this technical report

The usual definitions of the following words do not apply. The following definitions shall be read carefully.

Required: If a feature or parameter value is Required, it means that it shall be used between compliant implementations. Compliant implementations are required to implement the feature. Interoperability is not guaranteed if Required features are not implemented. Each Required feature will include a note that describes the condition(s) in which the feature shall be used.

Invocable: If a feature or parameter value is Invocable, it means that it may be used between compliant implementations. Compliant implementations are required to implement the feature. Invocable is different from Required in that an implementation may use the feature if needed, but is not required to use it. No discovery process is necessary prior to use of an Invocable feature.

Allowed: If a feature or parameter value is Allowed, it means that it may be used between compliant implementations. Compliant implementations are not required to implement the feature. Typically, the potential user of an Allowed feature may determine if an implementation supports it via an Invocable discovery process.

Prohibited: If a feature is Prohibited, it means that it shall not be used between compliant implementations. This report does not Prohibit the implementation of features, only their use between compliant implementations. Interoperability is not guaranteed if Prohibited features are used.

Table 1 summarizes the above definitions.

Table 1 – Summary of implementation and use of features

Term	Implementation	Use
Required	Shall	Shall
Invocable	Shall	May
Allowed	May	May
Prohibited	May	Shall not

Table 2 and Table A.1 list features described in the various Fibre Channel standards and technical reports. These tables indicate whether the features are Required, Prohibited, Invocable, or Allowed for compliance with this technical report; or whether a parameter is Required to be a particular value for compliance with this report. Features or parameters that are not listed do not affect the interoperability of FC-AE-RDMA devices.

The following legend is used for table entries in Clause 4:

- 'R' Required
- 'I' Invocable
- 'A' Allowed
- 'P' Prohibited
- 'n' the parameter shall be set to this value
- 'X' this parameter has no required value; any value is Allowed
- '-' this parameter or feature is not meaningful

4 FC-AE-RDMA protocol

4.1 Introduction to FC-AE-RDMA

This technical report defines an Upper Level Protocol (ULP), FC-AE-RDMA. FC-AE-RDMA is based on Fibre Channel Protocol for SCSI (FCP). This clause lists features described in ANSI/INCITS X3.269-1996 (FCP) and indicates whether the features are Required, Prohibited,

Allowed, or Invocable in FC-AE-RDMA. This clause also defines some new features for FC-AE-RDMA that are not defined in FCP.

FC-AE-RDMA follows the FCP standard in its definition of the services necessary to support low latency, low overhead communication between elements of a mission-critical avionics system. The key feature of FC-AE-RDMA is that it allows an Initiator to read data from or write data to a remote Target memory in peer-to-peer mode (similar to SCSI-3 processor device type) with lower latency. All FC-AE-RDMA compliant devices shall be capable of operating as a Initiator and a Target.

NOTE There is another protocol similar in operation to FC-AE-RDMA. It is the Fibre Channel Lightweight Protocol described in 4.5 of INCITS TR 31-2002 (FC-AE). While they are similar, there are sufficient differences to warrant their own protocols, especially in their use of the FCP_XFER_RDY IU and the Command Descriptor Block.

4.2 Remote direct memory access (RDMA) using FCP

4.2.1 FC-AE-RDMA modifications to FCP

FC-AE-RDMA places certain restrictions on FCP in order to improve support for low latency, real-time applications. For example, the FCP_XFER_RDY IU is not used and data immediately follows the FCP_CMND IU.

In FC-AE-RDMA, a Target shall predefine one or more regions of its local memory for RDMA use. An Initiator then reads or writes data directly into the Target memory. The contents of the FCP_CMND IU are not redefined, but usage of the Parameter field of the header is modified. The Parameter field in the FCP_CMND IU is used to indicate that the transfer is FC-AE-RDMA and to specify the base address that the Target uses as a starting point for the read or write operation.

Note that this protocol does not prohibit a target memory region from overlapping with any other target memory region. A system designer may choose to restrict memory region usage depending on the application.

4.2.2 Logical Unit

Logical Units may be used to identify separate memory regions within a Target. The FCP_LUN field of the FCP_CMND IU is used to identify a specific Logical Unit for each Message.

All Targets shall provide at least one RDMA memory region assigned to Logical Unit 0 (FCP_LUN = 0). Targets may optionally define additional memory regions assigned to Logical Units with FCP_LUN values of 1-255. The additional LUNs do not have to be contiguous. These Logical Units are designated in the LUN/Target field of FCP_LUN_0 (First Level) as defined in ISO/IEC 14776-411 (SAM). All other fields within FCP_LUN shall be 0 (FCP_LUN = hex '00 xx 00 00 00 00 00').

For avionics applications, provisions should be made for multiple target buffers as well as for limiting the range of the buffers in order to provide a memory protection scheme.

FC-AE-RDMA does not define an explicit method for an Initiator to discover Target device characteristics, including supported Logical Units, their assigned LUN values, memory region size, etc. System implementers shall use predetermined Target addressing or define implementation-specific discovery methods.

4.2.3 FC-AE-RDMA frame header

All fields shall follow the standard ISO/IEC 14165-251 (FC-FS) and ANSI/INCITS X3.269-1996 (FCP) definitions except for

- a) the RX_ID field for Write Data Transfers for FCP_DATA IUs (see 4.2.4.6) and

b) the Parameter field of the FCP_CMND IU defined in this clause.

Figure 1 defines the FC-AE-RDMA usage of Header Word 5, the Parameter field, for the FCP_CMND IU.

The Relative Offset Present bit in F_CTL field of the Frame Header of the FCP_CMND IU and the FCP_RSP IU shall be set to '0'b for FC-AE-RDMA. The Relative Offset Present bit in F_CTL field of the Frame Header of the FCP_DATA IU shall be set to '1'b for FC-AE-RDMA.



CODE hex 'A6' = FC-AE-RDMA with Target Notification
 CODE hex 'A7' = FC-AE-RDMA without Target Notification

Figure 1 – FC-AE-RDMA parameter field usage – FCP_CMND IU only

The upper 7 bits of the Code field (Bits 31-25 of the Parameter field) shall be used to indicate that the FCP_CMND IU is initiating an FC-AE-RDMA message. A specific non-zero value ('1010 011'b) is used to identify an FC-AE-RDMA message.

The low-order bit of the Code field (Bit 24 of the Parameter field) shall be used to indicate whether the Target application shall be notified that the FC-AE-RDMA message has been processed. In either case the FCP_RSP IU is transmitted. Use of this feature is implementation-specific.

If any value is received in the Code field other than hex 'A6' or hex 'A7' behaviour of the Target is undefined.

The Base Address field (Bits 23-0 of the Parameter field) shall be used to specify the Base Address for the FC-AE-RDMA message. The Target shall add the Base Address and the Relative Offset of each FCP_DATA IU to the beginning of the Logical Unit's memory region to calculate the actual address to read/write the FCP_DATA IU payload.

Note that in order to prevent ambiguous behaviour, connecting FCP-only devices and FC-AE-RDMA devices is administratively Prohibited.

4.2.4 FC-AE-RDMA Features

4.2.4.1 Overview

Table 2 defines the features for FC-AE-RDMA compliant nodes. Devices that are compliant with FC-AE-RDMA shall comply with the mandatory features defined in FCP unless noted herein. Table 2 identifies optional features that represent potential interoperability concerns and indicates whether they are Required, Invocable, Allowed, or Prohibited for FC-AE-RDMA compliance.

Table 2 – FC-AE-RDMA Features

FC-AE-RDMA Features	Nx_Port	Notes
Process Login Service Parameters Page:		
Type Code = hex '08'	R	
PRLI Flags		
Process Associator Validity Bits = 1	P	
Establish Image Pair = 1	R	
Service Parameters		
Data Overlay Allowed = 1	P	
Initiator Function = 1	R	
Target Function = 1	R	
Command/Data Mixed Allowed = 1	P	
Data/Response Mixed Allowed = 1	P	
Read XFER_RDY Disabled = 1	R	
Write XFER_RDY Disabled = 1	X	^a
Process Login Service Parameter Response Page:		
Type Code = hex '08'	R	
Process Associator Validity Bits = 1	P	
PRLI Response Flags		
Image Pair Established = 1	R	
Response Code = '0001'b	I	Request Executed
Service Parameter Response		
Data Overlay Allowed = 1	P	
Initiator Function = 1	R	
Target Function = 1	R	
Command/Data Mixed Allowed = 1	P	
Data/Response Mixed Allowed = 1	P	
Read XFER_RDY Disabled = 1	R	
Write XFER_RDY Disabled = 1	X	^a
FCP_CMND IU		
Header Word 5 Available for Protocol use	R	
FCP_LUN = 0	R	
FCP_LUN > 0	A	
FCP_CNTL		
Byte 1, Task Codes		
SIMPLE_Q = '000'b	I	
Others	A	
Byte 2, Task Management Flags = 0		
Byte 3, Execution Management Codes		
READ DATA = 1	I	^b
WRITE DATA = 1	I	^b
FCP_CDB	X	
FCP_DL	R	
FCP_XFER_RDY IU	P	

FC-AE-RDMA Features	Nx_Port	Notes
FCP_DATA IU	R	
Class 2 RX_ID for writes = hex '00 00' (Initiator)	A	See 4.2.4.6
Class 2 RX_ID for writes = hex '00 00' (Target)	I	See 4.2.4.6
Class 3 RX_ID for writes = hex '00 00'	R	See 4.2.4.6
FCP_RSP IU	R	
FCP_STATUS		
Byte 2, Validity and Status Indicators		
FCP_RESID_UNDER = 1	I	
FCP_RESID_OVER = 1	I	
FCP_SNS_LEN_VALID = 0	R	
FCP_RSP_LEN_VALID = 0	R	
Byte 3, SCSI Status Byte		
Status Code = hex '00'	I	GOOD
Status Code = hex '08'	I	BUSY
Status Code = hex '28' ^c	I	TASK SET FULL/QUEUE FULL
Status Code = hex '5A' ^d	A	LOGICAL UNIT NOT AVAILABLE
Status Code = hex '5B' ^d	A	WRITE PROTECT ERROR
Status Code = hex '5C' ^d	A	READ PROTECT ERROR
Status Code = hex '5D' ^d	A	BUFFER OVERRUN
FCP_RESID	I	^e
FCP_SNS_LEN = 0	R	
FCP_RSP_LEN = 0	R	
The instructions in the column Notes are NORMATIVE and are requirements of this technical report.		
Key		
A Allowed		
I Invocable		
P Prohibited		
R Required		
X this parameter has no required value; any value is Allowed		
^a The PRL Write XFER_RDY Disabled field does not have meaning for FC-AE-RDMA interfaces and shall be ignored. FCP_XFER_RDY IUs shall never be sent in FC-AE-RDMA. This behaviour is required to support early implementations. For future implementations, the bit should be set to '1'b.		
^b Execution Management Code Read Data shall be set to '1'b or Write Data shall be set to '1'b. Both set to '0'b or both set to '1'b is Prohibited for all FCP_CMD IUs.		
^c Defined as QUEUE FULL in SCSI-2. Defined as TASK SET FULL in SCSI-3.		
^d Unique to FC-AE-RDMA, reserved in SCSI-3.		
^e Only meaningful if FCP_RESID_OVER or FCP_RESID_UNDER set to '1'b		

4.2.4.2 Port Login service parameters

FC-AE-RDMA compliant ports shall perform Port Login (PLOGI) as described in FCP.

4.2.4.3 Process Login service parameters

FC-AE-RDMA compliant ports shall perform Process Login (PRLI) as described in FCP.

An FC-AE-RDMA port that receives a PRLI request with acceptable parameters (as shown in Table 2) shall respond with a valid PRLI Accept to complete the Process Login. A valid PRLI Accept shall contain the values shown in Table 2, including an Accept Response Code of Request Executed ('0001'b) and the Image Pair Established bit set to '1'b."

If an FC-AE-RDMA responder receives a PRLI Request with unexpected values it shall respond with an LS_RJT or a PRLI Accept with an Accept Response Code other than Request Executed.

If an FC-AE-RDMA originator receives a PRLI Accept with unexpected parameter values it should attempt to log out with the responder by sending a Process Logout command (PRLO).

4.2.4.4 FCP_CMND IU

4.2.4.4.1 Overview

In FC-AE-RDMA compliant devices, Header Word 5 of the FCP_CMND IU is used to denote an FC-AE-RDMA transfer, as described in 4.2.3.

For read transfers, the FCP_CMND IU shall be a T1 IU. For write transfers, the FCP_CMND IU shall be a T2 IU (see FCP).

4.2.4.4.2 Logical Unit Number (FCP_LUN)

FC-AE-RDMA supports SCSI Level 1 peripheral addressing, which defines the first two bytes of the FCP_LUN field, as described in 4.2.2. If a Target defines multiple target buffers, FCP_LUN shall be used to select the desired buffer.

4.2.4.4.3 Control Field (FCP_CNTL)

The Task Attribute Simple_Q is Invocable. Other Task Attributes are Allowed.

Initiators shall not set Task Management Flags to 1 in transmitted FCP_CMND IUs. If an FCP_CMND IU is received with one or more Task Management Flags set, Targets shall ignore the Flags and maintain or automatically resume normal Target operation.

For a read operation, the Initiator shall set the Execution Management Code READ DATA bit to '1'b. For a write operation, the Initiator shall set the Execution Management Code WRITE DATA bit to '1'b.

4.2.4.4.4 Command Descriptor Block (FCP_CDB)

The contents of the 16-byte FCP_CDB are not defined in FC-AE-RDMA and its use is beyond the scope of this technical report.

4.2.4.4.5 Data Length (FCP_DL)

The FCP_DL field shall contain the exact number of data bytes that will be transferred in the payload of the FCP_DATA IUs. The FCP_DL field may be set to zero. When FCP_DL is zero, no FCP_DATA IUs are transmitted.

4.2.4.5 FCP_XFER_RDY IU

FCP_XFER_RDY IUs are Prohibited in FC-AE-RDMA.

4.2.4.6 FCP_DATA IU

For read transfers, the FCP_DATA IU shall be a I3 IU when data is transferred (see ANSI/INCITS X3.269-1996 (FCP)).

For write transfers, the FCP_DATA IU shall use T6 or T7 IUs when data is transferred.

- When the Initiator transfers all of the data in one Sequence, a single T6 IU shall be used.
- When the Initiator transfers all of the data in more than one Sequence, all Sequences except the last Sequence shall be T7 IUs. The last IU shall be a T6 IU (see ANSI/INCITS X3.269-1996 (FCP)).

FCP_DATA IUs are used as defined in FCP with the following exception:

The Relative Offset (RLTV_OFF) Field value in the first frame of the first FCP_DATA IU of an exchange shall be set to zero. Subsequent frames shall be transmitted with RLTV_OFF set to the value of RLTV_OFF in the previous frame + the length of the payload of the previous frame.

FCP_DATA IU frame headers are used as defined in FC-FS with the following exceptions:

For Class 3 write data transfers, the Initiator shall set the RX_ID field in the header to hex '00 00'.

For Class 2 write data transfers, the Initiator may use the RX_ID value returned in the Acknowledgment frame (ACK) following the FCP_CMND IU, or the Initiator may set the RX_ID to hex '00 00'. The Target shall accept either value as valid.

Note that this non-standard behaviour of substituting hex '00 00' for the RX_ID in Classes 2 or 3 is required for continuing support of existing implementations.

4.2.4.7 FCP_RSP IU

Usage of fields in the FCP_RSP IU for FC-AE-RDMA are defined in Table 2.

For all transfers, the FCP_RSP IU shall be an I4 IU (see FCP).

FCP_RSP IUs shall be exactly 24 bytes in length as Length of the SCSI Sense Information (FCP_SNS_LEN) and Length of the FCP Response Information (FCP_RSP_LEN) shall be zero.

Residual Count (FCP_RESID) and status indicators FCP_RESID_UNDER and FCP_RESID_OVER shall be used to report mismatches between FCP_DL and the number of bytes transferred in associated FCP_DATA IUs. FCP_RESID is only meaningful when FCP_RESID_UNDER or FCP_RESID_OVER is set to 1. If either FCP_RESID_UNDER or FCP_RESID_OVER is set, data may be written to memory except that any data beyond the FCP_DL (in the case of an overrun) shall be discarded.

A mismatch between FCP_DL and the number of bytes actually transferred constitutes a protocol error and the receiver of the Data IU(s) associated with that Message should not assume that the content is valid.

FC-AE-RDMA defines three SCSI status codes (FCP_Status, Byte 3) that are defined in SCSI-3 (hex '00', '08' and '28') and four status bytes that are reserved in SCSI-3 (hex '5A', '5B', '5C' and '5D').

A SCSI status code of GOOD (hex '00') shall be used to indicate successful completion of the FC-AE-RDMA transfer.

SCSI status codes of BUSY (hex '08') or QUEUE FULL (hex '28') shall be used to indicate that the Target has memory or processing resource limitations that preclude an FC-AE-RDMA transfer from successfully completing.

A SCSI status code of LOGICAL UNIT NOT AVAILABLE (hex '5A') is optional and may be used to indicate that the Target Logical Unit is unavailable for an FC-AE-RDMA transfer.

A SCSI status code of WRITE PROTECT ERROR (hex '5B') is optional and may be used to indicate that the Target RDMA Logical Unit is write protected and therefore the FC-AE-RDMA transfer cannot occur.

A SCSI status code of READ PROTECT ERROR (hex '5C') is optional and may be used to indicate that the Target RDMA Logical Unit is read protected and therefore the FC-AE-RDMA transfer cannot occur.

A SCSI status code of BUFFER OVERRUN (hex '5D') is optional and may be used to indicate that the length of the transfer (plus offset) would exceed the Target RDMA Logical Unit and therefore the FC-AE-RDMA transfer cannot occur.

SCSI status bytes of BUSY and QUEUE FULL typically indicate a temporary Logical Unit resource limitation. It is likely that simply reinitiating the transfer would be successful. Systems should be designed with sufficient performance to minimize the occurrence of these status byte values during normal operation.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-314:2013

Annex A (informative)

Example FC-AE-RDMA avionics network profile

A.1 Summary

This annex contains a profile of the FC-FS and FC-AL-2 standards for an example avionics Fibre Channel network which uses FC-AE-RDMA. This profile was originally part of the main body of the FC-AE technical report in the FC-AE-RDMA section. It was moved to an annex in this technical report because the lower level features profiled herein are, for the most part, independent of the Upper Level Protocol that is used. Therefore, the use of the terms Required, Invocable, Allowed and Prohibited are only guidelines here.

A.2 FC-FS and FC-AL-2 features for FC-AE-RDMA

A.2.1 Overview

Table A.1 is intended primarily for private arbitrated loop and fabric topologies. Devices that are compliant with FC-AE-RDMA shall comply with the mandatory features defined in FC-FS and FC-AL-2, unless noted herein. Table A.1 identifies optional features that represent potential interoperability concerns and indicates whether they are Required, Invocable, Allowed, or Prohibited. In addition to interoperability concerns, this profile addresses certain features that are needed in order to achieve the performance necessary for real-time avionics systems. More information is provided after the table.

Because FC-AE-RDMA was developed for use in low-latency real-time environments, it is expected that most applications will require commensurate N_Port performance. Typical performance requirements would specify stringent limits on the time between IUs within an exchange (~1 ms to 2 ms), relatively short values of E_D_TOV (~5 ms to 10 ms) and high data transfer rates (~50 Mbytes/s). Developers of components and equipment should verify specific performance requirements with system integrators.

In many cases, these features have Login Parameters associated with them. For features that are Required or Invocable, the corresponding login parameters shall indicate that the feature is supported. For features that are Prohibited, the corresponding login parameters may indicate that the feature is supported, even though the feature shall not be used by compliant implementations. For features that are Allowed, the corresponding login parameters shall reflect whether or not the feature is supported by the implementation.

Table A.1 – FC-FS and FC-AL-2 Features for example avionics network

FC-FS/FC-AL-2 Features	Nx_Port	Fx_Port	Notes
Link Protocols (point-to-point links)			
Link Initialization	R	R	
Online to Offline	A	A	
Link Failure	I	I	
Link Reset	I	I	
Loop Port State Machine (LPSM) (arbitrated loop)	A	A	
Loop Protocols			
Loop Initialization	I	I	
Loop Port Bypass	I	I	
Loop Port Enable	I	I	

FC-FS/FC-AL-2 Features	Nx_Port	Fx_Port	Notes
Link Reset	I	I	
Loop Port Bypass Circuit	A	A	
Hard Addressing	I	I	
Loop Port Position Mapping	A	A	
Dynamic Half-Duplex	A	A	
Login BB_Credit ≥1	I	I	
Programmable Loop Tenancy	A	A	
Broadcast Open Replicate – OPN(fr)	A	A	
Old Port State	I	I	
Request Old Port State – REQ (old-port)	I	I	
Fabric Login			
Explicit Login	I	I	
S_ID = hex '00 00 00' for N_Ports S_ID = hex '00 00 AL_PA' for NL_Ports	R	-	
Implicit Login	A	A	
Fabric Login – Common Service Parameters			
BB_Credit ≥2	R	R	
Max BB Receive Data Field Size ≥2 048		I	
Alternate BB_Credit Management = 0 for Fabric Alternate BB_Credit Management = 1 for Fabric Attached Loop	R	R	
Multicast	P	P	
Broadcast	P	P	
Hunt Group	P	P	
FLOGI Payload Length = 0	R	R	116 bytes
R_A_TOV = 20ms	R	R	
E_D_TOV = 10ms	R	R	
R_T_TOV Value = 1	R	R	≤1ms
Fabric Login – Class Specific Service Parameters			
Class of Service			
Class 1	P	P	
Class 2	A	A	
Class 3	R	R	
Class 4	P	P	
Class 6	P	P	
Sequential Delivery = 1	R	R	
Priority = 1	A	I	
Clock Sync Primitive Capable	P	P	
Clock Sync ELS Capable	I	I	Server is in the fabric
N_Port Login			
Explicit Login	R	-	
Implicit Login	P	-	
N_Port Login – Common Service Parameters			
BB_Credit ≥ 2	R	-	point-to-point
Max BB Receive Data Field Size ≥2 048	I	-	

FC-FS/FC-AL-2 Features	Nx_Port	Fx_Port	Notes
Continuously Increasing Relative Offset = 1	R	–	
Random Relative Offset = 1	P	–	
(Continuously Increasing) SEQ_CNT	I	–	
Alternate BB_Credit Mgmt = 0 for Point-to-Point Alternate BB_Credit Mgmt = 1 for Arbitrated Loop	R	R	
PLOGI Payload Length = 0	R	–	116 bytes
Total Concurrent Sequences ≥ 16	I	–	
Relative Offset by Information Category	R	–	per FCP
E_D_TOV = 10 ms	R	R	
N_Port Login – Class Specific Service Parameters			
Class of Service			
Class 1	P	–	
Class 2	A	–	
Class 3	R	–	
Class 4	P	–	
Class 6	P	–	
N_Port Login – Class 2 Service Parameters			
Max BB Receive Data Field Size $\geq 2\ 048$		–	
Priority	A	I	
Initial Process Associator = '00'b	R	–	Process Associators not used in FCP or FC-AE-RDMA
Recipient X_ID Interlock	–	–	
ACK_0 Initiator/Recipient Capable	I	–	
ACK Generation Assistance	A	–	
Initiator Clock Sync ELS Capable	P	–	Server is in the fabric
Recipient Clock Sync ELS Capable	A	–	
Concurrent Sequences ≤ 16	I	–	
N_Port EE_Credit ≥ 2	I	–	
Open Sequences per Exchange ≥ 1	I	–	
N_Port Login – Class 3 Service Parameters			
Max BB Receive Data Field Size $\geq 2\ 048$	I	–	
Priority bit	A	I	
Initial Process Associator = '00'b	R	–	Process Associators not used in FCP or FC-AE-RDMA
Initiator Clock Sync ELS Capable	P	–	Server is in the fabric
Recipient Clock Sync ELS Capable	I	–	
Concurrent Sequences ≥ 16	I	–	
Open Sequences per Exchange ≥ 1	I	–	
Fabric Reject Reason Codes			
hex '01' Invalid D_ID	–	I	
hex '03' N_Port not available, temporarily	–	A	
hex '04' N_Port not available, permanently	–	I	
hex '05' Class of service not supported	–	I	
hex '16' Login required	–	I	

FC-FS/FC-AL-2 Features	Nx_Port	Fx_Port	Notes
Others	-	A	Within the bounds of Class 2
Fabric Busy Reason Codes			
hex '1' Fabric is Busy	-	A	
hex '3' D_ID busy with a Class 1 connection	-	P	
Port Reject Frames	I	-	Reason codes not specified
Port Busy Frames	I	-	Reason codes not specified
Well Known Address Support			
hex 'FF FF FF' (Broadcast)	A	A	
hex 'FF FF FE' (F_Port Server)	-	I	
hex 'FF FF FD' (Fabric Controller)	-	A	
hex 'FF FF FC' (Directory/Name Server)	A	A	
hex 'FF FF F6' (Clock Sync Server)	P	I	
Implicit N_Port login	R	-	
Multicast capable	I	I	For CSU
Others	P	-	
Class of Service to/from WKA			
Class 3	I	I	
Class 2	P	P	
others	P	P	
Basic Link Services			
BA_ACC	I	I	
BA_RJT	I	I	
ABTS	I	I	
Others	P	P	
Extended Link Services			
CSR	I	I	To the Clock Sync Server only
CSU	I	I	
FLOGI	I	I	
LOGO	I	I	
PLOGI	R	-	
PRLI	R	-	
PRLO	I	-	
RLS	I	I	
Others	A	A	
The instructions in the column Notes are NORMATIVE and are requirements of this technical report.			
Key A Allowed I Invocable P Prohibited R Required X this parameter has no required value; any value is Allowed - this parameter or feature is not meaningful			

A.2.2 Link protocols

Support of basic Link Initialization described in FC-FS is Required and will be used at power up or upon re-initialization on point-to-point links. Link Failure and Link Reset protocols are Invocable and will be used as needed.

Online-to-Offline protocol is Allowed but not required to be implemented because typically an Offline state is not needed in avionics systems. The equipment will operate until power is removed.

A.2.3 Arbitrated loop

This profile was written primarily for switched fabrics, which may include attached loops. Although there is no inherent restriction on the use of a Private loop with the FC-AE-RDMA profile that topology is not addressed specifically.

The Loop Port State Machine (LPSM) is Allowed. If the LPSM is implemented, then the port is an NL_Port or an FL_Port. Otherwise, the port is an N_Port or an F_Port. The items indented under LPSM are only applicable if the LPSM is implemented.

Loop Initialization is Invocable (but not Required) because this enables the System Designer to implement a system using “Implicit Loop Initialization” in order to achieve faster start-up times. Implicit Loop Initialization shall be implemented in such a way as to insure interoperability with devices that only support “Explicit Loop Initialization”. If Implicit Loop Initialization is used it is up to the System Designer to figure out how to provide all the necessary information to the nodes, such as AL_PA assignments.

If Explicit Loop Initialization is used, Loop Initialization Hard Addressing (LIHA) is Invocable during AL_PA Assignment.

NOTE 1 Uniquely defined Hard Addressing is very desirable in avionics networks because every node gets the same address every time the loop is initialized.

If Explicit Loop Initialization is used, it may be sped up by not building a Loop Port Position Map. Although Loop Port Position Mapping is Required in FC-AL-2, it is only Allowed here and is generally not needed and is discouraged for embedded avionics systems because all of the participating nodes and their positions are known ahead of time.

NOTE 2 The use of a centralized hub may be desirable, depending on the number of nodes involved. A hub provides a centralized location for wire routing and the ability to isolate faulty nodes or wiring without affecting the rest of the loop.

Support for Loop Port Bypass and Loop Port Enable Primitive Sequences, which control access to the loop, is Invocable. Support for the optional Bypass circuit is Allowed.

Dynamic Half Duplex is Allowed and encouraged because it improves bandwidth utilization.

Specifying Login BB_Credit ≥ 1 is Invocable. Login BB_Credit is used to set the available BB_Credit every time the loop circuit is opened. If login BB_Credit is ≥ 1 and the transmitting port remembers the value, the transmitting port doesn't have to wait for an R_RDY to begin sending frames.

Programmable Loop Tenancy is Allowed. This implementation feature gives the system designer the ability to set the maximum Loop Tenancy of a loop port in firmware/hardware if desired. Alternatively, the loop tenancy may be controlled in software.

Dynamic Half Duplex, Login BB_Credit and Loop Tenancy are only some of the parameters that affect loop performance. Other parameters that should be considered include:

- whether to allow outstanding R_RDYs when the loop circuit is closed (referred to as “unbalanced” BB_Credit),
- the amount of data to be transferred during each loop tenancy,
- the arbitration time,
- the round trip time and
- the time required to empty the buffers.

Broadcast Open Replicate is Allowed.

The capability to operate in OLD_PORT state is desirable in order for a 2-node pair to operate in point-to-point mode as opposed to operating as a 2-node loop. There are two ways to enter OLD_PORT state:

- a) one node in the pair doesn't support LPSM; or
- b) both nodes support LPSM, but OLD_PORT state is requested via the REQ (old-port) L_Port Control.

The ability to request Old Port state is Invocable.

A.2.4 Fabric Login

A.2.4.1 Overview

Explicit Fabric Login is Invocable (but not Required) because this enables the System Designer to implement a system using Implicit Login. One of them (Explicit or Implicit Fabric Login) shall be used. Explicit Fabric Login will normally be used. If Implicit Login is used it is up to the System Designer to figure out how to provide all the necessary information to the fabric and the attached nodes.

NOTE If attached Loops are not present; it is highly desirable to disable the attempt to perform Loop Initialization if possible in order to achieve faster start-up times.

The S_ID of the Nodes shall be hex '00 00 00' at login for N_Ports requesting Fabric Login. This will reduce the complexity of the N_Ports. The S_ID shall be hex '00 00 AL_PA' for NL_Ports.

Either Registered or Non-registered names are Allowed for Ports, Nodes and Fabrics. The only requirement is that all names in the system shall be unique by type.

A.2.4.2 Fabric Login – Common service parameters

BB_Credit shall be at least 2 in order to allow data frames to be received and processed simultaneously by N_Ports and F_Ports.

The Max BB Receive Data Field Size shall be at least 2 048 bytes.

Alternate BB_Credit Management is required and shall be used in Arbitrated Loops, including Fabric Attached Loops, but shall not be used in non-loop topologies.

The System Designer will determine the values of E_D_TOV and R_A_TOV for the network, but in this example network E_D_TOV shall be 10 ms and R_A_TOV shall be 20 ms.

The default value for R_T_TOV in commercial nodes is 100 ms. The short R_T_TOV value, which enables quicker detection of Loss of Sync, is recommended for avionics networks and is therefore Invocable. For purposes of this network example, the short R_T_TOV value is specified as ≤ 1 ms with a goal of achieving 100 μ s in the future. This is a deviation from FC-FS, which specifies that the short value will be 100 μ s.

FCP, and therefore FC-AE-RDMA, does not support the use of Multicast. However, Multicast may be used for certain ELS commands, so it is Allowed. If Multicast is used, pre-defined Multicast Group IDs (i.e. alias addresses) shall be used.

NOTE The use of pre-defined addresses saves time, reduces complexity in the nodes, eliminates the need for an Alias Server and guarantees that the alias addresses will always be the same from one mission to the next.

The use of Broadcast and Hunt Groups is Prohibited.

Extended length Fabric login payloads shall not be used.

A.2.4.3 Fabric Login – Class specific service parameters

Class 3 shall be supported and shall be used for most Exchanges between nodes and for all Exchanges between nodes and servers in the fabric. Class 2 is Allowed for Exchanges between nodes. Class 3 shall be used for F_Port Login, even when Class 2 is supported for Exchanges between nodes.

Sequential Delivery is Required for fabrics and nodes. Sequential Delivery simplifies re-assembly and error detection at the recipient node and is not difficult to achieve in avionics fabrics because they generally contain only 1 or 2 switches. This example fabric is no exception.

Support for Priority is Invocable for fabrics and Allowed for nodes. The fabric shall route frames with the Priority bit set ahead of frames without the bit set. System Designers may require the use of Priority for enhanced performance, but since they are not required to use it, this parameter is only Invocable in the fabric.

Clock Synchronization using the ELS Method is Invocable for both fabrics and nodes. In order to reduce complexity of compliant nodes, the Clock Sync Server is required to be within the fabric in this profile.

A.2.5 N_Port Login

A.2.5.1 Overview

Explicit N_Port Login is Required. See 4.2.4.2.

Either Registered or Non-registered names are Allowed for Ports and Nodes. The only requirement is that all names in the system shall be unique by type.

A.2.5.2 N_Port Login – Common service parameters

As in FLOGI, BB_Credit shall be at least 2. In N_Port Login, this applies to point-to-point topologies. This field is also used to set the Login BB_Credit for Loop devices operating on the same arbitrated loop. See A.2.3.

The maximum BB Receive Data Field Size shall be at least 2 048 bytes.

Continuously Increasing Relative Offset set to '1'b is Required.

Alternate BB_Credit Management is required in Arbitrated Loops, but shall not be used in non-loop topologies (i.e. Point-to-Point).

Nodes shall be capable of supporting ≥ 16 Total Concurrent Sequences. This requirement is included more for performance reasons rather than interoperability purposes.