

FC Frame Encapsulation WG Last Call

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC 2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as Reference material or to cite them other than as "work in progress".

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Abstract

This ips (IP Storage) working group draft contains all the working group last call comments received regarding:

draft-ietf-ips-fcencapsulation-06.txt

The proposed disposition of each comment also is recorded in this draft.

Summary of Comments

Technical:	17 Comments, resulting in about	11 Changes
Editorial:	60 Comments, resulting in about	42 Changes

Totals:	77 Comments, resulting in about	53 Changes

Table Of Contents

1. Comments from David Black 2

2. Comments from Mallikarjun Chadalapaka 19

3. Comments from Rob Elliott 23

4. Additional Changes Discovered After WG Last Call 32

1. Comments from David Black

=====

Comment 1

-- Section 1 - Scope

The organization responsible for the Fibre Channel standards (NCITS Technical Committee T11) has determined that some functions and modes of operation are not interoperable to the degree required by the IETF. This draft includes applicable T11 interoperability determinations in the form of restrictions on the use of this encapsulation mechanism.

[E] Is there an official citation for this statement? It really needs one to be published in an archival unchangeable format such as an RFC.

Accepted resulting in the following changes

Add reference to FC-MI. Change "NCITS" to "INCITS".

Comment 2

-- Section 2 - Encapsulation Concepts

FC frames have several possible lengths.

[E] Should read "variable length" or something like that - this implies several possible choices of fixed length, which is incorrect.

Accepted as follows

Replace the sentence with: "FC frames are variable length."

Comment 3

-- Section 2 - Encapsulation Concepts

To facilitate transporting FC frames over TCP the native FC frame needs to be contained in (encapsulated in) a slightly larger structure as shown in figure 1.

[E] The use of TCP in this context is overly restrictive. This encapsulation is in principle applicable to any means of transport over IP, including TCP, SCTP, UDP, and carrier pigeon :-), even though in practice all the initial uses will use TCP.

Accepted as follows

Change "over TCP" to "over an IP based transport such as TCP".

Comment 4

-- Section 2 - Encapsulation Concepts

The format and content of an FC frame is described in the FC

[E] "is" --> "are"

Accepted

Comment 5 Technical

-- Section 3 - The FC Encapsulation Header

-- Section 3.1 - FC Encapsulation Header Format

The values in the Protocol# field are assigned by IANA [8]. The following values are known to be in use:

- FCIP -- TO BE ASSIGNED by IANA
- iFCP -- TO BE ASSIGNED by IANA

[T] Delete the text starting with "The following values" and insert a forward reference to the IANA Consideration section.

Accepted as written.

Comment 6 Technical

- Section 3 - The FC Encapsulation Header
- Section 3.1 - FC Encapsulation Header Format

FC Encapsulation receivers may compare the Protocol# and -Protocol# fields as an additional verification that an FC Encapsulation Header is being processed.

FC Encapsulation receivers may compare the Version and -Version fields as an additional verification that an FC Encapsulation Header is being processed.

[T] Those "may"s are misleading. I think "SHOULD" is appropriate, but I could accept "SHOULD"s that only applied when the CRC is not valid.

Accepted as follows

Replace the two cited sentences with:

FC Encapsulation receivers SHOULD either validate the CRC or compare the Protocol# and -Protocol# fields to verify that an FC Encapsulation Header is being processed according to a policy defined by the encapsulating protocol.

FC Encapsulation receivers SHOULD either validate the CRC or compare the Version and -Version fields to verify that an FC Encapsulation Header is being processed according to a policy defined by the encapsulating protocol.

As per comment 8, sentences similar to those shown above will be added to the -Flags and -Frame Length descriptions.

Comment 7 Technical

- Section 3 - The FC Encapsulation Header
- Section 3.1 - FC Encapsulation Header Format

Flags (bits 31-26 in word 3): The Flags bits provide information about the usage of the FC Encapsulation Header as shown in figure 3.

Note: Implementers are advised to consult the specifications of protocols that use this header to determine how each individual protocol uses the bits in the Flags field.

[T] The "Note:" paragraph is part of the CRCV issue (see below), and probably needs to be deleted as part of resolving that issue. This paragraph also has the additional problem in that it implies that protocol specific uses of the reserved flags are allowed, which is not the case.

Accepted

Comment 8 Technical

- Section 3 - The FC Encapsulation Header
- Section 3.1 - FC Encapsulation Header Format

Reserved (Flags, bits 31-27 in word 3): These bits are reserved for use by future versions of the FC Encapsulation and SHALL be set to zero on send. Protocols employing this encapsulation MAY require checking for zero on receive, however doing so has the potential to create incompatibilities with future versions of this encapsulation.

[E] Second sentence is poorly worded. Suggested rewrite: Protocols employing this encapsulation SHOULD ignore the Reserved bits on receive in order to avoid creating incompatibilities with possible future versions of this encapsulation. I believe this change is editorial, and it also applies to the -Flags and -Frame Length fields.

Rejected

This change is not editorial. It is technical and specifically recommends against some of the validity checking described in FCIP. The working group argued this issue several times and (I thought) agreed that checking the version number was sufficient to know that the reserved flags must be zero.

The last sentence of the comment is a misnomer with respect to the rest of the comment, however it makes sense when applied to comment 6.

Comment 9 Technical

- Section 3 - The FC Encapsulation Header
- Section 3.1 - FC Encapsulation Header Format

CRCV (CRC Valid Flag, bit 26 in word 3): A CRCV bit value of one indicates that the contents of the CRC field are valid. A CRCV bit value of zero indicates that CRC are invalid. Some protocols may always check the CRC without regard for the state of this bit. The value of the CRCV bit SHALL be constant for all FC Encapsulation Headers sent on a given TCP connection.

[T] The "Some protocols may always check the CRC ..." is the CRCV issue that Mallikarjun also found and that has been problematic in the past. I believe that what's going on here is that all protocols have to check the Protocol#, and once that's been checked, the implementation knows whether there's supposed to be a CRC there and hence doesn't need to look at CRCV. In practice this won't cause problems, as including the CRC when it's not supposed to be there is harmless, and failing to include it when it should be there will almost certainly cause a CRC check failure.

I offer a proposal to resolve this by expanding the Protocol# registry that IANA will create so that each registered protocol must supply not only its name and an RFC reference, but also whether the protocol uses (Yes) or does not use (No) the CRC in this header. The above text could then be revised to make the CRCV check at the receiver OPTIONAL in all cases because its value can be inferred from the protocol#.

Rejected in principle, with some changes required

At the Nashua interim, someone wanted a client protocol to be free to use CRC or not according to operating environments (e.g., lab-local vs. network attachment). The proposed definition of usage by IANA based on protocol would eliminate this option.

Actions to be taken

It must be noted that the content of Annex A also conflicts with this result from the Nashua interim. To correct that, the following text from Annex A must be replaced:

"CRC

"Protocols employing this encapsulation SHALL either:

- "1) Require a valid CRC to be sent and the CRCV Flag bit to be sent as one, or
- "2) Require the CRC field to be sent as zero and the CRCV Flag bit to be sent as zero."

The Annex A CRC discussion (shown above) will be replaced with:

"Protocols employing this encapsulation SHALL define the procedures and policies necessary for verifying that an FC Encapsulation Header is being processed."

Also, make the change described in the response to comment 35.

Comment 10

- Section 3 - The FC Encapsulation Header
- Section 3.1 - FC Encapsulation Header Format

[E] Also need to generalize away from TCP connection to allow possible future use with other transports.

Accepted, resulting in the following changes

- 1) In the description of CRCV: change "TCP connection" to "connection";
- 2) In Section 4:
 - change "TCP-connected elements" to "IP-connected elements";
 - change "traverse the TCP connection" to "traverse the connection";
 - change "injected into a TCP connection" to "injected into a connection"; and
- 3) In Section 5.3: change "transmission over TCP" to "transmission over an IP Network";

Comment 11 Technical

-- Section 3 - The FC Encapsulation Header
-- Section 3.1 - FC Encapsulation Header Format

[T] Here or in the description of the Protocol Specific fields, a warning to implementers is needed says some sort of error checking redundancy (e.g., the ones complements found elsewhere in the header) SHOULD (or MUST) be used when the CRC is not used. This warning should be duplicated in Section 3.2.1. This is a technical comment, but should not be controversial.

Rejected

Specific statements of action have been added to each applicable field as per comment 6.

Comment 12 Technical

-- Section 3 - The FC Encapsulation Header
-- Section 3.1 - FC Encapsulation Header Format

Time Stamp [integer] and Time Stamp [fraction] (words 4 and 5): The two Time Stamp fields contain time at which the FC Encapsulated frame was sent as known to the sender. The format of integer and fraction Time Stamp word values is specified in Simple Network Time Protocol (SNTP) Version 4 [9]. The contents of the Time Stamp [integer] and Time Stamp [fraction] words SHALL be set as described in section 4.

[E] For convenience, it might be good to summarize those formats here with an indication that [9] is the normative authority. I don't feel strongly about this, though.

[T] We have a problem here - RFC 2030 is Informational, and hence can't be referenced in a normative fashion from a standards track document. I'll talk to Ralph offline about how to get around this.

Accepted resulting in the following changes

- 1) Copy the definitions of the two time stamp words from RFC 2030 to this draft (estimated to be 2 paragraphs);
- 2) Copy any necessary ancillary definition text from RFC 2030 to this draft (estimated to be no more than 5 paragraphs); and
- 3) Make the reference to RFC 2030 information, both in the body text and in a Informative References section (which will have to be added).

Comment 13

-- Section 3 - The FC Encapsulation Header
-- Section 3.1 - FC Encapsulation Header Format

CRC (word 6): When the CRCV Flag bit is zero, the CRC field SHALL contain zero. When the CRCV Flag bit is one, the CRC field SHALL contain a CRC for words 0 to 5 of the FC Encapsulation Header computed using the polynomial, initial value, and bit order defined for Fibre Channel in FC-FS [3]. Using this algorithm, the bit order of the resulting CRC corresponds to that of FC-1 layer. The CRC transmitted over the IP network shall correspond to the equivalent value converted to FC-2 format as specified in FC-FS.

[E] I realize that FC-FS is the latest and greatest version of the FC frame standard, BUT, referencing a project in progress for this sort of basic CRC mechanism is an invitation to procedural problems. Can this reference be changed to FC-PH accompanied by a note that FC-FS is supplanting FC-PH, but will make *no* changes in this area? Note that I'm comfortable with the earlier reference to FC-FS for frame contents.

Accepted (Partially)

T11 has clearly stated that FC-PI and FC-FS are the preferred documents over FC-PH. The statement takes the form of a refusal to process FC-PH for international standardization, with the preferred recourse being to process FC-PI and FC-FS when they are available.

Furthermore, referencing FC-PH really requires references to FC-PH-2 and FC-PH-3, since FC-PH is not a fully correct document in and of itself. In other words, there is a rats nest here.

It is not a matter of FC-FS being the latest and greatest. T11 has unambiguously stated that FC-PI and FC-FS are the only true path with very clear reasons for that decision.

Action to be taken:

In the References section, the note will be added following the FC-FS reference: "Note: The Fibre Channel frame structure and CRC features referenced by this draft, while formally described in FC-FS, are substantially unchanged from similar features described in Fibre Channel Physical and Signaling Interface (FC-PH), ANSI X3.290-1994, June 1, 1994."

[Comment 14 Technical](#)

-- Section 3.2.1

[T] The warning that the protocol-specific fields SHOULD (or MUST) be protected by redundancy needs to go here.

[Accepted.](#)

[Comment 15](#)

-- Section 3.2.1

Redundancy based header validation can be built from simple logic (e.g., XORs and comparisons). Header validation based on redundancy also is a step wise process in that the first word is validated, then the second, then the third and so on. A decision that a candidate header is not valid may be reached before the complete header is available.

[E] First sentence is superfluous and probably should be deleted as it's rather hardware-oriented.

[Accepted with the following results](#)

Replace the cited paragraph with: "Header validation based on redundancy is a step wise process in that the first word is validated, then the second, then the third and so on. A decision that a candidate header is not valid may be reached before the complete header is available."

[Comment 16](#)

-- Section 3.2.2

CRC based header validation employs a straight forward algorithm (e.g., compute the CRC for all bytes preceding the CRC word and compare the results to the CRC word's contents). The number of comparisons required to perform CRC validation is exactly one and the method for computing the CRC is well known with proven implementations.

[E] Last sentence is superfluous and probably should be deleted as it's rather hardware-oriented.

[Accepted with the following results](#)

Replace the cited paragraph with: "Header validation based on the CRC defined in section 3.1 requires computing the CRC for all bytes preceding the CRC word, and comparing the results to the CRC word's contents."

[Comment 17](#)

-- Section 4 - Measuring Fibre Channel frame transit time

To comply with FC-FS [3], an FC Fabric must specify and limit the lifetime of a frame.

[E] Same comment as before about referencing FC-FS. Can this be changed to reference FC-PH with a note that FC-FS won't change this ... or is FC-FS tinkering with things here?

[Rejected see response to comment 13.](#)

Comment 18 Technical

-- Section 4 - Measuring Fibre Channel frame transit time

When originating an encapsulated frame, an entity that does not support transit time calculation SHALL always set the Time Stamp [integer] and Time Stamp [fraction] fields to zero. When receiving an encapsulated frame, an entity that does not support transit time calculation SHALL ignore the contents of the Time Stamp words. The protocol SHALL specify whether or not implementation support is required.

[T] This is about "MUST/SHOULD/MAY implement". Need a similar requirement on the protocol to specify "MUST/SHOULD/MAY use" and under what conditions.

Accepted with the following results

Add the following sentence: "The protocol SHALL specify those conditions under which a received encapsulated frame MUST have its transit time checked before forwarding."

Comment 19 Technical

-- Section 4 - Measuring Fibre Channel frame transit time

The policy for processing frames while in the Unsynchronized state SHALL be defined by the protocol specification, including whether or not the entity may continue to send and receive frames from the IP network.

[T] On the receive side, this condition appears to be specified in the wrong direction. Receiving frames from the IP network cannot possibly cause problems, the issues are in forwarding (stale) frames into FC.

Accepted resulting in the following changes

- 1) Change "processing" to "forwarding"; and
- 2) Remove "including whether ..." to the end of the sentence.

[Comment 20](#)

-- Section 4 - Measuring Fibre Channel frame transit time

When de-encapsulating a frame, an entity in the Synchronized state SHALL:

[E] While the sub-bullets are correct, they leave a reader unfamiliar with FC somewhat high and dry. I would include a "for example" in both a) and b), along the lines of:

- a) For example, if a calculated transit time exceeds a value that could cause the frame to violate FC maximum time in transit limits (Time Out Values), the protocol may specify that the frame is to be discarded.
- b) For example, a protocol may specify that frames for which transit time cannot be determined are never to be forwarded over FC.

[Accepted with changes](#)

Everything except the phrase "(Time Out Values)" will be incorporated as written.

[Comment 21](#)

-- Section 4 - Measuring Fibre Channel frame transit time

[T] At the end of this section, it would be good to warn protocol designers that well-designed protocols are unlikely to accomplish useful communication when the communicating entities are in different states, and hence protocol designers need to consider how to coordinate state transitions, especially the Unsynchronized to Synchronized transition on startup and an unexpected Synchronized to Unsynchronized transition (e.g., caused by loss of contact with an external time service). This is related to some issues that Mallikarjun found.

[Accepted but only in principle](#)

The problem is not coordinating states between the two entities. The problem is keeping both entities in the Synchronized state as much of the time as possible.

Little, if anything, is accomplished by adding protocol overhead to coordinate state transitions.

This is not to say that the comment lacks merit.

Action to be taken:

Add the following note at the end of the section: "Note: For most purposes, communication between entities is possible only while in the Synchronized state."

Comment 22

- Section 5 - The FC frame
- Section 5.1 - FC frame content

As shown in figure 4, the FC frame content is defined as the data between the EOF and SOF delimiters (including the FC CRC) after conversion from FC-1 to FC-2 format as specified by FC-FS [3].

[E] This needs some more explanation. The important things that need to be said are:

- FC uses the same 8b/10b encoding as Gigabit Ethernet in which each 8 bit byte is transmitted using 10 bits on the wire for reasons that include redundancy and low level timing synchronization between sender and receiver.
- All discussion of FC frame content in this draft is at the 8b level prior to 8b->10b expansion on send or after 10b->8b reduction on receive.

The Gigabit Ethernet reference is particularly important in increasing accessibility of this document to a network-savvy, but new to FC audience.

Accepted but only in principle

The 8b/10b statement is no more accurate for Fibre Channel than it is for Ethernet. Ten Gigabit Fibre Channel will use 64b/66b encoding, the same as ten Gigabit Ethernet. Such a statement must be worded as an example.

Action to be taken:

Add the following paragraph at the end of the section: "In the conversion to the FC-0 physical transport, an encoding is applied to the FC frame content (e.g., 8b/10b encoding like that used in Gigabit Ethernet) for reasons that include redundancy and low level timing synchronization between sender and receiver. All discussion of FC frame content in this document is at the 8-bit byte level, prior to the application of any such encoding."

Comment 23

-- Section 5.3 - FC SOF and EOF

The FC frame content is composed of 8-bit bytes that can be translated directly for transmission over TCP. The FC SOF and EOF [3] require 8b/10b special characters that cannot be translated directly to 8-bit bytes, encoded values are required.

[E] I think this paragraph needs to be moved to Section 5.1, and replaced with a sentence here that refers back to it. One important editorial change is "8b/10b special characters that cannot be translated directly to 8-bit bytes" should be changed to "10b special characters that have no 8b equivalents" or something like that.

Accepted

Comment 24 Technical

-- Section 5.3 - FC SOF and EOF

SOF (bits 31-24 and bits 23-16 in word 0): The SOF fields contain the encoded SOF value selected from table 1.

[T] As we've learned from the Class 4 adventure, this table is subject to change/extension. IANA will need to manage it, and will need some sort of allocation guidelines to remain consistent with whatever mechanism produced this peculiar set of values. While we probably don't want to allow value recycling, we may want to write some text dealing with retiring values (making them no longer usable). This also applies to the EOF values in Table 2.

Rejected

The SOF/EOF codes are stable and have not changed for at least three years. They have been implemented in numerous products for tunneling FC over ATM and SONET. The only instability is the editor's lack of understanding about which SOF/EOF codes are required for FC Class 4 operation.

The codes are assigned by T11 and involving IANA in there assignment would constitute a conflict of jurisdictions.

Comment 25

-- Section 5.3 - FC SOF and EOF

Note: FC-BB-2 [6] lists SOF and EOF codes not shown in table 1 and table 2 (e.g., SOFi1 and SOFn1). However, FC-MI [7] identifies these codes as not interoperable, so they are not listed in this specification.

[T] There are a couple of problems here. If FC-BB-2 has assigned values to SOF and EOF encodings that MUST NOT be used with FCIP, then we need to instruct IANA to reserve and not allocate those values. As part of allocating future values in this table, we need to (1) instruct the author(s) of the draft/RFC doing the allocation to ensure that T11.3 review of the proposed allocation is obtained (2) that the IPS WG chair(s) (if the IPS WG still exists) and the Transport ADs are informed of this review, and (3) that IANA allocates the values approved by T11.3 as opposed to choosing values. Since Murali's been appointed as T11's official liaison to IETF, I think it's his responsibility to suggest a coordination process.

Accepted only because a platform is needed for an editorial change

As per the response to comment 25, IANA must not be involved in assigning SOF/EOF codes.

Actions to be taken

The SOF and EOF tables will be modified to add a column for "Class" and each SOF or EOF value will have one of the following entered in the new column: "2", "3", "2 & 3", "2, 3 & 4", or "4".

The new column will allow FCIP and iFCP to reference Class 2, Class 3, and Class 4 SOF/EOF values in their statements of what the protocol supports.

Comment 26

-- Section 7 - Normative References

I would really like to remove the normative reference to FC-FS, substituting FC-PH with a note that FC-FS will replace FC-PH. I don't object to an FC-FS reference where it's really needed, but the portions of FC-FS that this draft relies on are sufficiently basic and stable that an FC-PH reference will make their stability clear. The FC-BB-2 and FC-MI references for SOF and EOF codes need to become non-normative as part of setting up the IANA registry and management process. The FC-SW-2 reference may not need to be normative here.

Rejected see response to comment 13.

Comment 27

-- Section 7 - Normative References

RFC 1700 is almost certainly the wrong reference to instruct IANA on what procedures to follow. See RFC 2434 for guidance on this topic, although it may not be necessary to reference it.

Accepted

Comment 28

-- ANNEX A - Protocol Requirements

[E] I think this should be an Appendix, rather than an Annex. Some changes may be in order here based on the above comments.

Accepted resulting in the following changes

- 1) Change "Annex A" to "Appendix A" and "Annex B" to "Appendix B", while remembering to correct the Table of Contents too;
- 2) Add discussion of IANA assignment of CRC usage, per the resolution of comment 9; and
- 3) In item d) change "processing" to "forwarding".

[Comment 29](#)

-- ANNEX B - IANA Considerations

[T] This needs to be made somewhat more explicit and direct. IANA is looking for simple straightforward instructions roughly of the form "IANA is instructed to do <X>". in particular, the following sentence is a problem:

Standards action on this RFC should be accompanied by IANA assignment of the following two Protocol# values:

It should read something like:

In addition to creating the FC Encapsulation Protocol Number Registry, the standards action of this RFC allocates the following two values from this registry:

While one normally asks IANA to allocate values, the exception is that when creating a registry, one can instruct IANA as to what the initial contents are (i.e., a new registry does not have to be created empty).

[Accepted](#)

[Comment 30](#)

-- ANNEX B - IANA Considerations

[T] Also, earlier comments suggest that the Protocol# registry needs to be expanded with a CRC field (Yes/No) and that registries need to be created for the SOF and EOF values.

[Rejected See comment 9 and comment 24.](#)

Comment 31

-- ANNEX B - IANA Considerations

It is requested that the ips working group chairs or the Transport Services area directors be notified when any new Protocol# value assignment is requested.

[T] Given that an approved RFC is required, this sentence seems redundant. If the intent was notification of the IPS WG chairs and/or ADs when a an I-D draft is submitted that will cause a Protocol# assignment if/when approved as an RFC, the language needs to say that and should be rephrased to require notification of the IP Storage WG chairs (don't use WG acronyms here) and notification of the Transport ADs instead in the case that the IPS WG does not exist or is not active.

Accepted, delete the sentence

Comment 32

-- ANNEX B - IANA Considerations

[T] Also see previous comments about needing to set up an IANA registry for SOF and EOF values. I'll work with Ralph on crafting the right IANA instructions.

Rejected see comment 24.

2. Comments from Mallikarjun Chadalapaka

=====

Comment 33 Technical

- Section 3.1, page 4. For the Protocol# values for FCIP and iFCP, the Annex B should instead be referred.

Accepted see comment 5.

Comment 34 Technical

- Section 3.1, pages 4 & 5. I notice that all the ones complement fields (-Protocol#, -Version etc.) are described as "contains the ones complement" as opposed to "SHALL contain ones complement", whereas the corresponding non-1's complement fields have the SHALL wording. Any reasons for this distinction?

Accepted

Change -xxx descriptions to use SHALL.

Comment 35 Technical

- Section 3.1, page 5, CRCV bit description.
"Some protocols may always check the CRC without regard for the state of this bit."
I am troubled by the literal implication of this sentence. Why would that be so?
Would the encapsulating protocol not mandate CRCV to be set to valid always instead? It seems like the purpose of defining a common encapsulation format and associated semantics is watered down for no stated reason...

Accepted

Delete the cited sentence.

The following response is provided in response to the question asked in the last paragraph of the comment. FCIP mandates that CRCV be zero. iFCP mandates that CRCV be one.

Comment 36

- Section 3.2.1, page 7. S/b "step wise" w/ "stepwise".

Accepted

Comment 37

- Section 4, page 7.

"The protocol SHALL specify whether or not implementation support is required."

A general comment on usage of the term "protocol" here and in other areas - I would recommend using "encapsulating protocol" instead to make it easier (or perhaps use "Protocol" may be...) for the reader to follow the usage.

Accepted, use "encapsulating protocol"

Comment 38

- Section 4, page 8. Since there is no mention of a notification frame to announce an entity's transition into/out of the Synchronized state, I assume it's possible and even anticipated that there may be times when one of the two entities may be in Unsynchronized state even while the operational encapsulating protocol requires Synchronized operation. The expectation is that the state rules (and encapsulating protocol-specified rules) should cause this type of start-up/transient scenario to be correctly sorted out. Is that right?

Inquiry

Yes it is anticipated that one entity may be in the Unsynchronized state and thus discarding some FC frames. Presumably, the entity will return to the Synchronized state quickly or close the connection.

It is not clear that any requirements need to be state for interoperability. It is definitely undesirable to add protocol overhead to coordinate the synch/unsynch state without there being a well demonstrated need.

Comment 39 Technical

- I think it might be useful to add a statement in this section along the lines of - If the encapsulating protocol mandates Synchronized operation, the entity MUST NOT accept TCP connections on the well-known port(s) unless the entity is in the Synchronized state.

Rejected

Since encapsulating protocols are allowed to specify operation in the Unsynchronized state, specifying this level of detail about how Synchronized operation is handled over reaches the bounds of this specification.

Comment 40 Technical

- Section 4, page 9, Synchronized state rules. I think this should also address what is to be done in case there's a case of "bad synchronization" when Time Stamp words are valid. For ex., when the received value is smaller than the received entity's timebase, I assume it would result in arriving at a huge transit time. While the huge transit time may cause the frame to be discarded, it isn't clear to me what may cause the TCP connection drop and a re-synch.

Rejected

RFC 2030 leads to the belief that "bad synchronization" should be very low probability event. Therefore, the FC Frame Encapsulation draft has chosen to ignore it.

Comment 41

- Section 4, page9, Synchronized state rules.
"If both Time Stamp words have a value of zero, the receiving entity SHALL process the de-encapsulated frame without computing the transit time. The disposition of the frame and any other actions by the recipient SHALL be defined by the protocol specification."

I am rather perplexed by the usage of the words here. While saying that the frame shall be "processed", this also seems to leave it up to the encapsulating protocol to define if it needs to be processed (as reaffirmed by rule (e) in Annex A). One change that makes it clear to me is:

S/b "SHALL process the de-encapsulated frame"
w/ "SHALL de-encapsulate the frame".

Accepted

Comment 42

- Section 7
It may be useful to add informative references to FCIP and iFCP.

Accepted with the following results

Informative references will be added for FCIP and iFCP. The references will be to the Internet-Drafts and will include the words, "Work in Progress".

Comment 43

- Section 9, page 13. S/b "no long working" w/ "no longer working".

Accepted

Comment 44

- Annex B, page 15. It isn't clear to me from this sentence if the Protocol# values (1 & 2) are temporarily assigned by this draft for now.

"Standards action on this RFC should be accompanied by IANA assignment of the following two Protocol# values:"

Inquiry

It is said that IANA will not assign the Protocol# values until presented with an RFC (not an internet draft) that instructs them to do so. Therefore, the current Protocol# values are requests to IANA (or perhaps instructions to IANA).

3. Comments from Rob Elliott

=====

Comment 45

Title page

I assume change bars will be gone from the final version.

Inquiry

The change bars appear only in the PDF file. The normative document in IETF terms is the TXT file, where there are no change bars.

Comment 46

Abstract

"common encapsulation" s/b "common Fibre Channel encapsulation"

Accepted

Comment 47

1 Scope
"NCITS" s/b "INCITS"

Accepted

Comment 48

All section headers
Should each word in a section header be capitalized or not? This document has a mix.

(3.2.2 doesn't capitalize validation, 4 doesn't capitalized frame transmit time, etc.)

Accepted as follows

All section headers will be modified to follow the capitalization in the header for section 3.1.

Comment 49

2. Encapsulation Concepts
"The Fibre Channel frame has a CRC that provides error detection..."
s/b "The Fibre Channel frame includes a Cyclic Redundancy Check (CRC) code that provides error detection..."

Accepted

Comment 50

Figure 1 Title
"FC frame Encapsulation" s/b "Encapsulated FC Frame"

Response

This figure and the text describing it are concerned with how a FC frame is encapsulated. Thus the figure title. The fact that the result is an Encapsulated FC Frame should not be introduced in a figure title.

[Comment 51](#)

3.1 Figure 2

Where is word defined as 32 bits?

[Response](#)

This figure seems to define that concept clearly. It would be a shame to add a glossary to the draft simply to contain this one definition.

[Comment 52](#)

3.1 FC Encapsulation Header Format

"Protocol (bits 31-24 in word 0):" s/b "Protocol# (bits 31-24 in word 0):"

[Accepted](#)[Comment 53](#)

3.1 FC Encapsulation Header Format

RE: "TO BE ASSIGNED by IANA"

When are these protocol values filled in?

[Inquiry](#)

See the response to comment 44.

[Comment 54](#)

3.1 FC Encapsulation Header Format

"The Version field SHALL contain 0x1 ..." s/b "The Version field SHALL contain 0x01 ..." since Version is an 8-bit field.

[Accepted](#)[Comment 55](#)

All instances of "ones complement"

"ones complement" s/b "one's complement"

Google reports 10,000 "one's complement" vs 3700 "ones complement"

[Accepted](#)

Comment 56

3.1 FC Encapsulation Header Format

"...i.e., the usage of this word is defined..." s/b "...i.e., the usage of these words is defined..." because there are two protocol specific words.

Accepted

Comment 57

3.1 FC Encapsulation Header Format

Regarding: "Protocols employing this encapsulation MUST NOT make use of the Reserved Flags bits in any fashion other than that described here." "here" s/b "by this encapsulation". "Here" implies that future versions are excluded.

Accepted

Comment 58

3.1 FC Encapsulation Header Format

"A CRCV bit value of zero indicates that CRC are invalid." s/b "A CRCV bit value of zero indicates that the contents of the CRC field are invalid."

Accepted

Comment 59

3.1 FC Encapsulation Header Format

Frame Length is a greater than 1 byte quantity. Which bit is the MSB?

There's discussion later on page 9 inside the FC frame section, but endianness should be covered before the encapsulation header is described.

Rejected

Per <http://www.ietf.org/ID-nits.html>:

"Historically, RFCs have specified byte and bit order according to a US DoD rule which made byte zero be the first byte in an address range, and bit zero be the most significant bit in a word or field. For example, you will find drawings like this one (from RFC 791) in many RFCs: when you make drawings like it, you should follow the same rule. Label your bit positions,

bit zero is the most significant bit, and place the first addressable byte at the upper left-hand corner."

Observance of these rules is enforced for all RFCs. Therefore, additional specifics are unnecessary.

See also comment 77.

Comment 60

3.1 FC Encapsulation Header Format

Regarding: "...the FC Encapsulation Header SHALL always be word-aligned;..."

Replace "SHALL" with "is".SHALL doesn't seem right here. There's no option to not make it aligned, since the format is fixed length. We don't say the CRC field shall be one word long - it just is one word long.

Accepted

Comment 61

3.1 FC Encapsulation Header Format

"...contain time..." s/b "...contain the time..."

Accepted

Comment 62

3.1 FC Encapsulation Header Format

Should the field names in SNTF ("Seconds" and "Seconds fraction") be referenced? It's not immediately obvious which words correspond.

Accepted as follows

Change all occurrences of "[integer]" to "[Seconds]" and all occurrences of "[fraction]" to "[Seconds Fraction]".

Comment 63

3.1 FC Encapsulation Header Format

SNTF describes its Seconds field formats with bit 0 on the left as the MSB.

I assume FCencap wants to use bit 31 on the left as the MSB. How can this be made clearer?

Rejected

MSB is always on the left and the bit numbering will be changed to match SNTF as per comment 77.

Comment 64

3.1 FC Encapsulation Header Format

"...CRC for words 0 to 5 of the FC Encapsulation Header computed using the polynomial, initial value, and bit order defined for Fibre Channel in FC- FS..." s/b "...CRC for words 0 to 5 of the FC Encapsulation Header computed using the equations, polynomial, initial value, and bit order defined for Fibre Channel in FC-FS..."

As indicated by iSCSI CRC discussion, the FC "initial value" assumes a certain implementation. I think if you add the word "equations" it implies that the FC annex be followed more completely.

Accepted

Comment 65

3.2 FC Encapsulation Header Validation

I would include a hyphen in both Redundancy-based and CRC-based (since they act as modifiers to "mechanism")

Rejected

Editor's prerogative.

Comment 66

3.2.2 CRC based FC Encapsulation Header validation

Straightforward is one word

Accepted

Comment 67

4. Measuring Fibre Channel frame transit time

Is it worthy of a note that this field runs out of bits in 2036? SNTF mentions the problems of zero. Using zeros to indicate invalid at the FCencap level means any solution future SNTF's define will not work.

Rejected

The way to handle this problem is to change the Version field when the SNTF rollover issue is resolved.

Comment 68

4. Measuring Fibre Channel frame transit time

"...accordance with the applicable protocol specification;..." s/b
"...accordance with the protocol specification;..."

Rejected

I see nothing gained by the removal of the work "applicable".

Comment 69

5.2 Bit and Byte Ordering

As mentioned before, description of ordering for the Encapsulation Header needs to be before the Encapsulation Header section, not buried in the FC frame chapter.

Rejected

As mentioned in comment 77 there is an enforced IETF depiction of byte ordering that clarifies this issue by having all IETF documents follow the same rules.

Comment 70

5.3 FC SOF and EOF

Somewhere it should be noted that FC frame content is always 32-bit aligned. Otherwise, FC encap would need insert pad bytes to keep EOFs as shown in figure 5.

Accepted as follows:

Add a the following sentence after this one: "The number of 8-bit bytes in the FC frame content is always a multiple of four."

Comment 71

5.3 FC SOF and EOF
"...bytes," s/b "bytes;"

Accepted

Comment 72

5.3 FC SOF and EOF
How were the SOF and EOF codes chosen?

It seems like the SOF codes should be chosen to increase the Hamming distance between each other. 0x28 is only one bit different from 0x29, so four bit errors could change SOF into SOFi4 undetected. With only 8 values to encode in 32 bits, it seems like better Hamming distance could be provided.

Rejected

The SOF and EOF code values are defined in FC-BB. They are already implemented in products based on FC-BB and are not open to changes.

Comment 73

5.3 FC SOF and EOF
what is the rule about checking the redundant SOF, -SOF, EOF, and -EOF fields? Same as the FCencap rules or different?

Rejected

There is no rule.

[Comment 74](#)

7. Normative References

Is it ok to reference a draft standard in this "normative references" section?

A link to a page that points to how to buy a copy of the official standard would be appropriate for [5] and any other completed T11 standards.

[Accepted \(Partially\)](#)

Yes, it is okay to reference draft standards in normative references. This assurance was provided by a Transport Area Director during the Minneapolis IETF meeting.

A clone of the ANSI web URL from a T10 or T11 standard will be add.

[Comment 75](#)

9. Acknowledgements

"...no long..." s/b "...no longer..."

[Accepted](#)[Comment 76](#)

10. Full Copyright Statement

"2001" s/b "2002"

[Accepted](#)

4. Additional Changes Discovered After WG Last Call

=====

Comment 77 Technical

The draft fails to conform to the following requirement stated in http://www.ietf.org/ID-nits.html.

Historically, RFCs have specified byte and bit order according to a US DoD rule which made byte zero be the first byte in an address range, and bit zero be the most significant bit in a word or field. For example, you will find drawings like this one (from RFC 791) in many RFCs: when you make drawings like it, you should follow the same rule. Label your bit positions, bit zero is the most significant bit, and place the first addressable byte at the upper left-hand corner.

3.1. Internet Header Format

A summary of the contents of the Internet header follows:

0				1				2				3																																					
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Version				IHL				Type of Service				Total Length																																					
				Identification				Flags				Fragment Offset																																					
				Time to Live								Protocol								Header Checksum																													
				Source Address																																													
				Destination Address																																													
				Options												Padding																																	

Accepted with the following results

- 1) Add the following paragraph immediately following the Table Of Contents: "Warning to Readers Familiar With Fibre Channel: Both Fibre Channel and IETF standards use the same byte transmission order. However, the bit and byte numbering is different. See Appendix A for guidance."
- 2) Change figures 2, 3, and 5 to conform to the IETF bit and byte numbering;
- 3) Remove bit and byte numbers wherever they appear in the text; and
- 4) Insert Appendix A with the following text:

"Appendix A - Fibre Channel Bit and Byte Numbering Guidance

"Both Fibre Channel and IETF standards use the same byte transmission order. However, the bit and byte numbering is different.

"Fibre Channel bit and byte numbering can be observed if the data structure heading shown in figure 6, is cut and pasted at the top of figure 2 and figure 5.



Fig. 6 - Fibre Channel Data Structure Bit and Byte Numbering

"Fibre Channel bit numbering for the Flags field can be observed if the data structure heading shown in figure 7, is cut and pasted at the top of figure 3.

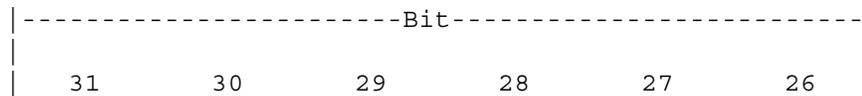


Fig. 7 - Fibre Channel Flags Bit Numbering"