

IPS Working Group
INTERNET-DRAFT
<draft-ietf-ips-fcip-wglc-02.txt>
(Expires November, 2002)

R. Weber

FCIP 1st 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-fcovertcpip-09.txt

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

Summary of Comments

Technical: 30 Comments, resulting in about 37 Changes
Editorial: 112 Comments, resulting in about 84 Changes

Totals: 142 Comments, resulting in about 121 Changes

Table Of Contents

1. Comments from David Black 2
 2. Comments from Mallikarjun Chadalapaka 47
 3. Comments from Don Fraser 70
 4. Comments from Murali Rajagopal 71
 5. Comments from Bob Snively 71
 6. Comments from Ralph Weber 72

1. Comments from David Black

=====

Comment 1

----- Title Page -----

E. Rodriguez, ips Liaison

[E] What sort of title is that? This looks like an invitation to IESG approval trouble.

Accepted with the following results

Change "Liaison" to "Co-chair".

Comment 2

-- Section 1 - Editors and Contributors

[E] A bit wordy, but basically ok. Please take out the Internet-Draft references.

Accepted with the following results

- 1) Change "draft-ietf-ips-fcip-slp-____.txt" to "the FCIP SLP internet-draft"; and
- 2) Change "draft-ietf-ips-fcip-mib-____.txt" to "the FCIP MIB internet-draft".

Comment 3

-- Section 1 - Editors and Contributors

Several T11 leaders supported this effort and advised the editors of this specification regarding appropriate interfaces to T11 documents.

[E] Is "leaders" the right T11 title? Let's make sure the right word is used.

Also I think "coordination with T11 documents and projects" might be better phrasing than "appropriate interfaces...".

Accepted (Partially)

I can think of no better word than "leaders". The only correct replacement I can think of is "chairs, vice-chairs, secretaries, facilitators, and document editors". I will make that change but only if mandated to do so.

Change "...regarding appropriate interfaces to T11 documents." to "...regarding coordination with T11 documents and projects."

Comment 4

-- Section 4 - Terminology
(really Section 2 - Purpose, Motivation and Objectives)

[E] Some of these [section 4] definitions implicitly exclude FC-AL, or at least the private (i.e., non-fabric) use of FC-AL across FC-IP. Section 2 would be a good place to make this exclusion explicit.

Accepted with the following results

Add the following new paragraph at the end of section 2: "The objectives of this document do not include using an IP Network as a replacement for the Fibre Channel Arbitrated Loop interconnect. No definition is provided for encapsulating loop primitive signals for transmission over an IP Network."

[Comment 5](#)

-- Section 3.2 - This Specification and Fibre Channel Standards

No attempt is being made to define a specific API between an FCIP Entity and an FC Entity at this time because doing so risks compromising the performance and efficacy of the resulting products. Current experience in this area is simply insufficient to guide definition of the interface appropriately.

[E] That's a bit negative. Please add some text indicating that the approach is to specify the functional interaction that is required, but allow implementers to choose how this will be realized.

[Accepted with the following results](#)

Replace the cited paragraph with: " No attempt is being made to define a specific API between an FCIP Entity and an FC Entity. The approach is to specify required functional interactions between an FCIP Entity and an FC Entity (both of which are required to forward FC frames across an IP Network), but allow implementers to choose how this will be realized."

[Comment 6](#)

-- Section 3.2 - This Specification and Fibre Channel Standards

The objectives and motivations of this specification are not impacted by the decision not to standardize a specific API between FCIP Entities and FC Entities because fully functional and compliant products can be built provided they contain both an FCIP Entity and an FC Entity. The only products that cannot be built are those that contain only one or the other.

[E] I don't like the product focus of this language. The basic point here is that in order to realize the full functionality of forwarding frames between FC and IP networks, both an FC Entity and an FCIP Entity are required. If someone wants to build something that only contains one of these, the fact that it won't have this functionality is their problem, not ours.

[Accepted with the following results](#)

This paragraph was intended to be the positive paragraph following on the negative paragraph cited in comment 5. Since the changes undertaken in response to comment 5 make that paragraph positive, the bulk of this paragraph is no longer needed.

A few of the critical thoughts from the comment have been added to the revised paragraph in the response to comment 5. With those changes in place, the cited paragraph will be deleted.

Comment 7

-- Section 4 - Terminology

Terms needed to clarify the concepts presented in FCIP are defined here.

[E] I don't like the usage of "clarify". How about "Terms used to describe FCIP concepts are defined in this section."?

Accepted, make change exactly as proposed.

Comment 8

-- Section 4 - Terminology

FC Entity - The Fibre Channel specific element that combines with an FCIP Entity to form an interface between an FC Fabric and an IP Network (see section 6.3).

[E] I think "element" is problematic in this definition, because it implies "fabric element" and leads to the sort of confusion that Mallikarjun got into. How about "functional component"?

Accepted, make change exactly as proposed.

Comment 9

-- Section 4 - Terminology

FC Receiver Portal - The access point through which an FC Frame and time stamp enters an FCIP Data Engine from the FC Entity.

FC Transmitter Portal - The access point through which a reconstituted FC Frame and time stamp leaves an FCIP Data Engine to the FC Entity.

[E] For clarity, shouldn't both of those terms be "FCIP" rather than "FC" to avoid any possible confusion with transmission and reception of FC frames on an FC fabric?

Accepted but only in principle

Change "FC" to "FC Frame" in these terms throughout the draft.

Inspection of the Portal names shows that "FC" is shorthand for "FC Frame", not for "FCIP".

Comment 10

-- Section 4 - Terminology

FCIP Entity - The principal FCIP interface point to the IP Network (see section 6.4).

[E] Definition needs to parallel FC Entity definition including "functional component" language (or whatever term/phrase is chosen).

Accepted, make change exactly as proposed.

Comment 11

-- Section 4 - Terminology

FCIP Frame - An FC Frame plus the FC Frame Encapsulation [27] header and encoded EOF that contains the FC Frame (see section 6.6.1).

[E] What about SOF?

Accepted

Add "encoded SOF" in the right place.

Comment 12

-- Section 4 - Terminology

FCIP Link Endpoint (FCIP_LEP) - The component of an FCIP Entity that contains one or more FCIP_DEs (see section 6.5).

[E] Needs to say something about its relationship to FCIP Links.

Accepted with the following results

Change "...that contains one..." to "...that handles a single FCIP Link and contains one..."

Comment 13

-- Section 4 - Terminology

Special Frame (SF) - A specially formatted FCIP frame containing information used by the FCIP protocol (see section 8).

[E] I suggest prefixing FCIP to this term for clarity.

Accepted with the following results

- 1) Change "Special Frame (SF)" to "FCIP Special Frame (FSF)" throughout the draft; and
- 2) Limit usage of spelled out FCIP Special Frame to once per section and sections titles.

Comment 14

-- Section 5 Protocol Summary

- 2) Viewed from the IP Network perspective, FCIP Entities are peers and communicate using TCP/IP. Each FCIP Entity is a TCP endpoint in the IP-based network.

[E] TCP endpoint is not the right language, as this implies endpoint of a single TCP connection. Probably needs to be described in terms of TCP ports.

Accepted with the following results

Change "...is a TCP endpoint..." to "...contains one or more TCP endpoints..."

Comment 15

-- Section 5 Protocol Summary

- 3) Viewed from the FC Fabric perspective, pairs of FCIP Entities, in combination with their associated FC Entities, serve as an FC Frame transmission component of the FC Fabric. The FC End Nodes are unaware of the existence of the FCIP Link.

[E] "FC frame transmission component" is a bit abstract. Can we say that it forwards frames between two FC ports, e.g., E_Ports? Note the use of the e.g., to avoid limiting this to E_Ports.

Accepted but only in principle

Change "...serve as an FC Frame transmission component of the FC Fabric." to "...forward FC Frames between FC Fabric elements."

Since comment 8 states that "element" has a well known meaning, use of that term here should not raise objections.

Meticulous effort has gone into avoiding use of E_Port and B_Port in this draft and a very good reason will have to be offered to get either term in.

Comment 16

-- Section 5 Protocol Summary

- 7) When multiple FCIP_LEPs with multiple FCIP_DEs are in use, selection of which FCIP_DE to use for encapsulating and transmitting a given FC Frame is outside the scope of this document. FCIP Entities do not actively participate in FC Frame routing.

[E] Does anything need to be said about requirements for or desirability of preserving the order in which frames are forwarded.

Response

No. The fact that TCP keeps all the sent on one TCP Connection in order is covered elsewhere here. All other frame ordering concerns are covered in FC-BB-2.

Comment 17

-- Section 5 Protocol Summary

- 8) The FCIP Control & Services function MAY use TCP/IP quality of service features (see section 11.2) to support Fibre Channel capabilities.

[E] This isn't quite right, as it seems to imply that FC has some features that map onto TCP/IP QoS features. Needs some editorial tweaking.

Accepted with the following results

Delete everything after the section reference.

Comment 18

-- Section 5 Protocol Summary

- 9) Each FCIP Entity is statically or dynamically configured with a list of IP addresses and TCP port numbers corresponding to participating FCIP Entities. If dynamic discovery of participating FCIP Entities is supported, the function SHALL be performed using the Service Location Protocol (SLPv2) [25]. It is outside the scope of this specification to describe any static configuration method for participating FCIP Entity discovery. Refer to section 9.1.2.2 for a detailed description of dynamic discovery of participating FCIP Entities using SLPv2.

[E] There's a requirement lurking in here. One way to express it would be to change the first "is" to "MUST be".

Accepted with the following results

Replace the first sentence in the cited text with: "It is necessary to statically or dynamically configure each FCIP entity with the IP addresses and TCP port numbers corresponding to FCIP Entities with which it is expected to initiate communication."

Comment 19

-- Section 5 Protocol Summary

- 10) Before creating a TCP Connection to a peer FCIP Entity, the FCIP Entity attempting to create the TCP connection SHALL statically or dynamically determine the IP address, TCP port, expected FC Fabric Entity World Wide Name, TCP Connection Parameters, and Quality of Service Information.

[E] Need to get the "configuration" word in here, as this is describing a requirement on the configuration mechanisms in 9), and consider rephrasing this to make the connection clearer.

Rejected

Item 9) in the list describes configuration. This item describes what is required in order to make a TCP connection and send the Special Frame.

Comment 20

-- Section 5 Protocol Summary

- 13) On a given TCP Connection, this specification relies on TCP/IP to deliver a byte stream in the same order that it was sent.

[E] "a given" --> "an individual" for clarity.

Accepted

[Comment 21](#)

-- Section 5 Protocol Summary

- 14) This specification defines only limited error detection and recovery mechanisms and relies on both TCP and FC to handle data loss and corruption within the IP Network.

[E] I'd rephrase this to talk about mechanisms that are complementary to the existing TCP/IP and FC mechanisms, and that this specification assumes the presence and requires the use of those existing mechanisms.

[Accepted with the following results](#)

Replace the cited paragraph with: "14) This specification assumes the presence of and requires the use of TCP and FC data loss and corruption mechanisms. The error detection and recovery features described in this specification complement and support these existing mechanisms."

[Comment 22](#)

-- Section 6.1 - FCIP Protocol Model

[E] Need to define every acronym in Figure 1 and make it clear that the FCIP Link is implemented via the IP Network Link. Consider using "Fibre Channel Fabric" instead of "Fibre Channel Environment" for consistency.

[Accepted \(Partially\)](#)

No changes will be made to "make it clear that the FCIP Link is implemented via the IP Network Link". This is a relatively standard network layering diagram. The notation is consistent already in place is 100% consistent with that concept.

[Actions to be taken:](#)

- 1) Add a key at the bottom of the figure; and
- 2) Change "Environment" to "Fabric".

Comment 23

-- Section 6.3 - FC Entity

A product that tunnels an FC Fabric through an IP Network MUST combine an FC Entity with an FCIP Entity (see section 6.4) to form a complete interface between the FC Fabric and IP Network as shown in figure 3.

[E] Again, I don't like the use of "product". How about "implementation"?

Accepted, make change exactly as proposed.

Comment 24

-- Section 6.3 - FC Entity

In general, the combination of an FCIP Link and FC and FCIP Entities is intended to replace a Fibre Channel defined connection between Fibre Channel components. For example, this combination can be used to replace a hard-wire connection between two Fibre Channel switches. There are limitations on the generally intended usage of the combination shown in figure 3. As another example, the combination cannot be used to replace cable connections in a Fibre Channel Arbitrated Loop because loop primitive signals cannot be encapsulated for transmission over TCP.

[E] I think "replace" is the wrong verb. For example, if the distance is well over 10km, it's not possible to have an FC connection to replace. I would rewrite this in terms of "function as" rather than "replace". I don't understand the "There are limitations ..." sentence. As noted earlier the absence of support for FC-AL should be stated in Section 2.

Accepted with the following results

- 1) Replace the first cited sentence with: "In general, the combination of an FCIP Link and FC/FCIP Entities is intended to provide a non- Fibre Channel backbone transport between Fibre Channel components.";
- 2) In the second cited sentence change "replace" to "function as"; and
- 3) Delete all text from "There are limitations..." to the end of the paragraph. (Note this works because comment 4 puts the FC-AL limitation in section 2.)

Comment 25

-- Section 6.3 - FC Entity

The interface between the FC and FCIP Entities is implementation specific. The minimum requirements placed on an FC Entity by this specification are listed in annex G.

[E] Suggest changing "minimal" to "functional".

Accepted

Comment 26

-- Section 6.4 - FCIP Entity

[E] In Figure 4, suggest changing "TCP connect request IP address and port" to "TCP port for incoming connections". The implicit need for an IP address should be obvious to the reader, or can be explained in the text.

Accepted

Comment 27

-- Section 6.4 - FCIP Entity

The FCIP Entity is the connection interface point for the IP Network and is the owner of the IP Address and Well Known Port used to form TCP Connections.

[E] That language isn't quite right. It owns the TCP port at that IP address, but does not need to exclusively own the IP address.

Accepted with the following results

Change "...is the owner of the IP Address and Well Known Port used to form..." to "...is the owner of the Well Known Port at an IP Address used to form...".

[Comment 28 Technical](#)

-- Section 6.4 - FCIP Entity

The interfaces to the IP Network features is implementation specific, however, to maintain interoperability, the notable TCP/IP mechanisms used are specified in this document as follows:

[E] I'd rephrase this to talk about "REQUIRED functional support" and remove the "to maintain interoperability" language.

[Accepted with the following results](#)

- 1) Change "is" to "are";
- 2) Replace everything from "however" to the end of the cited text with: "however, REQUIRED TCP/IP functional support is specified in this document, including:"

[Comment 29](#)

-- Section 6.5 - FCIP Link Endpoint (FCIP_LEP)

An FCIP_LEP MAY communicate with its FC Entity counterpart to coordinate flow control.

[E] Insert "local" before "FC Entity" to make it clear that this communication does not occur over the IP network.

[Accepted](#)

[Comment 30](#)

-- Section 6.6 - FCIP Data Engine (FCIP_DE)

Data flows through the FCIP_DE in the following seven steps:

[E] The text that follows this actually describes data flow through a pair of FCIP_DEs connected across an IP network - this sentence needs to be revised accordingly.

[Accepted with the following results](#)

Replace "...the FCIP_DE..." with "...a pair of IP Network connected FCIP_DEs..."

Comment 31

-- Section 6.6 - FCIP Data Engine (FCIP_DE)

Table 2 shows the SOF and EOF code values that are legal in FCIP Frames. This list may be a subset of the SOF and EOF codes listed in the FC Frame Encapsulation [27].

[T] This is a problem because these codes are being specified in more than one place. I think the FC Frame Encapsulation document is the right place for the normative specification of these codes (and see my comments against it on the need to get IANA involved). This would be ok as a list of codes that are currently valid, but the specification authority needs to be in one place.

Accepted (Partially) with the following results

Replace the cited text with: "In FCIP, the Class 2, Class 3, and Class 4 SOF and EOF codes listed in the FC Frame Encapsulation [27] are legal.

Note: This change depends on adding the Class column in the FC Frame Encapsulation draft.

Comment 32

-- Section 6.6.2.1 - TCP Assistance With Error Detection and Recovery

TCP [8] requires in order delivery, generation of TCP checksums, and checking of TCP checksums. Thus, the byte stream passed from TCP to the FCIP_LEP will be in order and free of errors detectable by the TCP checksum. If TCP did not perform these functions, the FCIP_LEP would have to.

[E] Strengthen the last sentence to indicate that the FCIP_LEP relies upon TCP performing these functions.

Accepted with the following results

Replace the cited last sentence with: "The FCIP_LEP relies on TCP performing these functions."

[Comment 33](#)

-- Section 6.6.2.2 - Errors in FCIP Headers and Discarding FCIP Frames

- 1) Length field validation -- $15 < \text{Length} < 545$;

[E] Explain where the 15 and 545 values come from.

[Accepted with the following results](#)

Add a note following the list that derives the values.

[Comment 34 Technical](#)

-- Section 6.6.2.2 - Errors in FCIP Headers and Discarding FCIP Frames

In addition to the tests above, the validity and positioning of the following FCIP Frame information SHOULD be used to detect encapsulation errors that may or may not affect synchronization:

- a) Protocol # field and its ones complement (2 tests);
- b) Version field and its ones complement (2 tests);

[... list continues, snipped ...]

[T] I think at least these two are MUSTs. At a minimum, the Protocol# and Version field must be checked against expected values - I might accept the checks against their ones complements being SHOULDs. Same comment applies to the Flags field and SOF. Someone MUST check the FC frame CRC before forwarding the frame, but that responsibility could be assigned to the FC Entity.

[Accepted \(Partially\) with the following results](#)

- 1) Add to 6.6.2.2 checking Protocol# and Version#. This addition will have to be in a separate paragraph before the current 1), 2), 3) list because it is not a synchronization issue;
- 2) Keep the one's complement tests in the SHOULD a), b), c) list, but reduce the number of tests in that list by 2 (1 in a and 1 in b); and
- 3) Change the count of optional tests required from "5 of 18" to "3 of 18".
- 4) Add the following after the tests are described: "Note: The process of selecting an 8b/10b encoding for the SOF by necessity includes some validation of the SOF fields."

Responses to other issues raised by comment

- a) The Flags field is not a MUST test because it provides no certain identification of the protocol beyond that provided by the Protocol# field; and
- b) Not even the Fibre Channel standards require that the FC CRC be validated by FC Fabric elements. FC Endnode validation of the FC CRC is sufficient.

Comment 35 Technical

-- Section 6.6.2.2 - Errors in FCIP Headers and Discarding FCIP Frames

At least 5 of the 18 tests listed above SHALL be performed. Failure of any of the above tests actually performed SHALL indicate an encapsulation error and the frame SHALL NOT be forwarded on to the FC Entity. Further, such errors SHOULD be considered carefully, since some may be synchronization errors.

[T] There aren't 18 tests, and I think some of the individual tests (or subsets) are MUSTs (see previous comment). This needs to be gone over carefully - in essence a MUST is only appropriate where failure to apply the test carries a non-negligible risk of forwarding a bad frame to FC.

The authors are the experts on this and need to do this analysis. I don't understand the last "SHOULD" - what is the (testable) requirement on an implementation?

No changes made

There are 18 tests: 2 + 2 + 1 + 2 + 2 + 1 + 4 + 1 + 2 + 1 = 18

What tests are MUSTs is covered by comment 34.

The authors have gone over the tests carefully and have concluded that no individual test or specific group of tests associates specifically with a non-negligible risk of forwarding a bad frame to FC. The requirement is that a suitable number (at least 5, or 12 when the 7 required tests are included) tests is necessary to reduce the risk of forwarding a bad frame to FC to a negligible level. The specific tests selected depends on the implementation, i.e., what test can be performed most efficiently in the implementation hardware.

The "SHOULD" statement is intended to guide implementations. Repeated failures of, for example, the CRC equal to zero test may

mean that synchronization has been lost. There are no hard and fast rules here.

Comment 36 Technical

-- Section 6.6.2.3 - Synchronization Failures

If the FCIP_DE attempts to recover synchronization, the resynchronization algorithm used SHALL meet the following requirements:

[T] One requirement is missing, which may be part of b):

- b) return to sending valid FC Frames only after synchronization has been verified; and

[T] Verification of synchronization MUST exclude any possibility of forwarding an FC frame that was not sent by the transmitting FCIP Entity. This includes the scenario in which a valid encapsulated FCIP frame occurs in the payload of an FC frame that is encapsulated and sent over FCIP; excluding this scenario is necessary but not sufficient to meet the requirement.

Accepted with the following results

Replace b) with: "return to forwarding FC Frames only after synchronization on the transmitted FCIP Frame stream has been verified; and"

Comment 37 Technical

-- Section 6.6.2.3 - Synchronization Failures

An example algorithm meeting these requirements can be found in annex C.

[T] That doesn't meet the missing requirement that my above comment wants to add. The problem is at step 2 of the algorithm description.

- 2) Use multiple strong candidate headers to locate a verified candidate header:

The Frame Length in one strong candidate header is used to skip incoming bytes until the expected location of the next strong candidate header is reached. Then the tests described in step 1) are applied to see if another strong candidate header has successfully been located.

The "skip incoming bytes" step makes it possible that a set of valid FC headers is interlaced in the payloads of FC frames in a fashion that causes all the real headers to be skipped. This is a rather unlikely, but not impossible scenario. This could be dealt with by searching for headers instead of skipping data and aborting if a header is found where data should have been or carrying on and aborting if an interlaced header chain scenario arises. The algorithm in Annex C does address the scenario of FCIP frames occurring in FC frame payloads, but it doesn't meet the "can't be fooled" requirement that I think should be added.

Unfortunately, this issue appears to not be resolvable within the WG. I have had heated and lengthy offline discussion with the FCIP Authors in which they have made clear their strong opposition to the "missing requirement" and the need to scan rather than skip data, and have argued that the algorithm in Annex C produces a long enough chain of headers that the odds of having followed a chain that is interlaced through frame payloads is small enough to be ignored. I will have to consult the Area Directors.

Accepted with the following comment

Modifications to either step 2 or step 3 will achieve the requested results. Because step 3 already includes complex steps such as verifying the FC CRC value, changes in response to the comment will be made in step 3.

Actions to be taken

1) Change the first paragraph of Annex C step 3 from:

"Incoming bytes are skipped and discarded until the next verified candidate header is reached. Each verified candidate header is tested against the full collection of tests listed in section 6.6.2.2 as would normally be the case."

to:

"Incoming bytes are inspected and discarded until the next verified candidate header is reached. Inspection of the incoming bytes includes testing for other candidate headers using the criteria described in step 1. Each verified candidate header is tested against the tests listed in section 6.6.2.2 as would normally be the case."

- 2) Change the second sentence in the third paragraph of Annex C step 3 from:

"If any verified candidate headers are invalid and fail to meet the tests for a strong candidate header, increment the retry counter and return to step 1."

to:

"If any verified candidate headers are invalid and fail to meet the tests for a strong candidate header or inspection of the bytes between verified candidate headers discovers any candidate headers, increment the retry counter and return to step 1."

Comment 38 Technical

Section 7 - Checking FC Frame Transit Times in the IP Network

The FC Entity MUST implement the measurement of Fibre Channel frame IP Network transit time as described in the FC Frame Encapsulation [27] specification.

[E] "MUST implement" --> "MUST implement and use" for clarity.

Accepted but not as the comment intended

The statement is misleading and needs to be revised.

Action to be taken

Replace the cited sentence with: "FC-BB-2 [4] defines how the measurement of IP Network transit time is performed, based on the requirements stated in the FC Frame Encapsulation [27] specification."

Additional note

See comment 40 for a discussion of why IP Network transit time checking is done by the FC Entity.

Comment 39

Section 7 - Checking FC Frame Transit Times in the IP Network

If no synchronized time stamp value is available to accompany the entering FC Frame a value of zero SHALL be supplied.

[E] "supplied" --> "used" for clarity.

Accepted

Comment 40 Technical

Section 7 - Checking FC Frame Transit Times in the IP Network

The FC Entity SHALL use suitable internal clocks and either Fibre Channel services or an SNTP Version 4 server [13] to establish and maintain the required synchronized time value. The FC Entity SHALL verify that the FC Entity it is communicating with on an FCIP Link is using the same synchronized time source as it is, either Fibre Channel services or SNTP server.

[T] I don't believe that this paragraph meets the requirements in the FC Frame Encapsulation specification for processing transit time information. Punting this up to the FC Entity is problematic - the minimum functional requirements on the FC Entity to meet the FC Frame Encapsulation requirements will need to be spelled out here in detail (i.e., a single sentence saying "must meet requirements in Section 4 of the FC Frame Encapsulation document" is probably not going to fly). Mallikarjun caught some issues in this area also.

Rejected

The decision to move time validation from FCIP to FC-BB-2 was made for sound technical reasons:

- 1) Having the FC Entity verify transit time makes the test more end-to-end;
- 2) Class F frames need not have their transit time verified. That decision is allowed by the FC Frame Encapsulation. Encoding that test in FCIP would necessitate a substantial increase in the FCIP reliance on FC specific knowledge, including but not limited to cracking FC Frames;
- 3) Allowing Class F frames to transit without transit time verification is required to allow the FC Time Service to be used as a source of synchronized time, a critical feature for the success of FCIP; and

- 4) The description of the interactions between the FC Entity and FCIP Entity for the purpose of maintaining a synchronized time based on the FC Time Service were getting impossible to verify for correctness when the requirements were stated in FCIP.

Having the FCIP Entity perform transit time tests was in the FCIP draft as recently as draft-ietf-ips-fcovertcpip-05.txt. The requested organization was tried and proved to be unworkable.

Comment 41

-- Section 8 - The FCIP Special Frame

[T] How is the FCIP Special Frame payload protected? I don't see the equivalent of an FC Frame CRC.

Inquiry

The Special Frame is protected by requiring that the connection be closed if the echoed SF differs from the transmitted SF. This is deemed to be a more rigorous test than any CRC.

Comment 42

-- Section 8 - The FCIP Special Frame

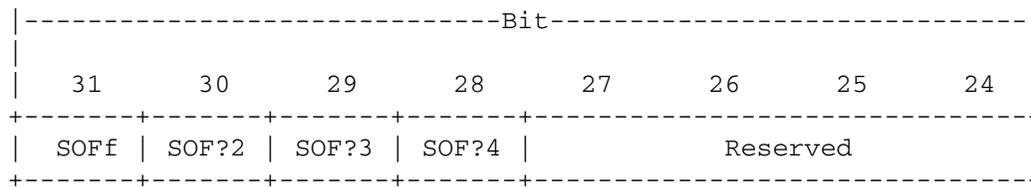


Fig. 10 Connection Usage Flags Field Format

[E] I don't think the "?"s were intended and suspect that MS Word or some other tool has been a little too helpful :-).

Inquiry

The question marks were intended to have the classic meaning of question mark in a file specification. SOF?2 == SOFi2 or SOFn2.

Comment 43 Technical

-- Section 8 - The FCIP Special Frame

The Connection Usage Code field contains Fibre Channel defined information regarding the intended usage of the connection as specified in FC-BB-2 [4].

[T] The authors need to talk to me about this, so that I can understand what's going on here, and whether we need another IANA registry, as is the case for the SOF and EOF values.

No changes made

The Connection Usage Code field allows pairs of FC Entities to communicate 16-bits of connection setup information in the Special Frame. It represents a request that the FC-BB-2 side of the house made on the FCIP side. Given that FC-BB-2 is defining a whole new SW_ILS to support a request made in the reverse direction, it is difficult to see why the presence of the Connection Usage Code field is an issue.

Comment 44 Technical

-- Section 8 - The FCIP Special Frame

[T] This section needs to describe the usage of the FCIP Special Frame, including the structure of the interaction between the two FCIP Entities, and how that establishes the security and connection usage properties that are desired. The descriptions in Section 9 contain a great deal of detail that mixes several mechanisms together - a high level guide to what's going on is necessary to understand them.

Accepted with the following results

It is considered desirable to leave the Special Frame open to additional uses in future versions of FCIP. This is why Section 8 lacks the requested overview.

Actions to be taken

- 1) Put the current Section 8 text in 8.1 "Special Frame Format";
- 2) Add 8.2 "Overview of Special Frame Usage in Connection Establishment"; and
- 3) Be sure to discuss the issues raised in comment 52, comment 53, and comment 109 in the new section.

The text for the new section is proposed to be the following:

8.2 Overview of FSF Usage in Connection Establishment

When a new TCP Connection is established, an FCIP Special Frame makes one round trip from the FCIP Entity initiating the TCP connect operation to the FCIP Entity receiving the TCP connect request and back. This FSF usage serves three functions:

- Identification of the FCIP Link endpoints
- Conveyance of a few critical parameters shared by the FC/FCIP Entity pairs involved in the FCIP Link
- Configuration discovery (used in place of SLP only when allowed by site security policies)

The specific requirements for this usage of the FSF are found in section 9.1. This section provides an overview of the FSF usage without stating requirements.

Because FCIP is only a tunnel for a Fibre Channel Fabric and because the Fabric has its own complex link setup algorithm that can be employed for many FCIP link setup needs, it is desirable to minimize the complexity of the FSF usage during TCP Connection setup. With this in mind, this FSF usage is not a login or parameter negotiation mechanism. A single FSF transits each newly established TCP connection as the first bytes sent in each direction.

Note: This usage of the FSF cannot be eliminated entirely because a newly created TCP Connection must be associated with the correct FCIP Link before FC Fabric initialization of the connection can commence.

The first bytes sent from the TCP connect request initiator to the receiver are an FSF identifying both the sender and who the sender thinks is the receiver. If the contents of this FSF are correct and acceptable to the receiver, the unchanged FSF is echoed back to the sender. This send/echo process is the only set of actions that allows the TCP Connection to be used to carry FC Fabric traffic. If the send and unchanged echo process does not occur, the algorithm followed at one or both ends of the TCP Connection results in the closure of the TCP Connection (see section 9.1 for specific algorithm requirements).

Note: Owing to the limited manner in which the FSF is used and the requirement that the FSF be echoed without changes before a TCP Connection is allowed to carry user data, no error checking beyond that provided by TCP is deemed necessary.

As described above, the primary purpose of the FSF usage during TCP

Connection setup is identifying the FCIP Link to which the new TCP Connection belongs. From these beginnings, it is only a small stretch to envision using the FSF as a simplified configuration discovery tool, and the mechanics of such a usage are described in section 9.1.

However, use of the FSF for configuration discovery lacks the broad range of capabilities provided by SLP v2 and most particularly lacks the security capabilities of SLP v2. For these reasons, using the FSF for configuration discovery is not appropriate for all environments. Thus the choice to use the FSF for discovery purposes is a policy choice to be included in the TCP Connection Establishment "shared" database described in section 9.1.1.

When FSF-based configuration discovery is enabled, the normal TCP Connection setup rules outlined above are modified as follows.

Normally, the algorithm executed by an FCIP Entity receiving an FSF includes verifying that its own identification information in the arriving FSF is correct and closing the TCP Connection if it is not. This can be viewed as requiring the initiator of a TCP connect request to know in advance the identity of the FCIP Entity that is the target of that request (using SLP, for example), and through the FSF effectively saying, "I think I'm talking to X." If the party at the other end of the TCP connect request is really Y, then it simply hangs up.

FSF-based discovery allows the "I think I'm talking to X" to be replaced with "Please tell me who I am talking to?", which is accomplished by replacing an explicit value in the Destination FC Fabric Entity World Wide Name field with zero.

If the policy at the receiving FCIP Entity allows FSF-based discovery, the zero is replaced with the correct Destination FC Fabric Entity World Wide Name value in the echoed FSF. This is still subject to the rules of sending with unchanged echo, and so closure of TCP Connection occurs after the echoed FSF is received by the TCP connect initiator.

Despite the TCP Connection closure, however, the TCP connect initiator now knows the correct Destination FC Fabric Entity World Wide Name identity of the FCIP Entity at a given IP Address and a subsequent TCP Connection setup sequence probably will be successful.

The Ch bit in the pFlags field (see section 6.6.1) allows for differentiation between changes in the FSF resulting from transmission errors and changes resulting from intentional acts by the FSF recipient.

[Comment 45](#)

-- Section 9.1.1 - Connection Establishment Model

What is important is the fact that the FCIP Entity MAY consult the database at anytime to determine its actions relative to TCP Connection establishment.

[E] "anytime" --> "any time"

[Accepted](#)

[Comment 46 Technical](#)

-- Section 9.1.2.1 - Non-Dynamic Creation of New TCP Connections

If the TCP connect request is rejected, the FCIP Entity SHALL act to limit unnecessary repetition of attempts to establish similar connections.

[T] That's a little vague. How about specifying a minimum time period that MUST elapse before retry?

[Accepted with the following results](#)

Add new sentence following cited text: "For example, the FCIP Entity might wait 60 seconds before trying to re-establish the connection."

Comment 47

-- Section 9.1.2.1 - Non-Dynamic Creation of New TCP Connections

It is recommended that an FCIP Entity not initiate TCP connect requests to another FCIP Entity if incoming TCP connect requests from that FCIP Entity have already been accepted.

[T] Needs an upper case "SHOULD" or "RECOMMENDED". This also needs a MUST or SHOULD on the configuration mechanism to indicate in which direction connections are to be established between a pair of FCIP Entities in order to deal with a problem that turns up near the end of Section 9.1.3. This is related to Mallikarjun's issue about handling of simultaneous opens.

Accepted (Partially)

Change "recommended" to "RECOMMENDED". See comment 124 for a discussion of the other issues cited here.

Comment 48

-- Section 9.1.2.2 - Dynamic Creation of New TCP Connections

[E] The information on connection establishment is basically common to this and Section 9.1.2.1. Consider reducing this section to focus on use of SLP, and referring to Section 9.1.2.1 for what to do after the configuration info is discovered via SLP. Also consider whether the SLP description should come before the connection establishment description.

Rejected

The information on connection establishment is indeed common to section 9.1.2.1, but its relationship to SLP is not. Section 9.1.2.2 already is less than one page long, that and a desire to have some amount of readable flow (as opposed to "do this, then do what it says to do over there, then do that, then do this other thing that is talked about in another section") seems like ample motivation for the current organization.

[Comment 49](#)

-- Section 9.1.2.3 - Connection Setup After a Successful TCP Connect Request

[T] This dives into the details of FCIP Special Frame handling without explaining the overall structure and goals of the use, and is unclear as a result. For example, For example, on p. 29, after constructing the FCIP Special Frame, the text says:

After the FCIP Special Frame bytes are sent on the newly formed connection, the FCIP Entity SHALL wait for the FCIP Special Frame to be echoed as the first bytes received on the newly formed connection.

But one has to wade all the way down to p.33 towards the end of Section 9.1.3 to discover that the other FCIP Entity is expected to perform validation actions before echoing the frame. The structural outline of the usage of the FCIP Special Frame (without the blow by blow details) needs to be described in Section 8.

[Duplicate comment, see response to comment 44.](#)

[Comment 50](#)

-- Section 9.1.2.3 - Connection Setup After a Successful TCP Connect Request

The remaining steps in this section SHALL be performed only if the echoed FCIP Special Frame bytes exactly match the FCIP Special Frame bytes sent (words 7 through 17 inclusive).

[E] There is a great deal of common text in the two descriptions that follow the above text. They should be combined into a single description, with the creation of the FCIP_LEP at step 2 being done only if necessary.

[Accepted \(Conditionally\)](#)

An attempt will be made to do this. If the cure is worse than the disease, the original text will be used.

[Comment 51](#)

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

If the requested connection is allowed, the FC Entity SHALL ensure that required IP security features are enabled and accept the TCP connect request.

[T] As written, this appears to require a dynamic interrogation of the IPsec Security Association Database, and possibly the IPsec Security Policy Database. I suspect that this is in excess of what was intended, and suspect a longer description is needed.

[Accepted with the following result](#)

Change "If the connection is allowed,..." to "If the connection is allowed by the "shared" database (see 9.1.1),..."

[Comment 52 Technical](#)

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

If the Destination FC Fabric Entity World Wide Name contains 0, the FCIP Entity SHALL take one of the following three actions:

- 1) Leave the Destination FC Fabric Entity World Wide Name field and Ch bit both 0;
- 2) Change the Destination FC Fabric Entity World Wide Name field to match FC Fabric Entity World Wide Name associated with the FCIP Entity that received the TCP connect request and change the Ch bit to 1; or
- 3) Close the TCP Connection without sending any response.

The choice between the above actions depends on the anticipated usage of the FCIP Entity and is outside the scope of this specification. The FCIP Entity may consult the "shared" database when choosing between the above actions.

[T] I'm not buying the "outside the scope" disclaimer. Some more description of why the three choices are available is in order even if the normative criteria for choosing among them are specified in FC-BB-2. If my assumption about FC-BB-2 is correct, the last sentence is almost certainly too weak - it needs to refer to consulting the FC Entity to determine what to do.

[Accepted but not as the comment intended](#)

Delete "... and is outside the scope of this specification"

Other actions to be taken

Note: Some non SLP implementations wish to use the SF as a configuration determination mechanism. The choice exists to allow that option.

Describe this in the new section added in response to comment 44.

Comment 53 Technical

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

If:

- a) The Destination FC Fabric Entity World Wide Name contains a non-zero value that does not match the FC Fabric Entity World Wide Name associated with the FCIP Entity that received the TCP connect request, or
- b) The contents of the Connection Usage Flags, and Connection Usage Code fields is not acceptable to the FCIP Entity that received the TCP connect request,

then the FCIP Entity SHALL take one of the following two actions:

- 1) Change the contents of the unacceptable fields to correct/acceptable values and set the Ch bit to 1; or
- 2) Close the TCP Connection without sending any response.

[T] 1) looks like a security hole that discloses information an attacker may find useful in establishing an undesired connection to FCIP.

What's the motivation/purpose for this?

No changes made

The motivation is to allow the SF to be used as a poor-man's SLP.

Option 1) is a security policy that is not a security hole because either IPsec or option 2) or both are available as security policy choices.

[Comment 54 Technical](#)

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

If the Source FC Fabric Entity World Wide Name and Source FC/FCIP Entity Identifier field values in the FCIP Special Frame match the Source FC Fabric Entity World Wide Name and Source FC/FCIP Entity Identifier associated with an existing FCIP_LEP, the FCIP Entity SHALL:

- 1) Request that the FC Entity authenticate the source of TCP connect request, providing the following information to the FC Entity for authentication purposes:

[T] Need to say more about what the FC Entity MUST do to "authenticate the source". I realize that the details are specified in FC-BB-2, but the functional requirements on the FC Entity can be specified here.

[Accept but only to a limited degree](#)

Requiring the FC Entity to do a specific thing in response to the request to authenticate goes substantially beyond the security policies that the IETF applies to itself (i.e., this would be MANDATORY to implement and MANDATORY to use).

[Action to be taken](#)

Replace the "warning" paragraph cited in comment 56 with: "The definition of the FC Entity SHALL include a mechanism for use in response to a TCP connect request source that communicates with the partner FC Entity on the FCIP Link using a previously authenticated TCP Connection to verify that the Connection Nonce has been sent in a yet to be completed TCP Connection setup. Failure of the partner FC Entity to verify the Connection Nonce SHALL be treated as an authentication failure."

[Comment 55 Technical](#)

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

The FCIP Entity SHALL wait indefinitely for the FC Entity to authenticate source of the TCP connect request and SHALL not use the new TCP Connection for any purpose until the FC Entity completes the authentication.

[T] "wait indefinitely" creates an obvious denial of service attack vulnerability. Try again.

[Accepted with the following results](#)

Change the cited text to: "The FCIP Entity SHALL not use the new TCP Connection for any purpose until the FC Entity authenticates the source of the TCP connect request."

[Comment 56](#)

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

Warning: The authentication mechanism described here and in FC-BB-2 [4] is not designed to thwart sophisticated security threats. The IP security mechanisms described in section 10 should be enabled in environments where security threats are suspected.

[E] Unfortunately, that's almost content-free. I suggest replacing this with a forward pointer to a more specific discussion of the threats involved in the Security Considerations section.

[Accepted with the following results](#)

- 1) Embed said pointer in the first paragraph of the step 1) text describing the process; and
- 2) Remove this paragraph entirely.

[Comment 57 Technical](#)

-- Section 9.1.3 - Processing Incoming TCP Connect Requests

Note that the originator of TCP connect requests uses IP Address and TCP Port to identify which TCP Connections belong to which FCIP_LEPs while the recipient of TCP connect requests uses the Source FC Fabric Entity World Wide Name, Source FC/FCIP Entity Identifier fields from the FCIP Special Frame to identify which TCP Connection belong to which FCIP_LEPs. For this reason, an FCIP Entity that both originates and receives TCP connect requests is unable to match the FCIP_LEPs associated with originated TCP connect requests to the FCIP_LEPs associated with received TCP connect requests.

[T] That's a problem. See comment against Section 9.1.2.1 for one suggestion for how to fix this, but some sort of fix appears necessary to me.

[Accepted with the following results](#)

Add a new section 9.1.4 titled "Simultaneous Connection Establishment" containing the following text.

"If two FCIP Entities perform simultaneous open operations, then two TCP Connections are formed and the SF originates at one end on one connection and at the other end on the other. Connection setup proceeds as described above on both connections, and the steps described above properly result in the formation of two FCIP Links between the same FCIP Entities.

"This is not an error. Fibre Channel is perfectly capable of handling to approximately equal connections between FC Fabric elements.

"The decision to setup pairs of FCIP Links in this manner is considered to be a site policy decision that can be covered in the "shared" database described in section 9.1.1."

[Comment 58](#)

-- Section 9.3 - TCP Connection Parameters

In order to provide efficient management of FCIP_LEP resources as well as FCIP Link resources, consideration of certain TCP Connection parameters is RECOMMENDED.

[E] As written, "RECOMMENDED" should be lower case, as it does not convey a requirement related to interoperability or good use of the network.

[Accepted](#)

[Comment 59](#)

-- Sections 9.3.2, 9.3.3, and 9.3.4

[E] Sections 9.3.2, 9.3.3, and 9.3.4 need references to the specifications of the TCP options being discussed.

[Accepted](#)

Good! We are being asked later to prune the normative references. This will allow them to be bulked up again.

[Comment 60 Technical](#)

-- Section 9.3.4 TCP_NODELAY Option

FCIP Entities SHALL set the TCP_NODELAY option to one. This will disable the Nagle Algorithm that is designed for usage in a telnet environment.

[T] I believe that "SHALL" should be a lower case "should". This is a local performance optimization that has no other effects.

[Accepted](#)

Comment 61 Technical

-- Section 9.5 - Flow Control Mapping between TCP and FC

Coordination of these flow control mechanisms one of which is credit based and the other of which is window based depends on painstaking design that is outside the scope of this specification.

[T] This is content-free. At least warn about some of the things that can go wrong, in particular BB-credit starvation and the potential to really screw up a Fibre Channel fabric if this is long-lived.

Rejected

This text was written in specific response to the decision of Orange County interim meeting that "The only statement that should appear in FCIP on the subject is, 'There be dragons here.'"

Comment 62 Technical

-- Section 10 Security

[T] Need to get this whole section aligned with the latest (currently -11) version of the IPS Security draft.

Accepted with the following results

Section 10 will be aligned with IPS Security Draft version 12. This will require a substantial number of changes, as detailed below. It would be desirable to avoid a "hairball" between FCIP and the IPS Security Draft. With this in mind, it is believed that changes in the IPS Security Draft will concern areas that do not directly impact FCIP. Of course, there are no guarantees.

In Section 10.1, add the following to the Threat Models: "8) An adversary may launch a variety of attacks against the discovery process [SLPv2, RFC2608]."

Note: This addition is placed in the list in such a way as to keep the FCIP-specific attack (the attack related to the Special Frame) last in the list.

Section 10.3.1, add the following to the end of the first paragraph: "When ESP is utilized, per-packet data origin authentication, integrity and replay protection must be used."

In addition to the changes in Section 10.3.2 (Key Management) described in comment 69, comment 70, comment 71, comment 72, and comment 73, make the following changes:

- 1) Add the following to the end of Paragraph 1:

"Conformant FCIP implementations MUST support peer authentication using pre-shared key and MAY support peer authentication using digital certificates. Peer authentication using public key encryption methods outlined in IKE [2409] Section 5.2 and 5.3 SHOULD NOT be used."

- 2) Change the last sentence of Paragraph 2 from:

"FCIP Entities MUST support "Main Mode" operation in Phase 1 and MAY support "Aggressive Mode" if identity protection is not required."

to:

"FCIP implementations MUST support IKE Main Mode and SHOULD support Aggressive Mode."

- 3) Add the following at the end of paragraph 3: "The Phase 2 Quick Mode exchanges MUST explicitly carry the Identity Payload fields (IDci and IDcr). Each Phase 2 payload SHOULD carry a single IP Address and a single non-zero port number and SHOULD NOT use the IP Subnet or IP Address Range formats. Other ID payload formats MUST NOT be used."
- 4) Add the following new paragraph after paragraph 3: "Since the number of Phase 2 SAs may be limited, Phase 2 delete messages may be sent for idle SAs. The receipt of a Phase 2 delete message SHOULD NOT be interpreted as a reason for tearing down an FCIP Link or any of its TCP connections. When there is new activity on that idle link, a new Phase 2 SA MUST be re-established."
- 5) In the paragraph that starts with "For the purposes of...", change the last sentence from:

"For the purposes of establishing IKE Phase 1 SA, static IP Addresses are typically used for identification."

to:

"Within IKE Phase 1, FCIP implementations must support the ID_IPV4_ADDR, ID_IPV6_ADDR (if the protocol stack supports IPv6) and ID_FQDN Identity Payloads. If FCIP Endpoint addresses are

dynamically assigned, it may be beneficial to use ID_FQDN, and for this reason, IP_FQDN Identity Payload MUST be supported. Other identity payloads (ID_USER_FQDN, ID_DER_ASN1_GN, ID_KEY_ID) SHOULD NOT be used.

Comment 63

-- Section 10.1 Threat Models

Using a general purpose, wide-area network such as an IP Network as a substitute for physical cabling introduces some security problems not normally encountered in Fibre Channel Fabrics. FC interconnect cabling typically is protected physically from outside access. Public IP Networks allow hostile parties to impact the security of the transport infrastructure.

[E] The intent is fine, but the "substitute" word has the same problems as the use of "replace" in Section 6.3. This is about "implementation of functionality that was originally specified only to use dedicated physical cabling" or something like that, as indicated in the latter two sentences.

Accepted with the following result

Replace "substitute" with "functional replacement".

Comment 64

-- Section 10.1 Threat Models

The general effect is that the security of the entire FC Fabric is only as good as the security of the entire IP Network through which it tunnels.

[E] "the entire" --> "any". This may be the first occurrence of "tunnel" in the document - that might be better rephrased to talk about "any IP Network" carrying FCIP links that are part of the fabric".

Accepted with the following result

Replace the cited text with: "The general effect is that the security of an FC Fabric is only as good as the security of the entire IP Network that carries the FCIP Links used by that FC Fabric."

Comment 65

-- Section 10.1 Threat Models

- 2) Unauthorized agents can monitor and manipulate Fibre Channel traffic flowing over physical media used by the IP Network and under control of the agent.

[E] "under control of" is too strong. "accessible to" would be better, especially for monitoring.

Accepted

Comment 66

-- Section 10.1 Threat Models

- 5) The payload of an FCIP Encapsulated frame may be altered or transformed in such a way that it preserves the TCP Checksum transform while altering content.

[E] Put a period after "transformed" and rephrase the rest to say that the TCP checksum, FCIP ones complement checks and FC frame CRC do not protect against this, as all of them can be modified or regenerated by a malicious and determined adversary.

Accepted

Comment 67

-- Section 10.3 - FCIP Security Components

FCIP Security compliant implementations MUST implement IPsec Protocol Suite based cryptographic authentication and data integrity [14], as well as confidentiality using algorithms and transforms as described in this section.

[E] Please insert the ESP acronym into this sentence.

[E] The official capitalization is "IPsec", not "IPSec".

Accepted with the following results

- 1) Replace "IPSec" with "ESP and IPsec"; and
- 2) Verify correct IPsec usage throughout.

[Comment 68 Technical](#)

-- Section 10.3.1 - IPSec ESP Authentication and Confidentiality

FCIP Entities MUST implement IPSec ESP [16] in Tunnel Mode for providing Data Integrity and Confidentiality. FCIP Entities MAY implement IPSec ESP in Transport Mode, if deployment considerations require use of Transport Mode.

[T] Those deployment considerations include RFC 2401 requirement for Transport mode because the IPsec implementation is a host implementation rather than a security gateway. I thought this was understood by the FCIP authors, and it needs to be explicit here including an appropriate reference to RFC 2401. I expect to be able to double-check this requirement with the IETF Security Area in the next few days.

Rejected

As per the consensus taken at the March 2002 IETF meeting, Transport Mode implementation is a MAY.

[Comment 69](#)

-- Section 10.3.2 - Key Management

FCIP Entities MUST support IKE [18] for peer authentication, negotiation of Security Associations (SA) and Key Management using the IPSec DOI [17]. Manual keying for establishing SA is not permitted since it does not provide the necessary elements for rekeying (see section 10.3.3).

[E] Reword last sentence to include "MUST NOT use", "SHALL NOT use" or equivalent.

Accepted

[Comment 70](#)

-- Section 10.3.2 - Key Management

Repeated rekeying using "Quick Mode" on the same shared secret will over time, reduce the cryptographic properties of that secret. To overcome this, Phase 1 MAY be invoked periodically to create a new set of IKE shared secrets and related security parameters.

[E] "MAY" --> "SHOULD" (I believe this captures the intention).

[Accepted](#)

[Comment 71 Technical](#)

-- Section 10.3.2 - Key Management

When pre-shared keys are used, IKE Aggressive Mode SHOULD be used and Main Mode SHOULD NOT be used.

[T] I think this is too strong and will cause problems. Pre-Shared keys are a "MUST", Aggressive Mode is a "MAY", so this "SHOULD" is inconsistent. The issue with Main Mode arises when dynamically assigned IP addresses are used (and hence Main Mode can't figure out which pre-shared key to use). The escape from this box appears to be that Aggressive Mode is a MUST (SHOULD?) when dynamic assignment of IP addresses to FCIP implementations is used, but support for dynamic assignment of IP addresses is NOT REQUIRED.

The problem with this approach is that one gets into trouble on one end of an FCIP Link when the *other* end has its IP address dynamically assigned. The obvious solutions to this issue are to forbid (MUST NOT) dynamic IP address assignment (which has no chance of making it through the IESG) or to REQUIRE Aggressive Mode (as iFCP does). In addition, the IPS Security draft appears to need some editing to allow Aggressive Mode to be a "MAY" for FCIP (and iFCP). Darn - I thought we had this issue closed in Huntington Beach - did I miss something?

[Accepted with the following results](#)

Replace the cited text with:

"When pre-shared keys are used, IKE Main Mode is usable only when both peers of an FCIP Link use statically assigned IP addresses. When support for dynamically assigned IP Addresses is attempted in conjunction with Main Mode, use of group pre-shared keys would be forced, and the use of group pre-shared keys in combination with

Main Mode is not recommended as it exposes the deployed environment for man-in-the-middle attacks. Therefore, if either peer of an FCIP Link uses dynamically assigned address, Aggressive Mode SHOULD be used and Main Mode SHOULD NOT be used."

Comment 72

-- Section 10.3.2 - Key Management

Support for IKE Oakley Groups is not required.

[T] What does this refer to? At a minimum, it needs a reference.

Accepted

The reference for IKE Oakley Groups is RFC 2412. A suitable informational reference will be added.

Comment 73

-- Section 10.3.2 - Key Management

For the purposes of establishing a secure FCIP Link, the two participating FCIP Entities consult a Security Policy Database (SPD).

[E] For this and the SAD, please add RFC 2401 references.

Accepted with the following results

- 1) Add a new normative reference for SPD to section 4.4.1 of RFC 2401; and
- 2) Add a new normative reference for SAD to section 4.4.3 of RFC 2401.

[Comment 74 Technical](#)

-- Section 10.4.2 - TCP Connection Security Associations (SAs)

For a TCP Connection establishment, IKE Phase 2 is employed, resulting in an SA, identified by an SPI. All IP datagrams of the TCP Connection MUST carry an ESP header with a valid SPI and Sequence Number to be accepted as valid by the receiving peer.

[T] The requirement for a phase 2 SA per TCP connection has been removed. This text (and possibly the rest of this section) need some editing to reflect that.

[Accepted with the following results](#)

- 1) Replace the cited text with: "Each TCP connection MUST be protected by an IKE Phase 2 SA. Traffic from one or more than one TCP connection may flow within each IPsec Phase 2 SA. While it is possible for a IKE Phase 2 SA to protect multiple TCP connections, all packets of a TCP connection is protected using only one IKE Phase 2 SA. FCIP implementations need not verify that the IP addresses and port numbers in the packet match any locally stored per-connection values, leaving this check to be performed by the IPsec layer."
- 2) Delete the last paragraph of the section, which currently reads: "When a TCP Connection is terminated or closed, all SAs associated with it MUST be removed from the local SAD."

[Comment 75](#)

-- Section 10.4.3 - Handling data integrity and confidentiality violations

An implementation MAY audit such events as a diagnostic aid.

[E] Almost but not quite. Auditing is about a lot more than "diagnostic aid". See Section 7 of RFC 2401, make sure the text is consistent with it, and refer to it.

[Accepted with the following results](#)

Replace the cited text with: "An implementation SHOULD follow guidelines for auditing all auditable ESP events per Section 7 of RFC2401."

[Comment 76 Technical](#)

-- Section 10.4.3 - Handling data integrity and confidentiality violations

Confidentiality checks MUST be performed if Confidentiality is enabled.

[T] And what would those be, please? Replace this with a prohibition on use of Confidentiality without Integrity.

[Accepted with the following results](#)

Replace the first "Confidentiality" with "Integrity".

[Comment 77 Technical](#)

-- Section 10.4.4 - Handling SA parameter mismatches

When SA parameters do not match, the TCP Connection may reach a point where no traffic moves, or there are excessive TCP retransmissions. In such a case, either side MAY take one of the following actions:

- a) Reestablish another set of SA parameters; or
- b) Close the TCP Connection and notify the FC Entity with the reason for the closure.

[T/E] Needs some more explanation, including an example of the sort of mismatch that could cause this problem. Recall that IKE negotiates SA parameters, and this fact needs to be included in the explanation and example.

[Accepted but perhaps not as intended](#)

The handling of SA parameter mismatches belongs in a security draft, not in FCIP. Therefore, section 10.4.4 will be removed.

Comment 78

-- Section 11.1 - Performance Considerations

Traditionally, the links between FC Fabric components have been characterized by low latency and high throughput. The purpose of FCIP is to replace some of these links with an IP Network,

[E] There's the "replace" work again. See comment against Section 6.3.

Accepted with the following result

Replace "...to replace some of these links with..." with "...to provide functionality equivalent to these links using..."

Comment 79

-- Section 11.2 - IP Quality of Service (QoS) Support

It is RECOMMENDED that some form of preferential QoS be used for FCIP traffic to minimize latency and drop precedence. No particular form of QoS is recommended.

[E] "drop precedence" --> "packet drops"

Accepted

Comment 80

-- Section 11.2 - IP Quality of Service (QoS) Support

If diffserv/PHB QoS is NOT implemented, the DSCP field for all IP packets SHALL be set to '000000'.

[T] Sorry, wrong answer. That's not consistent with RFC 2474. Best bet is to drop this sentence, but if the authors want to say something here, they should contact me directly to discuss/vet appropriate text.

Accepted with the following results

Replace the cited text with: "If no form of preferential QoS is implemented, the DSCP field SHOULD be set to '000000' to avoid negative impacts on other network components and services that may be caused by uncontrolled usage of non-zero values of the DSCP field."

Comment 81

-- Section 12 - Normative References

[E] This needs to be carefully checked to minimize normative references. [7] is definitely non-normative. Most of the QoS references are or can be non-normative, specifically [11], [22], [23], [24]). [22] is an Informational RFC and hence has to be referenced in a non-normative fashion, and I really want to see [23] and [24] be non-normative (else any FCIP implementation will have to implement both EF and AF, which is surely not what was intended). Need to add a non-normative MPLS reference.

Accepted with the following results

- 1) Remove reference 7 entirely;
- 2) Make the SNTP reference [13] informative;
- 3) Make all references from section 11 (Performance/QoS) informative (note: this covers all the other cited references); and
- 4) Add an informative reference to RFC 3031 for MPLS.

All other references appear to be necessary.

Comment 82

-- Section 14 - Acknowledgments

[E] This is a good place to thank everyone who's reviewed the document, commented on ideas in it, or helped in other ways.

Accepted

Acknowledge Mallikarjun Chadalapaka and David Black.

Comment 83

-- Section 15 - Contributors' Addresses

We'll try this structure of not separating the folks listed on p.1 from the other contributors and see if there are any procedural objections. It's unusual to say the least.

Unusual demands beget unusual responses.

[Comment 84 Technical](#)

-- Annex A - IANA Considerations

[T] Instruct IANA to change the authority for those port allocations to reference this RFC when it is approved. Add a sentence forbidding the use of UDP with FCIP even though IANA has allocated a port.

[Accepted](#)

[Comment 85](#)

-- Annex A - IANA Considerations

[T] Will need to reference the SOF/EOF registry that the FC Frame Encapsulation Draft will need to set up. Point to the Protocol# allocation done by that draft also. If a connection usage registry is needed (see comment against Section 8, details of that will have to go here).

[Accepted \(Partially\)](#)

Add discussion of Protocol# allocation.

The SOF/EOF registry will not happen.

[Comment 86](#)

-- Annex A - IANA Considerations

[E] Should the ANNEXes be APPENDIXes instead?

[Accepted, make exactly the proposed change](#)

[Comment 87](#)

-- Annex C - Example of synchronization recovery algorithm

[T] See comment on this Annex under Section 6.6.2.3 above.

[See response to comment 37.](#)

2. Comments from Mallikarjun Chadalapaka

=====

Comment 88

> 2. Purpose, Motivation and Objectives

.....

> Network. Since Fibre Channel and IP Networking technologies are
> compatible,

I am not sure what's implied by this sentence....

Generally, I would think that the motivation to extend an FC SAN using IP networks is based on the ubiquity of the IP networks.

No changes made

The cited phrase says that it is technologically possible to use IP Networks to extend FC SANs. That is, IP Networks are NOT tin cans and strings, but are in fact electromechanical signaling systems that are similar enough to FC to be useful in FC SANs.

Comment 89

> 2. Purpose, Motivation and Objectives

.....

> The fundamental assumption made in this specification is that the
> Fibre Channel traffic is carried over the IP Network in such a
> manner that the Fibre Channel Fabric and all Fibre Channel
> devices on the Fabric are unaware of the presence of the IP
> Network.

Can someone please comment on the protocol interactions that result in the failure to set up an FCIP Link if this latency expectation is not met?

Inquiry

All data frames will fail their transit time test and will be discarded. No data will move from application to application and the end users will riot.

Comment 90

> 2. Purpose, Motivation and Objectives

.....

> 2) apply the mechanism described in 1) to an FC Fabric using an
> IP network as an interconnect for two or more islands in an
> FC Fabric.

S/b "an" w/ "the"; S/b "for" w/ "between"

Accepted (Partially)

Change "for ... islands" to "between ... islands".

Do not change "an IP Network" to "the IP Network" because
this specification assumes that there can be more than one
IP Network.

Comment 91

> 3. Relationship to Fibre Channel Standards

.....

> FC is standardized under American National Standard for
> Information Systems of the National Committee for Information
> Technology Standards (ANSI-NCITS) in its T11 technical committee.

I believe ANSI stands for American National Standards Institute.
Also, NCITS had been changed to INCITS last year.

Accepted with the following result

Replace the cited text with: "FC is standardized as a family of
American National Standards developed by the T11 technical committee
of INCITS (InterNational Committee for Information Technology
Standards).

Comment 92

- > 3. Relationship to Fibre Channel Standards
-
- > FC-BB and FC-BB-2 describe the relationship between an FC Fabrics
- > and interconnect technologies not defined in by Fibre Channel
- > standards (e.g., ATM and SONET). FC-BB-2 is the natural Fibre
- > Channel home for describing relationships to TCP/IP and FCIP.

This is the first instance of FCIP's usage as a *protocol*. It would be best preceded by a definition of what FCIP means as a protocol. The previous text only describes the objectives of the document, but not the protocol.

Accepted but not as intended

The section in which the cited text appears is attempting to describe the relationship between standards documents. The second sentence in the cited text fails to maintain the purpose of the section.

Actions to be taken

Replace the second cited sentence with: "FC-BB-2 is the Fibre Channel document describing the relationships between FC and TCP/IP, including the FC use of FCIP.

Comment 93

- > 3.2 This Specification and Fibre Channel Standards
-
- > No attempt is being made to define a specific API between an FCIP
- > Entity and an FC Entity at this time because doing so risks
- > compromising the performance and efficacy of the resulting
- > products.

I am somewhat concerned that the names are implying an incorrect distribution of functionality between "FC Entity" and "FCIP Entity"....

From the name, I had assumed that FC Entity is simply a standard FC fabric element with no FCIP-isms. But further reading of the document made me realize that it in fact knows about the TCP connections and even actively participates in QoS and authentication decisions.

As a first time reader, it appears to me that retaining only the FC E_Port functionality perhaps additionally providing timestamp and flow control services to the FCIP Entity (and dropping everything else into

the FCIP Entity) may be a lot cleaner distribution of functionality. At any rate, I would like to understand the current distribution rationale.

BTW, can someone please clarify if the expected "role" of the FC Entity on the FC side is indeed an E_Port? It appears so from the requirement that FCIP be totally transparent to the FC fabric, but I don't see it called out in Annex G.

Inquiry

The process of developing FCIP has shown repeatedly that FCIP should contain only the Fibre Channel knowledge required to interact with TCP/IP. Attempting to include FC concepts such as R_A_TOV only leads to confusion in the IP Network community that in turn leads to proposals for changes in FCIP that are, in fact, proposals to change the definition of R_A_TOV in ways that are unacceptable to T11.

The dividing line between an FC Entity and an FCIP Entity is drawn along political not technological lines.

To understand the function of an FC Entity, you must read FC-BB-2, <http://www.t11.org/t11/docreg.nsf/ldl/fc-bb-2>.

Comment 94

```
> 3.2 This Specification and Fibre Channel Standards
> .....
> ....fully functional and compliant
>   products can be built provided they contain both an FCIP Entity
>   and an FC Entity. The only products that cannot be built are
>   those that contain only one or the other.
```

The last sentence seems to stress the obvious at first glance, but I would think that products with just the FC Entity should be able to be built (not that one would...) to act as regular fabric elements with no FCIP features?

Also, I take it that products with two FC Entities and one FCIP Entity for ex. are disallowed - but the last sentence seems to imply otherwise.

No changes made

This section and the cited paragraph are about the compliance relationship between FCIP and FC-BB-2. The Model section is where the numbers of FC Entities and FCIP Entities are discussed.

[Comment 95](#)

- > 4. Terminology
-
- > FCIP Entity - The principal FCIP interface point to the IP
- > Network (see section 6.4).

This doesn't sound right to me....this definition is more appropriate for FCIP_LEP. This can perhaps be described as "the entity responsible for the FCIP protocol exchanges on the IP Network and which encompasses FCIP_LEP(s) and FCIP Control & Services module".

[Accepted](#)

[Comment 96](#)

- > 5. Protocol Summary
- ...
- > 2) Viewed from the IP Network perspective, FCIP Entities are
- > peers and communicate using TCP/IP. Each FCIP Entity is a TCP
- > endpoint in the IP-based network.

The second sentence seems a little context-sensitive, since each FCIP Entity can be the TCP endpoint for several TCP connections.

[Changed as described in the response to comment 14.](#)

[Comment 97](#)

- > 5. Protocol Summary
- ...
- > 3) Viewed from the FC Fabric perspective, pairs of FCIP
- > Entities, in combination with their associated FC Entities,
- > serve as an FC Frame transmission component of the FC Fabric.
- > The FC End Nodes are unaware of the existence of the FCIP
- > Link.

Can a "FC Frame transmission component" be something other than an E_Port?

[Inquiry](#)

An FC Frame transmission component can also be a B_Port.

[Comment 98](#)

```
> 5. Protocol Summary
...
> 6) An FCIP Entity MAY contain multiple FCIP Link Endpoints,
> but...
```

I would add "each of which MAY service multiple TCP connections, " here...

[Rejected](#)

Such a change would detract from the focus on the point to be made.

[Comment 99](#)

```
> 5. Protocol Summary
...
> 7) When multiple FCIP_LEPs with multiple FCIP_DEs are in use,
> selection of which FCIP_DE to use for encapsulating and
> transmitting a given FC Frame is outside the scope of this
> document. FCIP Entities do not actively participate in FC
> Frame routing.
```

Couple of comments on this bullet (7) -

Let's consider the case of multiple FCIP_DEs in one FCIP_LEP. This draft does not specify how each inbound FC frame from the FC fabric is distributed onto one of these FCIP_DEs (TCP connections). Where is it specified in wrt routing on TCP connections? I take it that the regular FC fabric routing rules aren't quite applicable in this case.

To stress the obvious, I think we should have some standard covering this case - else we will end up with frames destined to the same D_ID being striped on multiple TCP connections, and thus ending up OOO.

[Accepted in principle](#)

The standard covering the questions raised by this comment is FC-BB-2.

[Action to be taken](#)

Replace "...is outside the scope of this document." with "...is covered in FC-BB-2 [4]."

Comment 100

> 5. Protocol Summary
...
> 13) On a given TCP Connection, this specification relies on
> TCP/IP to deliver a byte stream in the same order that it was
> sent.

Perhaps we should add - ", but allows confirmation of the same for each frame by checking the FCIP and FC CRCs."....

Rejected

1) FCIP has no CRC to check. 2) It is difficult to see how checking the FC CRC would confirm in order delivery.

Comment 101

> 6.3 FC Entity
....
> the combination shown in figure 3. As another example, the
> combination cannot be used to replace cable connections in a
> Fibre Channel Arbitrated Loop because loop primitive signals
> cannot be encapsulated for transmission over TCP.

I am not sure the last sentence is needed since figure 3 is explicit about the usage of fabrics.

Changed as described in the response to comment 24.

Comment 102

> 6.4 FCIP Entity
.....
> The FCIP Entity is the connection interface point for the IP
> Network

As commented earlier, the "connection interface point" doesn't sound totally correct - I would suggest "interfacing element" instead.

Accepted

[Comment 103](#)

> 6.4 FCIP Entity
.....
> and is the owner of the IP Address and Well Known Port used to
> form TCP Connections. An FC Fabric to IP Network interface
> product SHALL provide each FC Entity FCIP Entity pair contained
> in the product

May I suggest a new term "FC-FCIP Pair" in place of "FC Entity FCIP Entity pair".

I think it improves the general readability.

[Accepted with the following results](#)

Change all occurrences of "FC Entity FCIP Entity pair" to "FC/FCIP Entity pair" to match the terminology used in FC-BB-2.

Also add the following to section 4:

FC/FCIP Entity pair - The combination of one FC Entity and one FCIP entity.

[Comment 104](#)

> 6.4 FCIP Entity
.....
> FC Entity with an interface to key IP Network features. The
> interfaces to the IP Network features is implementation specific,
> however, to maintain interoperability, the notable TCP/IP
> mechanisms used are specified in this document as follows:

I think the last sentence is incorrectly implying interoperability issues around the (private) FC Entity-FCIP Entity interface.

I would suggest a rewrite for the last sentence as something like -

"The interfaces to the IP Network features are implementation-specific. To aid interoperability, this document specifies the notable TCP/IP Network features that are typically used."

[Changed as described in the response to comment 28.](#)

Comment 105

- > 6.5 FCIP Link Endpoint (FCIP_LEP)
-
- > An FCIP_LEP
- > MAY communicate with its FC Entity counterpart to coordinate
- > flow control.

Suggest adding the phrase "across the domains".

Rejected

Without a definition of "domains" such a change adds confusion, not reduced it.

Comment 106

- > 6.6 FCIP Data Engine (FCIP_DE)
-
- > 7) In the absence of errors, the de-encapsulated FC Frame and
- > time stamp SHALL be passed to the FC Transmitter Portal for
- > delivery to the FC Entity.

It is nice to add a sentence about the handling in the presence of errors. At a minimum, this should provide a cross-reference to the error detection section.

Accepted with the following results

Add the following at the end of the paragraph: "Error handling is discussed in section 6.6.2.2."

[Comment 107](#)

> 6.6.1 FCIP Encapsulation of FC Frames

....

> The FCIP Special Frame SHALL only be sent as the first bytes
> transmitted in each direction on a newly formed TCP Connection
> and only one FCIP Special Frame SHALL be transmitted in each
> direction at that time (see section 9.1). After that all FCIP
> Frames SHALL have the SF bit set to 0.

This para seems slightly out of context here, and perhaps should be moved to section 9.1. This usage semantics should ideally be preceded by what *is* a Special Frame and its purpose - though I could gather that from the usage and format descriptions in the entire document, I don't really find a place where its purpose is defined...

[Accepted, see response to comment 44](#)

[Comment 108](#)

> 6.6.1 FCIP Encapsulation of FC Frames

....

> Table 1 summarizes the usage of the pFlags SF and Ch bits.

- It may be useful to add a protocol-specific bit to distinguish originated vs echoed SF, so the parsing and validation can be self-contained.

- Also, I think a sentence should be added which says that the - pFlags field SHALL contain the ones-complement of the pFlags field.

[Accepted](#)

[Comment 109 Technical](#)

- > 6.6.1 FCIP Encapsulation of FC Frames
-
- > The CRCV (CRC Valid) Flag SHALL be set to 0.
- >
- > The CRC field SHALL be set to 0.

I am surprised that the FCIP encapsulated header is never protected by a CRC. The error detection section clearly acknowledges the possibility that TCP checksum could be inadequate for certain deployments - and even suggests checking the FC frame CRC (sort of like a data digest) on reception at the Encapsulated Frame Receiver Portal.

My recommendation is to require an FCIP Entity to employ the header CRC if the SF that it receives asks for CRC - IOW, similar to iSCSI's approach. So, I guess I am also advocating a new bit in the pFlags field to announce this. When this CRC expectation is announced, the FC CRC checking test in 6.6.2.2 should also be a mandatory test - from the "semi-optional" list it is currently in.

[Accepted with the following result](#)

Add the following to the new section created in response to comment 44: "Note: Owing to the limited manner in which the FSF is used and the requirement that the FSF be echoed without changes before a TCP connection is allowed to carry user data, no error checking beyond that provided by TCP is deemed necessary."

[Comment 110](#)

- > 6.6.2 FCIP Data Engine Error Detection and Recover

The last word in the above should be "Recovery".

[Accepted](#)

Comment 111

> 6.6.2.1 TCP Assistance With Error Detection and Recovery
.....
> Thus, the byte stream passed from TCP to
> the FCIP_LEP will be in order and free of errors detectable by
> the TCP checksum. If TCP did not perform these functions, the
> FCIP_LEP would have to.

Suggest rewording the last sentence (TCP always performs these functions to *its* satisfaction, the question is if FCIP feels the same; secondly, FCIP_LEP's error mgmt behavior is not dynamically contingent upon TCP's behavior as this sentence implies) -

"To address the possibility that TCP did not perform these functions adequately in a given FCIP deployment context, the FCIP_LEP verifies if these expectations are met."

Changed as described in the response to comment 32.

Comment 112

> 6.6.2.2 Errors in FCIP Headers and Discarding FCIP Frames
.....
> Further, some errors in the encapsulation will result in the
> FCIP_DE losing synchronization with the FCIP frames in the byte
> stream

I don't see "frames" here with the uppercase "F" used everywhere else.

Also, one observation is that FCEncap document uses "frames" consistently throughout, whereas this document chooses to use uppercase "F" (mostly).

Accepted with notes

"FCIP frame" s/b "FCIP Frame". That is a specific named construct that is defined and used in FCIP.

FC Frame Encapsulation uses "frame" (lowercase) when referring to a Fibre Channel frame because lowercase frame is the Fibre Channel term.

[Comment 113](#)

```
> 6.6.2.2 Errors in FCIP Headers and Discarding FCIP Frames
> .....
> 1) Length field validation -- 15 < Length < 545;
```

I assume "Frame Length" is meant by "Length" above.

[Accepted with the following results](#)

Change "Length" to "Frame Length".

[Comment 114](#)

```
> 6.6.2.3 Synchronization Failures
> .....
> If an FCIP_DE determines that it cannot find the next FCIP Frame
> header in the byte stream entering through the Encapsulated
> Frame Receiver Portal, the FCIP_DE SHALL either:
```

S/b "either" w/ "do one of the following".

[Accepted](#)

[Comment 115](#)

```
> 6.6.2.3 Synchronization Failures
> .....
> b) recover synchronization by searching the bytes delivered by
> the Encapsulated Frame Receiver Portal for a valid FCIP Frame
> header having the correct properties
```

Useful to refer to 6.6.2.2 here for "correct properties".

[Accepted](#)

[Comment 116](#)

```
> 7. Checking FC Frame Transit Times in the IP Network
> ....
> requirement in the FC Entity is based on a desire to include
> the Frame.
```

S/b "in" w/ "on"

[Accepted](#)

[Comment 117 Technical](#)

- > 7. Checking FC Frame Transit Times in the IP Network
-
- > ... If no synchronized time stamp value is available to
- > accompany the entering FC Frame a value of zero SHALL be
- > supplied.

From this, it isn't clear to me if FCIP *requires* only Synchronized operation. If so, the draft should also specify how implementations are expected to deal with "benign" transitions into and out of the Unsynchronized state (i.e. transitions happening when no Encapsulated Frames are being received).

Rejected

Class F frames can be transmitted and forwarded in the Unsynchronized state.

The requirements for transit time determinations are in FC-BB-2 for all the reasons described in the response to comment 40.

[Comment 118](#)

- > 7. Checking FC Frame Transit Times in the IP Network
-
- > The FC Entity SHALL
- > verify that the FC Entity it is communicating with on an FCIP
- > Link is using the same synchronized time source as it is, either
- > Fibre Channel services or SNTP server.

I see a chicken-and-egg problem for some implementations: Since Fibre Channel time services are not available until the FCIP Link is successfully set up and since timestamp is to be carried on every FC Frame (including those Fibre Channel time service exchanges) once the Link is set up, the only decent option for an implementation (assuming it supports only Synchronized operation) is to rely on SNTP server. Is this correct?

If Unsynchronized operation is intended to be disallowed (my earlier question), then it seems to me that SNTP is the only option.

Inquiry

Yes but Unsynchronized operation is allowed precisely so that Class F frames can be processed to setup the FC Time Services Time for use by the FC Entity in computing transit times.

Comment 119

> 8. The FCIP Special Frame

.....

> The FCIP Special Frame SHALL only be sent as the first bytes
> transmitted in each direction on a newly formed TCP Connection
> and only one FCIP Special Frame SHALL be transmitted in each
> direction.

A general comment on this wording (and others similar to this). I would suggest rewording (to be much stronger) to:

The FCIP Special Frame SHALL be the first application data exchanged on each newly formed TCP connection, and only one FCIP Special Frame SHALL be transmitted in each direction.

Rejected

The use of "application data" could cause confusion with application data generated by FC Endnodes.

See also comment 44 for more information on the SF and how it will be described in the next FCIP revision.

Comment 120

- > 8. The FCIP Special Frame
-
- > Note: The combination of the Source FC Entity World Wide Name and
- > Source FCIP Entity Identifier fields uniquely identifies every FC
- > Entity FCIP Entity pair in the IP Network.

- S/b "Source FCIP Entity Identifier" w/ "Source FC/FCIP Entity Identifier"

- Also I take it that the "FC/FCIP Entity Identifier" is unique only within the scope of the given FC Entity's WWN. So, does the model allow multiple FCIP Entities to be associated with the FC Fabric Entity WWN?

- From this statement, it implies to me that the Destination FC/FCIP Entity Identifier must be present in the special frame to ensure that the TCP connection is indeed established to the right "FC/FCIP Entity" under the scope of the given Destination FC Fabric Entity WWN. What am I missing?

Accepted (Partially) as follows

Make the change described in the first bullet.

Response to the second bullet: Yes, the "FC/FCIP Entity Identifier" is unique only within the scope of the given FC Entity's WWN. Yes, the model allows multiple FCIP Entities to be associated with the FC Fabric Entity WWN.

The following changes will be made to clarify this and other issues raised during the discussion of this comment on the ips list.

Add the following to section 4:

FC Fabric Entity - A Fibre Channel specific element containing one or more Interconnect_Ports (see [FC-SW-2 ref 5]) and one or more FC/FCIP Entity pairs. See FC-BB-2 [4] for a details about FC Fabric Entities.

Change the title of figure 3:

from: FC Entity and FCIP Entity Model
to: Model for Two Connected FC/FCIP Entity Pairs

Add the following sentence at the end of the paragraph preceding figure 3:

An FC Fabric Entity may contain multiple instances of the FC/FCIP Entity pair shown on either the right-hand or left-hand side of figure 3.

Change the first sentence following figure 4:

from:

The FCIP Entity is the connection interface point for the IP Network and is the owner of the IP Address and Well Known Port used to form TCP Connections.

to:

The FCIP Entity receives TCP connect requests on behalf of the FCIP_LEPs that it manages. In support of this, the FCIP Entity is the sole owner of at least one TCP port/IP Address combination used to form TCP Connections. The TCP port may be the FCIP well known port at a given IP Address.

Response to the third bullet: The motivation for permitting the "Destination FC Fabric Entity WWN" to be zero is to allow the SF to be used as a poor-man's SLP (a configuration discovery mechanism). This usage is described in the new section added in response to comment 44.

Comment 121

> 8. The FCIP Special Frame

.....

> The Destination FC Fabric Entity World Wide Name field MAY
> contain

Why isn't the above requirement a SHALL?

Inquiry - see the response to comment 120.

Comment 122

> 9.1.2.2 Dynamic Creation of New TCP Connections

...

- > - The expected Destination FC Fabric Entity World Wide Name of
- > the FC Entity FCIP Entity pair to which the TCP Connection is
- > being made
- > - TCP Connection Parameters (see section 9.3)
- > - Quality of Service Information (see section 11)
- >
- > Based on this information, the FCIP Entity SHALL generate a TCP
- > connect request [8] to the FCIP Well-Known Port of 3225 (or other
- > configuration specific port number) at the IP Address specified
- > by the service advertisement. If the TCP connect request is
- > rejected, act to limit unnecessary repetition of attempts to
- > establish similar connections. If the TCP connect request is
- > accepted, the FCIP Entity SHALL follow the steps described in
- > section 9.1.2.3 to complete the establishment of a new FCIP_DE.
- >
- > It is recommended that an FCIP Entity not initiate TCP connect
- > requests to another FCIP Entity if incoming TCP connect requests
- > from that FCIP Entity have already been accepted.

This entire text is duplicated from previous section 9.1.2.1. Seems like we could do with one instance (possibly in a new subsection).

Rejected

The information on connection establishment is indeed common to section 9.1.2.1, but its relationship to SLP is not. Section 9.1.2.2 already is less than one page long, that and a desire to have some amount of readable flow (as opposed to "do this, then do what it says to do over there, then do that, then do this other thing that is talked about in another section") seems like ample motivation for the current organization.

[Comment 123](#)

- > 9.1.2.3 Connection Setup After a Successful TCP Connect Request
-
- > All fields in the FCIP Special
- > Frame SHALL be filled in as described in section 8, particularly:

While the sentence above is unequivocally clear, I don't understand the need for all the text that follows "particularly".... It is confusingly repetitive since as far as I can tell, all these are covered in more or less the same language in section 8.

No changes made

The draft is structured so that the Special Frame can be used for more features than just initial connection setup. As of yet, additional uses have not been defined, but it seems prudent to simplify such enhancements to ensure that future versions remain correct with minimal editorial effort.

In keeping with this, section 8 describes the overall SF format and section 9.1... defines SF usage during connection setup.

[Comment 124](#)

- > 9.1.2.3 Connection Setup After a Successful TCP Connect Request
-
- > After the FCIP Special Frame bytes are sent on the newly formed
- > connection, the FCIP Entity SHALL wait for the FCIP Special Frame
- > to be echoed as the first bytes received on the newly formed
- > connection.

- S/b "bytes are" w/ "is"
- S/b "first bytes" w/ "first TCP segment data"
- I take it that the onus is on the FCIP Entity that did the TCP active open to send the SF. That leads me to: What if there's a TCP simultaneous open?
Would not each end up sending the SF and waiting for the echo?
Also, would not the earlier rule on only one frame being transferred in each direction be violated then?

Accepted (Partially) as follows

Make the change described in the first bullet.

Do not make the change described in the second bullet because it requires more knowledge of TCP than FCIP should have.

The last issue is resolved as per the response to comment 57.

Comment 125

```
> 9.1.2.3 Connection Setup After a Successful TCP Connect Request
> .....
>   If the echoed FCIP Special Frame bytes do not exactly match the
>   FCIP Special Frame bytes sent (words 7 through 17 inclusive), the
>   FCIP Entity SHALL close the TCP Connection and notify the FC
>   Entity with the reason for the closure.
```

Seems like all the 11 words are required to be compared. If so, what is the Ch bit being used for? IOW, why SHALL it be set by the responder?

Inquiry

The Ch bit serves to identify the difference between changes made intentionally by the echoing FCIP Entity and changes that result from transmission errors.

Comment 126

```
> 9.1.2.3 Connection Setup After a Successful TCP Connect Request
> .....
>   The FCIP Entity SHALL listen for new TCP Connection requests [8]
>   on the FCIP Well-Known Port (3225). An FCIP Entity MAY also
>   accept and establish TCP Connections to a TCP port number other
>   than the FCIP Well-Known Port, as configured by the network
>   administrator.
```

It may be useful to add that this usage is outside the scope of this document.

Accepted with the following results

Add "... in a manner outside the scope of this specification." at the end of the last cited sentence.

[Comment 127](#)

```
> 9.1.2.3 Connection Setup After a Successful TCP Connect Request
> .....
>   The FCIP Entity SHALL determine the following information about
>   the requested connection:
>
>   - Whether the requested connection is allowed
```

I take it that by means not specified in this spec? If so, it's useful to call it out.

[Accepted with the following results](#)

- 1) Change "Whether the requested connection is allowed" to "Whether the "shared" database (see section 9.1.1) allows the requested connection"; and
- 2) Ensure that this change is applied wherever it is needed.

[Comment 128 Technical](#)

```
> 9.1.3 Processing Incoming TCP Connect Requests.....
>   abort the TCP connect request [8]. If the requested connection is
```

I was told that "abort" isn't available on all implementations - perhaps "abort/close" would be better....

[Accepted with the following results](#)

Change "abort the TCP connect request [8]" to "reject the connect request using appropriate TCP means"

Comment 129

- > 9.1.2.3 Connection Setup After a Successful TCP Connect Request
.....
> The FCIP Entity MAY apply a timeout of not less than 90 seconds
> to the waiting for the FCIP Special Frame bytes and if the
> timeout expires the FCIP Entity SHALL close the TCP Connection
> and notify the FC Entity with the reason for the closure.

I am not clear why this notification is required, since the local FC Entity did not motivate the TCP connection establishment. If it is intended for logging, perhaps notifying the Control & Services module would be more appropriate.

- > 9.1.2.3 Connection Setup After a Successful TCP Connect Request
.....
> If the Connection Nonce field contains a value identical to the
> most recently received Connection Nonce from the same IP Address,
> the FCIP Entity SHALL close the TCP Connection and notify the FC
> Entity with the reason for the closure.

Same comment.

Rejected

Current plans call for the MIB interface to be in the FC Entity. Therefore, this notification is necessary for MIB logging. Also, requirements are being added to FC-BB-2 that depend on the requirement as stated in FCIP.

Comment 130

- > 9.1.2.3 Connection Setup After a Successful TCP Connect Request
.....
> 1) Leave the Destination FC Fabric Entity World Wide Name field
> and Ch bit both 0;

So allow the FCIP Link to be established? It is unclear to me how implementations adopting this option would be able to prevent unintended FCIP Link formation.

Accepted

A requirement needs to be added somewhere that the TCP Connection MUST be closed if the echoed Destination FC Fabric Entity World Wide Name field contains zero.

[Comment 131 Technical](#)

```
> 9.1.2.3 Connection Setup After a Successful TCP Connect Request
.....
> 2) Change the Destination FC Fabric Entity World Wide Name field
> to match FC Fabric Entity World Wide Name associated with the
> FCIP Entity that received the TCP connect request and change
> the Ch bit to 1; or
```

In effect, this is being indirectly allowed as a legal discovery process. Is there a DoS concern here? Also, would revealing the FC WWN be acceptable in confidentiality terms?

[Duplicate of comment 53.](#)

[Comment 132](#)

```
> 9.1.2.3 Connection Setup After a Successful TCP Connect Request
.....
> 3) Close the TCP Connection without sending any response.
```

I like this option best, :-)

[Inquiry](#)

See comment 53.

[Comment 133](#)

```
> 10.2 FC Fabric and IP Network Deployment Models
.....
> Entities in an equal manner. This varies from a true Client-
S/b "varies" w/ "differs".
```

[Accepted](#)

3. Comments from Don Fraser

=====

Comment 134

9.1.3 (page 31)

"If an FCIP Entity receives a duplicate FCIP Short Frame during the FCIP Link formation process,..."

"Short Frame" s/b "Special Frame"

Accepted

Comment 135

Annex G (page 63)

The reference to section 9.5, should refer to 9.4.

Accepted

Comment 136

Annex G

The table should have a number and title like the rest of the tables in the document. Also, do not put just the table header on a page by itself.

Accepted

4. Comments from Murali Rajagopal

=====

Comment 137

-- Section 5 Protocol Summary

- 8) The FCIP Control & Services function MAY use TCP/IP quality of service features (see section 11.2) to support Fibre Channel capabilities.

[E] "function" s/b "Module".

Accepted

5. Comments from Bob Snively

=====

Comment 138

-- 6.6.2.3 Synchronization Failures

I would suggest the following minor editorial correction:

- b) return to sending valid FC Frames only after synchronization has been verified; and

should be changed to:

- b) return to emitting frames through the FC Transmitter Portal only after synchronization has been verified; and

Accepted with changes

Make the change as written, except replace "emitting" with "forwarding".

6. Comments from Ralph Weber

=====

Comment 139

Annex C (first step in algorithm)

"f) Reserved field and its ones complement."

The "reserved" field is now pFlags.

Accepted

Comment 140 Technical

The bit/byte numbering resolution approved for the FC Frame Encapsulation draft must be replicated in this draft.

Accepted

Comment 141 Technical

In order to support the FC-BB-2 Link Keep Alive (LKA) switch internal link service, it is necessary for FC/FCIP Entities to communicate a time interval for transmission of the LKA. The T11 FC-BB-2 working group has asked that this 32-bit quantity be added to the Special Frame.

Accepted

Comment 142

Update contributors and acknowledgements.

Update contact information for Murali Rajagopal.

Move Jim Nelson from contributors to acknowledgements since Jim no longer is FC-FS editor and has provided no updated contact information.

Accepted