

APPROVED MINUTES OF THE 20TH FAIS AD HOC MEETING

The twentieth meeting of the T11.5 FAIS Ad Hoc group took place at the Embassy Suites in La Jolla, CA, on April 6, 2004.

The following people were present:

AAROHI	Ranga Bakthavathsalan	ranga@aarohi.net
ASTUTE NETWORKS	Robert Johnson	robert@astutenetworks.com
BROCADE	Ed McClanahan	edmc@brocade.com
CISCO	Claudio DeSanti	cds@cisco.com
CISCO	Maurilio Cometto	maurilio@cisco.com
EMC	Anshul Chadda	chadda_anshul@emc.com
EMC	Greg Golubski	golubski_greg@emc.com
EMULEX	Ken Hirata	ken.hirata@emulex.com
HITACHI	Hidehisa Shitomi	hidehisa.shitomi@hds.com
IBM	Ken Hallam	khallam@us.ibm.com
INTEL	Ramamurthy Krithivas	ramamurthy.krithivas@intel.com
IVIVITY	Zulfiqar Qazilbash	zqazilbash@ivivity.com
IVIVITY	Eddy Quicksall	eddyquicksall@ivivity.com
MCDATA	Larry Hofer	larry.hofer@mcddata.com
NETWORK APPLIANCE	Tanjore Suresh	suresht@netapp.com
QLOGIC	Craig W. Carlson	craig.carlson@qlogic.com
TROIKA NETWORKS	William Chow	wchow@troikanetworks.com
UNH-IOL	David Woolf	djwoolf@iol.unh.edu
VERITAS	Roger Cummings	roger.cummings@veritas.com
VERITAS	Amitava Guha	amitava.guha@veritas.com

A total of 20 people from 16 organizations.

1. Opening remarks and introductions

The Chairman of FAIS, Claudio DeSanti of Cisco, called the meeting to order and welcomed the attenders, and led a round of introductions.

2. Administrivia

2.1 Approval of minutes

The minutes of the February FAIS Ad Hoc meeting had been uploaded to www.t11.org by William Chow of Troika Networks as 05-171v0. Claudio noted several changes. Pursuant to these changes, Claudio moved to approve the minutes; Ed McClanahan seconded. The minutes were approved unanimously.

The minutes for the conference calls (05-168v0 and 05-252v0), these were deferred.

2.2 Approval of this agenda

Claudio had distributed a draft agenda for this meeting as 05-221v0 via the T11 Web site.

The following changes were noted:

- Added presentation by Roger Cummings for 05-228v0.
- Added discussion of feedback on statistics.
- Added review of draft 05-115v1.

Claudio moved to approve the agenda; Ed seconded. The agenda was approved unanimously.

3. Call for Patents

Claudio stated that this Working Group operates under the INCITS procedures, which are available on the web at <http://www.incits.org/>. Amongst these procedures are the ANSI patent policies. He directed persons to forward any patent related issue to the T11.5 and T11 Plenary meetings.

4. Review of old FAIS action items

Old action items were reviewed. Their status was updated as follows:

AI 9: Ed McClanahan will propose an approach for supporting failover in the event of CPP failure. (Opened 1/13/2004)

Open. EM said this is blocked behind 04-402 completion as he needs to understand the Provider state that must be managed during failover. Target: ??? meeting.

AI 22: Ed McClanahan will propose a draft for chapter 1 and 2. (Opened 5/8/2004)

Partially completed (only chapter 2). Still open. Target: ??.

- AI 27: Ronak Desai will write up API requirements for Client takeover and handling of a faulted I/O. (Opened 7/13/2004)
Closed per 05-091v1.
- AI 35: Anshul Chadda will send object model (drawn in FrameMaker) to editor. (Opened 10/6/2004)
Closed.
- AI 38: Ed McClanahan will incorporate the Xmap changes for single read/write Vdev (from 04-665v1) into the draft specification. (Opened 10/6/2004)
Closed per 05-115v1. New action for Ed to incorporate new object model from AI 38 into draft.
- AI 40: Ed McClanahan will incorporate the new Multipath Vdev (from 04-665v1) into the object model. (Opened 11/16/2004)
Open.
- AI 44: Roger Cummings to propose a separation of Reason & Action Codes, and an enumeration of Action Codes. (Opened 1/13/2005)
Closed per 05-228v0.
- AI 46: Ed McClanahan will incorporate 05-006v1 into the next FAIS draft. (Opened 2/3/2005)
Closed per 05-115v1.
- AI 47: Anshul Chadda will include support for Task Management Function responses in the SCSI command data structure. (Opened 2/3/2005)
Closed per 05-008v1.
- AI 48: All providers will give feedback for statistics, per 05-113v0 and 05-116v0. (Opened 2/3/2005)
Open.
- AI 49: Ed McClanahan to incorporate 05-114v0 and 05-131v0 into the draft. (Opened 2/3/2005)
Open. Partially completed (131v0 already incorporated into 115v1)
- AI 50: Anshul Chadda will update the FAIS_SCB to support task management functions. (Opened 3/4/2005)
Closed per 05-008v1.
- AI 51: Ed McClanahan and William Chow will incorporate FAIS_ISOLATOR support into the Port API proposal (05-101). (Opened 3/4/2005)
Open.
- AI 52: William Chow will propose modifications to enhance processing of front-end commands. (Opened 3/4/2005)
Open.
- AI 53: Anshul Chadda to incorporate changes from this call into a separate document. The SCSI Command event moves out of the API document. (Opened 3/14/2005)
Closed per 05-091v1.
- AI 54: Maurilio Cometto to propose how Control event for XMAP would be supported. (Opened 3/14/2005)
Open.

5. Old business

There was no old business at this meeting.

6. Scheduled FAIS business

6.1 FAIS API Proposal (05-008v0)

Anshul Chadda (EMC)

Anshul noted that IN/OUT are no longer comments.

Anshul noted that handles are all now passed by value.

Anshul noted that there is now a requirement for the provider must drain/complete outstanding I/Os for FITL_Action and VDEV_XMAP_MASK.

Ed noted that fais_ObjFITL needs to be changed to fais_ObjPFITL.

VDEV_XMAP_MASK:

Ed noted that there needs to be a prioritization indicated, i.e. which one takes precedence in the event of multiple qualifications. Maurilio's action item is expected to address this. William asked whether this is really a mask; Anshul said no, and that he'll rename to be a permission. Ed asked what happens when an I/O overlaps against multiple fault reasons. William asked whether these might also be faulted concurrently. Anshul will add a note stating that there can be only one outstanding Xmap fault for an given lba range, Xmap, and I/O. The remaining issues apply to the events and were tabled until events are discussed.

6.2 FAIS Draft (05-115v1)

Ed McClanahan (Brocade)

Ed asked that future submissions for incorporation into the draft need to be completely ready for simple inclusion (e.g. no side notes or overlapping changes, etc).

Ed indicated that he will add region object.

There was discussion about objects that are not created by the client (e.g. no C_handle). Ed indicated that such objects can be differentiated as they will not have a Chandle attribute in their object.

Ed noted that the FIT needs to be fixed as it doesn't have handles and the associations with FT and FITL are incorrect. William noted that the FIT needs a P_handle as the a client might want to logout a FIT. Tanjore also noted that the client might want to operate on a FIT for statistics.

Ed asked how the Xmap's NumberOfEntries attribute is used, and whether it was useful. At the very least, it should only be set by the provider. William said this is useful for queries. Ed said he will remove it as no operations currently apply to it (e.g. no query for xmap entries).

Ed asked about the Multipath Vdev. He noted it wasn't clear whether the BITLset remained. William stated that when the Multipath Vdev was approved, it was agreed that the BITLset was kept.

Ed was assigned an action to update the object model diagram. Anshul was assigned an action to update the API proposal.

6.3 Clause 2 of API proposal (05-008v1)

Anshul Chadda (EMC)

Anshul said he hasn't incorporated support for managing state transitions for a control-path command.

William asked how resources are generally released. Larry noted that some resources are inherently variable-length (e.g. sense data). Ed noted that a "resolve" call is already required to serve as the event-completion function, which allows the provider to release any resources associated with the event. Will noted that this means we can now adopt the general rule that variable-length resources can use pointers. Ed said that we should maintain the rule that if passing by value is possible, it should be used.

Anshul asked whether the MPA values should or should not be defined. William asked whether a provider can support a subset of the possible functionality. Claudio says the API may define functionality that a provider may not support, i.e. a provider is allowed to return "not supported" (e.g. a particular MPA policy is not implemented). Ed asked for a vendor-specific range so that vendors can use their own encoding/interpretations. Anshul said he will add this.

Ed noted that it is necessary to add a "buffered" flag into the Vdev-copy. William noted something similar is useful in the BITLset. Anshul said he will add both.

Claudio suggested that the Port API proposal needs to be integrated asap. Anshul was assigned an action to publish a revision with today's changes, and then a subsequent revision with the Port API. It was requested that the scope of the Port API integration shall be limited to just the functionality currently defined. Any areas that don't immediately overlap shall be deferred to a future revision.

William asked that we have separate events for control versus faults, where former is owned by Client and latter is still owned by Provider. He said this was desirable to make it explicit in the API between control path I/Os and data path I/Os, since each have different operational semantics. An action was noted for William and Anshul to support this.

7. Unscheduled FAIS business

Claudio proposed having letter ballot on a draft that incorporates chapter 1 of the API proposal. He said this would not include the events, and reducing the scope of the work for FAIS 1 will allow this group to finally produce a specification. William noted that this will effectively defer any currently outstanding work, such as the events. Ed suggested that it could be useful for certain types of development (e.g. management s/w). Claudio suggested that we can revisit this discussion in June. A straw poll was taken: yes=11, no=2, abstain=1

8. Review of FAIS action items

New action items were reviewed. The list was updated as follows:

- AI 9: Ed McClanahan will propose an approach for supporting failover in the event of CPP failure. (Opened 1/13/2004)
Defer to phase 2. EM said this is blocked behind 04-402 completion as he needs to understand the Provider state that must be managed during failover. Target: ??? meeting.
- AI 22: Ed McClanahan will propose a draft for chapter 1 and 2. (Opened 5/8/2004)
Partially completed (only chapter 2). Still open. Target: late april concall.
- AI 40: Ed McClanahan will incorporate the new Multipath Vdev (from 04-665v1) into the object model. (Opened 11/16/2004)
Open. Target: 1st draft.
- AI 48: All providers will give feedback for statistics, per 05-113v0 and 05-116v0. (Opened 2/3/2005)
Open, possibly defer to phase 2.
- AI 49: Ed McClanahan to incorporate 05-114v0 and 05-131v0 into the draft. (Opened 2/3/2005)
Open. Partially completed (131v0 already incorporated into 115v1). Target: 1st draft.
- AI 51: Ed McClanahan and William Chow and Anshul Chadda will incorporate Port API proposal (05-101) into API proposal. (Opened 3/4/2005)
Open.
- AI 52: William Chow and Anshul Chadda will propose modifications to enhance processing of front-end commands.(Opened 3/4/2005)
Defer to phase 2.
- AI 54: Maurillio Cometo to propose how Control event for XMAP would be supported. (Opened 3/14/2005)
Defer to phase 2.
- AI 55: Ed McClanahan will incorporate the latest object model diagram into the draft standard.
New.
- AI 56: Anshul Chadda will publish a revision of the API proposal with just changes discussed per this (4/6/2005) meeting (i.e. w/o Port API proposal).
New. Target: 4/18.
- AI 57: Anshul Chadda and Ed McClanahan and William Chow will incorporate the Port API proposal into a subsequent revision of the API proposal (i.e. based on that from AI 56).
New. Target: 4/25.
- AI 58: William Chow and Anshul Chadda will separate Client handling of control-path I/Os from faulted data-path I/Os.
New. Deferred to phase 2.

9. Next meeting schedule

Claudio will schedule 2 conference calls:

- Week of 4/25
- Week of 5/2.

He will post concall times next week (week of 4/11).

Claudio will ask for 5-6 hours at the june plenary.

10. Adjournment

A motion to adjourn the meeting was passed unanimously.