

DRAFT MINUTES OF THE 17TH FAIS AD HOC MEETING

The seventeenth meeting of the T11.5 FAIS Ad Hoc group took place at Cisco in San Jose, CA, on November 16, 2004.

The following people were present:

AAROHI	Parag Bhide	parag@aarohi.net
BROCADE	Bob Snively	rsnively@brocade.com
CISCO	Maurilio Cometto	maurilio@cisco.com
CISCO	Ronak Desai	rodesai@cisco.com
CISCO	Claudio DeSanti	cds@cisco.com
IVIVITY	Eddy Quicksall	eddyquicksall@ivivity.com
MCDATA	Gib Cardwell	gibc@mcddata.com
MCDATA	Larry Hofer	larry.hofer@mcddata.com
TROIKA NETWORKS	William Chow	wchow@troikanetworks.com
VERITAS	Roger Cummings	roger.cummings@veritas.com
VERITAS	Amitava Guha	amitava.guha@veritas.com
VERITAS	Prasanna Narayana	prasanna.narayana@veritas.com
VERITAS	Gopal Sharma	gopal.sharma@veritas.com

A total of 13 people from 7 organizations.

1. Opening remarks and introductions

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

2. Administrivia

2.1 Approval of minutes

The minutes of the October FAIS Ad Hoc meeting had been uploaded to www.t11.org by William Chow of Troika Networks as 04-739v0.

The minutes were approved in the absence of any objections.

2.2 Approval of this agenda

Claudio had distributed a draft agenda for this meeting as 04-745v0 via the T11 Web site, and he now asked for changes to that agenda.

The agenda was approved in the absence of any objections.

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 4: Amitava Guha to provide requirements for statistics gathering in FAIS and present proposals (Opened 9/10/2003)
Open: Roger Cummings stated that they are waiting for legal/SMI/SNIA issues to be resolved. Target: Feb meeting.
- AI 5: Ed McClanahan and William Chow will work on a draft of Chapter 5. (Opened 11/9/2003)
Open. Target: Dec meeting.
- AI 7: Tanjore Suresh will cross-reference names/types of capabilities against existing standards/prior work. (Opened 1/13/2004)
Open. Target: the next meeting after completion of action item #8.
- AI 8: John Kemeny/Brian Geisel will each provide a list of capabilities that they must be able to query support for. (Opened 1/13/2004)
Open. Roger said Veritas is not going to provide guidance. Target: Dec meeting.
- AI 9: Ed McClanahan will propose an approach for supporting failover in the event of CPP failure. (Opened 1/13/2004)
Open. Target: Dec meeting.
- AI 12: Elizabeth Rodriguez will propose support for journaling. (Opened 4/6/2004)
Open. Roger asked if any wants to take this action. Claudio stated we can make a decision to keep this in the next meeting.
- AI 20: Tanjore Suresh and Eddy Quicksall will propose chapter 6 text for iSCSI-specific port attributes. (Opened 5/8/2004)
Open. Target: Dec 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: Dec meeting.

- AI 23: Larry Hofer will propose a draft for chapter 3. (Opened 5/8/2004)
Closed per 04-393v1.
- AI 27: Ronak Desai will write up API requirements for Client takeover and handling of a faulted I/O. (Opened 7/13/2004)
Open. Target: Dec meeting.
- AI 29: Roger Cummings will propose object model modifications to support pass-thru of non-read/write I/Os. (Opened 7/13/2004)
Open. Target: Dec meeting.
- AI 31: Claudio DeSanti will notify the membership (via the reflector) of the remaining schedule of the working group upon completion of the initial draft of the spec. (Opened 9/14/2004)
Open. Claudio said a Feb target for last technical input seems unrealistic; need to review this AI upon completion of this meeting.
- AI 33: Eddy Quicksall will propose an object to represent a port in the object model. (Opened 10/6/2004)
Open. Eddy said he would like to review this AI upon completion of his presentation in this meeting.
- AI 34: Roger Cummings will provide a list of SCSI documents to cite in the draft specification. (Opened 10/6/2004)
Roger Cummings said emailed Ed McClanahan, but wants to keep it upon since he needs to update the list. Open.
- AI 35: John Kemeny will make the source for the object model available. (Opened 10/6/2004)
Open.
- AI 36: William Chow will post a revised version of the proposal to t11.org (as 04-665v1) for Xmap changes for single read/write Vdev, as agreed to at the meeting. (Opened 10/6/2004)
Closed per 04-665v1.
- AI 37: John Kemeny will incorporate the Xmap changes for single read/write Vdev (from 04-665v1) into his API proposal. (Opened 10/6/2004)
Closed per 04-402v2.
- 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)
Open.

5. Old business

There was no old business at this meeting.

6. Scheduled FAIS business

6.1 Xmap/Multipath Object Model Changes (04-665v1) William Chow

William described a scenario where a read coming into a multipathing object goes out on a selected path. He said that he wanted separation of read and write without needing an extent-

based resource. Claudio noted that the BITLSet is not hierarchical in the current object model, and William agreed that his proposal would involve renaming the object. He saw the relationship between two VDevs as being provider-specific, and said that a particular provider could enforce a BITL set limitation. Claudio felt that the issue of XMAP being too heavy was similar to the separate read/write case raised last time that had been solved with a hint, and he asked if consideration of such a thing could be deferred until after a later presentation.

Ronak asked if a provider could reject the adding of an entry to an Xmap, and William replied that this was why a hint was necessary. Parag asked how different path weights could be represented in an XMAP, and was told that it was not possible and that a new multipath VDev object would be necessary to do that.

Roger asked a general question about how we make the types of VDev extensible, both in terms of standard and vendor-unique definitions, but there was general agreement that we were not ready to address this in FAIS Rev 1. Claudio emphasized that what was being proposed was not just a textual change, but a change to the relationships represented in the object model, and Larry viewed the new VDev as acting as a multipath Vdev or a BITLSet, but not both. William said that if his Option A was selected it would be desirable for a Vdev to reference a BITL without going thru BITLSet. Parag expressed a preference for one object for both uses, and Ronak saw this as a big change in that it would involve a BITLSet having Vdev as a child where today only BITL is a child. William thought that was going too far, however and suggested we stick to Option A. Larry asked if a multipath object would be mandatory, and William replied that Silvano had said last time that we had agreed that all objects are mandatory for the provider.

Claudio stated that he would prefer a separate Vdev. Maurilio stated that this could be supported by a Mirrored Vdev, since it supports separate policy weights for reads and writes. William said this would be confusing unless the Mirrored Vdev is renamed. Claudio stated that there seems to be consensus for a new Vdev. William moves that 04-665v1 (slides 5 and 6) be incorporated into the object model. Roger seconded. No objections or abstentions; the motion carried.

6.2 Xmap Optimization (04-708v2)

Parag Bhide

William noted that the proposal doesn't state that the size needs to be passed, with the hint, in the Xmap creation function. Parag affirmed this. William said this could be added to the XMAP_VDEV_TYPE_PARAM (1.8.3.8).

Claudio stated that there seems to be consensus for this proposal. He suggested that Parag needed to follow-up with specific changes to the API proposal.

6.3 iSCSI As It Applies To FAIS (04-744v0)

Eddy Quicksall

Claudio commented that FAIS should not be transport-specific, and suggested that we stick to terms defined in SAM-3. Bob thought that this might be difficult in the area of discovery and correct buffer structures.

Claudio also asked if we needed a parameter such as a transport-specific name, and noted that a FAIS_Client will have to generate some WWNs as part of defining the front end nexus. Roger noted that a VPD Page 83 name will also have to be generated.

6.4 FAIS Chapter 3 (04-393v1)

Larry Hofer

Larry Hofer then reviewed the Chapter 3 definitions contained in 04-393v1. A number of changes were agreed, and Larry accepted an Action Item to post another revision by early the following week.

6.5 FAIS API Proposal (04-402v2)

Larry Hofer

Larry then presented the latest FAIS API document (04-402v2). He reported that this was the result of a small group getting together and work on the document. Clause 1 is believed to be in line with the object model, but he felt that Clause 2 was not ready for detailed review. A number of different editorial approaches for this document were discussed, and Claudio noted that we also need to discuss and define event handling and fault handling for inclusion in Chapter 2.

1.3

The need for FAIS Provider Info was questioned, and the necessity for both a FAIS_Client name or identifier was questioned - Eddy felt it would be better if the client obtained a handle from the provider that would be given back in subsequent calls. Inconsistencies in the size of the ID were also noted. An agreement of a Client_ID of 32 bits in length was reached. A change from an IEEE-derived identifier to a T10 Vendor ID was also agreed.

Roger asked why app name must be unique. Ronak stated that it would be useful for a Client to query the Provider for this; William suggested an example where there are multiple Clients. William asked how this proposal is coordinated with Chapter 5 (i.e. bootstrapping function addresses). Claudio said that this is an open issue and this proposal does not need to address it yet.

1.6.2

Claudio suggested that the login stuff isn't firm yet and should go into section 2 of the proposal. William asked whether there needs to be an "event services" section; does this suggest that there are common types/functions for all "async events"? He said this forces the Client to demux if there are different Clients handling different types of events.

1.6.2.3

Parag said "hold" and "control" are not in the object model chapter. Larry will add editor's note for this open issue.

1.6.2.4

Claudio asked why the port ID is 32 bits. William noted that this section doesn't address previously raised issues concerning mechanisms to manage transport-level identities. Claudio suggested William and Eddy take an action to propose object model and API changes to support virtual fais scsi ports.

1.6.4.1

William asked why the object is identified only by a client-specified handle. He suggested that we must have a provider-specified handle. Parag said the client-specified handle is useful for object failover b/t clients. Ronak said that the client-specified handle might also be unique and persistent. There was consensus for 2 handles, one for client and one for provider.

Much work was noted that still needed to be done to this document to bring it into line with agreed FAIS terminology, and all details of the object model. Larry agreed to post a new revision with agreed changes and some of the global edits by the end of the following week, and it was agreed that we would ask for comments to be made against that document by December 8.

7. **Unscheduled FAIS business**

Claudio said that it is clear that a lot of work remains to complete the API proposal. He asked whether we needed a conference call. William suggested that folks need to post feedback before we have another conversation. Larry said he could post an update by end of week (11/19). Claudio suggested that we could set a deadline for feedback by 12/8. He also suggested a conference call to resolve the issue on IDs in the API; he proposed either monday 12/22 or possibly friday 12/19, depending on the availability of EMC. Claudio is actioned to post a email request for feedback, after the call.

8. **Review of FAIS action items**

New action items were not reviewed. The list was updated as follows:

- AI 4: Amitava Guha to provide requirements for statistics gathering in FAIS and present proposals (Opened 9/10/2003)
Open: Roger Cummings stated that they are waiting for legal/SMI/SNIA issues to be resolved. Target: Feb meeting.
- AI 5: Ed McClanahan and William Chow will work on a draft of Chapter 5. (Opened 11/9/2003)
Open. Target: Dec meeting.
- AI 7: Tanjore Suresh will cross-reference names/types of capabilities against existing standards/prior work. (Opened 1/13/2004)
Open. Target: the next meeting after completion of action item #8.
- AI 8: John Kemeny/Brian Geisel will each provide a list of capabilities that they must be able to query support for. (Opened 1/13/2004)
Open. Roger said Veritas is not going to provide guidance. Target: Dec meeting.
- AI 9: Ed McClanahan will propose an approach for supporting failover in the event of CPP failure. (Opened 1/13/2004)
Open. Target: Dec meeting.
- AI 12: Elizabeth Rodriguez will propose support for journaling. (Opened 4/6/2004)
Open. Roger asked if any wants to take this action. Claudio stated we can make a decision to keep this in the next meeting.

- AI 20: Tanjore Suresh and Eddy Quicksall will propose chapter 6 text for iSCSI-specific port attributes. (Opened 5/8/2004)
Open. Target: Dec 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: Dec meeting.
- AI 23: Larry Hofer will propose a draft for chapter 3. (Opened 5/8/2004)
Closed per 04-393v1.
- AI 27: Ronak Desai will write up API requirements for Client takeover and handling of a faulted I/O. (Opened 7/13/2004)
Open. Target: Dec meeting.
- AI 29: Roger Cummings will propose object model modifications to support pass-thru of non-read/write I/Os. (Opened 7/13/2004)
Open. Target: Dec meeting.
- AI 31: Claudio DeSanti will notify the membership (via the reflector) of the remaining schedule of the working group upon completion of the initial draft of the spec. (Opened 9/14/2004)
Open.
- AI 33: Eddy Quicksall and William Chow will propose an object to represent a port in the object model. (Opened 10/6/2004)
Added William Chow. Open.
- AI 34: Roger Cummings will provide a list of SCSI documents to cite in the draft specification. (Opened 10/6/2004)
Open.
- AI 35: John Kemeny will make the source for the object model available. (Opened 10/6/2004)
Open.
- 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)
Open.
- AI 39: Larry Hofer will update the definitions in the Chapter 3 draft (04-393).
New.

9. Next meeting schedule

The next meeting of the FAIS group will be hosted by Broadcom in Indian Wells, CA, on 12/18.

10. Adjournment

A motion to adjourn the meeting was passed unanimously.