

ICSC ITWG minutes for 11/16/04

(Taking minutes: HP - JR)

Y HP Jim Hamrick (JH)
Y HP Jay Rosser (JR)
N HP Fred Worley (FW)
Y IBM Fredy Neeser (FN)
N NetApp Arkady Kanevksy (AK)
N Sun Matt Pearson (MP)

cascading ascii art attendance diagram. (if you have more than 1 minus visible, you are not eligible to vote.)

	hp	ibm	netapp	sun
	----	-----	-----	----
m-3	+	+	-	-
m-2	+	+	-	-
m-1	+	+	-	-
m-0	+	+	-	-

o Agenda bashing, approve minutes

o Minutes to approve:

o Minutes for 26. Oct. 2004, email "ICSC ITWG Draft minutes 10/26/04" sent by Fred Worley on 11. Nov.

APPROVED

o Minutes for 2. Nov. 2004, email "ICSC ITWG Draft minutes, 2. Nov. 2004" sent by Fredy Neeser on 3. Nov.

[approve in next meeting:]

APPROVED

o Action item review

The first three AIs related to the IRD/ORD Header (IOH) and DAT compatibility:

o AI (jr): Change reject code name to IT_IWARP_VERSION_MISMATCH

o AI (jr): Figure out how many bits will be needed for the RNIC type field and size the field appropriately.

o AI (jr): Specify that IRD/ORD from EP attributes is only used if IOH Header is used.

Responding to the three AIs above, JR sent out a revised IOH in email (subject "Revised: AI - Modifications to IOH and disable IOH", sent 11/10/04 6:14PM PT):

o JR notes that he has added a reject code for bad IRD and asks for feedback - JH/FN think it is useful.

o FN notes that there is a proposal for IRD/ORD informational RFC in the RDDP

§ FN - should we wait for the IETF to finish it?

- § JR - notes concern about time required for such an RFC to make it to last call
- § FN - can we introduce our proposal in IT-API and adopt the IETF later?
 - § JR/JH - requirement is that IETF use a header version number field in the exact same placement in the private data as the IT-API header
- § JH/JR - concern about how we define interoperability between the two versions of IOH header.
 - § Recognizing the version number is sufficient if the implementations support both versions - it is a problem is we have to ensure interoperability with a version that only supports version 0.
 - § JH - concerned that this increases the size of interoperability matrix<>
- § JH - much further discussion of the possible ways to achieve interoperability here ensues and is not captured.

-
- o AI (fn): Get update from RNIC-PI WG / John Carrier regarding interoperability and version number problem.
(Discuss this under iWARP CM issues)

Discussion of the above not captured.

-
- o AI (jr): Add text to diagram to say that TCP ACK could be piggybacked on the first FPDU.

Pending

-
- o AI (jr): Add discussion of how IRD/ORD negotiation is done into app usage section of it_ep_connect man page, and reference to it from app usage section of it_ep_accept man page.

Pending

-
- o AI (fn): Ask RNIC-PI WG how to unbind a narrow RMR after a QP has been disconnected.

FN - describes findings:

- FN - Seems to be the case that no WR may be processed if QP not in RTS state.
 - o FN - Describes CB's response on this issue - Bind behind Send on QP will be stuck.

- o FN - Implication is having lost a connection, cannot unbind a MW - must deallocate MW.
 - § Similarly for Fast Register, must deallocate the memory region.
- o JH notes that Consumer would be able, in theory, to connect a busted QP back up again in order to free MW/MR (where latter is Fast Register).
 - § JH notes that pragmatically, Consumers will simply free the memory resource.
- o JR notes that on InfiniBand 1.1, Consumer could use a separate EP (connected to another EP on the same IA) and do unbind on that EP (i.e. loopback).
- o **AI - FN must adjust current text must represent problem solution on iWARP.**

-
- o AI (fn): Find out if Bind, Fast Register, Unbind, and Invalidate are supported in anything other than the RTS state.

Closed - Covered by response on previous action item.

o iWARP CM issues

- o Discussion on RDMAP/DDP version number problem and solution via MPA Startup
- o Decide on how to proceed with IOH extensions / suppression of IOH and on following proposition

"The IT-API phase 2 implementation shall support suppression of the IOH on both the active and passive sides. The IOH shall support additional "type" and "downgrade capability" fields. There shall also be a defined IT_IWARP_VERSION_MISMATCH reject code for use in the IOH.",

taking into account the latest development on the RDDP WG reflector and John Carrier's feedback.

The above were discussed in action item review (above). No vote was taken.

- o Email thread started by Fredy Neeser on "it_listen_create connection qualifier"

FN - covers JH's response on this. Listen() on an ephemeral port - query would be used to figure out local port to convey to remote.

Nothing further in the agenda was discussed (out of time):

- o Discussion of draft man pages
- o MM issues
 - o Discussion based on MM Detailed Requirements v0.97a
 - o MM Detailed Requirements: Remote invalidation
 - o Continue discussion on handling of remotely detected errors
- o Next steps
 - o Focus on man page generation, occupy spare time in telecons with errata review, next round of detailed requirements to be prioritized.
- o Any other business

Meeting adjourns at the 2-hr mark.