

ICSC ITWG Meeting Minutes, Dec. 21, 2004

Taking minutes: HP (jr)

Meeting attendance diagram (if you have more than 1 minus within the last four ITWG meetings, you are not eligible to vote):

mtg	date	hp	ibm	netapp	sun
---	-----	--	---	-----	---
m-3		+	+	-	-
m-2		+	+	-	-
m-1		+	+	-	-
m-0	dec 21	+	+	-	-

Present:

jim hamrick (jh) [HP]
fredy neeser (fn) [IBM]
jay rosner (jr) [HP]

Minutes to approve: Minutes for 7. Dec. 2004, email "First draft of minutes from 12/07/04 ITWG meeting" sent on 8. Dec. by Jim Hamrick.
Approved with no objections.

o Action item review

o AI (FN): Does IT-API Consumer have a need to control MPA markers, given the RNIC Interop Draft ?

See email by FN, sent on 12/16/2004, subject

"AI: Does IT-API Consumer have a need to control MPA markers, given the RNIC Interop Draft ?"

JR - agrees with analysis and only concern is industry support for the interop draft.

FN - optimistic that draft will be supported.

FN - also goes through the 4 possible motivations for exposing marker control.

FN - first reason - interop (covered by draft)

FN - second reason - suppressing MPA startup (however implies enabling markers since it is implicit that you are talking to an RDMAC device)

FN - third reason - unusual ULP problem. FN not convinced of this and thinks at worst we could defer supporting this beyond version 2.0

FN - fourth issue - performance due to RNIC preference that needs be over-ridden by an application. FN - not convinced by this.

JR - remarks that closing on Marker control needed before proceeding with man pages.

FN - summary - okay to live without markers.

JR - okay as well

JH - okay with it.

FN - okay - explicit marker support to be dropped from API.

New AI - JR - drop marker support from CM man pages.

Existing AI Closed.

o AI (FN): Change MM requirements MM-15.1.D1.2.2.1 for `it_post_send_and_unlink` to simplify remote indication of unlink by adding only an `it_handle_t` to `it_dto_cmpl_event_t`.

FN - **closed**. Will appear in 0.99 version of MM req'ts.

o AI (FN): Get rid of MM-16.4.D2.1 and MM-16.4.D2.2

FN - **pending**.

o AI (FN): Document that an unlink operation that fails remotely will not directly manifest as an error locally. Also add text to Implementor's guide that they need to absorb the Async Error and not manifest it via the IT API.

FN - **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.

FN - Stresses that the IT API 1.0 specification is insufficiently clear on this issue and really does need the text modified.

FN - notes that his email on the subject includes a description of the typical algorithm that a Consumer should use to achieve compatible IRD/ORD. There may be other ways and this algorithm can just appear as an example.

FN - an additional issue with IRD/ORD handling: For 2-way connection establishment, the active side does not find out the passive side IRD/ORD values.

JH - correct.

FN - concerned that active side needs to know what the passive side has chosen

JH - not clear that the active side needs to know the passive side's choice

FN - what if active side wishes to be able to reduce ORD where it has offered a very large ORD that the passive cannot support?

JH - notes that this is legitimate reason but concerned about supporting it.

JR - we support changing ORD at any time in IT API 2.0 (for iWARP iSER support) and thus, if we expose IRD/ORD in the connection event on the active side, there is a mechanism for the Consumer to make use of ORD.

FN - do we want IRD/ORD to appear in the connection event on the passive side as well? They appeared originally in the connection request event on the passive. What do they mean if they show up in the connection established event? Should they be the same or zero or what?

JH - points out that the 3-way model is more compatible with the conveying of this information

JH - wonders whether we should consider supporting the 3-way model over iWARP?

JR - (points out that without the "RTR" state, the requirement for a first Send is actually a 3-way model anyhow).

JR - another consideration is tossing support for IRD/ORD entirely (admits this may be yet another pandora's box).

FN - thinks we could fix the asymmetry by simply returning the same IRD/ORD to the passive side in their connection event as in the connection request event.

JH - i.e. the events always contain the remote's IRD/ORD request.

JR - summarizes that we have a scenario where active side sends IRD/ORD, passive sees a connection request event containing the active side IRD/ORD, passive replies with an adjusted IRD/ORD which manifests in a connection established event on active side. Now passive side sees a connection established event with the active side original IRD/ORD even though the passive's response may have changed the active side setting.

JR - thinks that this can be documented in the man page that the active side connection event IRD/ORD is valid but the passive side connection event IRD/ORD is invalid

FN - wishes also to include an rationale in the man page as follows: the IRD/ORD seen in the passive side connection event is described as invalid due to the fact that the active side may have changed IRD/ORD in response to the IRD/ORD sent by the passive side accept call.

New AI - JR - change cm_msg page to reflect the IRD/ORD usage model as described above.

Existing AI Closed.

- o AI (JR): Generate a proposal for a mapping service to generate a PBL (with bus addresses) given a virtual address range and a virtual address space ID(virtual addresses)

- o Why a mapping/unmapping service?

- o See first discussion in Minutes for 23. Nov. 2004

FN - notes that a discussion was started in the RNIC-PI work group on this subject. Recommends folks look at the document uploaded by Bernard Metzler on the subject.

JR - **Pending.**

- o AI (JR): Produce requirements for RDDP WG informational draft to support IOH functionality.

JR - **Pending.**

- o AI (JR): Revise remaining man pages that erroneously state "logical OR" instead of "bitwise OR"

JR - **Closed.** Text will appear in next CM draft.

- o CM issues

- o Proposal to add support for disabling IOH to IT API

FN - There is still a problem for passive side to support connection requests with and without IOH.

JH - Mentions the problems with trying to use a field that "should" be present to identify whether IOH is valid or not.

JH - The implementation could mis-interpret Consumer's private data.

FN - okay with leaving the proposal as is and forcing the control of the use or not of IOH entirely into the Consumer's hands.

JH - Thinks there may be a middle ground - a passive side could disable IOH and the Consumer could make use of knowledge of a identifying field in IOH to determine it exists and toss it.

FN - Describes the need for an IOH-using active side to receive a valid IOH somehow generated by the IT-API passive side (i.e. the Consumer would need to construct the IOH).

<more to and fro on this>

JH/FN - conclude draft on disabling IOH as it stands is fine (no point in adding special identifier).

o Graceful closes during MPA Startup (RDMAC vs. Non-Permissive IETF RNIC). See email thread starting on 12/15/2004, subject "Graceful closes during MPA Startup (RDMAC vs. Non-Permissive IETF RNIC)"

FN - recalls JR in agreement on the thread.

AI - JR - implement the non-peer rejects as described in the email thread.

FN - one more comment - this can happen both with TDI and TII and needs to be represented in both places.

o Discussion of third draft of CM man pages

JR - some questions on FN's comments <not captured here>

AI - JR - follow up on further responses to FN via email.

o S-RQ issues

FN - in creating a S-RQ, you cannot specify the low water mark for S-RQ. Can only be specified later on by specifying modify_S-RQ. What is the motivation?

FN - notes that the verbs allow you to specify the low water mark on creation in iWARP.

JH - is it also the case for IB? (May be the rationale).

FN - On pg 563 of IB 1.2 - finds that low water mark not in creation verb so that explains rationale.

o MM issues

o Continue discussion based on MM Detailed Requirements v0.98

FN - proposes deferring the discussion on MM.

JH - asks about issue of improving readability of spec vs reducing number of change bars of the specification. Also the applicability section. There is a fundamental conflict between an existing consumer seeing a massively changed spec and a new consumer seeing a revised spec.

FN - thinks that a specification that is more clear is highly desirable.

JH - that implies we provide two specifications - one with change bars and another without.

FN - yes.

FN - describes motivation for an "applicability section" in each man page. Some man pages do not apply to a consumer's transport - for instance UD transport pages, and the new iWARP-only pages.

JH - Should the applicability section appear on every man page or just on those where there is a restriction? Thinks that the majority of man pages have no restrictions.

FN - thinks at least 20 man pages currently have restrictions. The new ones will apply to iWARP only.

JR - Don't the iWARP ones also apply to IB1.2?

FN - gives example of `it_post_rdma_read_to_rmr` - only applies to iWARP.

JH - such a call depends on an IA attribute.

FN - yes

JH - Thus applicability section would need to refer to such attributes to constrain them.

FN - Service type will also be another constraint.

JH - agreed.

FN - Transport is another

JH - Concerned about constraining to transport since that may impede creative implementors.

FN - what about the `it_socket_convert` man page? Isn't it specific to iWARP only?

FN - summarizes constraints are attributes, service types (e.g. UD), transports.

JH - thinking about `it_socket_convert` - today it is iWARP only but could foresee it being applicable to another transport in the future.

JH - thus thinks that `it_socket_convert` should be tied to an IA attribute.

FN - okay - this may simplify the specification. Then we would strictly use IA attribute or service type to simplify the man pages with the applicability section.

JH - adding an applicability section to every man page needs to be divided up amongst group.

JH - summarizes - applicability section is added to every man page identifying the service type where applicable and if an IA attribute is applicable, it will identify that as well.

JR - we need to make sure we have no conflict with OpenGroup style

AI - JR - research in OpenGroup style guide.

o Continue discussion on handling of remotely detected errors

FN - asks why the Error section is empty typically.

JH - thinks this was due to OpenGroup style - Error section was supposed to contain Errno values.

AI - JR - to send OpenGroup style guide to FN.

FN - has been adding completion errors into the Error section so that the errors associated with a particular call are easily identified.

<much discussion of this with no outcome defined>

o Any other business

o DAT Collaborative request for assignment of 256 IB Service IDs to DAT
o Basis and implications of this request?

JR - agrees with Fujitsu proposal to preserve the contiguity of the ICSC Service ID space.

FN - also agrees.

<some discussion of the construction of Service IDs and history of this issue>

o Next ITWG Meeting

JR - next meeting was scheduled for Jan 11th 2004

JH/FN - both have conflicts.

FN - proposes we meet next on the 18th and use the intervening time to clean up the man pages.

Meeting adjourns at the 2hr 25min mark.