

ICSC ITWG Meeting Minutes, 15 Feb 2005
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	feb 15	+	+	-	-

Present:

jim hamrick (jh) [HP] (after 40 mins)
fredy neeser (fn) [IBM]
jay rosner (jr) [HP]

Minutes to approve: Email from Fredy Neeser, subject "ICSC ITWG draft minutes for 1. Feb. 2005", sent 2/2/05 2:57AM PT - **Approved**.

Pending Action Items Summary:

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

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

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)

AI (JR): Lead further discussion to complete: Produce requirements for RDDP WG informational draft to support IOH functionality.

AI (JR/FN): Determine if keeping around socket descriptors for lifetime of converted connection is too resource-consuming

New Action Items Summary:

AI (JR) - Add text describing possible future detailed status code to it_dto_events in Future Enhancements section.

AI (JR) - Replace "type of net_addr" with "type of passed in source address" at line 11025 of it_api_2_0_021405_fn_rev.doc

AI (JR) - Revise text in Implementer's guide for it_ep_rc_create to make it clear that Modify QP for RDMA enables can be done in the it_ep_rc_create implementation or in the CM as the implementation sees fit - the former is preferable as it makes endpoint state consistent.

AI (JH) - Supply a definition for the Service Request Handle object

AI (JR) - look at resolving the bad formatting of the Figures section in the table of contents (at a low priority)

AI (FN) - Fix language of Sections 1.5 and 4.1

AI (JR) - Drop "IT-API" from "IT-API Connection Management" and fix any links.

AI (JR) - do some experimentation to fix how it_post_rdma_read_to_rmr appears in the "legal blackline" comparison document.

AI (FN) - Draft a sentence about local errors being described in `it_dto_status_t` as a cross reference in all DTO pages (should appear as a first sentence in ASYNCHRONOUS ERRORS section).

AI (JR) - Replace `ird_ord_support` flags in `it_ia_info_t`. Where `ird_ord_support` is `IT_TRUE`, add text stating that if `IRD_ORD` suppression is attempted on a transport that cannot support it (say `IB`), then an immediate error will be emitted (`IT_ERR_INVALID_FLAGS`). If the `ird_ord_support` is `IT_FALSE`, add text stating that `IRD_ORD_SUPPRESS` flag is ignored.

AI (JR) - Compile all the needed compability header file text and add to the BW Compat section (as well as to the header file).

AI (all) - Review Tentatively Done TODOs as per FN's email.

Minutes:

Action item review

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

Pending

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

Pending

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)

Pending

AI (JR): Lead further discussion to complete: Produce requirements for RDDP WG informational draft to support IOH functionality.

Pending

AI (JR): Add macro for `rdma_read_incoming` / `rdma_read_outgoing` field names to support `IRD` / `ORD` as alternate (additional) names; determine where this changes should be documented (includes `msg_events` page)

- Global edit

JR - Done

FN - Have we considered this as a backwards compatibility feature?

JR - Not yet (see later discussion on backwards compatibility too).

AI (JR/FN): Determine if keeping around socket descriptors for lifetime of converted connection is too resource-consuming

Pending

AI (FN): Evaluate "detailed status code" idea.

FN - How important is this? If we don't do it now will we never do it?

JR - Consider last question first - how would this be implemented and how significant a backward's compatibility issue would this be?

FN - Discusses using a changed completion Event data structure with another element for detailed code.

FN/JR - Agree that FN's approach is feasible and backwards compatible.

JR - Thinks that given this reasoning and the effort required to implement revised detailed status code at this point, we can defer this.

FN - Agrees and suggests we add a future enhancement note indicating the possibility to support a detailed status code in the future. Suggests doing so in `it_dto_events`.

JR - Agree.

AI (JR) - Add text describing possible future detailed status code to `it_dto_events` in Future Enhancements section.

Done

AI (ToDo list/JR) Reformat Reference page to use separate columns for reference tag and reference

JR - Added new References section with agreed upon format.

FN - Asks about the latest IETF draft references.

JR - Has revised IT-API spec to refer to the latest.

Done

AI (JR/FW, others?) Add this text to the man pages for all blocking calls

JR - JH had a comment about adding `IT_ERR_INTERRUPT` to `it_convert_netaddr` and latest draft adds that.

Done

AI (FN) will keep to-do list for edits that still need to be done; noted edits should be provided to FN for tracking

FN - list is up to date.

Done

AI (JR) Implement resolution for GN55=HP15 proposed by FN

Done

AI (JH) Implement resolution to GN23=SUN11

JR - fix to `it_convert_netaddr`

FN - notes that Functional Changes text uses the wrong argument name

AI (JR) - Replace "type of `net_addr`" with "type of passed in source address" at line 11025 of `it_api_2_0_021405_fn_rev.doc`

Done

AI (All) Review portion of implementers guide corresponding to owned man pages to see if transport specifics need to be called out

FN - this is largely done but we can discuss when we talk about TODO list.

Done

AI (ToDo list) Implement terminology changes described in mail from FN on 2/8/05 Subject: Re: Terminology

Done

AI (Add to ToDo list) Add text to Implementers Guide to specify that on IB one must modify the EP after creation (with `it_ep_rc_create`) to ensure that the endpoint attributes match those requested in endpoint creation

FN - questions whether the RDMA enables truly need to be done by CM or whether they could be done under the covers of the `it_ep_rc_create` call (i.e. by calling the `modify_QP` verb after the `create_QP` verb).

JR - notes that the above is possible

FN - this would keep the Endpoint attributes consistent with the Consumer's settings.

AI (JR) - Revise text in Implementer's guide for `it_ep_rc_create` to make it clear that Modify QP for RDMA enables can be done in the `it_ep_rc_create` implementation or in the CM as the implementation sees fit - the former is preferable as it makes endpoint state consistent.

Done

AI (ToDo list) Document the completion or async errors that can occur associated with RDMA enables being disabled

Done

AI (ToDo list) Document how to use the endpoint attributes in connection with RDMA Reads

Done

AI (FN) Review `it_rmr_link` to ensure existence of text discussing ramifications of stale STags.

FN - agrees with JR's comment that we indicate we can reduce the potential of problems not eliminate them

Done

AI (ToDo list) Remove proposed fix and reference to `INVALID_SGID/SLID` from latest draft of v2.0 document

JR - as agreed in telecon, fix will cause compile error for old code.

Done

AI (ToDo list) Update DTO flags

Done

Editorial issues:

Email thread started by Fredy Neeser, subject "List of TODOs for v2.0 specification", sent 2/8/05 3:33PM PT.

- Also "Re: `it_ia_query`" sent by FN 2/10/05 5:37AM

- Also "Latest TODOs for v2.0", sent by FN 2/15/05 9:25AM PT

FN - only the "Latest TODOs for v2.0" needs to be discussed

JR - agrees

FN/JR - discuss `it_ia_query` tangent

(JH joins)

<Discussion of the fact that `it_ia_info_t` is returned by implementation since it is a potentially changing size structure. This imposes the need for the implementation to allocate the data structure and track the allocations for freeing later.>

FN - returns to latest TODO list - we discuss below heading by heading through the TODO list attachment in his email.

FN - two major sections: Things that are not yet done and things that are tentatively done.

FN - clarifies that the class of "Tentatively Done" fixes are integrated in the latest draft emitted by FN today in private email.

FN - suggests we discuss what needs to be done first and discuss tentatively done if we still have time.

FN - Global Edits:

FN - Bad links are low priority

FN - Service Request Handle is a noun

JH - Service Request Handle is an encapsulation of a SIDR request from point of view of Consumer

AI (JH) - Supply a definition for the Service Request Handle object

FN - Varia:

FN - Bad formatting of Figures

AI (JR) - look at resolving the bad formatting of the Figures section in the table of contents (at a low priority)

FN - Introduction mostly revised

FN - Section 1.5 and Section 4.1 are slightly out of sync in language usage.

AI (FN) - Fix language of Sections 1.5 and 4.1

JR - meta issue - do we have time to fix this?

FN - proposes giving ourselves one more day

JR - by close of business Feb 16, 2005?

JH - asks if any more review should be done?

JR - recommends we do continue review prior to passing doc to OpenGroup. Wishes to catch as much as possible prior to handover.

FN - References needing updating

FN - already done by JR in latest draft

FN - Should we drop IT-API from "IT-API Connection Management"?

JR - Agrees

AI (JR) - Drop "IT-API" from "IT-API Connection Management" and fix any links.

FN - the document with change bars

AI (JR) - do some experimentation to fix how `it_post_rdma_read_to_rmr` appears in the "legal blackline" comparison document.

FN - "Man pages to be updated" section.

FN - API Calls - should we add text describing the locally detected errors to all DTO pages?

JR - What does this change require as an effort? Can we only have a cross reference to `it_dto_status_t`?

AI (FN) - Draft a sentence about local errors being described in `it_dto_status_t` as a cross reference in all DTO pages (should appear as a first sentence in ASYNCHRONOUS ERRORS section).

FN - `it_ep_connect` and `it_listen_create` - should we add an immediate error in IB if Consumer attempts to defeat IRD/ORD?

JH - notes that JR's latest text implies that it is not desired to support support disabling IRD/ORD on IB (and it may not even be possible).

JH - also notes that VIA cannot support IRD/ORD

FN/JR - suggest that we change the `it_ia_info_t` flags `ird_support` and `ord_support` into one new flag "ird_ord_support" (and add macros for backwards compatibility).

All - agree

FN - Now if `ird_ord_support` is false, then ignore the `IT_XXX_SUPPRESS_IRD_ORD` flag. If `ird_ord_support` is true, then omitting `IT_XXX_SUPPRESS_IRD_ORD` flag is okay. If transport cannot support suppression, and Consumer attempts to suppress IRD/ORD, should an immediate error be emitted (say on InfiniBand)?

All - agree

AI (JR) - Replace `ird_ord_support` flags in `it_ia_info_t`. Where `ird_ord_support` is `IT_TRUE`, add text stating that if `IRD_ORD` suppression is attempted on a transport that cannot support it (say IB), then an immediate error will be emitted (`IT_ERR_INVALID_FLAGS`). If the `ird_ord_support` is `IT_FALSE`, add text stating that `IRD_ORD_SUPPRESS` flag is ignored.

FN - "Data types"

FN - JR and FN agreed that two-level `dto_status` code can be postponed to V2.1 (see above in AI review section).

FN - this allows FN to defer discussion of IB1.2 MM errors and other DTO errors.

FN - "Header file"

FN - Proposes some changes needed.

JR - Concerned about losing the information added to the non-normative header file

JH - Suggests placing such information into the Backwards Compatibility document.

AI (JR) - Compile all the needed compatibility header file text and add to the BW Compat section (as well as to the header file).

FN - notes that there are classes of constants that are deprecated and there are others that are just gone.

FN - now on to the Tentatively Done TODOs:

FN - "Global Edits"

FN - covers various changes

JR - okay with me

FN - "Varia"

FN - Version number text in Preface - will this change be acceptable to OpenGroup?

JH - Let us see if they complain.

FN - recommends remainder of Tentatively Done TODOs be reviewed and commented by email if necessary

JR - okay.

AI (all) - Review Tentatively Done TODOs as per FN's email.

Email thread started by Jim Hamrick, subject "Propose we reject GN23=SUN11", sent 1/31/05 2:52PM PT

All - this is resolved.

Email thread started by Jim Hamrick, subject "Missing error code in `it_ep_free`", sent 2/9/05, 4:51PM PT

All - okay with revision.

Email thread started by Jim Hamrick, subject "AI to check for transport dependencies in the Implementer's Guide", sent 2/9/05 5:02PM PT

- And "More AI to check for transport dependencies in the Implementer's Guide", sent by JH 2/10/05 3:39PM PT
- And "AI to check for transport dependencies in the Implementer's Guide", sent by FW 2/10/05 7:25PM PT

FN - asks that we look closely at the changes for it_rmr_query text in TODOs (Tentatively Done)

FN - Also, need to replace "Query Memory Region" with "Query Memory Window" in new it_rmr_query text in Implementer's Guide

JR - (done)

Email thread started by Fredy Neeser, subject "it_post_rdma_read_to_rmr", sent 2/3/05 7:43AM PT

FN - resolved.

Email thread started by Jay Rosser, "Direct LMR Handle", sent 2/14/05 11:03AM PT

JR - Had removed all reference

FN - Okay.

Ending discussion:

JR - Will compile document by COB 2/16/05 and forward for final editing to OpenGroup.

FN - Cannot attend next two meetings but will be available sporadically by email.

Meeting adjourned 5-minutes over time.