

ICSC ITWG Meeting Minutes, 28 June 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	jun 28	+	+	-	-

Present:

jim hamrick (jh) [HP] (left after 1 hour)
fredy neeser (fn) [IBM]
fred worley (fw) [HP]
jay rosser (jr) [HP]

Minutes to approve:

- Email from Fred Worley, subject "ICSC ITWG DRAFT minutes for 6/7/05 and 6/14/05", sent 06/17/2005 - **Approved both emended versions sent 6/28/05 (revised as follows)**
 - AI (FW) - minutes of 6/7/05 require following corrections
 - pg 5 - replace "physical phase" with "physical memory"
 - **Missed a couple other fixes (they appear in emended minutes, though).**
 - AI (FW) - minutes of 6/14/05
 - pg 4 - remove "(using rmr_link)"
 - pg 5 - replace "If we don't have the address in lmr_link" with "If we don't have the address space in lmr_link".
 - pg 6 - replace "Error will occur if STag 0 requests comes in on non-priv SRQ" with "Error will occur if STag 0 DTO is dequeued onto a non-priv EP"
 - pg 8 - remove text on page 7/8 "-what if the consumer attempts to register a mismatched pair of LKEY of one page size and a PBL with a different page size-"
 - Add text below "Is this possible" stating that we believe this is not the case.

Action item review:

- AI (MM subgroup): Add text to implementer's guide clarifying proper behavior if VA on input and output differ
 - **Closed**
- AI (FN): Summarize it_lmr_create issue of different VAs on input and output and provide text to resolve; post to bugzilla
 - **Closed**
- AI (MM subgroup): Replace wording on line 4251 that currently reads "refers to the same physical memory pages as the new LMR" with "the identical set of physical memory pages as the new LMR" and add to the end of the sentence "and has the same length".
 - **Closed**

- AI (FN): Clarify definition of "same physical memory pages" on the `it_lmr_create()` man page; post to bugzilla
 - Closed
 - FN - asks if on `it_lmr_create` man page (line 4259-60 of IT-API 2.0), does ignoring "shared_id" make sense?
 - JR - agrees that we have confusion on use of shared flag
 - AI (FN) - summarize use of shared flag and start email discussion
- AI (JR): Determine the text in ITAPI 2.0 that needs to be clarified that defines relationship between # of outstanding work requests and the size of the work queue
 - JR - describes text in `it_ep_attributes` on line 8456-7 that states that there will be an error on overflowing DTO queues
 - FN - asks for text that defines the relationship between posting and reaping completions that shows the consumption of work queue space
 - JR - notes the possibility of implying a rigid (strict/conservative) interpretation of work queue size rather than exposing the "sloppy" semantic when adding such text to API.
 - FN - thinks it is worth describing the strict semantic and then word-smithing it to convey the sloppy semantic as agreed in our DTO-6.0 vote.
 - AI (JR) - emit such text
 - Pending
- AI (JR): Post summary of physical to bus address translation issue to reflector for discussion [See email by JR, subject "AI - recap physical to bus address translation discussion todate", sent on 06/15/2005]
 - Closed
 - Discussion pending on this
 - FN - translation of physical addresses to bus addresses is a heavy responsibility on Consumer (requiring DDK use)
 - AI (All) - reflect on this summary and collect info on various O/Ses to see how this is implemented. Discuss on reflector.
- AI (JR): Update proposal of a flag for address space selection in `it_lmr_create()`; also discuss applicability of such a flag for RMRs
 - Pending
 - FN - we discussed whether or not flag would be an attribute of an LMR and whether also an attribute of an RMR.
 - Use of the flag is only necessary when registering the address - thus it need not be an attribute of the LMR or RMR.
 - Once LMR is registered, it corresponds to an underlying PBL
 - Moreover, the LMR will be associated with the virtual address that the Consumer passed in
 - The responsibility of tracking the relationship is with the implementation.

- o FN - summary - thinks flag is useful for it_lmr_create but it is not needed as an attribute.
 - o JR - will ponder FN's thoughts and send a summary if in agreement.
- AI (MM subteam): Issue a revised set of detailed requirements with summary updates to prepare for votes
 - o FN - FN/JR are revising the MM requirements and are converging
 - it_ia_create detailed requirements being revised to maintain the IT-API 1.0 signature
 - Consumer can simply query the IA to find the actual device PBL mode
 - o FN - we will post an interim copy with the proviso that we are not quite ready to vote.
 - We will start to vote in two weeks on the fully resolved sections.
 - o Closed
 - o (JH leaves mtg)
- AI (MM subteam): Review MM-4.0.D3 to make sure the list of completion errors for it_lmr_link() (fast registration with PBL) is complete
 - o Closed
- AI (FN): Revisit the completion error where a privileged Receive WR on an S-RQ (i.e., a Receive using the Direct LMR Handle) ends up being consumed by a non-privileged Endpoint. (IB uses "MM Operation Error", while iWARP may use the "Invalid STag" or "QP not in Privileged Mode" completion error)
 - o FN - iWARP has a more specific error than IB. Hence FN recommends adding an additional IT-API status code.
 - o FN - sent ballot in email (subject "Completion error for privileged Receive WR consumed by non-privileged EP", sent 6/28/05).
 - JR - no objection
 - FN - **vote passes**, proposal accepted.
- AI (All, from June 7 meeting): Consider moving S-RQ handle from RC-only attributes to generic attributes
 - o FN - recommends deferring till JH available.

List of Work Requests

- Review status after voting on WR coalescing in RNIC-PI WG
 - o FN - recommends deferring till JH available.
 - o FN - notes that RNIC-PI has accepted single work request calls and coalesce flag so they have an efficient means to implement IT-API
 - JR - RNIC-PI did, however, retain also its list of work request interface
 - FN - Yes - two classes of function vectors - one with list, one without.
 - o FN - impact on IT-API is immediate error processing occurs at post time - very positive for IT-API.

Recent email threads

- Email from Fredy Neeser, subject "Atomics and access control settings for MRs and MWs", sent 6/13/05 6:21AM PT
 - JR - this corresponds to MM-24.0, right?
 - FN - this should cause a change to the `it_post_atomics` man page to define behavior if these flags are not set
 - **AI (JR)** - check to see if this is currently handled in the proposed atomics page and revise if not.
- Subject "Cleaning up Receive WRs when using an S-RQ with iWARP", started by Jim Hamrick on 06/23/2005, and impact on `it_ep_reset()`.
 - JR - thinks that JH finally recommended RNIC-PI adopt the "Last WQE flushed" event from IB into the RNIC-PI.
 - FN - recommends deferring till JH available.

Any other business

- No meeting next week (week of July 4th)
 - **AI (JR)** - post notice of cancelled mtg
- JR - FN/JR had private email discussion of immediate errors being emitted on use of `direct_lmr_handle` when not allowed.
 - JR - notes that he withdraws concern with costs of immediate checking given overhead of marshalling
 - FN - recommends if `direct_lmr_handle` not supported, then `direct_lmr_handle` MUST be an invalid LMR handle. This will allow IT-API implementation to discover use of the `direct_lmr_handle` on IB 1.1 without incurring an additional check on IB 1.2 or iWARP that is unnecessary.
 - FN - recommends using "IT_NULL_HANDLE" for this case.
 - JR - this seems fine.

Ending discussion:

Meeting adjourned at the 1hr 33 minute mark