

ICSC ITWG Meeting Minutes, 31. May 2005
Taking minutes: IBM (FN)

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	10/05	+	+	-	-
m-2	17/05	+	+	-	-
m-1	24/05	+	+	-	-
m-0	31/05	+	+	-	-

Present:

Jim Hamrick (JH) [HP]
Fredy Neeser (FN) [IBM]
Jay Rosser (JR) [HP]
Fred Worley (FW) [HP]

Action item summary after today's meeting

Pending AIs:

AI (FN): Report back on the semantics of IOVA in InfiniBand's Register Physical Memory Region verb.

AI (JH): Add text to the usage section of all posting calls to clarify appropriate usage to avoid WQ and CQ overflow.

AI (JR): Determine how DTO 6 high level requirement vote was resolved.

AI (JH): Provide additional description of what is meant in list of work requests proposal by # of work requests processed; include usage example.

AI (FN): Propose resolution for ambiguity of shared state of LMRs for LMRs that are a) created in a shared state and b) transition into a shared state.

New AIs:

AI (JR): Review status of MM Detailed Requirements and propose sections ready for voting.

AI (JR): Forward FN's email, subject "Semantics of IOVA in IB's Register Physical Memory Region verb", to HP's IBTA folks for comment.

AI (JR): Generate an Erratum for changing the names of it_lmr_sync_rdma_read and it_lmr_sync_rdma_write to something more

intuitive. Names proposed are `it_lmr_flush_to_mem` (performs writeback of dirty cache lines) and `it_lmr_refresh_from_mem` (performs invalidation of cache lines).

AI (All): See if there is any value in S-RQ support for UD (see also Erratum GN216 = IBM203).

Begin: Meeting Minutes

Agenda bashing

Additional agenda item: Bugzilla

Minutes to approve

Email from Fred Worley, subject "ICSC ITWG minutes for 5/17/05", sent 5/17/05 11:20AM PT. Minutes approved.

Action item review

AI (FN): Report back on the semantics of IOVA in InfiniBand's Register Physical Memory Region verb.

FN: See email "Semantics of IOVA in IB's Register Physical Memory Region verb". Why does InfiniBand's Register Physical Memory Region verb have a parameter "IOVA" (input and output), while iWARP's Register Non-Shared Memory Region verb has a parameter "Virtual Address" (input only). Is there a semantic difference between these verbs in terms of who performs the mapping from virtual to bus addresses?

FN: So far, I have not received any response from IBTA members.

AI left pending.

JR: AI for me to forward FN's email to HP's IBTA folks for comment.

AI (All) - consider whether we wish to rename `it_lmr_sync_rdma_read` and `it_lmr_sync_rdma_write` to more intuitive names and then deprecate the existing calls.

FN: Suggests `it_lmr_flush_to_mem` and `it_lmr_refresh_from_mem` to better indicate direction of synchronization.

FW: Extra hint in current name referring to RDMA ops is not sufficient justification for keeping the names confusing.

AI closed.

New AI assigned to JH to generate an Erratum for changing the names.

AI (JH): Add text to the usage section of all posting calls to clarify appropriate usage to avoid WQ and CQ overflow
AI pending.

AI (JR): Determine how DTO 6 high level requirement vote was resolved
AI pending. FN will also look at this.

AI (All): Consider if the IT-API documentation should elaborate on async errors pertaining to the data source in all work request man pages.

FN: It's a completeness argument, not a correctness argument.

All: Postpone to Phase 3.

AI (JH): Provide additional description of what is meant in list of work requests proposal by # of work requests processed; include usage example

AI pending.

AI (JH): Research waiting on an fd attached to an EVD in current ITAPI. See email from Jim Hamrick, subject "AI: Research waiting on an fd attached to an EVD in current", sent 5/26/05 6:15PM PT

JH: IT-API does not allow enabling multiple notification mechanisms for an EVD at the same time. Waiting on an EVD with an associated fd is an error.

Use of file descriptors with EVDs: SEVD with associated fd does not allow associated AEVD. AEVD with associated fd is legal.

JR: Recalls that the Consumer should not wait on an EVD that is associated with and enabled for Notification to higher-level objects.

AI closed.

AI (JH): Propose requirements such that when a callback routine is attached to an EVD, that is the only notification pathway that functions

AI closed.

AI (FN): Propose resolution for ambiguity of shared state of LMRs for LMRs that are a) created in a shared state and b) transition into a shared state

FN: Explains proposed resolution.

JR: Concern about shadow MRs and increase in resource usage.

FN: Create two shared LMRs that match. Results in one primary MR and two derived MRs.

JR: Pathological case. Single shared LMR. Results in one primary MR and one derived MR.

JR: Concerned that resource penalty could be heavy for ill-behaved applications.

FN: Overhead goes down asymptotically with number of shared LMRs.

AI pending.

Erratum GN216 = IBM203 - Missing S-RQ support for UD service.

FN: Proposal to move `srq_handle` to set of EP attributes common to RC and UD.

JH: What is your use model? - A single UD Endpoint can serve multiple remote UD Endpoints, so it's not clear what is gained by introducing S-RQ support for UD.

New AI (All): See if there is any value in S-RQ support for UD (see also Erratum GN216 = IBM203).

Erratum GN217 = IBM204 - Shared state of an LMR cannot be reliably inferred

See AI above regarding ambiguity of shared state of LMRs.

Callbacks

Email from Jim Hamrick, subject "Second draft of detailed requirements for callbacks", sent 5/26/05 5:58PM PT

AI (JH): Call for a vote on second draft for June 14.

General IT-API 2.1 discussion

New AI for JR to review status of MM Detailed Requirements and propose sections ready for voting.

Any other business

FN: Martin Kirk enabled Bugzilla tool for ICSC bug tracking. Seems convenient to use. Would this be an alternative to our spreadsheet?

JH: Open to take a look at the tool.

Call ends at the 70 minute mark.

End: Meeting Minutes