OPENFABRICS ALLIANCE

iWARP Update

Authors:

Felix Marti (felix@chelsio.com)

Terry Hulett (terry.v.hulett@intel.com)

Date: April 5th, 2011

Past Presentations, 2008-2010

- Memory Extension Verbs
  - Fast Register Non-Shared Memory Region
  - RDMA Read Invalidate Local STag
  - Invalidate Local STag
  - …
- Provide Applications and Upper Layer Protocols with various Memory Registration choices (with different tradeoffs)

Past Presentations (2)

- Peer 2 Peer Support
  - Client/Server Connection Setup
  - Implemented by OFED RNICs
  - Internet-Draft submitted to IETF STORM WG: ‘Enhanced RDMA Connection Establishment’
  - Draft co-authored by
    - VMware
    - Intel
    - Open Grid Computing

Past Presentations (3)

- ‘iWARP 2.0 – potential features list’
  - Close the gap between *different* RDMA transport Verbs *and* semantics – yes, OFED is more than IB
  - Goals
    - Remove or at least reduce transport aware code in Applications/Upper Layer Protocols
    - Prevent potentially hard to catch bugs due to differences in semantics

RDMA Protocol Extensions

- Internet-Draft submitted to IETF STORM WG: ‘RDMA Protocol Extensions’
- Extensions to RDMAP RFC5040 (aka iWARP)
- Draft co-authored by
  - Broadcom Corporation
  - Chelsio Communications, Inc
  - Intel Corporation

RDMA Protocol Extensions (2)

- Adds support for Atomic Operations
  - FetchAdd
  - Swap
  - CmpSwap
- Adds support for Immediate Data
  - RDMA Write w/Immediate
  - …
- Wire Protocol to implement missing Verbs and semantics

Next Steps

- RDMA Protocol Extensions are only half of the solution
- Exposed via Verbs
- Find a home for Verbs
  - RDMA Consortium is no longer active
  - IETF does not specify Verbs

Next Steps (2)

- Proposal was made to the OFA Board on March 14th for OFA to set up a Technical Working Group to update the iWARP Verbs specification
  - This proposal was retracted on March 17th
  - Fundamental concern was expanding OFA charter to developing specifications

- Before the next OFA Workshop the iWARP Verbs document will be updated.
  - If anyone would like to participate please contact Terry Hulett (terry.v.hulett@intel.com)

Conclusions

- iWARP delivers an IP routable implementation of the OFA API for the Enterprise Data Center
- iWARP is evolving to remove differences in functionality relative to IB as exposed through the OFA API