

RDMA Protocol Extensions

81st IETF – Quebec City, Canada

26th July 2011

draft-ietf-storm-rdmap-ext-01

Hemal V. Shah, Broadcom Corporation
hemal@broadcom.com

Co-authors:

Felix Marti, Wael Nouredine, Asgeir Eiriksson, Chelsio Communications Inc.
Robert Sharp, Intel Corporation

Overview

- RFC 5040 defines RDMA Protocol (RDMA)
- This draft extends RDMA
- Proposed extensions
 - Atomic operations on remote memory locations
 - Immediate Data messages
- History
 - Draft -00 was submitted in March 2011
 - Received valuable comments on the reflector
 - Incorporated feedback received on the reflector in -01
 - Draft -01 was submitted in July 2011

RDMA Atomic Operations

- Used to execute atomic ops at remote locations
- Draft defines two new RDMA Messages for executing a 64-bit atomic op at a specified destination address
 - Messages: Atomic Request and Atomic Response
 - Defined operations: FetchAdd, CmpSwap & Swap
- Motivation behind atomic operations
 - Enhance the usability of RDMAP
 - Important building block for distributed shared memory applications

Immediate Data

- Used in conjunction with an RDMA Operation
- Two new RDMA Messages are defined
 - Immediate Data
 - Immediate Data with Solicited Event (SE)
- 8B of immediate data is allowed
- Motivation
 - Improve ULP efficiency
 - Allow sender to provide a small amount of data

Why adopt this draft?

- Enhance RDMAP
- Increase usability of RDMAP
- Improve ULP efficiency
- Provide additional building blocks for distributed shared memory implementations