

Progress in Standardization of RDMA technology
Arkady Kanevsky, Network Appliance

The last several saw a tremendous progress in the development, commercialization, and standardization of the RDMA technology. >From early years of VIA to current InfiniBand commercial RDMA technology is poised to provide a huge performance leap and unify the fabric for networking, clustering and storage infrastructures.

While fledging InfiniBand products appear from several dozen companies, the work on the next round had already begun. IETF RDDL WG started its work on developing RDMA over IP protocols, while RDMA Consortium defines RDMA over TCP.

Taking the best ideas from the VIPL APIs, DAT Collaborative has defined user (uDAPL-1.0) and kernel (kDAPL-1.0) level APIs for RDMA protocols. The first reference implementations of DAPL started to appear on Source Forge. The Open Group, taking DAT APIs as a starting point, is working on their extensions for datagrams and other features. The Open Group also defines APIs for Socket extension protocol to take advantage of the RDMA performance for sockets. Storage community has defines new APIs and protocols, like DAFS and SRP, to take advantage of the RDMA technology with several products available already.

This presentation provides the highlights of various RDMA related standardization activities and its impact on HPEC and HPC communities.

Arkady Kanevsky
Network Appliance, Inc.
375 Totten Pond Rd.
3rd Floor
Waltham, MA 02451-2010

email: arkady@netapp.com
phone: 781-768-5395
fax: 781-895-1195
http: www.netapp.com
general phone: 781-768-5300