

RDMA over Converged Ethernet (RoCE) Fact Sheet

What is RoCE?

Remote Direct Memory Access (RDMA) is a critical technology at the heart of the fastest supercomputers and many of the largest data centers in the world. RDMA first became widely adopted in the Supercomputing space with InfiniBand but has expanded into enterprise markets and is now being widely adopted over Ethernet networks with RDMA over Converged Ethernet or RoCE (pronounced like “rocky”). RoCE is driving an advanced data center architecture that eliminates dedicated storage area networks and converges compute, networking and storage onto a single fabric. Leveraging the latest advances in reliable Ethernet and Data Center Bridging (DCB), efficient RDMA mechanisms in RoCE provide lower CPU overhead and increase mainstream data center application performance over Layer 2 and Layer 3 fabrics at 10, 40, and 100 Gigabit Ethernet speeds and beyond.

What is RDMA?

Remote Direct Memory Access (RDMA) is a remote memory management capability that allows server-to-server data movement directly between application memory without any CPU involvement. Offloading the chore of data movement from the CPU makes both the data movement and the execution of applications more efficient. RDMA delivers performance and efficiency gains that are not available from any other communications protocol; including: low latency, improved resource utilization, flexible resource allocation, scalability and fabric unification.

Data center architects running RDMA applications on an Ethernet infrastructure can expect to see application performance and efficiency improvements along with cost and power savings that come from the offloading of data movement and the higher availability of CPU resources to the application. Adopters of RoCE can make use of RDMA’s capabilities without leaving the familiar transport and network management system of Ethernet. In this way, adopters can upgrade their application performance without investing in alternative switching interconnect technologies.



InfiniBand Trade Association Administration
3855 SW 153rd Drive, Beaverton, OR 97006
503.619.0565

www.infinibandta.org

Who Benefits from RoCE?

Targeted applications and benefits include:

Financial services

- Sample applications: Low latency messaging applications include Tibco, Wombat/NYSE, IBM WebSphere MQ, Red Hat MRG, 29West/Informatica
- Benefits: High performance, scalable

Data warehousing

- Sample applications: Oracle RAC, IBM DB2 PureScale, Microsoft SQL
- Benefits: Significantly higher job operations per second, linear scaling with cluster size, maintains table scan time in the face of exponential growth in database table sizes

Data Storage

- Sample applications: Microsoft SMB-Direct,
- Benefits: iSCSI over RDMA delivers higher throughput

Cloud computing clustered solutions

- Sample applications: VMware, Red Hat KVM, Citrix Xen, Microsoft, Amazon EC2, Google App Engine
- Benefits: Improved service level agreements through deterministic performance, efficient clustering allowing for elastic/scale out computing and higher performance from virtualized servers through its support of SR-IOV

Web 2.0 clustered solutions

- Sample applications: Hadoop, Memcached, Eucalyptus, Cassandra
- Benefits: Minimize response time, Maximize jobs/second, enable to build highly scalable infrastructure

What is the Return on Investment?

By reducing the Ethernet network latency and offloading CPU overhead, RoCE increases performance in search, storage, database, financial and high transaction rate applications. By increasing CPU efficiency and improving application performance, RoCE can reduce the number of servers needed which produces energy savings and reduces the footprint of Ethernet-based data centers.

Who supports RoCE?

RoCE has received broad industry support from such hardware, software and system vendors as Arista, Broadcom, Cisco, Emulex, IBM, Mellanox, Microsoft, QLogic and Red Hat, end users including NYSE Euronext and industry organizations including the OpenFabrics Alliance and the Ethernet Alliance.

RoCE-enabled drivers are downloadable today in the latest OpenFabrics Enterprise Distribution (OFED) stack. Many Linux distributions, which include OpenFabrics Software, support a wide and rich range of middleware and application solutions such as IPC, sockets, messaging, virtualization, SAN, NAS, file systems and databases. RoCE is also supported on Windows Servers with SMB-Direct for Storage Applications. RoCE can therefore deliver all three dimensions of unified networking on Ethernet – IPC, NAS and SAN.

On the hardware side, RoCE-supported network interface cards (NICs) and switches are available today from many hardware vendors.



The InfiniBand Trade Association was founded in 1999 and is chartered with maintaining and furthering the InfiniBand specification. The IBTA is led by a distinguished steering committee that includes Cray, Emulex, HP, IBM, Intel, Mellanox, Microsoft, Oracle and QLogic. Other members of the IBTA represent leading enterprise IT vendors who are actively contributing to the advancement of the InfiniBand specification. The IBTA markets and promotes InfiniBand from an industry perspective through online, marketing and public relations engagements, and unites the industry through IBTA-sponsored technical events and resources. For more information on the IBTA, visit: www.infinibandta.org

©2015 InfiniBand Trade Association. All rights reserved.