

RDMA over Converged Ethernet Delivers Efficient Lightweight Transport

[Electronic Design](#)

[Bill Lee](#)

Fri, 2015-10-09 09:43



The ever-growing data deluge isn't showing any signs of slowing down. The rise of IoT, Big Data, mobile video, and the plethora of smartphones and tablets create a need for faster application performance to support all of these devices in real-time. CIOs, enterprise data-center architects, and solutions engineers are seeking an efficient and cost-effective way to improve application performance and data-center productivity.

This is where RoCE steps in. RDMA over Converged Ethernet, or RoCE (pronounced "Rocky"), is one approach that allows CIOs and data-center managers to improve performance, while leveraging their investment in Ethernet-based data centers. The RoCE specification has been defined by the InfiniBand Trade Association (IBTA) to enable efficient data movement around the data center. So how does this all work?


Related

[Chips Make Software Defined Networking Work](#)

[Break Through The TCP/IP Bottleneck With iWARP](#)

[Ethernet Everywhere: Accelerating Data-Center Migration to the Cloud](#)

Let's start with RDMA. RDMA stands for remote direct memory access. It enables faster movement of data between servers and between servers and storage. In one scenario, data movement utilizes traditional transport technologies such as TCP/IP, which generates multiple copies and significant CPU overhead in order to get data on or off a server. RDMA, on the other hand, utilizes hardware offloads to move data faster with less overhead, and it bypasses the TCP/IP stack altogether. This frees up the CPU to do the work that it's meant to do—run applications and process massive amounts of data (*see the figure*).



RDMA delivers a number of benefits to cloud, storage, virtualization, and hyper-converged infrastructures:

low latency and CPU overhead

- Higher network efficiency due to reduced server bottlenecks
- Efficient data transfer; OS bypass enables fastest access to remote data
- Efficient computing reduces power, cooling, and space requirements
- Support for message passing, sockets, and storage protocols
- Supported by all major operating systems

RoCE brings the benefits of RDMA to the Ethernet. It delivers more efficient network utilization with lower latency and improved CPU efficiency, while reducing overall financial investment by increasing server productivity and leveraging existing Ethernet technology. It transports data across Layer 2 and Layer 3 networks, thus improving traffic isolation and enabling hyperscale data-center deployments. RoCE features include:

- Efficient, lightweight transport, layered directly over Ethernet
- Lowest latency in the Ethernet industry; 1- μ s server-to-server RDMA latency
- Doubles storage throughput compared with traditional software transport
- IBTA standard, supported in OpenFabrics Software stack

RoCE is already seeing traction in the industry with an ecosystem of vendors, with multiple server and storage OEM vendors shipping RoCE products and multiple adapter vendors supporting the technology. Millions of RoCE adapters are deployed in hyperscale data centers today.

Microsoft is one vendor supporting the technology. At Microsoft Ignite 2015, RoCE demonstrations showed that the technology is able to deliver remote flash-storage access at the same throughput and latency performance as locally attached flash SSDs. **Check out the video below:**

Impact of RDMA Networking in an SMB3 workload over 10...



With RoCE enabled, the SMB3 server achieved twice the throughput, half the latency, and 33% less CPU overhead than alternative solutions.

Introducing the RoCE Initiative

Recognizing the need for education and resources around RoCE, IBTA launched the [RoCE Initiative](#) to serve as the go-to source for information on RDMA over Ethernet solutions, as well as advance the RoCE multi-vendor ecosystem. The initiative's goals are to accelerate adoption and development of RoCE applications, and enable CIOs, enterprise data-center architects, and solutions engineers to deploy fabrics that will significantly boost application performance and data-center productivity. IBTA member companies are contributing education and technical resources, including case studies, white papers, and training webinars, in support of this initiative.

To learn more about RoCE and the RoCE Initiative, check out the IBTA's "What is RoCE?" video:

What is RDMA over Converged Ethernet (RoCE)?



Content Classification: Curated

Source URL: <http://electronicdesign.com/interconnects/rdma-over-converged-ethernet-delivers-efficient-lightweight-transport>