



Intel Ethernet 800 Series Network Adapters in Dell EMC PowerEdge R740xd servers offer near-local storage performance over the network*

Deploy the Intel Ethernet 800 Series using either RoCEv2 or iWARP RDMA protocols to meet business needs

New storage technologies require more advanced networking technologies to keep latency in check and deliver the fast direct I/O experience that customers have come to expect. Dell EMC™ PowerEdge™ R740xd servers now offer Intel® Ethernet 800 Series Network Adapters as another option to meet these needs.

In the Principled Technologies data center, a Dell EMC PowerEdge R740xd server with a 100GbE 800 Series network adapter and eight NVMe™ drives offered storage performance over the network that approached performance levels of direct attached storage. We used the Flexible I/O (fio) tool to test the NVMe over Fabrics performance of a 100GbE Intel Ethernet Network Adapter E810-CQDA2 using two protocols for remote direct memory access (RDMA): RDMA over Converged Ethernet v2 (RoCEv2) and iWARP.

If your organization requires new technologies that can support workloads such as machine learning, analytics, and supercomputing, know that Dell EMC PowerEdge R740xd servers with 800 Series adapters can deliver strong network performance and offer the flexibility to choose from multiple RDMA protocols.



- Near-local storage performance**
 - up to **4M IOPS** (iWARP)
 - up to **3.9M IOPS** (RoCEv2)vs. 4M IOPS locally on a 50/50 read/write workload
- Supported NVMe over Fabrics using both iWARP and RoCEv2 RDMA protocols**

*when running a 4K random 50/50 read/write workload with 100GbE Intel Ethernet Network Adapter E810-CQDA2 in a Dell EMC PowerEdge R740xd server in our tests performed in Sept 2020

Demanding workloads need a step up in networking

Traditional TCP/IP protocol networking is adequate for many workloads, but it comes with its share of limitations, including CPU overhead and higher than desired round-trip latency. The RDMA protocol reduces CPU overhead and latency by skipping through layers of the stack; this can help with workloads that need the quickest possible data transfers, such as business analytics, artificial intelligence, and high-performance computing. Two network transports, iWARP and RoCEv2, implement RDMA and the entire network stack in the network adapter hardware, offloading nearly all the work from the CPU to offer faster networking.

iWARP implements RDMA over IP networks using TCP, making it ideal for organizations that wish to use RDMA over their existing IP network infrastructure without any specialized hardware.

RoCEv2 relies on an Ethernet network configured to use Layer 2 Priority Flow Control (PFC) or Layer 3 DSCP PFC to minimize congestive packet loss.



Intel Ethernet 800 network adapters with RDMA support both iWARP and RoCEv2 protocols concurrently. Whether you house your Dell EMC PowerEdge R740xd servers in a single data center or across multiple locations, your organization can get the low latency and CPU overhead that demanding workloads require.

About Intel Ethernet 800 Series Network Adapters

According to Intel, the new Intel Ethernet 800 Series offers technologies that “address a variety of workloads used in Cloud, Communications, and Enterprise market segments.”¹ Other features of the 800 Series Network Adapters include ADQ (Application Device Queues), which can help deliver greater SLA consistency and more administrator control over pipeline programming. To learn more about the Intel Ethernet 800 Series, visit <https://www.intel.com/content/www/us/en/architecture-and-technology/ethernet/introducing-800-series.html>.



Dell EMC PowerEdge R740xd with Intel Ethernet 800 Series Network Adapters support NVMe over Fabrics with two RDMA protocols

NVMe over Fabrics is a feature that enables you to use NVMe commands to transfer data between a server and a target storage device over network fabric. Both NVMe drives and RDMA protocols offer speed that aims to decrease latencies and improve performance. NVMe over Fabrics hooks the two together and allows NVMe devices on different servers to communicate quickly via RDMA for large, fast storage pools.

In our tests, the 800 Series Network Adapters allowed the client server to communicate with the eight NVMe devices on the target server using both RoCEv2 and iWARP RDMA protocols. This functionality has the potential to improve performance and drive down latency.

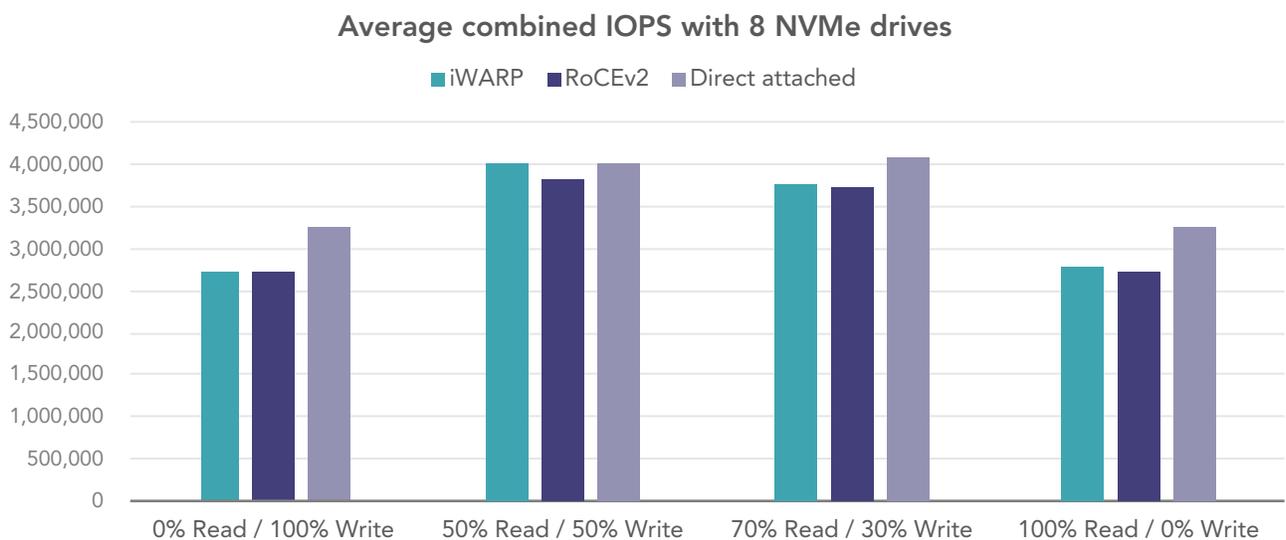


Figure 1: Combined IOPS for each read/write ratio we tested on the configuration with eight NVMe drives. Higher is better. Source: Principled Technologies

Client

Dell EMC PowerEdge R740 server



Target

Dell EMC PowerEdge R740xd server with 8x Intel® Optane™ SSD DC P4800X drives

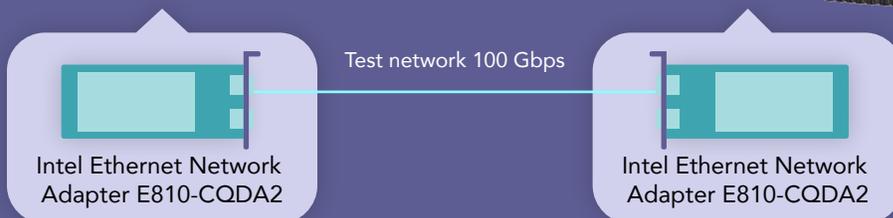


Figure 2: How we set up RDMA functionality over a 100 GbE network. Source: Principled Technologies

Get near-local storage performance over the network with Intel Ethernet 800 Series Network Adapters on the Dell EMC PowerEdge R740xd

Every network adapter has a theoretical maximum performance that it can achieve, though system overhead ensures that performance will approach but not reach that maximum. Using local, direct-attached storage by definition has the least overhead, but that isn't always an option. 800 Series Network Adapters on the Dell EMC PowerEdge R740xd with eight NVMe drives provided strong IOPS performance in our tests using two RDMA protocols to route fio test traffic over the network, as Figure 2 shows. Particularly for workloads that were a mix of reads and writes, modeling some typical database workloads, IOPS performance over the network was close to the same test workload running locally. This shows that equipping your Dell EMC PowerEdge R740xd servers with 800 Series Network Adapters can provide your organization with a high-performing way to link your infrastructure across multiple sites.

In some cases, the IOPS were limited in the eight-drive configuration by the maximum throughput of the 100Gbps network connection. We conducted further testing using four and six NVMe drives, respectively, to show that IOPS over the network were virtually identical compared to direct attached storage when the physical network throughput limit was no longer a bottleneck. To learn more, read the [science behind the report](#).

Average bandwidth (MB/s) with 8 NVMe drives

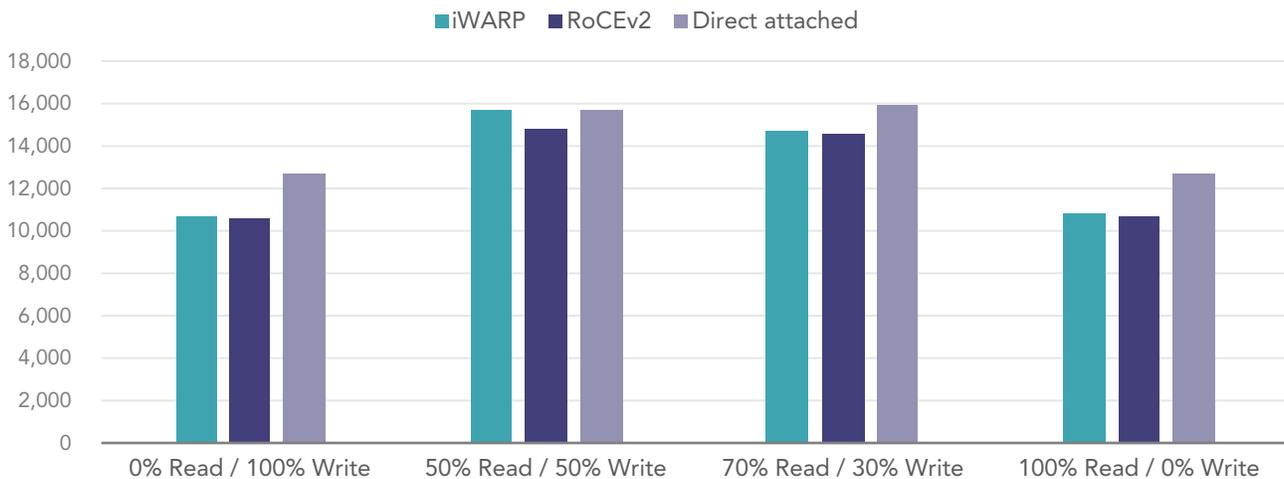


Figure 3: Average bandwidth (MB/s) for each read/write ratio we tested on the configuration with eight NVMe drives. Higher is better. Source: Principled Technologies

About 2nd Generation Intel Xeon Scalable processors

The 2nd Generation Intel Xeon® Scalable processor platform features a wide range of processor types to support the workloads you run, including Bronze, Silver, Gold, and Platinum. According to Intel, the 2nd Generation Intel Xeon Scalable platform can handle a variety of workloads, including enterprise, cloud, HPC, storage, and communications. The new processor line also supports a new memory and storage technology to further accelerate workloads, Intel Optane DC persistent memory. To learn more about the Intel Xeon Scalable processor family, visit <https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html>.



About the Dell EMC PowerEdge R740xd

What does the “xd” stand for? Extra drives—with up to 24 NVMe drives or 32 x 2.5” or 18 x 3.5” drives in just 2U of rack space, the Dell EMC PowerEdge R740xd promotes flexibility through several drive options and density that allows you to use software-defined storage solutions. The two-socket Dell EMC PowerEdge R740xd rack server features 2nd Generation Intel Xeon® Scalable processors, embedded Dell EMC iDRAC9, and Dell EMC OpenManage™ software for management. To learn more about the Dell EMC PowerEdge R740xd, visit <https://www.dell.com/en-us/work/shop/povw/poweredge-r740xd>.

In conclusion

Demanding workloads such as high-performance computing, analytics, and machine learning require fast networking to accomplish tasks in a timely manner. To that end, the Intel Ethernet 800 Series Network Adapter is now available in select Dell EMC PowerEdge R740xd server configurations. We verified that the server with the 800 Series Network Adapter delivered strong IOPS performance with eight NVMe drives, both over the network and locally using direct attached storage. Additionally, the server-and-network adapter solution supported NVMe over Fabrics functionality using both RoCEv2 and iWARP RDMA protocols, allowing admins the choice of whichever protocol best suits their organization’s needs.

-
- 1 Intel, “Move Data Faster,” accessed September 24, 2020, <https://www.intel.com/content/www/us/en/architecture-and-technology/ethernet.html>.

Read the science behind this report at <http://facts.pt/olx740z> ►



Facts matter.®

This project was commissioned by Dell EMC.

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.