



## Get the server I/O capacity you need for heavy mixed, write, and read workloads

### Configure your Intel Xeon Scalable processor-powered Dell EMC PowerEdge R740xd with more NVMe PCIe SSDs to increase IOPS and throughput for your workloads


Choosing the right server means deciding on the right balance of I/O and compute power for your workloads. When you need a tremendous amount of raw I/O power, you may want to consider a configuration with NVMe PCIe solid-state drives (SSDs). These SSDs connect directly to the processors, bringing storage close to compute and providing fast performance. The Intel® Xeon® Scalable processor-powered Dell EMC™ PowerEdge™ R740xd has the compute and I/O scalability to handle four, eight, or twelve NVMe PCIe SSDs.

We used the Iometer tool to measure the storage capabilities of a single PowerEdge R740xd powered by Intel with varying numbers of NVMe SSDs. We configured Iometer to mimic several different types of workloads, including mixed, write-heavy, and read-heavy workloads. For all profiles we tested, input/output operations per second (IOPS) and throughput rose dramatically as we added more NVMe SSDs, while latency remained consistently under 1 ms. (For more on the importance of throughput and latency, see page 2.)

The test results in this report highlight the raw IOPS capacity of the PowerEdge R740xd in three configurations. Whether you need the I/O performance of four, eight, or twelve NVMe SSDs, the Intel Xeon Scalable processor-powered Dell EMC PowerEdge R740xd delivers.



Choose the NVMe PCIe SSD configuration that matches your needs



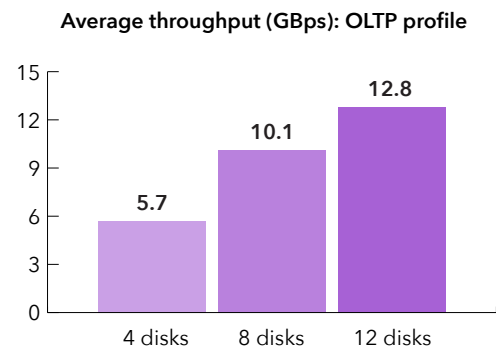

## Increase your server's I/O capacity

We tested each configuration of the Intel processor-powered Dell EMC PowerEdge R740xd with three different Iometer profiles. In each set of profile testing, more NVMe PCIe SSDs translated into higher IOPS and higher throughput. Not only that, but the solution delivered fast response times throughout testing. Ten years ago, solutions delivering sub-1ms latency for enterprise workloads were not widely available—but in our testing, the PowerEdge R740xd with NVMe PCIe SSDs consistently kept response times below 1 ms.

### Mixed workload

Our first Iometer profile, which imitated an online transaction processing (OLTP) workload, had an 8K block size and was 70 percent read, 30 percent write, and 100 percent random. OLTP workloads include a variety of database workloads and match up well with companies that run ecommerce websites or frequently utilize large databases of customer or employee information.

Running this profile, the solution sustained 75 percent more IOPS and 75 percent higher throughput when we moved from four to eight NVMe PCIe SSDs. Upgrading to twelve SSDs enabled the solution to handle over 27 percent more IOPS and 27 percent more throughput than it could with eight SSDs. The PowerEdge R740xd provided sub-1ms response times at each configuration.



### The relationship among IOPS, throughput, and latency

Though IOPS is a useful measurement metric, it does not paint a complete picture of a solution's storage performance. Throughput and latency are additional metrics that tell us more about what a solution can do for you. Throughput takes block size into account, measuring the total number of GB per second the solution can handle, while latency measures how long the system waits between processing orders. A solution with high throughput and low latency would be able to process a lot of data at a time while still responding to users quickly.


Together, the three metrics paint a more holistic picture of a solution's capabilities. If a solution has high IOPS, high throughput, and low latency, it's a good sign.



## Write-heavy workload

At 20 percent read, 80 percent write, 80 percent random, and 20 percent sequential with a 4K block size, we designed our second lometer profile to mimic a virtual desktop infrastructure (VDI) workload. Reference the results for this workload if you're planning on implementing VDI in your organization, which could be useful if you have many users who work at desktops in similar programs.

When we upgraded from four to eight NVMe PCIe SSDs, we saw 91 percent more IOPS with 91 percent higher throughput. Performance continued to rise as we added four more SSDs, with the twelve-SSD configuration delivering 32 percent more IOPS and 32 percent more throughput than the configuration with eight SSDs. At all three configurations, the solution's average latency remained under 1 ms.



## Testing with lometer

lometer is an industry-standard tool for measuring raw I/O capacity, or the maximum number of IOPS that a solution can provide.<sup>1</sup> Because lometer stresses exclusively storage, its results do not indicate how many IOPS the solution can deliver while running a real-world workload, which will also require work from the rest of the components in the hardware.

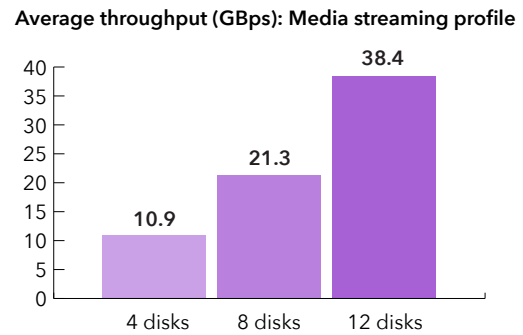

However, lometer does permit us to create different profiles by customizing three factors:

- **Block size:** The size of the file being inputted or outputted. Setting the block size larger imitates a workload that might have to deal with more complex operations and larger files, while setting a smaller block size will get you closer to a workload with many simple operations.
- **Read/write ratio:** The ratio between read operations, which simply review existing data, and write operations, which add new data or change existing data.
- **Random/sequential ratio:** The ratio between random reads or writes (the system chooses specific, random packets of data with no relation to each other) and sequential reads or writes (the system chooses the data packet that immediately follows the packet it has just chosen).

Together, these factors combine to mimic different real-world workloads. If you're trying to decide on the right storage configuration for your work, you may wish to review the results for the lometer profile closest to your workload.

## Read-heavy workload

Our final Iometer profile was 98 percent read, 2 percent write, and 100 percent sequential with a 64K block size. This profile imitates a media streaming workload, ideal if your server is sustaining many users watching videos or listening to music. Upgrading from four to eight NVMe SSDs yielded 95 percent more IOPS with 95 percent more throughput. When we added the next four SSDs, the solution delivered 81 percent more IOPS and 81 percent more throughput than it could with eight SSDs. At four, eight, and twelve disks, the solution consistently provided average response times below 1 ms.



### The Dell EMC PowerEdge R740xd

The 14th-generation Dell EMC PowerEdge R740xd is available in many configurations, so when you order your PowerEdge R740xd, you can choose the number of NVMe PCIe SSDs that's right for you. According to Dell EMC, the 2U PowerEdge R740xd—powered by the latest Intel Xeon Scalable processors—offers strong storage performance and includes new management capabilities with iDRAC9, SupportAssist for proactive customer support, and an out-of-band management architecture.<sup>2</sup>





## Choose a configuration with more NVMe PCIe SSDs, get more I/O performance with consistently low latencies

Whether you measure storage performance by IOPS, throughput, or both, the results of our Iometer testing were clear. If you choose a configuration of the Intel Xeon Scalable processor-powered PowerEdge R740xd with eight NVMe SSDs as opposed to four, or twelve as opposed to eight, you can gain substantially more raw I/O performance—all while maintaining latencies under 1 ms. We tested with several different workloads so you can compare your own I/O needs to the results of our tests and find the solution that will be right for you.

---

1 "Iometer," accessed June 26, 2017, <http://www.iometer.org/>

2 "PowerEdge R740/xd servers", accessed June 21, 2017, <https://dellservervr.dell.com/poweredge-r740xd/>.



On June 9, 2017, we finalized the hardware and software configurations we tested. Updates for current and recently released hardware and software appear often, so unavoidably these configurations may not represent the latest versions available when this report appears. For older systems, we chose configurations representative of typical purchases of those systems. We concluded hands-on testing on June 16, 2017.

All of the hardware we used in this study, including both the server and the NVMe SSDs, is pre-production and not currently available in the market. Dell EMC provided us with this hardware for our testing.

## Appendix A: System configuration information

Server configuration information	Dell EMC PowerEdge R740xd
BIOS name and version	Dell 1.0.2
Non-default BIOS settings	N/A
Operating system name and version/build number	Windows Server® 2016 Datacenter x64
Date of last OS updates/patches applied	06/01/2017
Power management policy	Performance
Processor	
Number of processors	2
Vendor and model	Intel Xeon Gold 6134 CPU
Core count (per processor)	8
Core frequency (GHz)	3.20
Stepping	4
Memory module(s)	
Total memory in system (GB)	192
Number of memory modules	12
Vendor and model	SK Hynix HMA82GR7AFR8N-VK
Size (GB)	16
Type	PC4-2666V
Speed (MHz)	2,666



Server configuration information		Dell EMC PowerEdge R740xd
Storage controller		
Controller #1		
Vendor and model	Dell PERC H730P	
Cache size	2 GB	
Firmware version	25.5.0.0018	
Driver version	10.0.14393.0	
Controller # 2		
Vendor and model	UCEA-200	
Cache size (GB)	N/A	
Firmware version	1.0.0	
Driver version	10.0.14393.0	
Local storage (type A)		
Number of drives	2	
Drive vendor and model	Dell ST300MM0006	
Drive size	300 GB	
Drive information (speed, interface, type)	10K, 6Gb SAS, HDD	
Local storage (type B)		
Number of drives	12	
Drive vendor and model	Dell P/N 0JDMHM (PM1725a)	
Drive size	3.2 TB	
Drive information (speed, interface, type)	NVMe	
Network adapter		
Vendor and model	Dell P/N 0G8RPD (BCM57800)	
Number and type of ports	2 x 10 GbE, 2 x 1 GbE	
Driver version	6.604.6.0	
Cooling fans		
Vendor and model	Nidec 4VXP3-X30	
Number of cooling fans	6	
Power supplies		
Vendor and model	Dell E1100E-S0	
Number of power supplies	4	
Wattage of each (W)	1,100	



## Appendix B: How we tested

### Installing Windows Server 2016 Datacenter Edition

1. Boot the target server.
2. When requested, press any key to boot to the installation media.
3. On the Windows Server 2016 screen, click Next.
4. Click Install Now.
5. On the Activate Windows Screen, click I don't have a product key.
6. Select Windows Server 2016 Datacenter (Desktop Experience), and click Next.
7. Agree to the License.
8. Select Custom Installation.
9. Select the target volume for the Windows installation, and click Next.
10. Wait for the installation to complete.
11. At the Customize settings screen, type and confirm a password, and click Finish.

### Configuring Windows Server 2016

1. At the login screen, enter your password, and press enter.
2. Open Windows Update, and select Check for updates.
3. Allow the updates to install.
4. Open Windows Firewall with Advanced Security.
5. Click Windows Firewall Properties.
6. On the Domain Profile, Private Profile, and Public Profile tabs, set the Firewall state to off.
7. Open System Properties.
8. In the Remote tab, select Allow remote connections to this computer. Click to deselect the Allow Connections only.... checkbox, and click OK.
9. Using Disk Management, create a simple volume using all defaults for each disk. There will be 12 total disks.

### Installing Iometer

1. Download Iometer 1.1.0 from [http://sourceforge.net/projects/iometer/files/iometer-stable/1.1.0/iometer-1.1.0-win64.x86\\_64-bin.zip/download](http://sourceforge.net/projects/iometer/files/iometer-stable/1.1.0/iometer-1.1.0-win64.x86_64-bin.zip/download).
2. Unzip the downloaded files.
3. In the iometer-1.1.0-win64.x86\_64-bin folder, run iometer.exe.
4. On the Intel Open Source License window, click I agree.

We set up Access Specifications for three profiles: media streaming, VDI, and OLTP. For each Access Specification, we ran Iometer for 30 minutes on 4, 8, and 12 disks. See Figures 1, 2, and 3 for the settings we used in each profile.





## Setting specifications for Access Profiles

See Figures 1, 2, and 3 to learn how we configured the Access Specifications for the OLTP, VDI, and media streaming profiles.

**Edit Access Specification**

Name: OLTP      Default Assignment: None

Size	% Ac...	% R...	% Ran...	Del...	B...	Alignment	Reply
0 MiB 8 KiB 0 B	100	70	100	0	1	0 MiB 8 KiB ...	none

Buttons: Insert Before, Insert After, Delete

**Transfer Request Size:** 0 Megabytes, 8 Kilobytes, 0 Bytes

**Percent of Access Specification:** 100 Percent

**Percent Read/Write Distribution:** 30% Write, 70% Read

**Percent Random/Sequential Distribution:** 0% Sequential, 100% Random

**Burstiness:** Transfer Delay: 0 ms, Burst Length: 1 I/Os

**Align I/Os on:**  Request Size Boundaries,  Sector Boundaries

**Reply Size:**  No Reply, 0 Megabytes, 8 Kilobytes, 0 Bytes

Buttons: OK, Cancel

Figure 1: OLTP profile access specifications

**Edit Access Specification**

Name: VDI      Default Assignment: None

Size	% Ac...	% R...	% Ran...	Del...	B...	Alignment	Reply
0 MiB 4 KiB 0 B	100	20	100	0	1	0 MiB 4 KiB ...	none

Buttons: Insert Before, Insert After, Delete

**Transfer Request Size:** 0 Megabytes, 4 Kilobytes, 0 Bytes

**Percent of Access Specification:** 100 Percent

**Percent Read/Write Distribution:** 80% Write, 20% Read

**Percent Random/Sequential Distribution:** 0% Sequential, 100% Random

**Burstiness:** Transfer Delay: 0 ms, Burst Length: 1 I/Os

**Align I/Os on:**  Request Size Boundaries,  Sector Boundaries

**Reply Size:**  No Reply, 0 Megabytes, 4 Kilobytes, 0 Bytes

Buttons: OK, Cancel

Figure 2: VDI profile access specifications

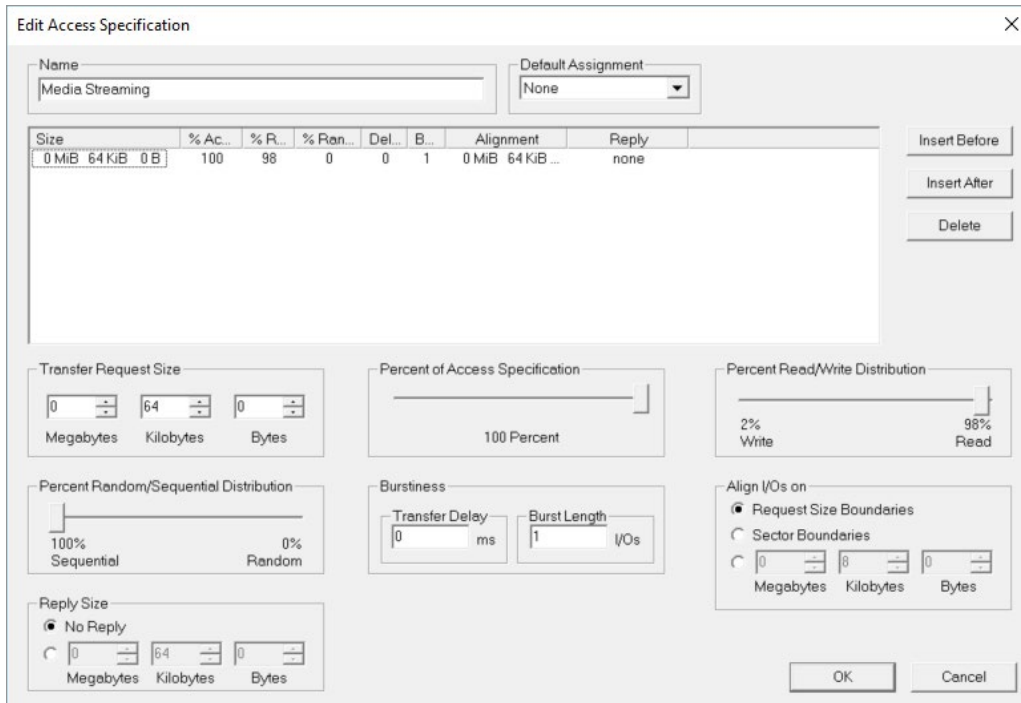


Figure 3: Media streaming profile access specifications

## Configuring general settings

We used the settings listed below for all tests. If a setting is not listed, we used the default.

- Maximum disk size: 1073741824 sectors
- Run time: 30 minutes
- Ramp up time: 300 seconds

## Configuring test-specific settings

The access profile, number of disks, workers per disk, and outstanding I/Os per target (queue depth) varied. See Table 1 for additional details.

	Access profile	Disks	Workers per disk	Outstanding I/Os per target
OLTP 4 disks	OLTP	4	2	48
OLTP 8 disks	OLTP	8	2	48
OLTP 12 disks	OLTP	12	2	48
VDI 4 disks	VDI	4	4	4
VDI 8 disks	VDI	8	4	4
VDI 12 disks	VDI	12	4	4
Media streaming 4 disks	Media streaming	4	1	32
Media streaming 8 disks	Media streaming	8	1	32
Media streaming 12 disks	Media streaming	12	1	32

Table 1: Test-specific iometer settings

## Appendix C: Complete test results

Test name	IO profile	Number of disks	Average IOPS	Average throughput (GBps)	Average latency (ms)	Host CPU utilization
OLTP 4 disks	OLTP	4	718,000	5.7	0.53	21.87%
OLTP 8 disks	OLTP	8	1,257,000	10.1	0.61	46.74%
OLTP 12 disks	OLTP	12	1,597,000	12.8	0.72	72.27%
VDI 4 disks	VDI	4	640,000	2.6	0.10	26.89%
VDI 8 disks	VDI	8	1,224,000	4.9	0.10	63.94%
VDI 12 disks	VDI	12	1,610,000	6.4	0.12	98.26%
Media streaming 4 disks	Media	4	170,000	10.9	0.75	7.75%
Media streaming 8 disks	Media	8	332,000	21.3	0.77	15.94%
Media streaming 12 disks	Media	12	600,000	38.4	0.64	25.17%

Table 2: Complete test results

This project was commissioned by Dell Technologies.



**Facts matter.®**

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners.

**DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:**

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.