



The science behind the report:

Get better database analytics performance at a lower cost with Dell EMC PowerEdge R840 servers equipped with value SAS and data center NVMe SSDs from Toshiba Memory

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report [Get better database analytics performance at a lower cost with Dell EMC PowerEdge R840 servers equipped with value SAS and data center NVMe SSDs from Toshiba Memory](#).

We concluded our hands-on testing on April 3, 2019. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on March 29, 2019, or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

The tables below presents our findings in detail.

Drive type	Operations per second (OPS)	Percent win (OPS)	Average latency (ms)	Percent win (latency)
Enterprise SATA SSDs	22,204		8.64	
Value SAS SSDs from Toshiba Memory	45,902	106.73%	4.17	51.66%
Data center NVMe™ SSDs from Toshiba Memory	52,791	137.76%	3.63	57.97%

Cost analysis results

To calculate the cost of the Dell EMC™ PowerEdge™ R840 system we tested, we used list pricing information obtained from Dell EMC on April 11, 2019. These numbers, while recent at the time of testing, are subject to change. To calculate the OPS-per-dollar metric, we divided the total OPS by the total solution cost of each configuration.

Configuration	Total OPS	Total solution cost	OPS per dollar	Percent win
Dell EMC PowerEdge R840 with enterprise SATA SSDs	22,204	\$57,791.01	0.3841	
Dell EMC PowerEdge R840 with RM5 Series value SAS SSDs	45,902	\$57,391.01	0.7997	108.20%
Dell EMC PowerEdge R840 with CD5 Series data center NVMe SSDs	52,791	\$58,991.01	0.8948	132.96%

System configuration information

The table below presents detailed information on the system we tested.

Server configuration information	Dell EMC™ PowerEdge™ R840
BIOS name and version	SMBIOS 3.0.0 (1.3.10)
BIOS settings	Default
Operating system name and version/build number	CentOS Linux® release 7.6.1810 6 kernel 3.10.0-957.10.1.el7.x86_64
Date of last OS updates/patches applied	04/03/2019
Power management policy	Performance
Processor	
Number of processors	4
Vendor and model	Intel® Xeon® Platinum 8180M CPU @ 2.50GHz
Core count (per processor)	28
Core frequency (GHz)	2.5
Stepping	4
Memory module(s)	
Total memory in system (GB)	256
Number of memory modules	8
Vendor and model	Samsung® M393A4K40BB1-CRC
Size (GB)	32
Type	DDR-4
Speed (MHz)	2,400
Speed running in the server (MHz)	2,400
Local storage OS	
Number of drives	2 (RAID 1)
Drive vendor and model	Dell MZ7LM1T9H MJPD3
Drive size (GB)	1,800
Drive information (speed, interface, type)	6Gbps SATA SSD
Network adapter	
Vendor and model	Broadcom Gigabit Ethernet BCM5720
Number and type of ports	1 x 1GbE
SATA drives	
Number of drives	4 (RAID 10)
Drive vendor and model	Intel SSDSC2KB960G8R
Drive size (GB)	960
Drive information (speed, interface, type)	6Gbps SATA SSD
RAID controller	PERC H740P v50.5.0-1750

Server configuration information		Dell EMC™ PowerEdge™ R840
SAS drives		
Number of drives		4 (RAID 10)
Drive vendor and model		Toshiba Memory KRM5XRUG960G
Drive size (GB)		960
Drive information (speed, interface, type)		12 Gbps SAS 2.5" SSD
RAID controller		PERC H740P v50.5.0-1750
NVMe drives		
Number of drives		4 (RAID 10)
Drive vendor and model		Toshiba Memory KCD5XLUG960G
Drive size (GB)		960
Drive information (speed, interface, type)		PCIe Gen 3 x4 2.5" NVMe SSD
RAID controller		mdadm v4.1-rc1
Cooling fans		
Vendor and model		Dell
Number of cooling fans		6
Power supplies		
Vendor and model		Dell PWR SPLY, RDNT, DELTA
Number of power supplies		2
Wattage of each (W)		1,600

How we tested

We installed and configured CentOS Linux® (release 7.6.1810) and MongoDB® (version 4.0.8) on the server under test, a Dell EMC PowerEdge R840. We configured the server with 256 GB of RAM and four Intel Xeon Platinum 8180M processors. We installed the OS on a pair of 1.8TB SATA SSDs set up in a RAID 1 configuration. We set up and managed the SATA and SAS drives using PERC H740P (hardware RAID 10), and set up and managed the NVMe drives using the Linux utility mdadm (software RAID 10).

We also deployed and configured a Dell EMC PowerEdge R740xd infrastructure server with VMware vSphere™ 6.7 and three YCSB load-generating/driver VMs, each with CentOS 7 as the operating system, 8 vCPUs, and 16 GB of RAM. We set up testbed networking using a 1GbE top-of-rack network switch and dual 1GbE ports on the server under test and infrastructure server.

Configuring the PERC RAID devices

We used the PERC RAID controller to build the RAID devices on the SATA and SAS drives.

1. Log onto the server's iDRAC console.
2. Redirect the screen output to a browser window.
3. Boot into the Dell™ Lifecycle Controller.
4. Configure the SATA and SAS RAID devices in the RAID Controller Configuration Utility with the following changes from default settings:
 - a. RAID 10
 - b. No Read Ahead
 - c. Write Through
5. Save settings and reboot the system to the OS installation disk in preparation of OS installation.

Installing CentOS 7

1. Insert a CentOS 7.4 bootable USB stick into the server or mount an ISO image as a virtual CD, and boot to it.
2. Select Install or upgrade an existing system.
3. Choose the language you wish to use, and click Continue.
4. Select Installation Destination.
5. Select the desired disk for the OS.
6. Under Other Storage Options, select I will configure partitioning.
7. Click Done.
8. Select Click here to create them automatically.
9. Click Done.
10. Click Accept Changes.
11. Select Network and Hostname.
12. Enter the desired hostname for the VM.
13. Turn on the desired network port, and click Configure.
14. On the General tab, select Automatically connect to this network when it is available.
15. On the IPv4 Settings tab, under Method, select Manual.
16. Under Addresses, click Add, and enter the desired static IP information for the server.
17. Enter the desired DNS information.
18. Click Save, and click Done.
19. Select Date and Time, and ensure you have set the correct date, time, and time zone.
20. Click the cog next to the Network Time On/Off switch to add your NTP server.
21. Add the IP address of your NTP server, and click +.
22. Uncheck all other NTP servers.
23. Click OK.
24. Click Done.
25. Click Software Selection, and select Server with GUI.
26. In the Add-Ons for Select Environment section, select Compatibility Libraries and Development Tools.
27. Click Done.
28. Click Begin Installation.
29. Select Root Password.
30. Enter the desired root password, and click Done.
31. When the installation completes, to restart the server, select Reboot.

Installing and configuring the OS on the server under test and driver VMs

1. Log into the server as the root user.
2. Stop and disable the firewall service by typing the following commands:

```
systemctl stop firewalld  
systemctl disable firewalld
```
3. Enter the SELinux configuration file by typing the following command:

```
vim /etc/sysconfig/selinux
```
4. Disable Hugepages:

```
echo 'never' > /sys/kernel/mm/transparent_hugepage/enabled  
echo 'never' > /sys/kernel/mm/transparent_hugepage/defrag
```
5. Perform system updates:

```
sudo yum update -y
```

Configuring the mdadm RAID device

1. Log into the server as the root user.
2. Install mdadm:

```
yum install -y mdadm
```
3. Create NVMe RAID volume (assumes drives nvme{0,1,2,3}n1):

```
mdadm --zero-superblock /dev/nvme{0,1,2,3}n1  
mdadm -C /dev/md/md-nvme -l10 -n4 /dev/nvme{0,1,2,3}n1
```

Preparing the RAID devices and installing MongoDB on the server under test

For the purpose of the testing, we labeled and mounted the arrays as follows.

Raw RAID devices:

- SATA: /dev/sdb
- SAS: /dev/sdc
- NVMe: /dev/md127, which is system created soft link to /dev/md/md-nvme

For documentation purposes, we refer to these as /dev/<raw>.

Formatted RAID devices:

- SATA: /dev/sdb1 mounted at /data/md-sata
- SAS: /dev/sdc1 mounted at /data/md-sas
- NVMe: /dev/md127p1 mounted at /data/md-nvme

For documentation purposes, they will be referred to as /dev/<drive> and /data/<mount>.

Installing MongoDB on the server under test

1. Add the Mongo community database to your yum repo and install MongoDB to your VM by typing the following commands:

```
vim /etc/yum.repos.d/mongodb-org-4.0.repo  
yum install -y mongodb-org
```
2. Format and build the file system, and mount the database drive by typing the following commands:

```
echo -e "o\nn\np\nl\n\n\nw" | fdisk /dev/<raw>  
mkfs.xfs /dev/<drive>  
mount /dev/<drive> /data/<mount>  
mkdir /data/<mount>/DB  
mkdir /data/<mount>/LOGS
```
3. Open the mongod configuration file by typing the following command:

```
vim /etc/mongod.conf
```

4. Inside the mongod configuration file, make the following changes:


```
systemLog:
  path: /data/<mount>/LOG/mongod.log
storage:
  dbPath: /data/<mount>/DB
  journal:
    enabled: true
    engine: "wiredTiger"
net:
  port: 27017
  bindIp: [your server's IP],localhost
```
5. Save the mongod configuration file.
6. To start your mongod server, type the following command:


```
systemctl start mongod
```
7. To verify that the mongod server has started correctly, type the following command:


```
systemctl status mongod
```

Installing YCSB on the server under test

1. Install Java with the following command:


```
sudo yum install java-devel
```
2. Install Maven with the following commands:


```
wget http://ftp.heanet.ie/mirrors/www.apache.org/dist/maven/maven-3/3.1.1/binaries/apache-maven-3.1.1-bin.tar.gz
sudo tar xzf apache-maven*-bin.tar.gz -C /usr/local
cd /usr/local
sudo ln -s apache-maven-* maven
echo export M2_HOME=/usr/local/maven >> /etc/profile.d/maven.sh
echo export PATH=${M2_HOME}/bin:${PATH} >> /etc/profile.d/maven.sh
```
3. Download and unpack YCSB into your driver with the following commands:


```
run("curl -O --location https://github.com/brianfrankcooper/YCSB/releases/download/0.12.0/ycsb-0.12.0.tar.gz")
run("tar xfvz ycsb-0.12.0.tar.gz")
```

Installing YCSB on the load-generating/driver VMs

1. Add the Mongo community database to your yum repo and install MongoDB to your VM by typing the following commands:


```
vim /etc/yum.repos.d/mongodb-org-4.0.repo
yum install -y mongodb-org
```
2. Add the Mongo community database to your yum repo and Install Java with the following command:


```
sudo yum install java-devel
```
3. Install Maven with the following commands:


```
wget "http://apache.cs.utah.edu/maven/maven-3/3.5.2/binaries/apache-maven-3.5.2-bin.tar.gz"
sudo tar xzf apache-maven*-bin.tar.gz -C /usr/local
cd /usr/local
sudo ln -s apache-maven-* maven
echo export M2_HOME=/usr/local/maven >> /etc/profile.d/maven.sh
echo export PATH=${M2_HOME}/bin:${PATH} >> /etc/profile.d/maven.sh
```
4. Download and unpack YCSB into your driver with the following commands:


```
run("curl -O --location https://github.com/brianfrankcooper/YCSB/releases/download/0.12.0/ycsb-0.12.0.tar.gz")
run("tar xfvz ycsb-0.12.0.tar.gz")
```

Creating the YCSB database

1. Log into a YCSB VM.
2. Run the following command to create a roughly 1.2TB database for MongoDB:

```
/root/ycsb-0.12.0/bin/ycsb load mongodb -s -P /root/ycsb-0.12.0/workloads/workloadc -threads 64 -p mongodb
.url=mongodb://<server-ip>:27017/ycsb?w=1 -p recordcount=100000000
```

Running the YCSB test

1. Log into the three YCSB driver VMs as root.
2. On the MongoDB Server, run the following command to clear all caches:

```
echo 3 > /proc/sys/vm/drop_caches
```
3. On all three VMs, run the following command to warm up the database:

```
/root/ycsb-0.12.0/bin/ycsb run mongodb -s -P /root/ycsb-0.12.0/workloads/workloadc -threads 64 -p mongo
db.url=mongodb://<server-ip>:27017/ycsb?w=1 -p operationcount=30000000 -p mongodb.upsert=true > ~/output/ycsbtest_64_threads_30000000_operations_workloadc_[your_ycsb_hostname]_runX.warmup
```
4. Once IO output stabilizes to indicate that the database has warmed up, run the following command:

```
/root/ycsb-0.12.0/bin/ycsb run mongodb -s -P /root/ycsb-0.12.0/workloads/workloadc -threads 64 -p mongo
db.url=mongodb://<server-ip>:27017/ycsb?w=1 -p operationcount=30000000 -p mongodb.upsert=true > ~/output/ycsbtest_64_threads_30000000_operations_workloadc_[your_ycsb_hostname]_runX.txt
```
5. Record the results from the test.

Read the report at <http://facts.pt/medhjiv> ►

This project was commissioned by Toshiba Memory America, Inc.



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.