

The science behind the report:

Enable greater data reduction and storage performance with Dell EMC PowerStore 7000 series storage arrays

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 [Enable greater data reduction and storage performance with Dell EMC PowerStore 7000 series storage arrays](#).

We concluded our hands-on testing on June 24, 2020. 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 June 4, 2020 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

Table 1: Results of our testing on the Dell EMC™ PowerStore™ 7000T and HPE Primera A670.

	Dell EMC PowerStore 7000T	HPE Primera A670	Win
Vdbench testing			
Data reduction ratio			
50% 128KB and 50% 256KB write workload	7.1:1 (reduced 800 GB of data to 111 GB)	2.0:1 (reduced 800 GB of data to 401 GB)	3.5x the data reduction
Input/output operations per second (IOPS)			
8KB random 100% write workload (data reduction enabled)	232,602	75,160	209% more IOPS
8KB random 100% write workload (data reduction disabled)	N/A	171,772	N/A
32KB 70/30 read/write workload	222,865	111,026	101% more IOPS
OLTP-like workload	182,030	105,687	72% more IOPS
Bandwidth (MB/s)			
Random read workload	23,239	10,763	116% more bandwidth
Sequential read workload	23,238	9,877	135% more bandwidth
Latency (ms)			
OLTP-like workload	0.56	2.01	72% lower latency
Accessing snapshot restore data			
Quiescing the application VM			
Time (seconds)	11	11	
Steps	5	5	
Removing the LUNs from the ESXi environment			
Time (seconds)	N/A	55	
Steps	N/A	3	
Unexporting the virtual volume set from the ESXi host			
Time (seconds)	N/A	32	
Steps	N/A	6	
Performing the snapshot restore			
Time (seconds)	52	991	
Steps	10	4	
Exporting LUNs back to the ESXi environment			
Time (seconds)	N/A	25	
Steps	N/A	5	

	Dell EMC PowerStore 7000T	HPE Primera A670	Win
Adding LUNs back to the virtualized ESXi environment			
Time (seconds)	N/A	7	
Steps	N/A	2	
Adding LUNs back to the VM as RDM volumes			
Time (seconds)	N/A	92	
Steps	N/A	5	
Restarting the application VM			
Time (seconds)	4	4	
Steps	1	1	
Total			
Time (seconds)	67	1,217	17.1x faster
Steps	16	31	15 fewer steps

Table 2: Results of our testing on the Dell EMC PowerStore 7000X and HPE Primera A670.

	Dell EMC PowerStore 7000X	HPE Primera A670	Win
Data reduction ratio			
50% 128KB and 50% 256KB write workload	7.1:1 (reduced 800 GB of data to 111 GB)	2.0:1 (reduced 800 GB of data to 401 GB)	3.5x the data reduction
Scenario 1: Testing storage performance on the Dell EMC PowerStore 7000X (hosting VMs internally) vs. the HPE Primera A670 (hosting VMs externally)			
4KB random write Vdbench workload (IOPS)	162,931	87,856	85% more IOPS
Scenarios 2-4: Testing database performance on the Dell EMC PowerStore 7000X (simultaneously hosting database VMs internally and VMs running Vdbench simulations externally) vs. the HPE Primera A670 (only hosting VMs running Vdbench simulations externally)			
Scenario 2: Dell EMC PowerStore 7000X internally hosting MongoDB VMs running a YCSB workload			
Database operations per second**	188,320	N/A	N/A
Read application latency (ms)**	0.33	N/A	N/A
Update (write) application latency (ms)**	0.41	N/A	N/A
Scenario 3: Dell EMC PowerStore 7000X and HPE Primera A670 externally hosting VMs running a Vdbench workload (bandwidth)			
256KB Vdbench workload (MB/s)**	13,704	10,450	31% more bandwidth
Scenario 4: Dell EMC PowerStore 7000X and HPE Primera A670 externally hosting VMs running a Vdbench workload (IOPS)			
8KB random write Vdbench workload (IOPS)**	129,909	78,775	65% more IOPS

**The PowerStore 7000X achieved these results while running Vdbench and MongoDB simultaneously. The HPE Primera A670 was running only Vdbench.

	Dell EMC PowerStore 7000X	HPE Primera A670	Win
Out-of-the-box VM deployment			
Configuring Fibre Channel switch zoning			
Time (seconds)	N/A	317	
Steps	N/A	7	
Mapping hosts			
Time (seconds)	N/A	57	
Steps	N/A	8	
Creating one virtual volume			
Time (seconds)	N/A	45	
Steps	N/A	6	
Adding LUNs to a virtualized ESXi environment			
Time (seconds)	N/A	38	
Steps	N/A	7	
Creating a single VMFS datastore from a LUN			
Time (seconds)	N/A	30	
Steps	N/A	6	
Deploying a VM in a VMware vSphere environment			
Time (seconds)	53	47	
Steps	12	9	
Total			
Time (seconds)	53	534	9.0x faster
Steps	12	43	31 fewer steps

System configuration information

Table 3: Detailed information on the servers we tested.

Server configuration information	4x Dell EMC PowerEdge™ R740	4x Dell EMC PowerEdge R740	4x Dell EMC PowerEdge R740
Testbed	Dell EMC PowerStore 7000X	Dell EMC PowerStore 7000T	HPE Primera A670
BIOS name and version	Dell EMC PowerEdge R740 2.3.10	Dell EMC PowerEdge R740 2.2.11	Dell EMC PowerEdge R740 2.2.11
Non-default BIOS settings	Virtualization enabled	Virtualization enabled	Virtualization enabled
Operating system name and version/build number	VMware ESXi™ 6.7.0 Update 3 Build 14320388	VMware ESXi 6.7.0 Update 3 Build 14320388	VMware ESXi 6.7.0 Update 3 Build 14320388
Power management policy	Performance	Performance	Performance
Processor			
Number of processors	2	2	2
Vendor and model	Intel® Xeon® Gold 6230	Intel Xeon Gold 6230	Intel Xeon Gold 6230
Core count (per processor)	20	20	20
Core frequency (GHz)	2.10	2.10	2.10
Memory module(s)			
Total memory in system (GB)	256	256	256
Number of memory modules	16	16	16
Vendor and model	Hynix® HMA82GR7CJR8N-WM	Hynix HMA82GR7CJR8N-WM	Hynix HMA82GR7CJR8N-WM
Size (GB)	16	16	16
Type	PC4-2666	PC4-2666	PC4-2666
Speed (MHz)	2,666	2,666	2,666
Speed running in the server (MHz)	2,666	2,666	2,666
Storage controller			
Vendor and model	Dell PERC H330	Dell PERC H330	Dell PERC H330
Firmware version	25.5.6.0009	25.5.6.0009	25.5.6.0009
Driver version	7.708.07.00	7.708.07.00	7.708.07.00
Local storage			
Number of drives	2	2	2
Drive vendor and model	Samsung® MZ-7LH480A	Micron® MTFDDAK480TDC	Micron MTFDDAK480TDC
Drive size (GB)	480	480	480
Drive information (speed, interface, type)	6Gbps, SATA, SSD	6Gbps, SATA, SSD	6Gbps, SATA, SSD
Network adapter			
Vendor and model	Broadcom® Gigabit Ethernet BCM5720	Broadcom Gigabit Ethernet BCM5720	Broadcom Gigabit Ethernet BCM5720
Number and type of ports	2x 1Gb 2x 10Gb	2x 1Gb 2x 10Gb	2x 1Gb 2x 10Gb
Driver version	21.40.25.31	21.40.22.20	21.40.22.20

Server configuration information	4x Dell EMC PowerEdge™ R740	4x Dell EMC PowerEdge R740	4x Dell EMC PowerEdge R740
Storage adapter			
Vendor and model	Emulex® LPe35002-M2-D	Emulex LPe35002-M2-D	Emulex LPe35002-M2-D
Number and type of ports	2x 2-port 32 Gb Fibre Channel	2x 2-port 32 Gb Fibre Channel	2x 2-port 32 Gb Fibre Channel
Firmware version	12.2.212.10	12.2.212.10	12.2.212.10
Power supplies			
Vendor and model	Dell 0PJMDNA01	Dell 05RHVVA03	Dell 05RHVVA03
Number of power supplies	2	2	2
Wattage of each (W)	750	750	750

Table 4: Detailed information on the storage we tested.

Storage configuration information	Dell EMC PowerStore 7000X	Dell EMC PowerStore 7000T	HPE Primera A670
Software version	1.0.0.0.5.109	1.0.0.0.5.109	4.1.0.27
Number of storage nodes	2	2	2
Number of data drives	21	21	24
Drive part number	118000740	118000740	P02434-002
Drive size (TB)	1.92	1.92	1.92

Detailed testing procedure

We received three separate testbeds from Dell EMC: one for testing the Dell EMC PowerStore 7000T, one for testing the PowerStore 7000X, and one for testing the HPE Primera A670. In addition to the storage arrays we tested, each testbed had four Dell EMC PowerEdge R740 servers equipped with two dual-port 32GB Emulex Fibre Channel adapters and VMware ESXi 6.7. With the exception of the internal VMs running on the PowerStore 7000X array, we performed all testing using Fibre Channel storage protocol and configured the arrays to use 16 ports connected to a 48-port Fibre Channel switch. We used 10GbE switches for testbed management and VM traffic. Where possible, we verified that the setups on the testbeds were identical, and we configured the hosts, multipathing, and arrays as closely as possible following best practices published by each storage vendor. For all of our testing, we used CentOS 7 VMs with Vdbench 5.04.07. We configured each VM with eight vCPUs and 24 GB of RAM.

We performed storage efficiency tests on the HPE Primera A670 array using 16KB deduplication unit sizes, as HPE recommends (<https://h20195.www2.hpe.com/v2/getdocument.aspx?docname=a50000209enw>). Both Dell EMC PowerStore arrays can deduplicate 4KB deduplication unit sizes, so we tested with this deduplication unit size on the PowerStore 7000X and 7000T arrays.

Data reduction testing

Dell EMC PowerStore 7000T and 7000X vs. HPE Primera A670

We created four 200GB volumes on all three arrays and mapped the volumes to one host in each testbed. We then added the four volumes to a single VM on each server host as RDM disks. After creating and mapping the volumes and adding the RDM disks, we used Vdbench to fill the volumes. Next, we checked each array's GUI to see by how much the array had reduced the data. After recording the data reduction rate, we deleted the RDM disks and volumes and repeated the volume provisioning and testing process twice more to get three runs per array. We present the median results.

We used a 50/50 mix of 128KB and 256KB sequential writes with a single thread to fill the LUNs with data:

```
compratio=3
dedupratio=2
dedupunit=16k (4k for PowerStore arrays)
dedupsets=50%
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=vdb-001
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
sd=sd3,hd=hd1,lun=/dev/sdd
sd=sd4,hd=hd1,lun=/dev/sde
wd=wd_default,sd=*
wd=wd_prefill,sd=sd*,xfersize=(128k,50,256k,50),seekpct=eof,rdpct=0
rd=default
rd=fill_1,wd=wd_prefill,elapsed=10h,interval=10,iorate=max,forthreads=(1)
rd=fill_2,wd=wd_prefill,elapsed=10h,interval=10,iorate=max,forthreads=(1)
```

Performance testing

Dell EMC PowerStore 7000T vs. HPE Primera A670

We created 32 800GB volumes and mapped these to all four hosts in each testbed. After rescanning the hosts HBAs, we attached the volumes as RDM disks to 16 Vdbench test VMs, each of which contained two RDMs. We used Vdbench to prefill the volumes and run the test. Once the prefilling had completed, we ran through a series of Vdbench I/O simulations and collected the results. After completing the tests, we unmapped the RDMs and volumes from the VMs and hosts and deleted the volumes on the array. We waited until the number of uninitialized chunklets on the HPE Primera A670 array returned to zero before we repeated the volume provisioning and testing process twice to get three runs per array. We present the median results.

Prefilling the data

We used 256KB sequential writes with a single thread to fill all 32 volumes with data:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
dedupsets=5%
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_prefill,sd=sd*,xfersize=256k,seekpct=eof,rdpct=0
rd=default
rd=rd_prefill,wd=wd_prefill,elapsed=10h,interval=10,iorate=max,forthreads=(1)
```

Running the random reads workloads

For these tests, we configured Vdbench to run at 256KB random reads with four threads:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_256k,sd=sd*,xfersize=256k,seekpct=100
rd=default
rd=read256k_test,wd=wd_256k,iorate=max,warmup=30,interval=10,forrdpct=(100),elapsed=120,forthreads=(4)
```

Running the sequential reads workloads

For these tests, we configured Vdbench to run 256KB sequential reads with four threads:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
```



```

sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_256k,sd=sd*,xfersize=256k,seekpct=seqnz
rd=default
rd=read256k_test,wd=wd_256k,iorate=max,warmup=30,interval=10,forrdpct=(100),elapsed=120,forthreads=(4)

```

Running the OLTP-like workloads

For these tests, we configured Vdbench to run an OLTP-like workload with a mix of 8KB reads and writes and 128KB reads and writes with eight threads:

```

compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_OLTP2A_RRH,sd=*,rhpct=100,rdpct=100,xfersize=8K,skew=20,range=(0,6m)
wd=wd_OLTP2A_RM,sd=*,rdpct=100,xfersize=8k,skew=45,range=(0,100)
wd=wd_OLTP2A_RW,sd=*,rdpct=0,xfersize=8K,skew=15,range=(0,100)
wd=wd_OLTP2A_SR,sd=*,rdpct=100,seekpct=seqnz,range=(0,100),xfersize=128K,skew=10
wd=wd_OLTP2A_SW,sd=*,rdpct=0,seekpct=seqnz,range=(0,100),xfersize=128K,skew=10
rd=default
rd=rd_OLTP2A,wd=wd_OLTP2A_*,iorate=max,pause=120,elapsed=120,interval=10,warmup=30,forthreads=8

```

Testing the 32KB 70/30 read/write mix workloads

For these tests, we configured Vdbench to run a 32KB mixed 70/30 read/write workload with 16 threads:

```

compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_32k,sd=sd*,xfersize=32k,seekpct=100
rd=default
rd=read32k_
test,wd=wd_32k,iorate=max,pause=120,warmup=30,interval=10,forrdpct=(70),elapsed=120,forthreads=(16)

```

Running the 8KB random writes workloads

For these tests, we configured Vdbench to run an 8KB random writes workload with 16 threads:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_8k,sd=sd*,xfersize=8k,seekpct=100
rd=default
rd=read8k_
test,wd=wd_8k,iorate=max,pause=120,warmup=30,interval=10,forrdpct=(0),elapsed=120,forthreads=(16)
```

Running the paced OLTP-like workloads

For these tests, we configured Vdbench to generate a 107,000 IOPS OLTP-like workload using a mix of 8KB reads and writes and 128KB reads and writes with eight threads. When the arrays reached the target IOPS, we captured latency metrics for both arrays:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_OLTP2A_RRH,sd=*,rhpct=100,rdpct=100,xfersize=8K,skew=20,range=(0,6m)
wd=wd_OLTP2A_RM,sd=*,rdpct=100,xfersize=8k,skew=45,range=(0,100)
wd=wd_OLTP2A_RW,sd=*,rdpct=0,xfersize=8K,skew=15,range=(0,100)
wd=wd_OLTP2A_SR,sd=*,rdpct=100,seekpct=seqnz,range=(0,100),xfersize=128K,skew=10
wd=wd_OLTP2A_SW,sd=*,rdpct=0,seekpct=seqnz,range=(0,100),xfersize=128K,skew=10
rd=default
rd=rd_OLTP2A,wd=wd_OLTP2A_*,pause=120,iorate=(107000),elapsed=120,interval=10,warmup=60,forthreads=(8)
```

After completing these tests, we repeated the 8KB random write test on the HPE Primera A670 with deduplication and compression disabled on 32x 250GB volumes. We repeated the volume provisioning and testing process twice more to get three test runs. We present the median results.

Dell EMC PowerStore 7000X with internal VMs vs. HPE Primera A670 with external VMs

The Dell EMC PowerStore 7000X array can allow VMs to consume compute and storage resources within the array without the need for external server hosts. We refer to VMs using this internal storage deployment model as "internal VMs" in this report. We provisioned four internal Vdbench test VMs on the Dell EMC PowerStore 7000X and four external VMs on the HPE Primera A670 testbed. On the Dell EMC PowerStore 7000X, we split these four VMs evenly across the two-node VMware PowerStore cluster. On the Primera array, we distributed the four VMs evenly between two Dell EMC PowerEdge R740 hosts. We created 16 800GB volumes on both arrays and assigned four volumes to each VM. We used Vdbench to prefill the volumes and run the test. Once Vdbench completed the prefilling, we ran through a series of Vdbench I/O simulations and collected the results. After we completed the tests, we unmapped the RDMS and volumes from the VMs and hosts and deleted the volumes on the array. We waited until the number of uninitialized chunklets on the Primera array returned to zero before we repeated the volume provisioning and testing process twice more to get three runs per array. We present the median results.

Prefilling the LUNS with data

To fill the LUNs with data, we used 256KB sequential writes with a single thread:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore arrays)
dedupsets=5%
messagescan=no
hd=default, shell=ssh, user=root, jvms=1
hd=hd1, system=vdb-001
hd=hd2, system=vdb-002
hd=hd3, system=vdb-005
hd=hd4, system=vdb-006
sd=default, openflags=o_direct
sd=sd1, hd=hd1, lun=/dev/sdb
sd=sd2, hd=hd1, lun=/dev/sdc
...
sd=sd15, hd=hd4, lun=/dev/sdd
sd=sd16, hd=hd4, lun=/dev/sde
wd=default, sd=*
wd=wd_prefill, sd=sd*, xfersize=256k, seekpct=eof, rdpcct=0
rd=default
rd=rd_prefill, wd=wd_prefill, elapsed=10h, interval=10, iorate=max, forthreads=(1)
```

Running the 4KB random write workload

For these tests, we configured Vdbench to run at 4KB random writes with 32 threads:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore arrays)
messagescan=no
hd=default, shell=ssh, user=root
hd=hd1, system=vdb-001
hd=hd2, system=vdb-002
...
hd=hd15, system=vdb-015
hd=hd16, system=vdb-016
sd=default, openflags=o_direct
sd=sd1, hd=hd1, lun=/dev/sdb
sd=sd2, hd=hd1, lun=/dev/sdc
...
sd=sd31, hd=hd16, lun=/dev/sdb
sd=sd32, hd=hd16, lun=/dev/sdc
wd=default, sd=*
wd=wd_4k, sd=sd*, xfersize=4k, seekpct=100
rd=default
rd=read4k_
test, wd=wd_4k, iorate=max, pause=120, warmup=30, interval=10, forrdpct=(0), elapsed=120, forthreads=(32)
```

Dell EMC PowerStore 7000X with internal MongoDB database VMs and external Vdbench VMs vs. HPE Primera A670 with external Vdbench VMs

We set up a MongoDB replication set consisting of three internal MongoDB VMs on the Dell EMC PowerStore 7000X cluster. We used a separate YCSB VM running on one of the Dell EMC PowerEdge R740 hosts in the PowerStore 7000X testbed to drive the MongoDB workload. Simultaneously, we created 32 400GB volumes on each array and mapped them to four Dell EMC PowerEdge R740 hosts on each testbed. We created RDM disks out of the 32 mapped volumes and mounted two RDMs per VM to 16 external VMs. We distributed all 16 VMs evenly between four Dell EMC PowerEdge R740 hosts. We used Vdbench to prefill the volumes and run the I/O test. Once Vdbench completed the prefilling, we initialized the replica set on the MongoDB VMs and used the external YCSB VM to load the MongoDB database with 100 million records. The 16 external VMs connected on the PowerStore 7000X and HPE Primera testbeds ran through 256KB and 8KB Vdbench I/O workloads. While the Vdbench tests were running on the PowerStore 7000X, the three internal VMs were also running a workload of 50 million operations driven by the external YCSB VM, using four separate YCSB processes and 64 total threads. We collected Vdbench results for the HPE Primera A670 and both Vdbench and MongoDB results for Dell EMC PowerStore 7000X. Below, Figure 1 demonstrates how we simultaneously ran workload on internally and externally hosted VMs on the Dell EMC PowerStore 7000X. Because the HPE Primera A670 lacks the ability to host VMs internally, we only ran the external Vdbench workloads on it. After the tests completed, we deleted all the volumes and the MongoDB databases and allowed the number of uninitialized chunklets on the HPE Primera array return to zero. We repeated the volume provisioning and testing process twice more to get three runs per array. We present the median results.




Figure 1: Time (minutes) it took to run two Vdbench workloads on externally hosted VMs while simultaneously running a workload on internally hosted VMs on the Dell EMC PowerStore 7000X. Source: Principled Technologies.

Prefilling the data

We used 256KB sequential writes with a single thread to fill the LUNs with data:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore arrays)
dedupsets=5%
messagescan=no
hd=default, shell=ssh, user=root, jvms=1
hd=hd1, system=vdb-001
hd=hd2, system=vdb-002
...
hd=hd15, system=vdb-015
hd=hd16, system=vdb-016

sd=default, openflags=o_direct
sd=sd1, hd=hd1, lun=/dev/sdb
sd=sd2, hd=hd1, lun=/dev/sdc
...
sd=sd31, hd=hd16, lun=/dev/sdb
sd=sd32, hd=hd16, lun=/dev/sdc
wd=default, sd=*
wd=wd_prefill, sd=sd*, xfersize=256k, seekpct=eof, rdpcnt=0
rd=default
rd=rd_prefill, wd=wd_prefill, elapsed=10h, interval=10, iorate=max, forthreads=(1)
```

Running the sequential reads workloads

For these tests, we configured Vdbench to run 256KB sequential reads with three threads:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore arrays)
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_256k,sd=sd*,xfersize=256k,seekpct=seqnz
rd=default
rd=read256k_test,wd=wd_256k,iorate=max,warmup=30,interval=10,forrdpct=(100),elapsed=300,forthreads=(3)
```

Running the 8KB random writes workloads

For these tests, we configured Vdbench to run at 8KB random writes with 16 threads:

```
compratio=3
dedupratio=3
dedupunit=16k (4k for PowerStore array)
messagescan=no
hd=default,shell=ssh,user=root
hd=hd1,system=vdb-001
hd=hd2,system=vdb-002
...
hd=hd15,system=vdb-015
hd=hd16,system=vdb-016
sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
...
sd=sd31,hd=hd16,lun=/dev/sdb
sd=sd32,hd=hd16,lun=/dev/sdc
wd=default,sd=*
wd=wd_8k,sd=sd*,xfersize=8k,seekpct=100
rd=default
rd=read8k_test,wd=wd_8k,iorate=max,pause=300,warmup=30,interval=10,forrdpct=(0),elapsed=300,forthreads=(16)
```

Prefilling the MongoDB data (Dell EMC PowerStore 7000X only)

We used an externally hosted YCSB VM to load our MongoDB replica set with 100 million records using four YCSB processes each loading 25 million records:

```
(./bin/ycsb load mongodb-async -s -P ./workloads/workloadb -p mongodb.url="mongodb://MongoD_
vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p recordcount=100000000 -p insertstart=0 -p
insertcount=25000000) &
(./bin/ycsb load mongodb-async -s -P ./workloads/workloadb -p mongodb.url="mongodb://
MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p recordcount=100000000 -p
insertstart=25000000 -p insertcount=25000000) &
(./bin/ycsb load mongodb-async -s -P ./workloads/workloadb -p mongodb.url="mongodb://
MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p recordcount=100000000 -p
insertstart=50000000 -p insertcount=25000000) &
(./bin/ycsb load mongodb-async -s -P ./workloads/workloadb -p mongodb.url="mongodb://
MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p recordcount=100000000 -p
insertstart=75000000 -p insertcount=25000000)
```

Running the MongoDB test (Dell EMC PowerStore 7000X only)

We used an externally hosted YCSB VM to run 50 million operations against our MongoDB replica set using four separate YCSB processes:

```
(./bin/ycsb run mongodb-async -s -P ./workloads/workloadb -threads 16 -p mongodb.url="mongodb://MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p operationcount=50000000 2>&1 |tee driver1.out) &
(./bin/ycsb run mongodb-async -s -P ./workloads/workloadb -threads 16 -p mongodb.url="mongodb://MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p operationcount=50000000 2>&1 |tee driver2.out) &
(./bin/ycsb run mongodb-async -s -P ./workloads/workloadb -threads 16 -p mongodb.url="mongodb://MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p operationcount=50000000 2>&1 |tee driver3.out) &
(./bin/ycsb run mongodb-async -s -P ./workloads/workloadb -threads 16 -p mongodb.url="mongodb://MongoD_vm1:27017,MongoD_vm2:27107,MongoD_vm3:27017/ycsb?w=1" -p operationcount=50000000 2>&1 |tee driver4.out)
```

Usability testing

Dell EMC PowerStore 7000X vs. HPE Primera A670

Out-of-the-box VM deployment

This testing scenario assumed that the user had initialized both arrays but not performed Fibre Channel zoning or host mapping. Our scenario also assumed that the HPE Primera A670 could use an available target VMware ESXi host and that the user had deployed a VMware vCenter® VM on the Dell EMC PowerStore and HPE Primera testbeds.

Deploying VMs in a VMware vSphere® environment on the PowerStore 7000X

1. Open a browser, and navigate to vCenter's vSphere client.
2. Enter credentials, and click Login.
3. Expand the PS7000X Data Center.
4. Expand the PS7000X Cluster.
5. Right-click an ESXi host, and select New Virtual Machine.
6. Select Create a new virtual machine, and click Next.
7. Enter a virtual machine name, and click Next.
8. On Select Storage screen, select PowerStore PS7000X storage, and click Next.
9. On compatibility screen, select ESXi 6.7 and later, and click Next.
10. Select appropriate OS details, and click Next.
11. On Customize hardware screen, click Next.
12. On Ready to complete screen, click Finish. vSphere will create the VM.

Configuring Fibre Channel switch zoning on the HPE Primera A670

1. Log into the Fibre Channel switch UI.
2. Click Configure→Zone Admin.
3. Click Zone tab→New Zone.
4. Create new Fibre Channel zones with the appropriate host's HBAs and storage ports.
5. Click Zone Config tab, and add all the new zones to Zone Config Members.
6. Click Save Config.
7. Click Enable config, and select the appropriate configuration to enable.

Mapping hosts on the HPE Primera A670

1. Open a browser, and navigate to the HPE Primera and 3PAR StoreServ Management Console (SSMC).
2. Enter credentials, and click Login.
3. Under Primera and 3PAR SSMC dropdown, select Hosts.
4. Click Create host.
5. On General screen, enter a host name, and select a host OS.
6. Under Paths, click Add FC.
7. On Add FC screen, select all host initiators, and click Add.
8. On Create Host screen, click Create.

Creating one virtual volume on the HPE Primera A670

1. Under Primera and 3PAR SSMC dropdown, select Virtual Volumes.
2. Click Create Virtual Volume.
3. On General screen, enter a virtual volume name, switch dedup and compression to Yes. For volume size, enter 1 TiB.
4. Expand Export screen, and click Add.
5. On Add screen, select the appropriate host, and click Add.
6. On Create Virtual Volume screen, click Create.

Adding LUNs to a virtualized ESXi environment on the HPE Primera A670

1. Open a browser, and navigate to the vCenter vSphere client.
2. Enter credentials, and click Login.
3. If necessary, expand Datacenter and cluster.
4. Select the target host.
5. Click the Configure tab.
6. To discover newly added storage, under Storage→Storage adapters, click Rescan Storage.
7. Select Scan for new Storage Devices, and click OK.

Creating a single VMFS datastore from a LUN on the HPE Primera A670

1. Right-click the target host, and select Storage→New Datastore.
2. On the Type screen, select VMFS, and click Next.
3. On the Name and device selection screen, enter a datastore name, and select a LUN for provisioning the datastore. Click Next.
4. On the VMFS version screen, select VMFS 6, and click Next.
5. On the Partition configuration screen, use default settings, and click Next.
6. On the Ready to complete screen, click Finish.

Deploying a VM in a VMware vSphere environment on the HPE Primera A670

1. On vCenter UI, expand the Data Center.
2. Right-click the ESXi host, and select New Virtual Machine.
3. Select Create a new virtual machine, and click Next.
4. Enter a virtual machine name, and click Next.
5. On Select Storage screen, select the newly created VMFS datastore, and click Next.
6. On compatibility screen, select ESXi 6.7 and later, and click Next.
7. Select appropriate OS details, and click Next.
8. On Customize hardware screen, click Next.
9. On Ready to complete screen, click Finish.

Dell EMC PowerStore 7000T vs. HPE Primera A670

Accessing snapshot restore data

This scenario assumed that both arrays provided storage resources to a single host and that the administrator had deployed 10 100GB volumes and mounted them as RDMs onto a single CentOS application VM on each testbed. Prior to the baseline manual snapshot, we filled all RDM volumes to 60 percent capacity using Vdbench with 256KB sequential writes. After executing the baseline snapshot, we ran a mixed I/O workload for one hour before attempting a restore from the snapshot procedure on 10 RDMs on each testbed.

Prefilling RDM data on the PowerStore 7000T and HPE Primera A670

For these tests, we used Vdbench to prefill RDM data:

```
compratio=2
dedupratio=2
dedupunit=16k (4k for PowerStore array)
dedupsets=5%
messagescan=no
hd=default, shell=ssh, user=root, jvms=1
hd=hd1, system=CB_command
```

```

sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
sd=sd3,hd=hd1,lun=/dev/sdd
sd=sd4,hd=hd1,lun=/dev/sde
sd=sd5,hd=hd1,lun=/dev/sdf
sd=sd6,hd=hd1,lun=/dev/sdg
sd=sd7,hd=hd1,lun=/dev/sdh
sd=sd8,hd=hd1,lun=/dev/sdi
sd=sd9,hd=hd1,lun=/dev/sdj
sd=sd10,hd=hd1,lun=/dev/sdk

wd=default,sd=*
wd=wd_prefill,sd=sd*,xfersize=256k,range=(0,60),seekpct=eof,rdpct=0

rd=default
rd=rd_prefill,wd=wd_prefill,elapsed=1h,interval=10,iorate=max,forthreads=(1)

```

Taking a manual snapshot of RDM volumes on the Dell EMC PowerStore 7000T

1. In the PowerStore Manager UI, choose Storage→Volume Groups.
2. Select the name of the volume group containing the RDM volumes.
3. On the volume group properties screen, click the Protection tab.
4. Under the Protection tab, click Take Snapshot.
5. On the Take Snapshot screen, keep the defaults, and click Take Snapshot.

Taking a manual snapshot of RDM volumes on the HPE Primera A670

1. In the Primera and 3PAR StoreServ Management Console (SSMC), under the Primera and 3PAR SSMC dropdown, click Show All.
2. Click Virtual Volume Sets.
3. Select the Virtual Volume Set.
4. Click the Actions button, and click Create Snapshot.
5. Keep the defaults, and click Create.

Running steady state I/O to RDM volumes on the PowerStore 7000T and HPE Primera A670

For these tests, we used Vdbench to run steady state I/O to RDM volumes:

```

compratio=2
dedupratio=2
dedupunit=16k (4k for PowerStore array)
dedupsets=5%
messagescan=no
hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=CB_command

sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sdb
sd=sd2,hd=hd1,lun=/dev/sdc
sd=sd3,hd=hd1,lun=/dev/sdd
sd=sd4,hd=hd1,lun=/dev/sde
sd=sd5,hd=hd1,lun=/dev/sdf
sd=sd6,hd=hd1,lun=/dev/sdg
sd=sd7,hd=hd1,lun=/dev/sdh
sd=sd8,hd=hd1,lun=/dev/sdi
sd=sd9,hd=hd1,lun=/dev/sdj
sd=sd10,hd=hd1,lun=/dev/sdk

wd=default,sd=*
wd=wd_OLTP2A_RRH,sd=*,rhpct=100,rdpct=100,xfersize=8K,skew=20,range=(0,6m)
wd=wd_OLTP2A_RM,sd=*,rdpct=100,xfersize=8k,skew=45,range=(0,30)
wd=wd_OLTP2A_RW,sd=*,rdpct=0,xfersize=8K,skew=15,range=(0,30)
wd=wd_OLTP2A_SR,sd=*,rdpct=100,seekpct=seqnz,range=(0,30),xfersize=128K,skew=10
wd=wd_OLTP2A_SW,sd=*,rdpct=0,seekpct=seqnz,range=(0,30),xfersize=128K,skew=10

rd=default
rd=rd_OLTP2A,wd=wd_OLTP2A_*,iorate=max,elapsed=1h,interval=10,warmup=30,forthreads=(1)

```


Quiescing the CentOS application VM on the Dell EMC PowerStore 7000T and HPE Primera A670

1. Open a browser, and navigate to the vCenter vSphere client.
2. Enter credentials, and click Login.
3. Expand the data center.
4. Expand the cluster.
5. Right-click the VM, and choose Power→Power Off.

Performing a snapshot restore on the Dell EMC PowerStore 7000T

1. Open a browser, and navigate to the PowerStore Manager UI.
2. Enter credentials, and click Login.
3. Click Storage→Volume Groups.
4. Click the name of the volume group containing the RDM volumes.
5. On the Volume group properties screen, click the Protection tab.
6. Check the snapshot timestamp of the snapshot you will restore.
7. Click on More Actions→Restore Using Snapshot.
8. On the Restore Using Snapshot screen, uncheck Take a backup snapshot..., and click Restore.
9. Confirm the restore process by clicking Restore.
10. On the volume group's Protection tab, review the snapshot restore success message.
11. In the vCenter UI, right-click the CentOS application VM and choose Power→Power On.

The HPE Primera A670 cannot promote snapshots if the user exports the target virtual volumes to a host (see https://support.hpe.com/hpesc/public/docDisplay?docLocale=en_US&docId=emr_na-a00085217en_us). The HPE Primera A670 requires the user complete the following steps to restore a snapshot successfully:

Removing the LUNs from the ESXi environment on the HPE Primera A670

1. In the vCenter UI, right-click the VM, and choose Edit Settings.
2. Expand Hard Disks.
3. Remove each RDM from the VM, choose to delete files from the datastore, and click OK.

Unexporting the Virtual Volume Set from the ESXi host on the HPE Primera A670

1. In Primera and 3PAR StoreServ Management Console (SSMC), under the Primera and 3PAR SSMC dropdown, click Show All.
2. Click Virtual Volume Sets.
3. Select the virtual volume set containing the base volumes.
4. Click Actions, and choose Unexport.
5. On the Unexport screen, select the volume set, and click Unexport.
6. Check the I have read and understand... checkbox, and click Yes, unexport.

Performing a snapshot restore on the HPE Primera A670

1. In Primera and 3PAR SSMC, select the virtual volume set containing the snapshot volumes.
2. Click Actions, and click Promote snapshot.
3. On the promote screen, select the Time Stamp of the snapshots you are restoring, and click Promote.
4. To confirm the snapshot promotion, click Yes, promote.

Exporting LUNs back to the ESXi environment on the HPE Primera A670

1. In Primera and 3PAR SSMC, select the virtual volume set containing the base volumes.
2. Click Actions, and choose Export.
3. On the Export screen, in the Export To section, click Add.
4. On the Add screen, choose the ESXi host.
5. On the Export screen, click Export.

Adding LUNs back to the virtualized ESXi environment on the HPE Primera A670

1. In the vCenter UI, if necessary, expand the data center and cluster.
2. To discover newly added storage, right-click the target host, and select Storage→Rescan Storage.

Adding LUNs back to the VM as RDM volumes on the HPE Primera A670

1. In the vCenter UI, right-click the VM, and choose Edit Settings.
2. Click Add New Device, and click RDM Disk.
3. Select one of the 100GB LUNS, and click OK.
4. Repeat step 2 until you have added all 10 LUNs as RDMs.
5. On Edit Settings screen, click OK.
6. Right-click the CentOS application VM, and choose Power→Power On.

Read the report at <http://facts.pt/nb98lrd> ▶

This project was commissioned by Dell EMC.



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.