The science behind the report:

Scale up your storage with higher-performing Dell APEX Block Storage for AWS

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 Scale up your storage with higher-performing Dell APEX Block Storage for AWS.

We concluded our hands-on testing on January 18, 2024. 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 December 13, 2023 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to http://facts.pt/calculating-and-highlighting-wins. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: The raw totals and per-node IOPS and throughput data from our Vdbench testing with 10 runs of each IO profile using the EBS-backed configuration of both solutions. Higher is better.

	Total		Per node	
EBS configuration	**Dell APEX Block Storage for AWS**	**Vendor A**	**Dell APEX Block Storage for AWS**	**Vendor A**
Random read IOPS over 10 runs at 128 threads				
Run 1	119,261.50	43,565.30	39,753.83	21,782.65
Run 2	119,227.50	43,589.70	39,742.50	21,794.85
Run 3	119,411.30	43,579.30	39,803.76	21,789.65
Run 4	119,342.10	43,481.20	39,780.70	21,740.60
Run 5	119,278.70	43,478.50	39,759.56	21,739.25
Run 6	119,255.50	43,418.30	39,751.83	21,709.15
Run 7	119,249.00	43,438.50	39,749.66	21,719.25
Run 8	119,260.00	43,379.10	39,753.33	21,689.55
Run 9	119,243.20	43,378.10	39,747.73	21,689.05
Run 10	119,232.40	43,308.00	39,744.13	21,654.00

	Total		Per node	
EBS configuration	**Dell APEX Block Storage for AWS**	**Vendor A**	**Dell APEX Block Storage for AWS**	**Vendor A**
Random write IOPS over 10 runs at 128 threads				
Run 1	57,090.10	49,001.80	19,030.03	24,500.90
Run 2	57,552.10	49,135.80	19,184.03	24,567.90
Run 3	57,568.90	48,623.90	19,189.63	24,311.95
Run 4	57,289.10	48,580.30	19,096.36	24,290.15
Run 5	57,371.20	49,592.30	19,123.73	24,796.15
Run 6	57,410.00	49,221.70	19,136.66	24,610.85
Run 7	57,566.20	48,543.80	19,188.73	24,271.90
Run 8	57,535.10	49,045.80	19,178.36	24,522.90
Run 9	57,558.90	47,786.70	19,186.30	23,893.35
Run 10	57,533.30	48,905.50	19,177.76	24,452.75
Sequential read throughput (MB/s) over 10 runs at 24 threads				
Run 1	3,306.55	1,983.61	1,653.27	991.80
Run 2	3,309.77	1,527.81	1,103.25	763.90
Run 3	3,305.86	1,354.98	1,101.95	677.49
Run 4	3,311.63	1,188.61	1,103.87	594.30
Run 5	3,306.60	1,109.11	1,102.20	554.55
Run 6	3,307.93	1,041.28	1,102.64	520.64
Run 7	3,302.82	971.63	1,100.94	485.81
Run 8	3,301.13	1,052.18	1,100.37	526.09
Run 9	3,299.41	839.75	1,099.80	419.87
Run 10	3,296.02	959.61	1,098.67	479.80
Sequential write throughput (MB/s) over 10 runs at 32 threads				
Run 1	1,622.95	991.37	540.98	495.68
Run 2	1,622.62	987.02	540.87	493.51
Run 3	1,622.33	981.85	540.77	490.92
Run 4	1,621.60	983.22	540.53	491.61
Run 5	1,619.71	975.13	539.90	487.56
Run 6	1,622.26	982.36	540.75	491.18
Run 7	1,622.63	969.62	540.87	484.81
Run 8	1,623.05	972.11	541.01	486.05
Run 9	1,620.00	962.19	540.00	481.09
Run 10	1,619.94	949.56	539.98	474.78

	Total		Per node	
EBS configuration	**Dell APEX Block Storage for AWS**	**Vendor A**	**Dell APEX Block Storage for AWS**	**Vendor A**
OLTP2 IOPS over 10 runs at 128 threads				
Run 1	51,079.50	50,804.80	17,026.50	25,402.40
Run 2	51,195.40	49,305.60	17,065.13	24,652.80
Run 3	50,586.00	49,623.50	16,862.00	24,811.75
Run 4	51,141.10	47,754.50	17,047.03	23,877.25
Run 5	48,642.20	47,395.00	16,214.06	23,697.50
Run 6	49,234.90	48,807.60	16,411.63	24,403.80
Run 7	50,642.60	49,119.30	16,880.86	24,559.65
Run 8	50,363.20	47,121.70	16,787.73	23,560.85
Run 9	50,020.10	45,669.50	16,673.36	22,834.75
Run 10	49,842.60	46,079.00	16,614.20	23,039.50

Table 2: The raw totals and per-node IOPS and throughput data from our Vdbench testing with 10 runs of each IO profile using the NVMe-supported configuration of both solutions. Higher is better.

	Total		Per node	
NVME-supported configuration	**Dell APEX Block Storage for AWS**	**Vendor A**	**Dell APEX Block Storage for AWS**	**Vendor A**
Random read IOPS over 10 runs at 128 threads				
Run 1	716,836.80	179,967.30	238,945.60	89,983.65
Run 2	714,792.50	164,035.20	238,264.16	82,017.60
Run 3	715,149.80	145,854.70	238,383.26	72,927.35
Run 4	713,904.70	133,650.50	237,968.23	66,825.25
Run 5	714,816.50	119,450.60	238,272.16	59,725.30
Run 6	712,507.30	110,250.60	237,502.43	55,125.30
Run 7	715,281.80	105,054.90	238,427.26	52,527.45
Run 8	712,219.30	101,469.10	237,406.43	50,734.55
Run 9	710,125.60	99,474.20	236,708.53	49,737.10
Run 10	716,445.40	98,245.00	238,815.13	49,122.50

	Total		Per node	
NVME-supported configuration	**Dell APEX Block Storage for AWS**	**Vendor A**	**Dell APEX Block Storage for AWS**	**Vendor A**
Random write IOPS over 10 runs at 96 threads				
Run 1	219,206.90	51,769.60	73,068.96	25,884.80
Run 2	219,332.70	52,220.10	73,110.90	26,110.05
Run 3	218,497.90	51,914.70	72,832.63	25,957.35
Run 4	218,892.70	51,577.00	72,964.23	25,788.50
Run 5	217,464.00	52,144.30	72,488.00	26,072.15
Run 6	218,857.10	50,831.00	72,952.36	25,415.50
Run 7	218,925.50	52,298.00	72,975.16	26,149.00
Run 8	217,675.90	51,005.20	72,558.63	25,502.60
Run 9	218,425.10	52,009.70	72,808.36	26,004.85
Run 10	218,594.50	52,013.30	72,864.83	26,006.65
Sequential read throughput (MB/s) over 10 runs at 32 threads				
Run 1	15,854.63	5,006.65	5,284.87	2,503.32
Run 2	15,667.77	4,658.13	5,222.59	2,329.06
Run 3	15,758.79	3,068.13	5,252.93	1,534.06
Run 4	15,629.31	2,497.25	5,209.77	1,248.62
Run 5	15,815.93	2,469.66	5,271.97	1,234.83
Run 6	15,753.10	2,260.57	5,251.03	1,130.28
Run 7	15,607.56	2,031.84	5,202.52	1,015.92
Run 8	15,735.47	2,055.17	5,245.15	1,027.58
Run 9	15,732.21	1,860.64	5,244.07	930.32
Run 10	15,345.17	2,018.69	5,115.05	1,009.34
Sequential write throughput (MB/s) over 10 runs at 24 threads				
Run 1	4,220.92	1,383.61	1,406.97	691.80
Run 2	4,223.72	1,365.82	1,407.90	682.91
Run 3	4,231.78	1,358.78	1,410.59	679.39
Run 4	4,234.40	1,370.53	1,411.46	685.26
Run 5	4,205.40	1,374.65	1,401.80	687.32
Run 6	4,222.04	1,358.55	1,407.34	679.27
Run 7	4,342.15	1,367.28	1,447.38	683.64
Run 8	4,303.89	1,319.25	1,434.63	659.62
Run 9	4,130.19	1,239.93	1,376.73	619.96
Run 10	4,220.32	1,310.99	1,406.77	655.49

	Total		Per node	
NVME-supported configuration	**Dell APEX Block Storage for AWS**	**Vendor A**	**Dell APEX Block Storage for AWS**	**Vendor A**
OLTP2 IOPS over 10 runs at 128 threads				
Run 1	301,905.50	87,719.70	100,635.16	43,859.85
Run 2	301,732.00	87,802.60	100,577.33	43,901.30
Run 3	300,776.90	89,943.20	100,258.96	44,971.60
Run 4	300,200.60	87,879.10	100,066.86	43,939.55
Run 5	301,336.30	87,129.50	100,445.43	43,564.75
Run 6	299,218.50	86,972.50	99,739.50	43,486.25
Run 7	298,048.00	86,776.40	99,349.33	43,388.20
Run 8	301,530.70	81,467.90	100,510.23	40,733.95
Run 9	302,080.70	79,306.80	100,693.56	39,653.40
Run 10	300,734.30	80,102.80	100,244.76	40,051.40

Table 3: The median IOPS and throughput results per node from runs 8 through 10 using the EBS-backed configuration of both solutions. Higher is better.

	Dell APEX Block Storage for AWS (per node)		Vendor A (per node)		
	IOPS	**MB/s**	**IOPS**	**MB/s**	**% increase**
Random read	39,747.73	-	21,689.05	-	83.26%
Random write	19,178.36	-	24,452.75	-	-27.50%
Sequential read	-	1,099.80	-	479.80	129.22%
Sequential write	-	540.00	-	482.42	11.94%
OLTP2	16,673.36	-	23,039.50	-	-38.18%

Table 4: The median per-node IOPS and throughput results per node from runs 8 through 10 using the NVMe-supported configuration of both solutions. Higher is better.

	Dell APEX Block Storage for AWS (per node)		Vendor A (per node)		
	IOPS	**MB/s**	**IOPS**	**MB/s**	**% increase**
Random read	237,406.43	-	49,737.10	-	377.32%
Random write	72,808.36	-	26,004.85	-	179.98%
Sequential read	-	5,244.07	-	1,009.34	419.55%
Sequential write	-	1,406.77	-	655.49	114.61%
OLTP2	100,510.23	-	40,051.40	-	150.95%

Table 5: The median per-node sequential read throughput results from the Dell APEX Block Storage for AWS scale testing.

Sequential reads 32 threads	Sequential reads - MB/s total	Sequential reads - MB/s per node	% increase over 3 nodes
3	15,511.91	5,170.63	-
12	69,537.06	5,794.75	12.07%
24	140,073.88	5,836.41	12.88%

How we tested

We conducted all testing using Dell-controlled cloud accounts and resources. We did not perform the setup of the solutions, but we worked closely with a Dell engineer to observe the setup and verified the configurations were fair. Dell engineers used default settings for both environments, with one exception: They disabled storage efficiency for the Vendor A solution and created a second storage pool with ownership assigned to the second node to allow it to be active-active at the node level (but not at the volume level). While engineers deployed the Dell APEX Block Storage for AWS solution manually, they used the same settings and configurations as if the instance were deployed using APEX Navigator.

Engineers set up both environments for high availability, distributing Dell APEX Block Storage for AWS across three availability zones, and Vendor A across two availability zones.

After setup, Dell gave us control of the instances for Vdbench testing. Because we did not perform the setup, this methodology provides only instance and storage type information, as well as the scripts we used to run Vdbench against the solution.

We tested Dell APEX Block Storage for AWS against a Vendor A solution, using AWS infrastructure in multiple configurations to measure relative performance using EBS gp3 and NVMe storage. We also recorded performance for APEX Block Storage for AWS scalability for up to 24 nodes. We ran a set of Vdbench workloads on both platforms using gp3 EBS storage with a two-node Vendor A configuration (the maximum node count for high availability (HA)) and a three-node APEX Block Storage deployment (the minimum node count available) using gp3 EBS storage. We then repeated these tests leveraging NVMe as storage for APEX and as a dedicated read cache for Vendor A.

Dell APEX Block Storage for AWS: EBS configuration

	Specifications	Comments
Controller host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network / 9,500 MB/s EBS	1 per availability zone
Storage	1.5TB gp3 14,000 IOPS / 125 MB/s	10 per controller host
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s Network	4 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Vendor A solution: EBS configuration

	Specifications	Comments
Controller host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network / 9,500 MB/s EBS	1 per availability zone
Storage	3.65TB gp3 16,000 IOPS / 1,000 MB/s	4 per aggregate + 4 for mirroring
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network	4 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Dell APEX Block Storage for AWS: NVMe configuration

	Specifications	Comments
Controller host	i3en.12xlarge 48 vCPU / 384 GB RAM / 50 GB/s network / 9,500 MB/s EBS	1 per availability zone
NVMe storage	7,500 GB	4 per controller
EBS storage	N/A	N/A
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network	4 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Vendor A solution: NVMe-supported configuration (NVMe/EBS)

	Specifications	Comments
Controller host	m5dn.24xlarge 96 vCPU / 384 GB RAM / 100 GB/s network / 19,000 MB/s EBS	1 per availability zone
NVMe storage	900 GB	4 per controller (used as extended read cache)
EBS storage	3.65 TB gp3 16,000 IOPS / 1,000 MB/s	4 per aggregate + 4 for mirroring
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network	4 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Dell APEX Block Storage for AWS: Scalability configuration - three-node NVMe storage

	Specifications	Comments
Controller host	i3en.12xlarge 48 vCPU / 384 GB RAM / 50 GB/s network / 9,500 MB/s EBS	1 per availability zone
NVMe storage	7,500 GB	4 per controller
EBS storage	N/A	N/A
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network	2 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Dell APEX Block Storage for AWS: Scalability configuration - 12-node NVMe storage

	Specifications	Comments
Controller host	i3en.12xlarge 48 vCPU / 384 GB RAM / 50 GB/s network / 9,500 MB/s EBS	4 per availability zone
NVMe storage	7,500 GB	4 per controller
EBS storage	N/A	N/A
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network	2 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Dell APEX Block Storage for AWS: Scalability configuration - 24-node NVMe storage

	Specifications	Comments
Controller host	i3en.12xlarge 48 vCPU / 384 GB RAM / 50 GB/s network / 9,500 MB/s EBS	8 per availability zone
NVMe storage	7,500 GB	4 per controller
EBS storage	N/A	N/A
Application host	c5n.9xlarge 36 vCPU / 96 GB RAM / 50 GB/s network	2 application hosts per controller node
Application host volumes	500 GiB	2 per host, no data reduction

Performance testing using Vdbench

We ran the following tests on two Dell APEX Block Storage for AWS configurations (EBS and NVMe) and two Vendor A configurations (EBS and NVMe/EBS).

Test 0: Prefill

Run a sequential write workload across all volumes to fully allocate them before testing:

```
/home/ec2-user/vdbench # ./vdbench -f tests/prefill -o tests/results/prefill
```

Test 1: Random read workload

100% random read workload across 100% of each volume – 4 KB IO size.

Measure IOPS and response time for gradually increasing thread counts.

```
./vdbench -f tests/random_read -o tests/results/random_read
```

Test 2: Random write workload

100% random write workload across 100% of each volume – 4 KB IO size.

Measure IOPS and response time for gradually increasing thread counts.

```
./vdbench -f tests/random_write -o tests/results/random_write
```

Test 3: Sequential read workload

100% sequential read workload across 100% of each volume – 256 KB IO size.

Measure bandwidth (MB/s) for gradually increasing thread counts.

```
./vdbench -f tests/seq_read -o tests/results/seq_read
```

Test 4: Sequential write workload

100% sequential write workload across 100% of each volume – 256 KB IO size.

Measure bandwidth (MB/s) for gradually increasing thread counts.

```
./vdbench -f tests/seq_write -o tests/results/seq_write
```

Test5: OLTP2-type workload

OLTP2-type workload across 100% of each volume. This workload is a combination of 8KB random read, 8KB read hits, 8KB random writes, 64KB sequential reads, and 64KB sequential writes.

Measure IOPS and response time for gradually increasing thread counts.

```
./vdbench -f tests/OLTP2 -o tests/results/OLTP2
```

Performance scaling with Dell APEX Block Storage for AWS

Test 1.0: Prefill

Run sequential write workload across all volumes to fully allocate them.

```
./vdbench -f tests/prefill_3node -o tests/results/prefill_3node
```

Test 1.1: Random read workload on three-node cluster

100% random read workload across 100% of each volume – 4 KB IO size.

Measure IOPS.

```
./vdbench -f tests/random_read_3node -o tests/results/random_read_3node
```

Test 1.2: Sequential read workload on three-node cluster

100% sequential read workload across 100% of each volume – 256 KB IO size

Measure MB/s.

```
./vdbench -f tests/seq_read_3node -o tests/results/seq_read_3node
```

Test 2.0: Prefill

Run sequential write workload across all volumes to fully allocate them.

```
./vdbench -f tests/prefill_12node -o tests/results/prefill_12node
```

Test 2.1: Random read workload on 12-node cluster

100% random read workload across 100% of each volume – 4 KB IO size.

Measure IOPS.

```
./vdbench -f tests/random_read_12node -o tests/results/random_read_12node
```

Test 2.2: Sequential read workload on 12-node cluster

100% sequential read workload across 100% of each volume – 256 KB IO size

Measure MB/s.

```
./vdbench -f tests/seq_read_12node -o tests/results/seq_read_12node
```

Test 3.0: Prefill

Run sequential write workload across all volumes to fully allocate them.

```
./vdbench -f tests/prefill_24node -o tests/results/prefill_24node
```

Test 3.1: Random read workload on 24-node cluster

100% random read workload across 100% of each volume – 4 KB IO size.

Measure IOPS.

```
./vdbench -f tests/random_read_24node -o tests/results/random_read_24node
```

Test 3.2: Sequential read workload on 24-node cluster

100% Sequential Read workload across 100% of each volume – 256 KB IO size

Measure MB/s

```
./vdbench -f tests/seq_read_24node -o tests/results/seq_read_24node
```

Configuration and scripts

Dell APEX Block Storage for AWS hosts/volumes

```
#compratio=1
#dedupratio=1
#dedupunit=4096
#dedupsets=50%

messagescan=no

hd=default,shell=ssh,user=root,jvms=1
hd=hd1,system=ip-172-31-11-156
hd=hd2,system=ip-172-31-11-117
hd=hd3,system=ip-172-31-12-53
hd=hd4,system=ip-172-31-12-114
hd=hd5,system=ip-172-31-13-79
hd=hd6,system=ip-172-31-13-243
hd=hd7,system=ip-172-31-11-168
hd=hd8,system=ip-172-31-12-101
hd=hd9,system=ip-172-31-13-70
hd=hd10,system=ip-172-31-11-172
hd=hd11,system=ip-172-31-12-145
hd=hd12,system=ip-172-31-13-213

sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/scinia
sd=sd2,hd=hd1,lun=/dev/scinib
sd=sd3,hd=hd2,lun=/dev/scinia
sd=sd4,hd=hd2,lun=/dev/scinib
sd=sd5,hd=hd3,lun=/dev/scinia
sd=sd6,hd=hd3,lun=/dev/scinib
sd=sd7,hd=hd4,lun=/dev/scinia
sd=sd8,hd=hd4,lun=/dev/scinib
sd=sd9,hd=hd5,lun=/dev/scinia
sd=sd10,hd=hd5,lun=/dev/scinib
sd=sd11,hd=hd6,lun=/dev/scinia
sd=sd12,hd=hd6,lun=/dev/scinib
sd=sd13,hd=hd7,lun=/dev/scinia
sd=sd14,hd=hd7,lun=/dev/scinib
sd=sd15,hd=hd8,lun=/dev/scinia
sd=sd16,hd=hd8,lun=/dev/scinib
sd=sd17,hd=hd9,lun=/dev/scinia
sd=sd18,hd=hd9,lun=/dev/scinib
sd=sd19,hd=hd10,lun=/dev/scinia
sd=sd20,hd=hd10,lun=/dev/scinib
sd=sd21,hd=hd11,lun=/dev/scinia
sd=sd22,hd=hd11,lun=/dev/scinib
sd=sd23,hd=hd12,lun=/dev/scinia
sd=sd24,hd=hd12,lun=/dev/scinib
```

Vendor A solution hosts/volumes

```
#compratio=1
#dedupratio=1
#dedupunit=4096
#dedupsets=50%

messagescan=no

hd=default,shell=ssh,user=root,jvms=2
hd=hd1,system=ip-172-31-19-40
hd=hd2,system=ip-172-31-18-173
hd=hd3,system=ip-172-31-42-65
hd=hd4,system=ip-172-31-43-35
hd=hd5,system=ip-172-31-31-127
hd=hd6,system=ip-172-31-38-125
hd=hd7,system=ip-172-31-23-23
hd=hd8,system=ip-172-31-40-130

sd=default,openflags=o_direct
sd=sd1,hd=hd1,lun=/dev/sda
sd=sd2,hd=hd1,lun=/dev/sdb
sd=sd3,hd=hd2,lun=/dev/sda
sd=sd4,hd=hd2,lun=/dev/sdb
sd=sd5,hd=hd3,lun=/dev/sda
sd=sd6,hd=hd3,lun=/dev/sdb
sd=sd7,hd=hd4,lun=/dev/sda
sd=sd8,hd=hd4,lun=/dev/sdb
sd=sd9,hd=hd5,lun=/dev/sda
sd=sd10,hd=hd5,lun=/dev/sdb
sd=sd11,hd=hd6,lun=/dev/sda
sd=sd12,hd=hd6,lun=/dev/sdb
sd=sd13,hd=hd7,lun=/dev/sda
sd=sd14,hd=hd7,lun=/dev/sdb
sd=sd15,hd=hd8,lun=/dev/sda
sd=sd16,hd=hd8,lun=/dev/sdb
```

Prefill

```
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)
```

4KB random read

```
wd=default,sd=*
wd=wd_4k,sd=sd*,xfersize=4k,seekpct=100
rd=default
rd=read4k_test,wd=wd_4k,iorate=max,warmup=30,interval=10,forrdpct=(100),elapsed=120,forthrea
ds=(1,2,4,8,16,24,32,64,96,128)
```

4KB random write

```
wd=default,sd=*
wd=wd_4k,sd=sd*,xfersize=4k,seekpct=100

rd=default
rd=read4k_test,wd=wd_4k,iorate=max,warmup=30,interval=10,forrdpct=(0),elapsed=120,forthrea
ds=(1,2,4,8,16,24,32,64,96,128)
```

256 KB sequential read

```
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,forthrea
ds=(1,2,4,8,16,24,32)
```

256 KB sequential write

```
wd=wd_256k,sd=sd*,xfersize=256k,seekpct=seqnz

rd=default
rd=read256k_test,wd=wd_256k,iorate=max,warmup=30,interval=10,forrdpct=(0),elapsed=120,forthrea
ds=(1,2,4,8,16,24,32)
```

OLTP2

```
wd=wd_OLTP2_RRH,sd=*,rhpct=100,rdpct=100,xfersize=8K,skew=20,range=(10m,30m)
wd=wd_OLTP2_RM1,sd=*,rdpct=100,xfersize=8k,skew=40,range=(73,80)
wd=wd_OLTP2_RM2,sd=*,rdpct=100,xfersize=8k,skew=5,range=(11,72)
wd=wd_OLTP2_RW1,sd=*,rdpct=0,xfersize=8K,skew=13,range=(73,80)
wd=wd_OLTP2_RW2,sd=*,rdpct=0,xfersize=8K,skew=2,range=(11,72)
wd=wd_OLTP2_SR1,sd=*,rdpct=100,seekpct=seqnz,range=(73,80),xfersize=64K,skew=9
wd=wd_OLTP2_SR2,sd=*,rdpct=100,seekpct=seqnz,range=(11,72),xfersize=64K,skew=1
wd=wd_OLTP2_SW1,sd=*,rdpct=0,seekpct=seqnz,range=(73,80),xfersize=64K,skew=9wd=wd_OLTP2_SW2,sd=*,rdpct=0
,seekpct=seqnz,range=(11,72),xfersize=64K,skew=1
rd=default
rd=rd_OLTP2STS,wd=wd_OLTP2_*,iorate=max,elapsed=120,interval=10,warmup=30,forthrea
ds=(1,2,4,8,16,24,32,64,96,128)
```

Read the report at https://facts.pt/bWT2Lbu ▶

This project was commissioned by Dell Technologies.

Principled 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.