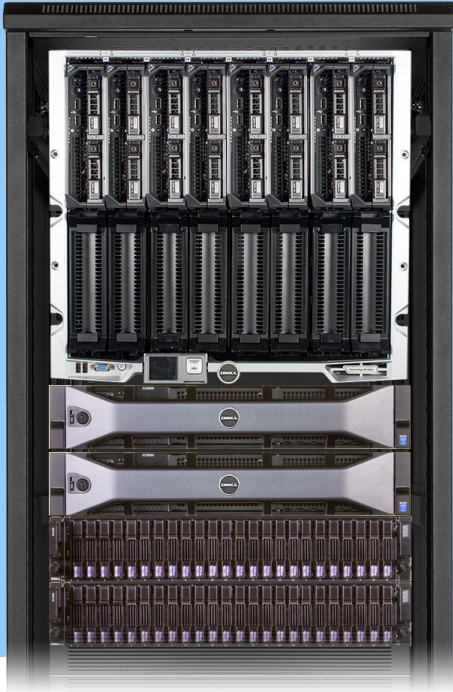



DATABASE PERFORMANCE IN BLADE ENVIRONMENTS: DELL POWEREDGE M1000E VS CISCO UCS 5108

Deploy easily and gain throughput with the 16Gb-enabled Dell™ PowerEdge™ M1000e and Storage SC9000 All Flash Array



Up to **91.7%** more storage throughput for better database performance*

Up to **36.9%** lower response time*

 ASM simplified deployment

*versus a comparably configured Cisco UCS® 5108 blade environment

When selecting a blade environment, strong performance is essential: It must support the increasing demands of your database applications and its users. When a blade environment also streamlines the task of deploying new blades and makes efficient use of precious datacenter rack space, it becomes even more attractive.

In the Principled Technologies datacenter, we set up two VMware vSphere 6 environments: one on a Dell PowerEdge M1000e enclosure with Brocade M6505 16Gb Fibre Channel switches and Dell PowerEdge M630 servers powered by Intel® Xeon® processors E5-2690 v3 and Qlogic QLE2662 16Gb Fibre Channel mezzanine cards and the other on a Cisco UCS 5108 enclosure and Cisco UCS B200 M4 servers, both using a Dell Storage SC9000 solution as a SAN target. We then measured the total storage IOPS, storage throughput, and latency of the two environments. As we scaled from one to eight blades, average performance per blade remained consistent in the 16Gb-enabled Dell environment, while it deteriorated in the Cisco environment. Response time, an important metric for gauging end-user experience, increased at a greater rate in the Cisco environment.

Thanks to Dell Active System Manager (ASM), deploying a PowerEdge M630 blade with a Dell Storage SC9000 solution can be automated, potentially saving your business valuable administration time. The design of the Dell chassis also makes for



greater space efficiency when scaling up in the future—the solution we tested had room for eight additional blades while the Cisco solution was already at capacity.

You might expect the better-performing, denser, and easier-to-manage Dell solution to come at a higher cost, but the prices for the Dell and Cisco environments are nearly equal (see [Appendix B](#) for pricing information). This means that with the Dell environment your business will get better value for your purchase price.

BETTER PERFORMANCE AND EASIER DEPLOYMENT

In our testing, we compared the database performance of the following similarly priced solutions in the datacenter environment:

- **A Dell PowerEdge M1000e blade environment** with Dell PowerEdge M630 servers powered by the Intel Xeon processor E5-2690 v3 and 16Gb Fibre Channel switches and mezzanine cards
- **A Cisco UCS 5108 blade environment** with up to eight Cisco UCS B200 M4 blades

Both solutions used the 16Gb-enabled Dell Storage SC9000 Array Controller, and two SC420 12Gb SAS Enclosures All Flash - SAN and VMware vSphere 6.

We first attached the Dell PowerEdge M1000e chassis to the Dell Storage SC9000 Array Controller and SC420 all-flash array enclosures. We then gathered baseline performance by installing a single Dell M630 blade with vSphere 6, with one VM running a Silly Little Oracle Benchmark (SLOB) 2.2 OLTP database workload tuned to a 70/30 percent read/write IO profile. Once we established this baseline, we scaled to two blades, then four blades, and finally to eight blades. At each blade count, we measured total input/output per second (IOPS) and SAN throughput.

We then removed the M1000e from the storage and attached the Cisco UCS 5108 solution to the storage via Cisco UCS 6248 Fabric Interconnects. We installed the Cisco B200 M4 blades and repeated the database testing on the Cisco UCS 5108 blade environment.¹ Figure 1 shows how we connected our two blade environments.

¹ For more information on the Dell hardware and the SLOB benchmark, see [Appendix A](#). For detailed configuration and pricing information, see [Appendix B](#). For details on how we tested, see [Appendix C](#).


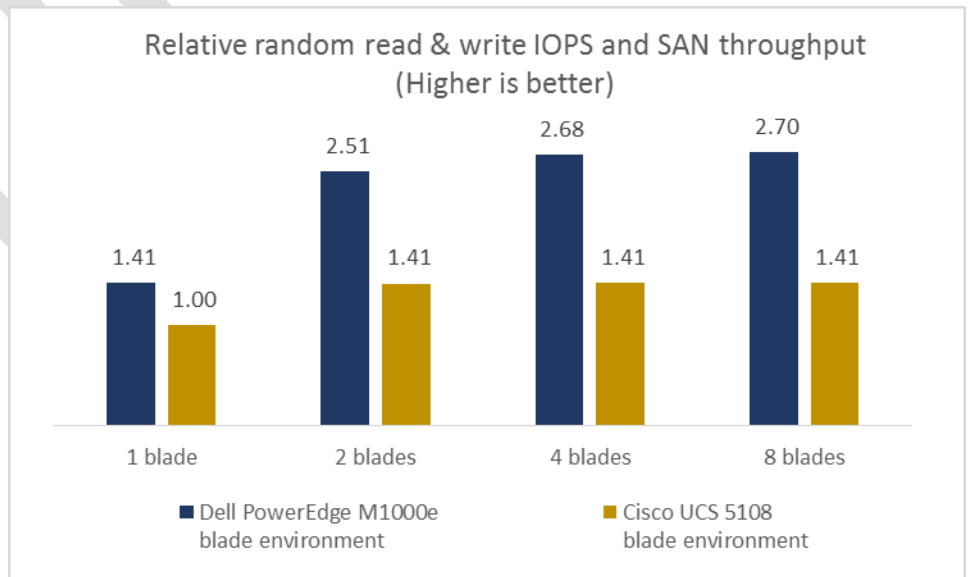


Figure 1: How we connected the two blade environments that we tested.

Greater IOPS and SAN throughput

Database workloads can demand large amounts of disk throughput to achieve solid performance. We ran a 70/30 percent read/write workload and compared total input/output per second (IOPS) with each blade count for our two test environments. Figure 2 shows how IOPS scaled for the two blade environments. The Dell PowerEdge M1000e blade environment delivered over 250,000 IOPS, 91.7 percent greater total IOPS than the Cisco UCS 5108 blade environment and we observed no major performance degradation on any single VM. Note that we normalized the scores to that of the lowest-performing configuration.

Figure 2: In the IOPS and SAN throughput tests, the Dell PowerEdge M1000e blade environment delivered better performance per blade with each blade addition.




SAN throughput is an additional metric for determining performance scalability in a given environment. Because throughput is a function of IOPS, normalized charts for the two metrics are identical even though the values themselves differ. For that reason, a normalized chart depicting the SAN throughput of the two solutions is identical to a normalized chart depicting IOPS. Just as the Dell PowerEdge M1000e blade environment delivered 91.7 percent greater IOPS than the Cisco UCS 5108 blade environment with eight blades, it also delivered 91.7 percent greater throughput.

Faster response time

Throughout the duration of the test, the 16Gb-enabled Dell PowerEdge M1000e blade environment read and write response time increased only slightly as we increased the blade count. In contrast, the response time of the Cisco UCS 5108 blade environment increased dramatically as we added blades (see Figure 3). Additionally, response time was lower for the Dell PowerEdge M1000e blade environment at each scale point. This means that users of the Dell PowerEdge M1000e blade environment would continue to get acceptable response times as the number of blades scaled. Users of the Cisco blade environment, on the other hand, would experience increasing lags that would require them to wait longer for their data as more blades joined the environment.

Figure 3: With each blade addition, the response time of the Dell PowerEdge M1000e blade environment increased only slightly compared to the Cisco UCS solution.



Streamlined blade deployment

When IT administrators can spend less time and effort deploying a blade, they have more time for other critical tasks. Both Dell and Cisco offer software products to help automate blade deployment: Dell Active System Manager (ASM) and Cisco UCS Director. In a previous comparison using a different hardware configuration, we found that using Dell ASM was faster and required fewer overall steps than UCS Director for deployment setup and design tasks.² This means that with Dell ASM, your business could finish deployment design tasks faster and save valuable administrative time for more pressing work.

More efficient use of datacenter rack space

In addition to maintaining an acceptable level of performance for database users and applications, companies are always looking for ways to maximize physical space in their datacenters. The Dell PowerEdge M1000e blade chassis, which occupies 10U of rack space, can hold up to 16 Dell PowerEdge M630 blades. That means each server actually occupies less than 1U. The 16 Dell PowerEdge M630 blades offer up to 32 Intel Xeon E5-2690 v3 processors with 12 cores each. This kind of density can pack an enormous amount of computing power into a small space.

The Cisco blade environment, which occupies 6U of rack space, holds eight blades. It also requires 2U of rack space to house fabric interconnects, which the PowerEdge M1000e blade chassis does not.³ This means that to match the 16-blade compute capacity of the Dell PowerEdge M1000e blade environment, your business would need two Cisco UCS 5108 chassis. That's 14U of space, or 4U more, in your datacenter to match the full capacity of the Dell PowerEdge M1000e chassis. While our test compared only eight blade servers in each solution, the Cisco UCS solution was already full. This means the moment your business needs to add a ninth blade server, the Cisco UCS solution would require another 6U of datacenter space. In contrast, the Dell PowerEdge M1000e blade environment would stay the same, resulting in 28.6 percent less rack space requirement than the Cisco solution—a factor that can translate to future savings.

² Another Principled Technologies study found deployment advantages of using Dell ASM vs. Cisco UCS Director. For more information, visit www.principledtechnologies.com/Dell/Dell_ASM_simplify_deployment_0415.pdf

³ On the Dell PowerEdge M1000e blade chassis, the fabric interconnects are rear-mounted modules and use no rack space.

CONCLUSION

Scalable and reliable compute performance, in conjunction with your 16Gb Fibre Channel capability with SAN-based storage, means you don't have to worry about your application VMs losing performance as your business grows. In our tests, the Dell PowerEdge M1000e blade environment delivered up to 91.7 percent greater throughput while running a virtualized OLTP database workload than the similarly priced Cisco UCS 5108 blade environment. More importantly, the Dell M1000e blade environment delivered consistent scaling as we added blade servers, in contrast to the performance degradation we observed in the Cisco UCS 5108 blade environment. The Dell M1000e also left room to grow within the 10U, with eight additional half-height slots available for compute nodes in the future, whereas the Cisco UCS 5108 blade environment would require a second enclosure to house additional blades. Finally, both solutions were similarly priced, so there is no additional cost associated with the advantages of the Dell environment. For applications demanding dense compute environments, high throughput, and lower response rates, companies will find the Dell PowerEdge M1000e blade environment better able to scale to meet predictable performance requirements as their business grows.

APPENDIX A – ABOUT THE COMPONENTS

About the Dell PowerEdge M1000e

Dell designed the PowerEdge M1000e blade enclosure to help reduce the cost and complexity of managing blade resources. According to Dell, this enclosure boasts the following:

- Powerful integrated management tools.
- Flexible remote management via the Chassis Management Controller, which lets you control entire blade infrastructure, in a single data center or around the world.
- Outstanding efficiency that leads to increased capacity, lower operating costs, and better performance per watt.
- A flexible and scalable system that lets you expand easily
- Advanced input/output (I/O) functionality
- Simplified product integration
- Enhanced systems management

For more information about the Dell PowerEdge M1000e blade enclosure, visit www.dell.com/us/business/p/poweredge-m1000e/pd

About the Dell PowerEdge M630 server nodes

According to Dell, the PowerEdge M630 blade server “offers flexibility to optimize in-server storage and IO performance, allowing you to tailor your servers to your workloads today and adapt as your needs change over time.”

The optional 12Gb PowerEdge RAID Controller (PERC) doubles the throughput and cache of previous generations and can dramatically boost IO for applications. The M630 supports up to four 1.8-inch SSDs or up to two 2.5-inch HDDs or Express Flash PCIe SSDs. Dell Select Network Adapters with 16Gb Fibre Channel capability offer quick access to the storage resources your server or cloud requires, via either NAS or SAN.

For more information about the Dell PowerEdge M630 blade server, visit www.dell.com/us/business/p/poweredge-m630/pd.

About the Intel Xeon processor E5-2600 v3 product family

According to Intel, the Intel Xeon processor E5-2600 v3 product family “helps IT address the growing demands placed on infrastructure, from supporting business growth to enabling new services faster, delivering new applications in the enterprise, technical computing, communications, storage, and cloud.” It also delivers benefits in performance, power efficiency, virtualization, and security.

The E5-2600 v3 product family has up to 50 percent more cores and cache than processors from the previous generation. Other features include the following:

- Intel Advanced Vector Extensions 2 (AVX2)
- Intel Quick Path Interconnect link
- Up to 18 cores and 36 threads per socket
- Up to 45 MB of last-level cache
- Next-generation DDR4 memory support
- Intel Integrated I/O providing up to 80 PCIe lanes per two-socket server

- Intel AES-NI data encryption/decryption

The Intel Xeon processor E5-2600 v3 product family also uses Intel Intelligent Power technology and Per-core P states to maximize energy efficiency.

Learn more at www.intel.com/content/www/us/en/processors/xeon/xeon-e5-brief.html.

About the Dell SC9000 storage array

Dell designed the SC9000 as an ideal solution for large-scale storage, high-end workloads, and distributed enterprise environments. According to Dell, this latest flagship array offers all the capabilities of the SC Series line in a fully upgraded, more scalable system leveraging Dell's 13G server platform.

- 16Gb Fibre Channel capability provides a massive network pipe for achieving the throughput needed to take full advantage of the all-flash performance.
- Quick, modular expansion to over 3PB raw SAN and/or NAS storage capacity per array, and the ability to be interconnected with other SC Series arrays in larger federated systems under unified management.
- Superior flash and hybrid performance. Comprehensive SSD integration plus a wide range of HDDs and ultra-dense enclosures provide flexible performance.
- Always available storage with automated cost savings. Ideal for private cloud deployments, the SC Series Live Volume feature now offers seamless disaster recovery with transparent auto-failover to fully synchronized backup volumes on another SC series array.
- Live Volume with auto-failover keeps workloads running continuously during unplanned outages and repairs your high-availability environment automatically when an array comes back online. No extra hardware or software is required.
- Application Protection Manager Suite ensures server-side data integrity in Oracle, VMware and Microsoft database environments. Optional FIPS-certified self-encrypting SSDs and HDDs protect against data theft, loss or unauthorized access.
- Built-in efficiencies include new block-level compression on all-flash arrays and pervasive "thin" methods. Patented auto-tiering leverages the best attributes of diverse drive types to optimize your data throughout its lifecycle.
- Redundant, hot-swappable components make the SC9000 a resilient, easy-to-deploy and highly available solution. Add capacity quickly with any of five supported expansion enclosures. Customize I/O connections to multiprotocol SANs (FC, iSCSI and FCoE) and support file storage from the same pool via the optional FS8600 NAS appliance.

Learn more at www.dell.com/us/business/p/storage-sc9000/pd

About the SLOB 2.2 benchmark

The Silly Little Oracle Benchmark (SLOB) can assess Oracle random physical I/O capability on a given platform in preparation for potential OLTP/ERP-style workloads by measuring IOPS capacity. The benchmark helps evaluate performance of server hardware and software, storage system hardware and firmware, and storage networking hardware and firmware.

SLOB contains simple PL/SQL and offers the ability to test the following:

1. Oracle logical read (SGA buffer gets) scaling
2. Physical random single-block reads (db file sequential read)
3. Random single block writes (DBWR flushing capacity)
4. Extreme REDO logging I/O

SLOB is free of application contention yet is an SGA-intensive benchmark. According to SLOB's creator Kevin Closson, SLOB can also offer more than testing IOPS capability such as studying host characteristics via NUMA and processor threading. For more information on SLOB, links to information on version 2.2, and links to download the benchmark, visit kevinclosson.net/2012/02/06/introducing-slob-the-silly-little-oracle-benchmark/.

DRAFT

APPENDIX B – SYSTEM CONFIGURATION AND PRICING INFORMATION

Figures 4 through 7 provide detailed configuration information for the test systems. Figure 8 provides detailed configuration information for the storage array.

| System | Dell PowerEdge M1000e blade enclosure |
|-----------------------------|---------------------------------------|
| Power supplies | |
| Number of power supplies | 6 |
| Vendor and model number | Dell C2700A-S0 |
| Wattage of each (W) | 2,700 |
| Cooling fans | |
| Total number of fan modules | 8 |
| Vendor and model number | Dell YK776 Rev. A00 |
| Chassis firmware | |
| Primary Firmware version | 5.11 |
| Hardware Version | A00 |
| I/O modules | |
| Switch | Dell Networking MXL 10/40GbE |
| Occupied bay | A1, A2 |
| Fibre Channel Switch | Brocade M6505 |
| Occupied bay | B1, B2 |

Figure 4: Configuration information for the Dell PowerEdge M1000e blade enclosure.

| System | Dell PowerEdge M630 server |
|-------------------------------------|----------------------------|
| Enclosure | |
| Blade enclosure | Dell PowerEdge M1000e |
| General | |
| Number of processor packages | 2 |
| Number of cores per processor | 12 |
| Number of hardware threads per core | 2 |
| System power management policy | Default |
| CPU | |
| Vendor | Intel |
| Name | Xeon |
| Model number | E5-2690 v3 |
| Stepping | M1 |
| Socket type | FCLGA2011-3 |
| Core frequency (GHz) | 2.6 |
| Bus frequency (GHz) | 9.6 |
| L1 cache | 768 KB |
| L2 cache | 3,072 KB |
| L3 cache | 30 MB |

| System | Dell PowerEdge M630 server |
|-----------------------------------|--|
| Platform | |
| Vendor and model number | Dell PowerEdge M630 |
| Motherboard model number | OJXJPT |
| BIOS version | 1.4.5 |
| Mezzanine card | Dell QLE 2662 Dual Port Fibre Channel HBA |
| Integrated LOM | Broadcom 57810S |
| Memory module(s) | |
| Total RAM in system (GB) | 64 |
| Vendor and model number | Hynix Semiconductor Inc.(r) HMA41GR7MFR8N-TF |
| Type | PC4-17000 |
| Speed (MHz) | 2,133 |
| Speed running in the system (MHz) | 2,133 |
| Size (GB) | 16 |
| Number of RAM module(s) | 4 |
| Chip organization | Double-sided |
| Rank | Dual |
| RAID controller | |
| Vendor and model number | Dell PERC H730 Mini |
| Firmware version | 25.2.1.0037 |
| Firmware | |
| Lifecycle Controller / iDRAC | 2.21.21.21 |
| Driver Pack | 14.05.04, X04 |

Figure 5: Configuration information for the Dell PowerEdge M630 blade server.

| System | Cisco UCS 5108 blade enclosure |
|-----------------------------|--------------------------------------|
| Power supplies | |
| Number of power supplies | 4 |
| Vendor and model number | Cisco Systems Inc. UCSB-PSU-2500ACPL |
| Wattage of each (W) | 2,500 |
| Cooling fans | |
| Total number of fan modules | 8 |
| Vendor and model number | Cisco Systems Inc. N20-FAN5 |
| Chassis firmware | |
| Board Controller | 5.0 |
| CIMC Controller | 2.2(3c) |
| IOM firmware | 2.2(3g) |
| I/O modules | |
| Switch | Cisco UCS 2208XP |
| Occupied bay | 1, 2 |

Figure 6: Configuration information for the Cisco UCS 5108 blade enclosure.

| System | Cisco UCS B200 M4 server |
|---|---|
| Enclosure | |
| Blade enclosure | Cisco UCS 5108 |
| General | |
| Number of processor packages | 2 |
| Number of cores per processor | 12 |
| Number of hardware threads per core | 2 |
| System power management policy | Default |
| CPU | |
| Vendor | Intel |
| Name | Xeon |
| Model number | E5-2690 v3 |
| Stepping | M1 |
| Socket type | FCLGA2011-3 |
| Core frequency (GHz) | 2.6 |
| Bus frequency (GHz) | 9.6 |
| L1 cache | 768 KB |
| L2 cache | 3072 KB |
| L3 cache | 30 MB |
| Platform | |
| Vendor and model number | Cisco UCS B200 M4 |
| Motherboard model number | UCSB-B200-M4 |
| BIOS version | B200M4.2.2.3c.0.101420141352 |
| Modular LAN on Motherboard (mLOM) adapter | Cisco UCS VIC 1340 (UCSB-MLOM-40G-03) |
| Memory module(s) | |
| Total RAM in system (GB) | 64 |
| Vendor and model number | Hynix Semiconductor Inc. UCS-MR-1X162RU-A |
| Type | PC4-17000 |
| Speed (MHz) | 2,133 |
| Speed running in the system (MHz) | 2,133 |
| Size (GB) | 16 |
| Number of RAM module(s) | 4 |
| Chip organization | Double-sided |
| Rank | Dual |
| RAID controller | |
| Vendor and model number | Cisco Systems Inc. UCSB-MRAID12G |
| Firmware version | 24.5.0-0021 |
| Firmware | |
| CIMC | 2.2(3C) |
| Board Controller | 5.0 |
| BIOS | B200M4.2.2.3c.0.101420141352 |
| Cisco UCS VIC 1340 | 4.0(1c) |

Figure 7: Configuration information for the blade server.

| System | Dell Storage SC9000 |
|----------------------------------|---|
| Storage Array | |
| Vendor and model number | Dell Storage SC9000 |
| OS Version | Storage Center 6.7.5 |
| Connection type | Fibre Channel |
| Number of Fibre Channel cards | 1 per controller (2 total) |
| Fibre Channel card model | QLogic® QLE2662 16G Fibre Channel Adapter |
| Number of SAS cards | 2 per controller (4 total) |
| SAS card model | SAS9300-8E 12G SAS Dual-Ext Port Adapter |
| Number of expansion enclosures | 2 |
| Expansion enclosure model number | SC420 |
| Number of disks per shelf | 24 |
| Type of disk | SSD |
| Disk capacity | 400GB |

Figure 8: Configuration information for the Dell Storage SC9000 storage array.

Figure 9 provides pricing information for 8 and 16 blades in both environments. While we tested only the 8-blade configurations, we provide the pricing for the 16-blade configurations to show the cost difference and potential savings you would see in the Dell environment when expanding the configuration, which would require a second chassis for the Cisco solution.

| Dell | | 8 blades (as tested) | | 16 blades | |
|-------------------------------------|---------------|----------------------|--------------------|-----------|---------------------|
| Component | Cost per unit | Quantity | Total | Quantity | Total |
| Dell Networking S4810 switch | \$5,414.86 | 2 | \$10,829.72 | 2 | \$10,829.72 |
| Brocade 6505 12-24 Port FC16 Switch | \$8,625.00 | 2 | \$17,250.00 | 2 | \$17,250.00 |
| Dell PowerEdge M630 | \$7,010.00 | 8 | \$56,080.00 | 16 | \$112,160.00 |
| Dell ASM | \$159.00 | 8 | \$1,272.00 | 16 | \$2,544.00 |
| Dell PowerEdge M1000E chassis | \$9,836.00 | 1 | \$9,836.00 | 1 | \$9,836.00 |
| Total | | | \$95,267.72 | | \$152,619.72 |

| Cisco | | 8 blades (as tested) | | 16 blades | |
|--|---------------|----------------------|--------------------|-----------|---------------------|
| Component | Cost per unit | Quantity | Total | Quantity | Total |
| Cisco UCS 5108 Chassis + 8 Cisco UCS B200 M4 blades | \$59,967.65 | 1 | \$59,967.65 | 2 | \$119,935.30 |
| Cisco UCS 6248UP Fabric Interconnect with two Cisco UCS 2208XP FEX modules | \$18,560.01 | 2 | \$37,120.02 | 2 | \$37,120.02 |
| Total | | | \$97,087.67 | | \$157,055.32 |

Figure 9: Pricing information for 8 and 16 blades in the two environments.

APPENDIX C – HOW WE TESTED

Installing VMware vSphere onto the blade servers

Before installing VMware vSphere onto the blades, we set the System Profile to Performance-per-watt (OS) on all Dell blades. We installed the VMware vSphere 6.0 hypervisor onto the blade servers' local hard drives, which we placed in a RAID1 pair.

1. Using the virtual KVM console, connect to the first blade server and attach the VMware vSphere 6.0 installation ISO to the virtual optical drive.
2. Boot the blade server to the VMware vSphere installer.
3. Press F11 to begin the installation wizard.
4. Select the RAID1 pair as the installation location.
5. Select US English as the keyboard type and language.
6. Enter the root password for the host.
7. Press Enter to start the installation.
8. When the installation completes press Enter to reboot the host.
9. After the host reboots, set the management network IP address.
10. Add the host to vCenter and enable SSH and ESXi shell.
11. Navigate to the host settings, and select Power Management.
12. Change the Power Management policy to High Performance.

Configuring the Fibre Channel networking

On the Dell M1000e chassis, we connected the 16Gb-enabled Dell Storage SC9000 to the Brocade M6505 16Gb Fibre Channel switches in the rear of the chassis. We left the switch ports at factory settings and created zones for traffic isolation. On the Cisco UCS chassis, we connected the storage directly to the Fabric Interconnects and placed the Fabric Interconnects in switch mode.

Configuring zoning on the Brocade M6505 switches

1. In a web browser, connect to the first Brocade M6505 switch and log in as the admin user.
2. Click Configure→Zone Admin.
3. In the Zone tab, create the following zones:
 - a. All visible physical ports on the SC9000
 - b. All visible virtual ports on the SC9000
 - c. Server 1 HBA 1 and all four visible virtual ports
 - i. Repeat Step C for remaining servers.
4. In the Zone Config tab, click New Zone Config and give it a name.
5. Add the newly created Zones to the Zone Config.
6. Click Save Config.
7. Click Enable Config.
8. Repeat the zoning process on the second Brocade M6505 for HBA 2.

Configuring the storage

After initial cabling and configuration of the Dell Storage SC9000, we disabled the read and write cache and created six volumes. Each volume contained the respective VMDKs for each VM. For example, the OS volume contained all the OS VMDKs. We had one OS volume, four Data volumes, and one volume for Logs. The data remained in the highest-performance RAID10 tier throughout testing.

Creating a server cluster and connecting the volumes

1. In a web browser, connect to the Storage Center virtual IP address and log in as the Admin user.
2. Click Storage Management→Servers→Create Server. Select the two HBA WWNs that correspond to the first blade server. Give the server a name and click continue.
3. Repeat Step 2 for the remaining blade servers.
4. Click Storage Management→Servers→Create Server Cluster. Click Add Existing Servers to Cluster. Select all servers and click Continue.
5. Click Storage Management→Volumes→Create Volume. Give the volume a size. Use 3.85 TB for the OS and Data volumes, and 2.71 TB for the Logs volume.
6. Repeat Step 5 for the remaining volumes (one OS, four Data, and one Logs).
7. Click Storage Management→Volume→Map volume to server. Select the first volume and click Continue. Select the server cluster and click Continue.
8. Using VMware vCenter, ensure that all volumes are showing as connected to each server. Rebooting the servers and chassis may be necessary to propagate all changes.

Configuring the host HBA settings

Run the following commands on each Dell PowerEdge M630 blade server to adjust the HBA queue depth and round robin pathing settings:

```
esxcli system module parameters set -m qlnativefc -p "ql2xmaxqdepth=255
ql2xloginretrycount=60 qlport_down_retry=60"

esxcli storage nmp device set --device=naa.* --psp=VMW_PSP_RR

esxcli storage nmp psp roundrobin deviceconfig set --device naa.* --
type=iops --iops=3
```

Run the following commands on each UCS B200 M4 blade server to adjust the HBA queue depth and round robin pathing settings:

```
esxcli system module parameters set -m fnic -p "fnic_max_qdepth=128"

esxcli storage nmp device set --device=naa.* --psp=VMW_PSP_RR

esxcli storage nmp psp roundrobin deviceconfig set --device naa.* --
type=iops --iops=3
```

Complete the following steps to adjust the IO Throttle Count on each UCS B200 M4 blade server:

1. In UCS Manager, click the Servers tab.
2. Expand Policies→Root→Adapter policies and select the active policy or create a new policy.
3. In the Options section, set the IO Throttle Count to 1024.
4. Click Save Changes.

Creating the workload virtual machines – SLOB

Creating the VM

1. In VMware vCenter, navigate to Virtual Machines.
2. Click the icon to create a new VM.
3. Leave Create a new virtual machine selected, and click Next.
4. Enter a name for the virtual machine, and click Next.
5. Place the VM on a host with available CPUs, and click Next.
6. Select the OS datastore for the 50GB OS VMDK, and click next.
7. Click Next.
8. Select the guest OS as Oracle Enterprise Linux 6, and click Next.
9. In the Customize Hardware section, make the following changes:
 - a. Increase the vCPUs to 8.
 - b. Increase the memory to 32GB.
 - c. Add a 50GB VMDK for Oracle data and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in one of the DATA datastores.
 - d. Repeat the previous step to add a 50GB VMDK to each DATA datastore, for four data VMDKs.
 - e. Add a 35GB VMDK for Oracle logs and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in the LOGS datastore.
10. Connect the VM to the test network.
11. Click Next.
12. Click Finish.
13. Follow the instructions below to install the guest OS.

Installing Oracle Enterprise Linux 6.5

1. Attach the installation ISO to the VM, and boot to it.
2. Select Install or upgrade an existing system.
3. If you are unsure of the fidelity of the installation disk, select OK to test the installation media; otherwise, select Skip.
4. In the opening splash screen, select Next.
5. Choose the language you wish to use, and click Next.
6. Select the keyboard layout, and click Next.
7. Select Basic Storage Devices, and click Next.
8. Select Fresh Installation, and click Next.
9. Insert the hostname, and select Configure Network.
10. In the Network Connections menu, configure network connections.

11. After configuring the network connections, click Close.
12. Click Next.
13. Select the nearest city in your time zone, and click Next.
14. Enter the root password, and click Next.
15. Select Use All Space, and click Next.
16. When the installation prompts you to confirm that you are writing changes to the disk, select Write changes to disk.
17. Select Software Basic Server, and click Next. Oracle Enterprise Linux installation begins.
18. When the installation completes, select Reboot to restart the server.

Installing VMware Tools

1. Right-click the VM in the Web Client, and select Install/Upgrade VMware Tools.
2. Log on to the guest as `root`.
3. Mount the CD ROM device:

```
mount /dev/cdrom /mnt
```
4. Untar VMware Tools into a temporary directory:

```
tar -C /tmp -zxf /mnt/VMwareTools-9.4.0-1280544.tar.gz
```
5. Run the install script and accept the defaults:

```
/tmp/vmware-tools-distrib/vmware-install.pl
```
6. Follow the prompts to configure and install VMware tools.
7. The installer will automatically load the NIC drivers, create a new `initrd`, and unmount the CD.
8. Reboot the VM.

Initial configuration tasks

Complete the following steps to provide the functionality that Oracle Database requires. We performed all of these tasks as `root`.

1. Disable firewall services. In the command line (as `root`), type:

```
service iptables stop  
chkconfig iptables off  
service ip6tables stop  
chkconfig ip6tables off
```
2. Edit `/etc/selinux/config`:

```
SELINUX=permissive
```
3. Modify `/etc/hosts` to include the IP address of the internal IP and the hostname.
4. Edit `/etc/security/limits.d/90-nproc.conf`:
Change this:

```
* soft nproc 1024
```


To this:

```
* - nproc 16384
```
5. Install 12c RPM packages, resolve package dependencies, and modify kernel parameters:

```
yum install oracle-rdbms-server-12cR1-preinstall -y
```
6. Install automatic system tuning for database storage through yum:

```
yum install tuned
```

```
chkconfig tuned on
tuned-adm profile enterprise-storage
```

7. Using yum, install the following prerequisite packages for Oracle Database:

```
yum install elfutils-libelf-devel
yum install xhost
yum install unixODBC
yum install unixODBC-devel
yum install oracleasm-support oracleasm lib oracleasm
```

8. Create the oracle user account and groups and password:

```
groupadd -g 1003 oper
groupadd -g 1004 asmadmin
groupadd -g 1005 asmdba
groupadd -g 1006 asmoper
usermod -G dba,oper,asmadmin,asmdba,asmoper oracle
passwd oracle
```

9. Create the /u01 directory for Oracle inventory and software and give it to the oracle user:

```
mkdir -p /u01/app/oracle/product/12.1.0/grid_1
mkdir -p /u01/app/oracle/product/12.1.0/dbhome_1
chown -R oracle:oinstall /u01
chmod -R 775 /u01
```

10. Edit bash profiles to set up user environments:

```
# vim /home/oracle/.bash_profile
# Oracle Settings
export TMP=/tmp
export TMPDIR=$TMP
export ORACLE_HOSTNAME=localhost.localdomain
export ORACLE_BASE=/u01/app/oracle
export GRID_HOME=$ORACLE_BASE/product/12.1.0/grid_1
export DB_HOME=$ORACLE_BASE/product/12.1.0/dbhome_1
export ORACLE_HOME=$DB_HOME
export ORACLE_SID=orcl
export ORACLE_TERM=xterm
export BASE_PATH=/usr/sbin:$PATH
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export
  CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/j
  lib
alias grid_env='. /home/oracle/grid_env'
alias db_env='. /home/oracle/db_env'

# vim /home/oracle/grid_env

export ORACLE_SID=+ASM1
export ORACLE_HOME=$GRID_HOME
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
```

```

export
    CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/j
    lib

# vim /home/oracle/db_env

export ORACLE_SID=orcl
export ORACLE_HOME=$DB_HOME
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export
    CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/j
    lib

```

Adding the log and four data drives to the VM and modifying the vmx file

1. Power off the VM.
2. Right-click the VM in the vSphere Web Client, and choose Edit Settings...
3. Click on the VM Options tab, and expand the Advanced menu option.
4. Choose Edit Configuration...
5. Click Add Row, and enter `disk.EnableUUID` in the parameter field and `TRUE` in the value field.
6. Go back to the Virtual Hardware tab.
7. Click the drop-down menu for New device, and choose New Hard Disk.
8. Name the Hard Disk and choose the size that you want it to be.
9. Repeat steps 7 and 8 for all remaining drives.
10. Click OK.
11. Power the VM back on.

Configuring disks for ASM

1. For each of the five shared disks, create a GPT label, and create one partition. For example, see the following shell script:

```

for disk in sdc sdd sde sdf sdg; do
    parted /dev/$disk mklabel gpt
    parted /dev/$disk mkpart primary "1 -1"
done

```
2. If desired, label the disk's partition with its Oracle function. For example:

```

parted /dev/sdc name 1 DATA1
parted /dev/sdd name 1 DATA2
parted /dev/sde name 1 DATA3
parted /dev/sdf name 1 DATA4
parted /dev/sdg name 1 LOG1

```
3. Initialize Oracle ASM on each server by executing the following commands as root on each node.

```

oracleasm init
oracleasm configure -e -u grid -g oinstall -s y -x sda

```
4. Label each shared disk-partition with an appropriate ASM name. For example, following the OS partition names created above, execute the following commands on one system:

```
oracleasm createdisk DATA1 /dev/sdc1
oracleasm createdisk DATA2 /dev/sdd1
oracleasm createdisk DATA3 /dev/sde1
oracleasm createdisk DATA4 /dev/sdf1
oracleasm createdisk LOG1 /dev/sdg1
```

5. On each server, scan the disks to make the disks immediately available to Oracle ASM.

```
oracleasm scandisks
oracleasm listdisks
```

Installing Oracle Grid Infrastructure 12c

1. Log in as the `oracle` user.
2. Unzip `linuxamd64_12c_grid_1of2.zip` and `linuxamd64_12c_grid_2of2.zip`
3. Open a terminal to the unzipped database directory.
4. Type `grid_env` to set the Oracle grid environment.
5. To start the installer, type `./runInstaller`
6. At the Updates screen, select Skip updates.
7. In the Select Installation Option screen, select Install and Configure Grid Infrastructure for a Standalone Server, and click Next.
8. Choose the language, and click Next.
9. In the Create ASM Disk Group screen, choose the Disk Group Name, and change redundancy to External.
10. Select the four disks that you are planning to use for the database, and click Next.
11. In the Specify ASM Password screen, choose Use same password for these accounts, write the passwords for the ASM users, and click Next.
12. Leave the default Operating System Groups, and click Next.
13. Leave the default installation, and click Next.
14. Leave the default inventory location, and click Next.
15. Under Root script execution, select Automatically run configuration scripts and enter root credentials.
16. In the Prerequisite Checks screen, make sure that there are no errors.
17. In the Summary screen, verify that everything is correct, and click Finish to install Oracle Grid Infrastructure.
18. At one point during the installation, the installation prompts you to execute two configuration scripts as root. Follow the instructions to run the scripts.
19. At the Finish screen, click Close.
20. To run the ASM Configuration Assistant, type `asmca`.
21. In the ASM Configuration Assistant, click Create.
22. In the Create Disk Group window, name the new disk group `log`, choose redundancy External (None), and select the log disk for redo logs.
23. Click Advanced Options, and type 12.1.0.0.0 in ASM Compatibility and Database Compatibility. Click OK.
24. Right-click the DATA drive, and choose Edit Attributes. Make sure both ASM and Database Compatibility fields list 12.1.0.0.0, and click OK.
25. Exit the ASM Configuration Assistant.

Installing Oracle Database 12c

1. Unzip `linuxamd64_12c_database_1_of_2.zip` and `linuxamd64_12c_database_2_of_2.zip`.
2. Open a terminal to the unzipped database directory.
3. Type `db_env` to set the Oracle database environment.
4. Run `./runInstaller.sh`.
5. Wait until the GUI installer loads.
6. On the Configure Security Updates screen, enter the credentials for My Oracle Support. If you do not have an account, uncheck the box I wish to receive security updates via My Oracle Support, and click Next.
7. At the warning, click Yes.
8. On the Download Software Updates screen, enter the desired update option, and click Next.
9. On the Select Installation Option screen, select Install database software only, and click Next.
10. On the Grid Installation Options screen, select Single instance database installation, and click Next.
11. On the Select Product Languages screen, leave the default setting of English, and click Next.
12. On the Select Database Edition screen, select Enterprise Edition, and click Next.
13. On the Specify Installation Location, leave the defaults, and click Next.
14. On the Create Inventory screen, leave the default settings, and click Next.
15. On the Privileged Operating System groups screen, keep the defaults, and click Next.
16. Allow the prerequisite checker to complete.
17. On the Summary screen, click Install.
18. Once the Execute Configuration scripts prompt appears, ssh into the server as `root`, and run the following command:

```
/home/oracle/app/oracle/product/12.1.0/dbhome_1/root.sh
```
19. Return to the prompt, and click OK.
20. Once the installer completes, click Close.

Adding hugepages

In Putty, edit `/etc/sysctl.conf` by adding the following lines:

```
vm.nr_hugepages = 13312  
vm.hugetlb_shm_group = 1001
```

Creating and configuring the database

1. Using Putty with X11 forwarding enabled, SSH to the VM.
2. Type `dbca`, and press Enter to open the Database configuration assistant.
3. At the Database Operation screen, select Create Database, and click Next.
4. Under Creation Mode, select Advanced Mode, and click Next.
5. At the Select Template screen, select General Purpose or Transaction Processing. Click Next.
6. Enter a Global database name and the appropriate SID.
7. At the Management Options screen, select Configure Enterprise Manager (EM) Database Express. Click Next.
8. At the Database Credentials screen, select Use the Same Administrative Password for All Accounts. Enter a password, and click Next.

9. At the Network Configuration screen, click Next.
10. At the Storage Configuration screen, select Automatic Storage Management, and select +DATA as the Database.
11. At the Database Options screen, click Next.
12. At the Initialization Parameters screen, click use Automatic Memory Management. Choose a database block size of 32KB.
13. At the Creation Options screen, select Create Database, and click Next.
14. At the summary screen, click Finish.
15. Close the Database Configuration Assistant.
16. In a Web browser, browse to `https://vm.ip.address:5500/em` to open the database manager.
17. Log in as system with the password you specified.
18. Go to Storage→Tablespaces.
19. Click Create.
20. Enter SLOB as the Name, and check the Set As Default box. Click OK.
21. Go to Storage→Redo Log Groups.
22. Click Actions→Switch file... until you get one of the groups to go inactive.
23. Highlight the inactive group, and click Actions→Drop group.
24. In Putty, enter sqlplus and type `ALTER DATABASE ADD LOGFILE GROUP 3 ('+logs/orcl/online/redo/redo_x.log') SIZE 9G BLOCKSIZE 4k;` where x is the number of the log file (1, 2, etc.).
25. Repeat steps 21 – 24 until all default logs are dropped, and two new ones exist.
26. Click on Configuration→Initialization Parameters
27. Change the `use_large_pages` parameter to Only.

Installing SLOB and populating the database

1. Download the SLOB kit from www.kevinclosson.net/slob/.
2. Copy and untar the files to `/home/oracle/dev/`.
3. Type `./setup.sh SLOB 128` to start the data population to the SLOB tablespace we created earlier.
4. When the setup is complete, the database is populated.

Running the performance test

To run the SLOB workload, we first rebooted all hosts and VMs and waited 20 minutes to allow the servers to reach full idle. We ran the production-level workload with `./runit.sh 128,0ms think time and 0 update percent.`

ABOUT PRINCIPLED TECHNOLOGIES



Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.

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.
