Migrate VMs faster with a new Dell EMC PowerEdge MX solution

Scalable SmartFabric boosted throughput and reduced latency vs. HPE Synergy and Cisco UCS solutions

Dell EMC™ PowerEdge™ MX platform powered by Intel® Xeon® Scalable processors*

Investing in a new modular server solution is a capital expenditure that requires careful consideration to ensure the greatest long-term return. It is important to assess things like chassis networking speed, application requirements, and platform longevity to make sure the solution can handle your most demanding workloads today and into the future. Efficient networking at the chassis level ensures multiple chassis and compute nodes talk to each other imperceptibly to the end user, functioning as a true unit for a more seamless experience. To test networking capabilities, we created multi-chassis solutions from popular vendors in the Principled Technologies datacenter and moved virtual machines from one chassis to another using VMware vSphere® vMotion®.

The modular Dell EMC PowerEdge MX solution we tested was more responsive than comparable HPE Synergy and Cisco® UCS™ blade solutions, moving VMs from one server to another in less time. By moving workloads quickly from chassis to chassis, you can minimize the time it takes to do routine maintenance or complete hardware updates. The Dell EMC PowerEdge MX solution also delivered more throughput and lower latency than the other solutions, showing the impressive bandwidth and responsiveness of the MX scalable SmartFabric. This first-generation kinetic infrastructure from Dell EMC introduces the next generation of modular infrastructure (replacing the longstanding M1000e architecture), so organizations can expect to see new, compatible server, storage, and networking options for the PowerEdge MX platform over the years as new technologies emerge.
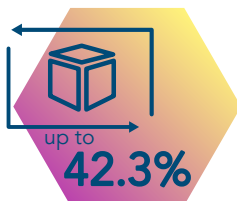
up to
42.3%
Move VMs
in up to 42.3% less time
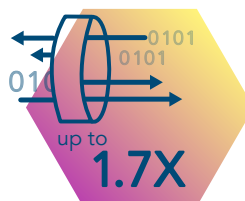
up to
1.7X
Move more data per second over the network
with up to 1.7X the throughput
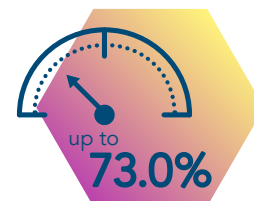
up to
73.0%
Increase responsiveness for vMotion
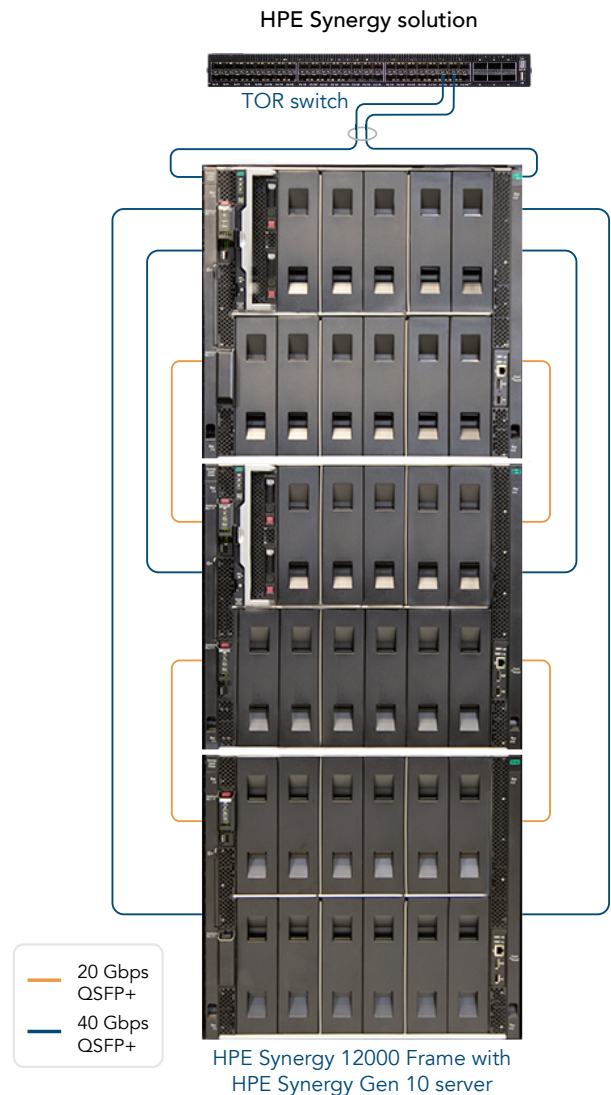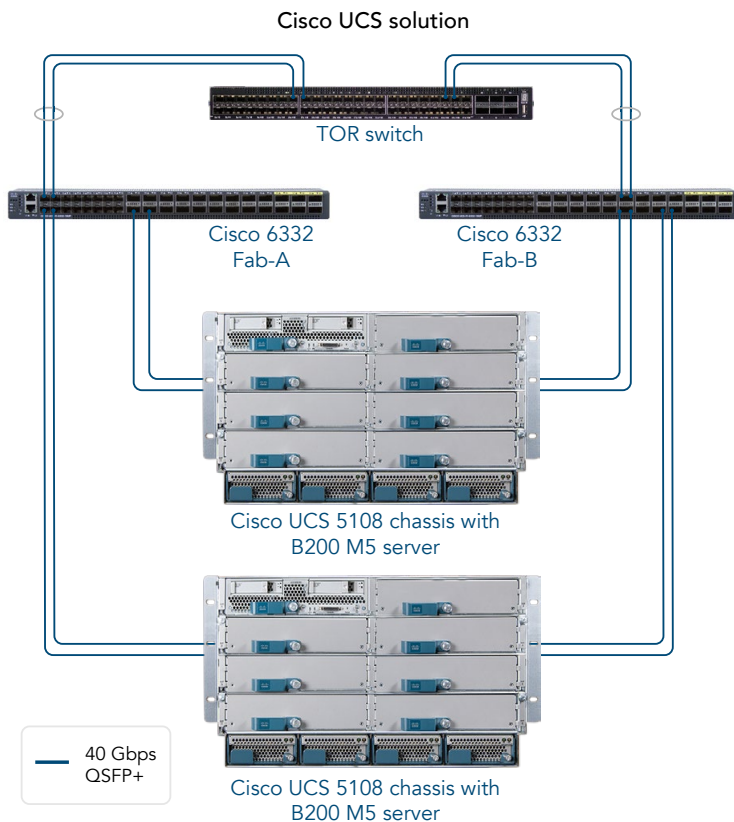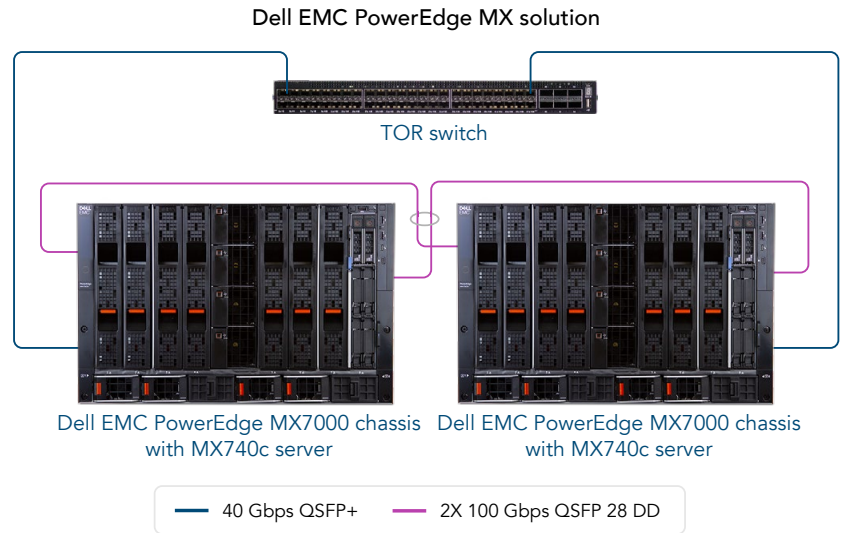with up to 73.0% lower latency

*Image provided by Dell EMC.

Assessing network performance in our labs

Using VMware vMotion, we evacuated VMs from one chassis to another on three leading blade solutions—configuring them with as similar specifications as possible:

- Dell EMC PowerEdge MX solution (two Dell EMC PowerEdge MX740c servers powered by Intel Xeon Platinum 8164 processors and Intel Ethernet Network Adapters XXV710 in 2x MX7000 chassis)

- HPE Synergy solution (two similarly configured HPE Synergy 480 Gen10 compute modules in 2x HPE Synergy 12000 Frame)

- Cisco UCS solution (two similarly configured Cisco UCS B200 M5 blade servers in 2x Cisco UCS 5108 Blade Chassis)

In our tests, we verified MTU 9000 on the network, set up two compute nodes in different chassis, created four virtual machines on one compute node, and used Testlimit64.exe to fill the memory of each VM to 78GB. We then ran a vMotion job to evacuate the VMs to another chassis. The maximum bandwidth available was 20Gbps for the HPE and Cisco solutions and 25Gbps for the Dell EMC solution utilizing Intel XXV710 technology. We measured time to migrate, network throughput, and latency during the evacuation.

Dell EMC PowerEdge MX solution

TOR switch

Dell EMC PowerEdge MX7000 chassis with MX740c server

Dell EMC PowerEdge MX7000 chassis with MX740c server

— 40 Gbps QSFP+   — 2X 100 Gbps QSFP 28 DD

HPE Synergy solution

TOR switch

— 20 Gbps QSFP+
— 40 Gbps QSFP+

HPE Synergy 12000 Frame with HPE Synergy Gen 10 server

Cisco UCS solution

TOR switch

Cisco 6332 Fab-A

Cisco 6332 Fab-B

Cisco UCS 5108 chassis with B200 M5 server

Cisco UCS 5108 chassis with B200 M5 server

— 40 Gbps QSFP+

Why it's important to move VMs quickly across chassis

Datacenter virtualization is everywhere—in fact, a 2016 Gartner report states that "most firms report 75% or higher virtualization of their datacenter."[1] One challenge to a successful virtualized infrastructure is ensuring your multi-chassis modular or blade solution's networking is fast enough that workloads dispersed across multiple chassis can work together with little delay. This includes the ability to move VMs quickly from one compute node to another, which admins must do when performing maintenance, balancing VM workloads for optimal resource distribution, or as a routine part of private cloud management. Business can't stop just because hardware must go offline, so evacuating workloads on a chassis quickly is a key part of maximizing workload uptime. Performing live migrations with VMware vMotion quickly and without disrupting service to end users is a benefit in and of itself. Though workloads and networks vary, the speed with which a solution can complete a vMotion live migration and the available bandwidth a solution provides can be a key indicator of inter-chassis networking performance.

Dell EMC PowerEdge MX scalable SmartFabric integrates the Intel Ethernet Network Adapter XXV710, which is a 25GbE adapter that can deliver the high networking performance of 25GbE while retaining backwards compatibility with 1/10GbE networks. To learn more, visit www.intel.com/content/www/us/en/ethernet-products/network-adapters/ethernet-xxv710-brief.html.

Move VMs in up to 42.3% less time

Migrate faster
Less time is better | minutes:seconds

Dell EMC PowerEdge MX solution
1:52

HPE Synergy solution
2:20

Cisco UCS solution
3:14

About the Dell EMC PowerEdge MX solution

The Dell EMC PowerEdge MX solution is a flexible, kinetic (meaning movable), modular infrastructure with eight bays in front that enables admins to add or remove servers and storage as business needs change and grow. According to Dell EMC, you can expand from two chassis to ten in a single domain quickly and easily with flexible SmartFabric Services. The gear we tested included:

• 7U modular Dell EMC PowerEdge MX7000 chassis, with two distinct redundant PowerEdge MX Scalable Fabrics, a third redundant storvage fabric, and a fault-tolerant interlinked Management Services module to manage compute, storage, and networking fabrics

• Dell EMC PowerEdge MX740c compute sled, which packs two Intel Xeon Scalable processors, up to 24 DDR4 DIMM slots, 3TB memory, and up to six 2.5" drives into one single-width sled

To learn more, visit https://www.dellemc.com/en-us/servers/modular-infrastructure.htm and https://searchconvergedinfrastructure.techtarget.com/news/252447223/Dell-EMC-PowerEdge-MX-teases-kinetic-compute-and-storage.

*Image provided by Dell EMC.

Dell EMC PowerEdge MX platform powered by Intel Xeon Scalable processors*

Evacuate VMs quickly with the kinetic infrastructure of the Dell EMC PowerEdge MX platform

The Dell EMC PowerEdge MX solution moved its virtual machines from a server in chassis 1 to a server in chassis 2 faster than the HPE Synergy and Cisco UCS solutions did, in part thanks to the next-generation modular (NGM) networking of Dell EMC Scalable SmartFabric. Taking less time to evacuate VMs running active workloads across chassis means that admins can complete routine maintenance more quickly, which can have a positive impact on particularly large datacenters with rows and rows of chassis that need updates.

The Dell EMC PowerEdge MX solution with scalable SmartFabric integrating Intel Ethernet Network Adapters XXV710 achieved an average of up to 72.8 percent greater (1.7x) Gbps throughput with as much as 73.0 percent lower average latency than the HPE Synergy and Cisco UCS solutions. Transferring more data per second through the networking pipe—achieving greater throughput—can mean that your network could transfer data faster, and lower network latency means that the servers and storage across your multiple chassis work together as a near-seamless unit so the network won't cause delays.

How are bandwidth, throughput, and latency related?

Envision a two-lane highway. The road can handle a certain number of cars comfortably traveling the speed limit. Drivers are happily moving along. If there's too much traffic, things start to slow down and drivers may get frustrated.

In this analogy, the highway is the network bandwidth—the wider the highway, the more cars you can fit. The number of cars that travel in a given time are the throughput. If there is too much work for a network to handle, or too many stops in a network's data path— a toll both in our highway analogy—there's a traffic jam. At this point, drivers—or in computing, the end users—start to feel a delay, which is latency.

So, more bandwidth, greater throughput, and lower latency are all ideal for computing— and they all play a role in how much work your servers can do and the experience end users have.

Up to
1.7X
the throughput

Transfer more data
More is better

Dell EMC PowerEdge MX solution

19,544 Mb/s

HPE Synergy solution

15,643 Mb/s

Cisco UCS solution

11,312 Mb/s

Reduce network latency
Less is better

Dell EMC PowerEdge MX solution

.61 ms

HPE Synergy solution

.82 ms

Cisco UCS solution

2.27 ms

Up to
73.0%
lower latency
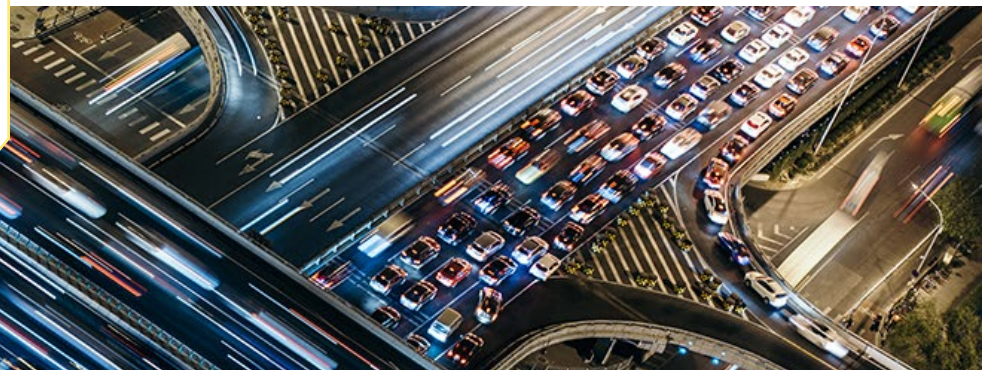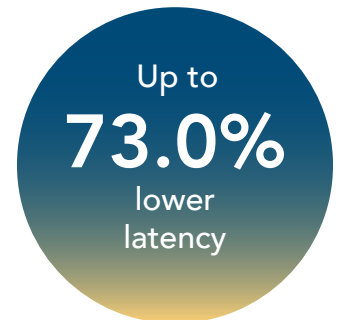
Enjoy the benefits of fast networking with the Dell EMC PowerEdge MX solution

In business, every second counts. Your workforce must have constant access to the information they need to do their jobs. Customers may need to access your databases to make purchases. Whatever the industry, the networking capabilities of your blade solution play a part in determining just how fast they can do these things, so choosing future-forward solution with better network responsiveness is a wise business decision.

With its kinetic infrastructure design, multi-chassis deployment capabilities, and scalable SmartFabric, the Dell EMC PowerEdge MX solution powered by Intel Xeon Platinum 8164 processors evacuated VMs from one chassis to another in less time than similar solutions from HPE and a Cisco did, while also delivering greater throughput and lower network latencies. This means that the Dell EMC PowerEdge MX solution can make routine datacenter maintenance less of a headache for admins by ensuring that virtual workloads move quickly and easily from chassis to chassis without incurring perceptible performance slowdowns.

Plus, the fact that the Dell EMC PowerEdge MX solution is a brand-new blade platform indicates that the line is likely to continue and expand over time. This means that organizations can continue to add new servers and storage as their needs change over the coming years, without the need to replace underlying chassis infrastructure to get all the benefits of the latest and greatest technology.

To find out more about Dell EMC PowerEdge MX, visit https://www.dellemc.com/en-us/servers/modular-infrastructure.htm

---

1   "Market Trends: x86 Server Virtualization, Worldwide, 2016," accessed August 2, 2018, https://www.gartner.com/newsroom/id/3315817

On July 20, 2018, we finalized the hardware and software configurations we tested. Updates for current and recently released hardware and software appear often, so unavoidably these configurations may not represent the latest versions available when this report appears. For older systems, we chose configurations representative of typical purchases of those systems. The Cisco and HPE hardware we selected for this study was currently available at time of writing. Cisco announced the 6300 series Fabric Interconnect in 2016. HPE announced the Synergy platform in late 2015, but the hardware was not generally available until late 2016 and 2017. We concluded hands-on testing on July 24, 2018.
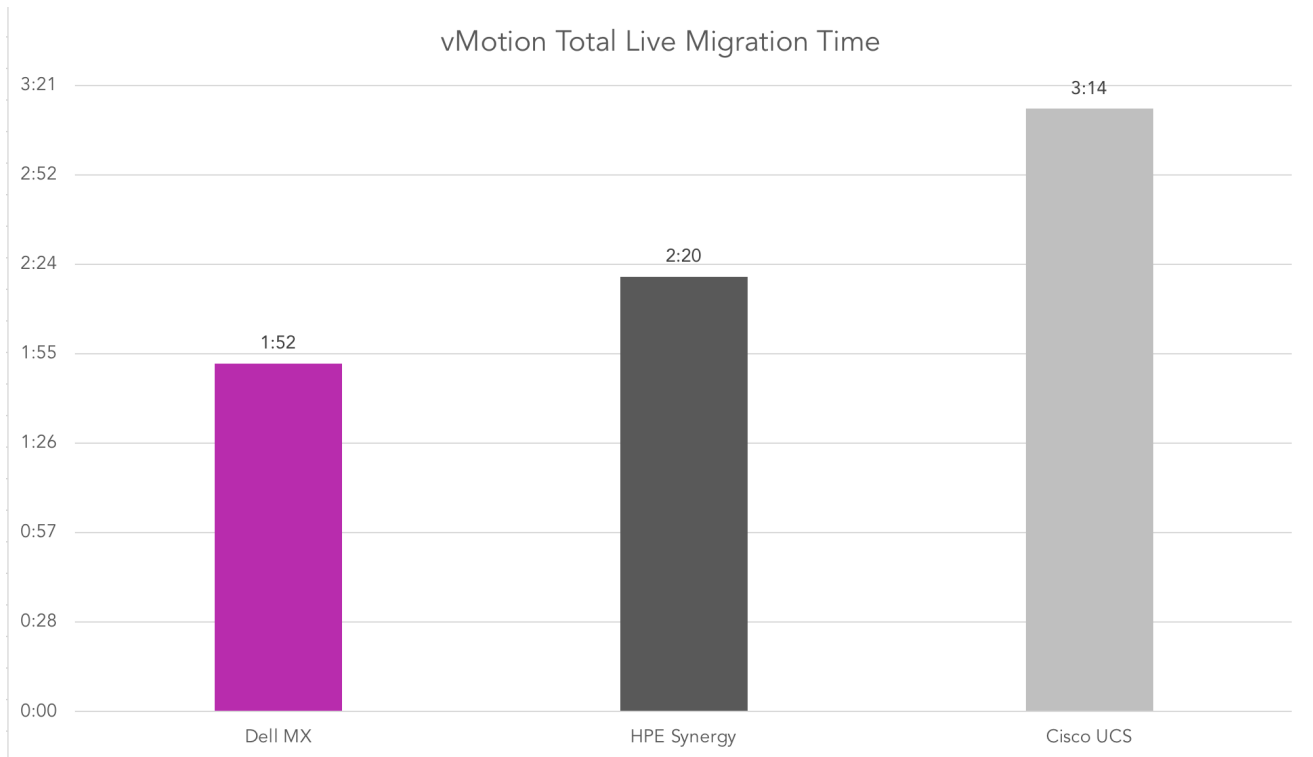
Appendix A

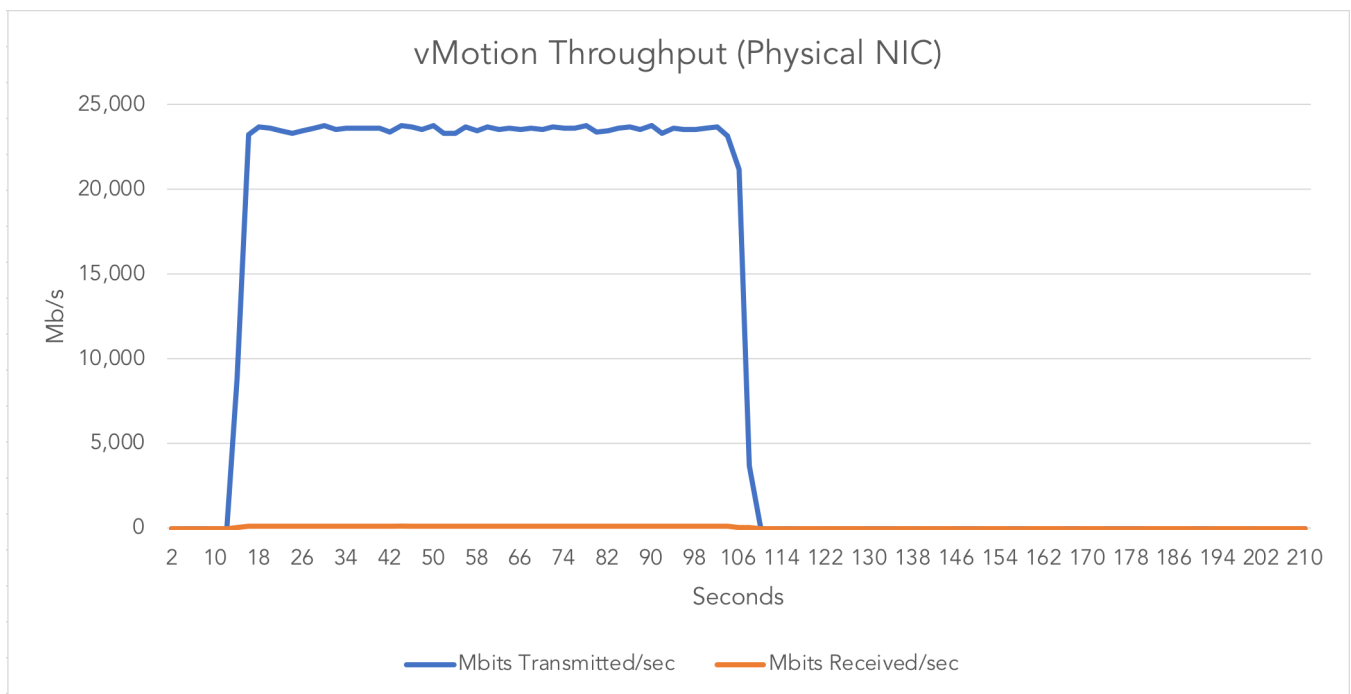## Our results

The table below presents our findings in detail.

Dell EMC PowerEdge MX solution	HPE Synergy solution	Cisco UCS solution
vMotion Live Migration time (min:sec)		
1:52	2:20	3:14
Avg network latency during vMotion (ms)		
0.61	0.82	2.27
Avg base network latency before vMotion (ms)		
0.24	0.47	0.09
Avg throughput during vMotion (Mb/s)		
19,544.39	15,643.11	11,312.07
Peak throughput (Mb/s)		
23,809.63	19,129.24	13,987.01

Time to migrate

vMotion Total Live Migration Time

	Dell MX	HPE Synergy	Cisco UCS
	1:52	2:20	3:14

Y-axis values: 0:00, 0:28, 0:57, 1:26, 1:55, 2:24, 2:52, 3:21

Throughput

Dell EMC PowerEdge MX solution

vMotion Throughput (Physical NIC)

Y-axis (Mb/s): 0, 5,000, 10,000, 15,000, 20,000, 25,000

X-axis (Seconds): 2 10 18 26 34 42 50 58 66 74 82 90 98 106 114 122 130 138 146 154 162 170 178 186 194 202 210

— Mbits Transmitted/sec   — Mbits Received/sec

HPE Synergy solution

vMotion Throughput (Physical NIC)

Mbits Transmitted/sec   Mbits Received/sec

Cisco UCS solution

vMotion Throughput (Physical NIC)

Mbits Transmitted/sec   Mbits Received/sec

Latency
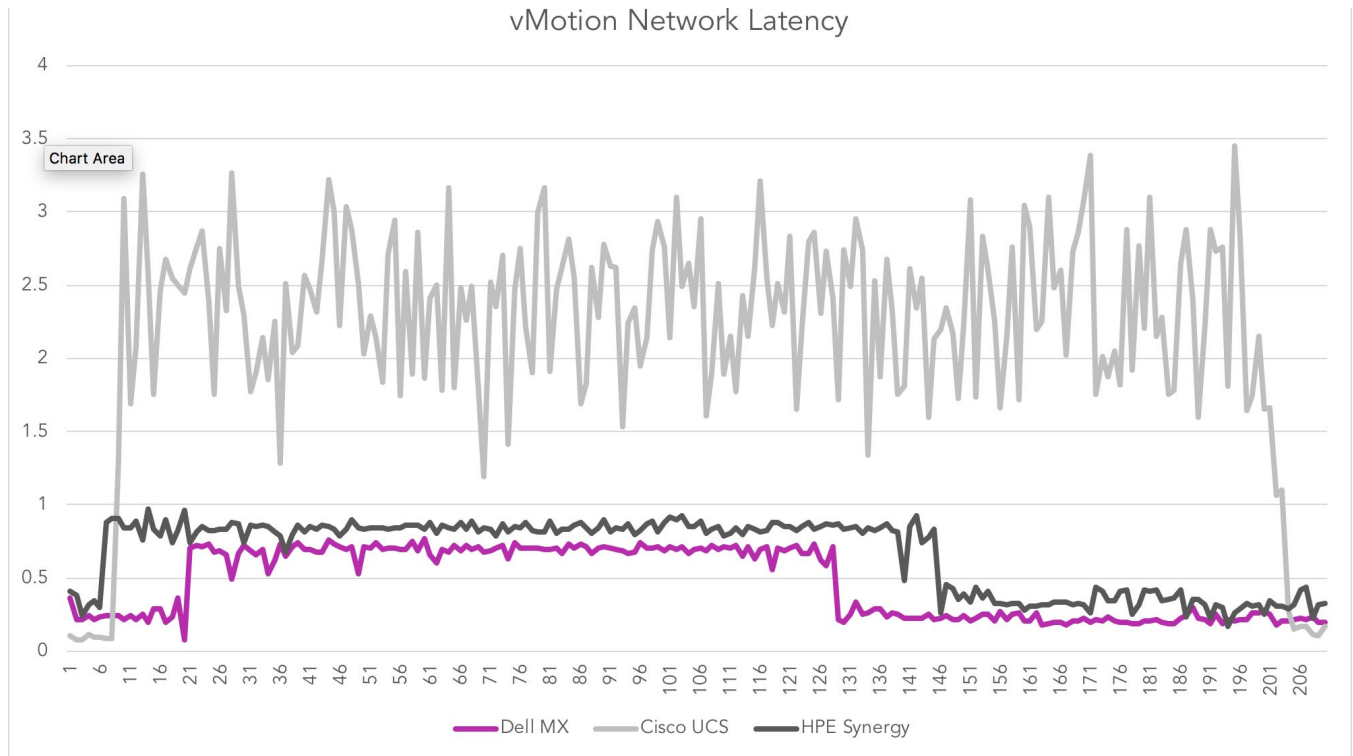

vMotion Network Latency

Appendix B

System configuration information

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

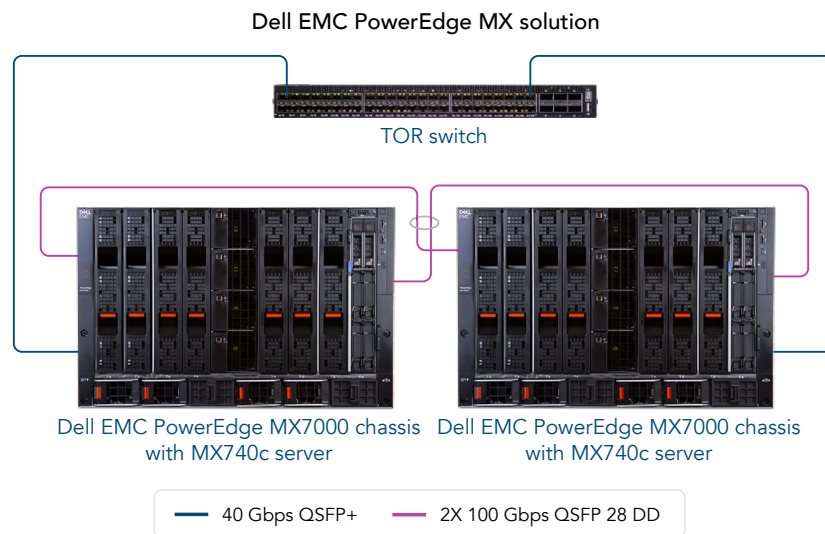| Server enclosure configuration information	Dell EMC PowerEdge MX7000	HPE Synergy Chassis	Cisco UCS 5108 blade server chassis
Power supplies			
Number of power supplies	6	6	4
Vendor and model number	Dell EMC	HPE 2650W AC Titanium Hot Plug Power Supply	Cisco Systems Inc N20-PAC5-2500W
Wattage of each (W)	3,000	2,650	2,500
Cooling fans			
Total number of fan modules	9	10	8
Vendor and model number	Dell EMC	HPE Synergy Fan Module	Cisco Systems Inc N20-FAN5
I/O modules			
Model	Dell EMC MX7116n Fabric Expander Module, MX9116n Fabric Engine	HPE Synergy Frame Link Module	Cisco UCS 2304
Occupied bay	1, 2	1, 2	1, 2

Server configuration information	Dell EMC PowerEdge MX740c	HPE Synergy 480c Gen10	Cisco UCS B200 M5
BIOS name and version	Dell 0.4.3	HPE I42 v1.36	Cisco B200M5.3.2.3c.0.0307181316
Non-default BIOS settings	None	None	None
Operating system name and version/build number	VMware ESXi 6.7 (VMKernel Release Build 8169922)	VMware ESXi 6.7 (VMKernel Release Build 8169922)	VMware ESXi 6.7 (VMKernel Release Build 8169922)
Date of last OS updates/patches applied	6/11/18	6/11/18	6/11/18
Power management policy	Performance	Performance	Performance
Processor			
Number of processors	2	2	2
Vendor and model	Intel Xeon Platinum 8164	Intel Xeon Platinum 8164	Intel Xeon Platinum 8164
Core count (per processor)	26	26	26
Core frequency (GHz)	2.0	2.0	2.0
Stepping	H0	H0	H0
Memory module(s)			
Total memory in system (GB)	384	384	384
Number of memory modules	12	24	12
Vendor and model	Hynix HMA84GR7AFR4N-VK	HPE SmartMemory 840757-091	UCS-MR-X32G2RS-H
Size (GB)	32	16	32

Server configuration information	Dell EMC PowerEdge MX740c	HPE Synergy 480c Gen10	Cisco UCS B200 M5
Type	DDR4-2666	DDR4-2666	DDR4-2666
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 H730P MX	HPE Smart Array P416ie-m SR G10	Cisco UCSB-LSTOR-PT
Cache size (GB)	2	2	1
Firmware version	25.5.3.0005	1.34	3.2(3a)B
Local storage (type A)			
Number of drives	6	2	2
Drive vendor and model	Samsung Express Flash PM1725a	HP 480GB 6G SATA MU-2 SFF SC SSD	HGST HUSMR7632BDP301
Drive size (GB)	3200	480	3200
Drive information (speed, interface, type)	PCIe SSD	SATA SSD	NVMe SSD
Local storage (type B)			
Number of drives	NA	6	NA
Drive vendor and model	NA	HP EO000800JWDKQ SAS	NA
Drive size (GB)	NA	800	NA
Drive information (speed, interface, type)	NA	12Gbps SAS SSD	NA
Network adapter			
Vendor and model	Intel Ethernet Adapter XXV710	Synergy 3820C 10/20 Gb CNA	Cisco UCS VIC 1340

Appendix C

How we tested

Dell EMC PowerEdge MX solution

TOR switch

Dell EMC PowerEdge MX7000 chassis with MX740c server   Dell EMC PowerEdge MX7000 chassis with MX740c server

— 40 Gbps QSFP+   — 2X 100 Gbps QSFP 28 DD

We installed a Dell EMC PowerEdge MX7000 chassis with a Dell EMC MX9116n Fabric Engine in slot A1 and a MX7116n Fabric Expander Module in A2, we installed a second Dell EMC PowerEdge MX7000 chassis with a Dell EMC MX7116n Fabric Expander Module in A1 and a Dell EMC MX9116n Fabric Engine in slot A2, we connected the two chassis via QSFP28DD cables and connected a TOR switch to each chassis via two 40 Gbps QSPF+ connections. In each MX7000 chassis we installed one PowerEdge MX740c compute node.
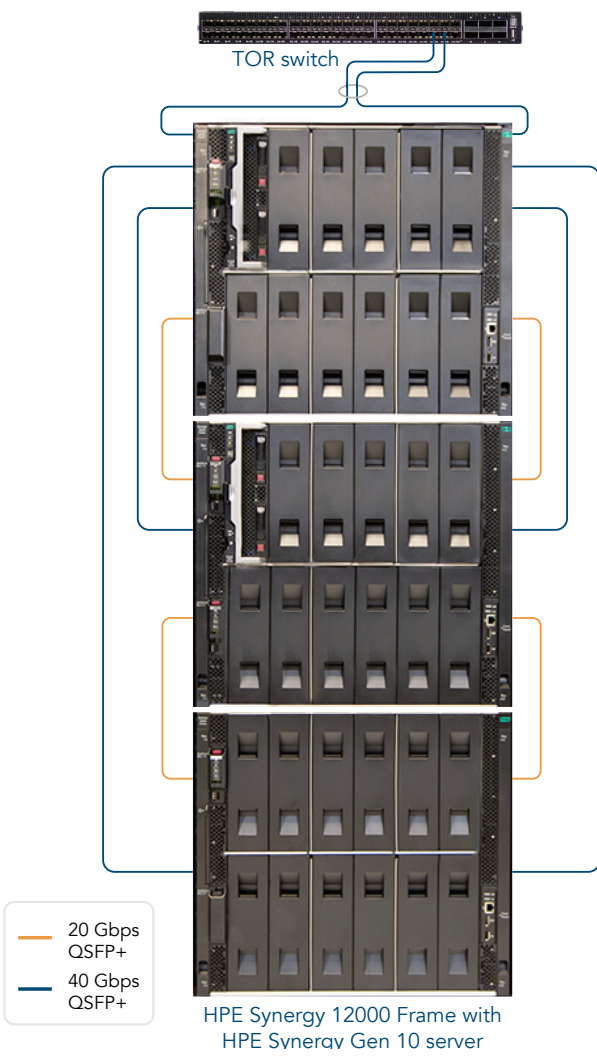
We installed each compute node with VMware ESXI 6.7 and configured the A1 NIC with two IP addresses. We used one for VMware management traffic and the other for vMotion. We also enabled MTU 9000 and ensured functionality by running vmkping between the two hosts.

We contacted HPE professional services to install our Synergy system. It consisted of three Synergy 12000 frames. We installed two Synergy 480 Gen10 compute node into two separate frames We installed each compute node with VMware ESXI 6.7 and configured a NIC with two IP addresses. We used one for VMware management traffic and the other for vMotion. We also enabled MTU 9000 and ensured functionality by running vmkping between the two hosts.
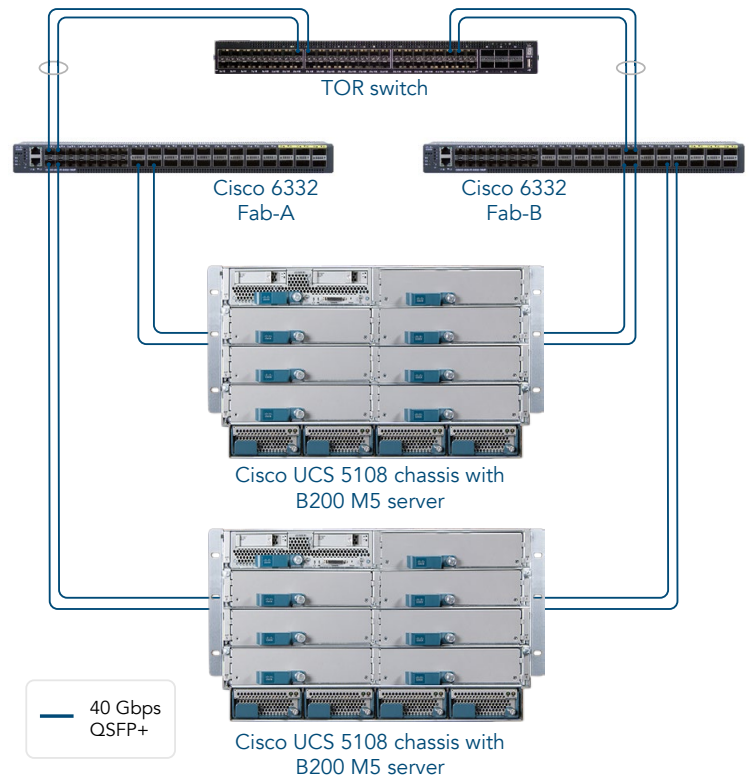
We set up two Cisco UCS 6332 16UP Fabric interconnects into a standard UCS cluster. For each Fabric Interconnect we connected two 40 Gbps QSFP+ cables to a Cisco UCS 5108 chassis via a pair of Cisco UCS 2304 IO modules, thus enabling a total of 80 Gbps of bandwidth from each Cisco UCS 6332 16UP Fabric interconnect to each chassis. We also connected each Cisco UCS 6332 16UP Fabric interconnect to a network TOR (Top Of Rack) switch via two 40 Gbps QSFP+ and configured an LACP port channel allowing a total of 80 Gbps from each Cisco UCS 6332 16UP Fabric interconnect to the TOR.

We installed one Cisco UCS B200 M5 in each of the two Cisco UCS 5108 chassis. We created a Cisco UCS service profile and applied it to each of the respective B200 M5s. The service profile had three vNICs configured, we forced all three vNICS on to Fabric A. Once the service profiles were applied we installed VMware ESXI 6.7 on both servers and configured each vNIC with an IP address. We used one for VMware management traffic and the other two vNICs for vMotion. We chose two vNICs for vMotion because Cisco UCS divides the physical 1340 VIC mLOM adapter into aggregated two 10 Gbps backplane ports, so with only one vNIC the highest achievable speed was 10 Gbps. We also enabled MTU 9000 on our TOR and via QoS rules via UCS manager, we ensured functionality by running vmkping between the two servers.

HPE Synergy solution

TOR switch

20 Gbps
QSFP+

40 Gbps
QSFP+

HPE Synergy 12000 Frame with
HPE Synergy Gen 10 server

Cisco UCS solution

TOR switch

Cisco 6332
Fab-A

Cisco 6332
Fab-B

Cisco UCS 5108 chassis with
B200 M5 server

40 Gbps
QSFP+

Cisco UCS 5108 chassis with
B200 M5 server

Installing VMware ESXi 6.7

1. Attach the installation media.
2. Boot the server.
3. At the VMware Installer screen, press Enter.
4. At the EULA screen, to Accept and Continue, press F11.
5. Under Storage Devices, select the appropriate virtual disk, and press Enter.
6. As the keyboard layout, select US, and press Enter.
7. Enter the root password twice, and press Enter.
8. To start installation, press F11.
9. After the server reboots, press F2, and enter root credentials
10. Select Configure Management Network, and press Enter.
11. Select the appropriate network adapter, and select OK.
12. Select IPv4 settings, and enter the desired IP address, subnet mask, and gateway for the server.
13. Select OK, and restart the management network.
14. Repeat steps 1 through 13 on the rest of the servers.

Deploying the VMware vCenter Server 6.7

1. On a Windows server or VM, locate the VMware-VCSA installer image.
2. Mount the image, navigate to the vcsa-ui-installer folder, and double-click win32.
3. Double-click installer.exe.
4. Click Install.
5. Click Next.
6. Accept the terms of the license agreement, and click Next.
7. Leave the default vCenter Server with an Embedded Platform Services Controller selected, and click Next.
8. Enter the FQDN or IP address of the host onto which the vCenter Server Appliance will be deployed.
9. Provide the server's username and password, and click Next.
10. Accept the certificate of the host you chose to connect to by clicking Yes.
11. Provide a name and password for the vCenter Appliance, and click Next.
12. Set an appropriate Appliance Size, and click Next.
13. Select the appropriate datastore, and click Next.
14. At the Configure Network Settings page, configure the network settings as appropriate for your environment, and click Next.
15. Review your settings, and click Finish.
16. When the deployment completes, click Next.
17. At the Introduction page, click Next
18. At the Appliance configuration page, select the time synchronization mode and SSH access settings, and click Next.
19. Select Create a new SSO domain.
20. Provide a password, and confirm it.
21. Provide an SSO Domain name and SSO Site name, and click Next.
22. At the CEIP page, click Next
23. At the Ready to complete page, click Finish.
24. When installation completes, click Close.
25. Using the vSphere web client, log into the vCenter server using the credentials previously provided.

Creating a Datacenter and adding the hosts to VMware vCenter

1. Once logged into the vCenter, navigate to Hosts and Clusters.
2. Select the primary site management vCenter.
3. Right-click the vCenter object, and select New Datacenter...
4. Enter a name for the new datacenter, and click OK.
5. Right-click the new datacenter, and click Add Host.
6. Enter the FQDN for the first server, and click Next.

7. Enter the root credentials for the server, and click Next.
8. To accept the server's certificate, click Yes.
9. Review the server details, and click Next.
10. Assign the desired license, and click Next.
11. Disable Lockdown mode, and click Next.
12. Click Finish.
13. Repeat steps 5 through 12 for the remaining servers in the environment.

Configuring the VMware vSphere Standard switch

1. Navigate to the Networking tab.
2. Expand the vCenter for the environment.
3. Select the Host, and select Networking, and click Virtual switches.
4. Select the vSwitch.
5. Under Properties, for the MTU setting, type `9000`
6. Click OK.
7. Select VMkernel adapters, and click Add Networking.
8. Select VMkernel network adapter, and click Next.
9. Choose Select an existing standard switch, and click Browse.
10. Choose the vSwitch, and click OK.
11. Click Next, and enter vMotion as the name of the new port group.
12. Set a VLAN to for vMotion.
13. Set the MTU to 9000.
14. Select the vMotion TCP/IP stack, and click Next.
15. Select Use static IPv4 settings, and set the IP address and subnet info for the VMkernel adapter.
16. Review the port group settings, and click Finish.
17. Browse to Virtual switches, and select the vSwitch.
18. In the Port Groups view, select the vMotion port group, and click Edit.
19. Select Teaming and Failover, and choose Override.
20. Assign the vmnic to be used for vMotion to Active, and the remaining vmnics to Unused.
21. Click OK.
22. For each host, repeat steps 3 through 21.

Installing Windows Server 2016

We used the following steps to install and configure Windows Server VMs:

1. Attach the Windows Server 2016 ISO to the virtual machine.
2. Open the VM console, and start the VM.
3. When prompted to boot from DVD, press any key.
4. When the installation screen appears, leave language, time/currency format, and input method as default, and click Next.
5. Click Install now.
6. When the installation prompts you, enter the product key.
7. Select Windows Server 2016 Datacenter Edition (Server with a GUI), and click Next.
8. Check I accept the license terms, and click Next.
9. Click Custom: Install Windows only (advanced).
10. Select Drive 0 Unallocated Space, and click Next. This starts Windows automatically, and Windows will restart automatically after completing.
11. When the Settings page appears, fill in the Password and Reenter Password fields with the same password. Log in with the password you set up previously.
12. Install VMware Tools in the VMs hosted on the ESXi servers.
13. From Server Manager, disable Windows Firewall.
14. Run Windows Updates.

Testing vMotion performance

We used the following steps to configure the test environment and capture data for the vMotion of four Windows VMs:

1. Configuring Windows memory usage to 78GB:
    a. On each VM, download TestLimit64 from SysInternals at https://live.sysinternals.com/windowsinternals/, and copy to the desktop.
    b. Create a shortcut to the EXE, and set the shortcut to run with the following command: `TestLimit64.exe -d -c 78000`
    c. Run the shortcut.
2. Capturing performance data from ESXi:
    a. Start two SSH sessions on the ESX host that will vMotion the VMs.
    b. On the first session, enter the command, but do not yet run: `esxtop -a -b -n 105 -d 2 > PLATFORMNAME-TESTNUMBER.csv`
    c. On the second session, enter the command, but do not yet run: `vmkping -I vmk1 -v -S vmotion -c 210 DESTINATION_IP > PLATFORM-ping-NUMBER.txt`
3. Preparing the vMotion:
    a. Select the four VMs to vMotion, and proceed through the wizard to create a high priority vMotion to the target host.
    b. On the final page of the wizard, do not yet click Finish.
4. Running the test:
    a. Press enter on the two SSH sessions to start capturing esxtop performance data, and begin pinging the VMkernel interface on the destination host.
    b. On the vMotion wizard, click Finish.
    c. When the vMotion is complete, record the elapsed time for the VMs to complete the vMotion.

This project was commissioned by Dell EMC.

Principled Technologies®

Facts matter.®