The science behind the report:

Cut server migration times by upgrading to Dell PowerEdge MX from legacy Cisco UCS

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 Cut server migration times by upgrading to Dell PowerEdge MX from legacy Cisco UCS.

We concluded our hands-on testing on January 13, 2023. 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 January 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: Results of our testing.

	Dell™ PowerEdge™ MX		Cisco UCS® X-series	
Initial configuration (hh:mm:ss)	Time	Steps	Time	Steps
Initial setup	0:04:09	15	0:22:10	28
(For Dell) Chassis configuration / (For Cisco) Required Firmware upload and update for Intersight management (v 4.1.3i)	0:20:51	15	1:09:49	29
Clear and reconfigure FIs	-	-	0:28:03	33
Enroll FI in Intersight	-	-	0:07:53	20
Domain template	-	-	0:27:05	69
Chassis template	-	-	0:02:57	19
Total initial configuration	**0:25:00**	**30**	**2:37:57**	**198**

	Dell™ PowerEdge™ MX		Cisco UCS® X-series	
Configuring and deploying server profiles				
Server template creation (one-time)	0:03:39	31	0:13:44	115
Profile deployment per server	0:08:42	6	0:06:03	3
Total configure and deploy server profiles (for three-node cluster)	**0:29:45**	**49**	**0:31:53**	**124**
Performing migration of workload cluster				
SAN zoning per server	0:05:08	9	0:07:23	9
Storage host per server	0:01:58	8	0:01:58	8
Migration per server	0:09:43	16	0:09:40	17
Total performing migration of workload cluster (for three-node cluster)	**0:50:27**	**99**	**0:57:03**	**102**
Total configuration and workload cluster migration	**1:45:12**	**178**	**4:06:53**	**424**

System configuration information

Table 2: Configuration information for the Dell PowerEdge MX7000 server enclosure.

System configuration information	Dell PowerEdge MX7000
Number of management modules	2
Management module firmware revision	2.00.00
First type of I/O module	
Vendor and model number	Dell MX9116n Fabric Engine
I/O module firmware revision	10.5.4.6.33 3.35.1.1-12
Number of modules	2
Occupied bay(s)	A1, A2
Second type of I/O module	
Vendor and model number	Dell MXG610s FC Switch 2SFP
I/O module firmware revision	8.1.0_lnx2
Number of modules	1
Occupied bay(s)	C1
Power supplies	
Vendor and model number	Dell
Number of power supplies	6
Wattage of each (W)	3,000
Cooling fans	
Vendor and model number	Dell
Number of fans	9

Table 3: Detailed configuration information for the Dell server under test.

System configuration information	Dell PowerEdge MX750c
BIOS name and version	Dell 1.8.2
Non-default BIOS settings	Intel Turbo Boost enabled, Virtualization enabled
Operating system name and version/build number	VMware ESXi, 7.0.3, 19482537
Date of last OS updates/patches applied	12/15/2022
Power management policy	Performance
Processor	
Number of processors	2
Vendor and model	Intel® Xeon® Gold 6338
Core count (per processor)	32
Core frequency (GHz)	2.00

System configuration information	Dell PowerEdge MX750c
Memory module(s)	
Total memory in system (GB)	512
Number of memory modules	8
Vendor and model	Hynix® HMAA8GR7AJR4N-XN
Size (GB)	64
Type	PC4-25600R (DDR4-3200)
Speed (MHz)	3,200
Speed running in the server (MHz)	3,200
Network adapter	
Vendor and model	Broadcom® Adv Quad 25Gb Ethernet
Number and type of ports	4 x 25GbE
Driver version	Dell FW 21.80.15.60
Vendor and model	QLogic QME2742 FC HBA
Number and type of ports	2 x 32Gb
Driver version	Dell FW 16.10.00

Table 4: Configuration information for the server enclosure.

System configuration information	Cisco UCSX-9508
Number of management modules	2 x UCS-FI-6454
Management module firmware revision	4.2(2c)
NX-OS Version	9.3(5)I42(2c)
First type of I/O module	
Vendor and model number	Cisco UCSX-I-9108-25G
I/O module firmware revision	4.2(2c)
Number of modules	2
Power supplies	
Vendor and model number	Cisco UCSX-PSU-2800AC
Number of power supplies	6
Wattage of each (W)	2,800
Cooling fans	
Vendor and model number	Cisco UCSX-9508-FAN
Number of fans	4

Table 5: Detailed configuration information for the Cisco UCSX server under test.

System configuration information	Cisco UCSX-201C-M6
BIOS name and version	Cisco X210M6.5.0.1d.0.0816211754
Non-default BIOS settings	Intel Turbo Boost enabled, Virtualization enabled
Operating system name and version/build number	VMware ESXi, 7.0.3, 19482537
Date of last OS updates/patches applied	12/15/2022
Power management policy	Performance
Processor	
Number of processors	2
Vendor and model	Intel Xeon Gold 6330
Core count (per processor)	28
Core frequency (GHz)	2.00
Memory module(s)	
Total memory in system (GB)	256
Number of memory modules	4
Vendor and model	Cisco UCSX-MR-X64G2RW (Samsung®)
Size (GB)	64
Type	PC4-25600R (DDR4-3200)
Speed (MHz)	3,200
Speed running in the server (MHz)	3,200
Network adapter	
Vendor and model	UCSX-V4-Q25GML
Number and type of ports	4 x 25GbE
Driver version	Cisco FW 5.2(1b

Table 6: Configuration information for the server enclosure.

System configuration information	Cisco UCS 5108
Number of management modules	2 x Cisco UCS® 6332 16UP
Management module firmware revision	5.0(3)N2(3.23b)
First type of I/O module	
Vendor and model number	Cisco UCS-IOM-2304
I/O module firmware revision	3.2(3b)
Number of modules	2
Occupied bay(s)	A, B

System configuration information	Cisco UCS 5108
Power supplies	
Vendor and model number	Cisco N20-PAC5-2500W
Number of power supplies	4
Wattage of each (W)	2500
Cooling fans	
Vendor and model number	Cisco
Number of fans	8

Table 7: Detailed configuration information for the Cisco UCSB server under test.

System configuration information	Cisco UCSB-B200-M5
BIOS name and version	Cisco B200M5.4.0.1b.0911180940
Non-default BIOS settings	Intel Turbo Boost enabled, Virtualization enabled
Operating system name and version/build number	VMware ESXi, 7.0.3, 19482537
Processor	
Number of processors	2
Vendor and model	Intel Xeon Silver 4114
Core count (per processor)	10
Core frequency (GHz)	2.20
Memory module(s)	
Total memory in system (GB)	64
Number of memory modules	2
Vendor and model	Cisco UCS-MR-X32G2RW
Size (GB)	32
Type	DDR4-3200
Speed (MHz)	3,200
Speed running in the server (MHz)	2,400
Network adapter	
Vendor and model	Cisco UCS VIC 1340
Number and type of ports	4 x 40 GbE/FCoE
Driver version	Cisco FW 4.2(3a)

Table 8: Detailed configuration information for the storage solution.

System configuration information	Dell Unity 650F
Storage Software Version	5.0.2.0.5.009
Number of storage controllers	2
Number of storage shelves	2
Number of drives per shelf	24
Drive vendor and model number	EMC™ 005052556
Drive size (GB)	7.68 TB
Drive information (speed, interface, type)	12 Gbps SAS Flash 4

Table 9: Detailed configuration information for the storage switches.

Storage switch configuration information	2 x Brocade DS_6620A
Firemware revision	8.0.1b
Number and type of ports	64 x 32Gb FC
Number and type of ports used in test	10 x 32Gb FC

Table 10: Detailed configuration information for the network switches.

Network switch configuration information	2 x Brocade DS_6620A
Switch software revision	OS 9.9(0.0P4)
Number and type of ports	48 x SFP+ 10GbE, 6 x QSFP 40GbE
Number and type of ports used in test	11 x SFP+ 10GbE, 5 x QSFP 40 GbE
Non-default settings used	Jumbo frames enabled, LACP LAGs

How we tested

Comparing Dell PowerEdge MX to Cisco UCS X-Series in a workload migration strategy from a legacy UCS environment

To compare the effort of migrating a diskless workload from an existing Cisco UCS environment to either a new Dell PowerEdge MX7000 environment or a new Cisco UCS X-Series environment, we set up a cluster of nodes leveraging boot from SAN (BFS). The legacy environment consisted of three Cisco UCS nodes, each running an up-to-date version of VMware ESXi containing the slipstreamed Cisco drivers.

Setting up the physical environment

For the legacy environment, we connected the UCS chassis and three UCS-B nodes via multiple SFP connections to a pair of clustered Fabric Interconnects (FIs). The upstream network connections from the FIs were 40Gb QSFP+ connections to a pair of Dell S4048-ON switches, which we configured as partners in a VLT pair. Additionally, we configured some ports on the FI for use as Fibre Channel (FC) uplink ports. The upstream storage connections were 16Gb FC connections to a Brocade DS_6620A storage switch. A Dell Unity 650F array served as the block storage provider for the environment via FC.

Each UCS-B 200 M5 node ran an instance of VMware ESXi v7.0.3d, and was clustered into a virtual cluster. We presented each node a boot LUN and one additional shared VM data LUN from the Dell Unity array.

The goal of this study was to migrate the active workload from the legacy Cisco UCS-B nodes to newer blade servers located within either a Dell PowerEdge MX7000 chassis or a Cisco UCS-X chassis, and compare the administrative time and effort required to achieve a functional end point with both solutions while minimizing workload downtime.

Configuring the Dell PowerEdge MX solution

Initial setup

1. Open a web browser, and enter the IP address of the embedded OpenManage Enterprise Modular (OMEM). For username, enter `root`, use the default password, and click Login.
2. In the Chassis Deployment Wizard, click Next to skip importing a profile.
3. Check the box for Configure Time settings. Check the box for Use NTP, and enter the IP address or FQDN of the time source you want to use for synchronization. Click Next.
4. To skip Activity and Alerts, click Next.
5. Check the box for Configure iDRAC Quick Deploy settings. Ensure the IPv4 Enabled box is checked and DHCP is the selected IPv4 Network Type. Click Next.
6. Check the box for configure I/O Module Quick Deploy Settings. Ensure the IPv4 Enabled box is checked and DHCP is the selected IPv4 Network Type. Click Next.
7. Check the box for Configure all devices to use the following catalog. Click New Catalog.
8. Provide a name, and select Latest component versions on Dell.com. Set the Update Catalog method to Automatically, set the Update Frequency to Weekly, and choose a time to update the catalog. Click Finish.
9. Use the drop-down menu to select the new catalog. Click Next.
10. To skip Proxy Configuration, click Next.
11. To skip Group Definition, click Next.
12. Click Submit, and click Yes to confirm.
13. Click Application Settings→Users.
14. Select the root user, and click Edit.
15. Enter and confirm a new password. Click Finish.

Configuring the chassis

1. Click Configuration→VLANs.
2. Click Define.
3. Enter a name, a brief description, and a VLAN ID, use the drop-down menu to select the type of VLAN you want to define, and click Finish. We defined the following VLANs for this test:
   - VLAN1000, PRIV, 1000, General Purpose (Platinum)
   - VLAN2000, vMotion, 2000, VM Migration
   - VLAN3000, vSAN, 3000, General Purpose (Platinum)
4. Click Devices→Fabric.

5. Click Add Fabric.
6. Provide a name and a description for the fabric. Click Next.
7. Use the drop-down menu to select the design type 2xMX9116n Fabric Switching Engines in same chassis. Use the drop-down menu to select the current chassis. Use the second drop-down menu to select Slot-IOM-A1 for Switch-A, and the third drop-down menu to select Slot-IOM-A2 for Switch-B. Click Next.
8. Click Finish. The fabric creation makes configuration changes to both MX9116n modules, and creates the interconnection networks.
9. After the fabric creation is complete, click the name of the fabric.
10. Click Add Uplink.
11. Provide a name and a description for the uplink. From the drop-down menu, select Ethernet. Click Next.
12. Select the ports on each IOM you have connected to the upstream switch pair. Check the boxes beside VLAN1000, VLAN2000 and VLAN3000 for Tagged networks. Click Finish.

Building and deploying server templates and profiles

Creating a template

1. Open a web browser, and connect to the OMEM login page. Enter a username and password, and click Login.
2. Select Configuration→Templates.
3. Click Create Template→From Reference Device.
4. Provide a name and a description, and click Next.
5. Click Select Device. Highlight the chassis, and select one of the servers you want configure from the list. Click Finish.
6. Accept the defaults, and click Finish.
7. When the status changes to completed, check the box beside the newly created template, and click Edit → Edit Template.
8. To leave the template information unchanged, click Next.
9. Click Advanced View. Expand FC → FC.Mezzanine.1C-1 → Fibre Channel Target Configuration. Make the following edits:
   - Clear the box beside Boot Scan Selection.
   - Check the First FC Target LUN checkbox. Enter 0 for the value.
   - Check the First FC Target World Wide Port Name checkbox. Enter a WWPN from the A side Storage Processor for the target array.
   - Check the Second FC Target LUN checkbox. Enter 0 for the value.
   - Check the Second FC Target LUN checkbox. Enter a WWPN from the B side Storage Processor for the target array.
10. Expand FC → FC.Mezzanine.1C-2 → Fibre Channel Target Configuration. Make the following edits:
   - Check the Boot Scan Selection checkbox.
   - Check the First FC Target LUN checkbox. Enter 0 for the value.
   - Check the First FC Target World Wide Port Name checkbox. Enter a WWPN from the A side Storage Processor for the target array.
   - Check the Second FC Target LUN checkbox. Enter 0 for the value.
   - Check the Second FC Target LUN checkbox. Enter a WWPN from the B side Storage Processor for the target array.
   - To accept the defaults in IO Pool Assignment, click Next.
   - To accept the defaults in Bandwidth, click Next.
11. In Configuration→Templates, click Edit Network.
12. Click Next.
13. For each adapter, click Untagged Network.
14. Select VLAN1, and click Finish.
15. For each adapter, click Tagged Network.
16. Check the boxes for VLAN1000, VLAN2000 and VLAN3000, and click the >> button to add it to the selected VLANs panel. Click Finish.
17. Click Next.
18. Click Finish.
19. To confirm, click Yes.

Deploying the template

1. Click Deploy Template.
2. To continue, click Yes.
3. Select Deploy to Devices. Click Select Sleds.
4. Select the chassis, and check the boxes for the servers you want to deploy the template. Click Finish.
5. Click Finish.
6. To confirm, click Yes.

Performing the migration

Configuring SAN zoning

Note: This procedure requires the server to have completed a POST operation and to have attempted logging into the storage switch. This means at least one reboot after the profile deployment has completed is necessary so the HBA cards can connect and scan for boot LUNs.

1. Open a supported web browser, and enter http://[ip_address_of_storage_switch]/switchExplorer_installed.html. Accept all security prompts to launch the Java-based SAN switch manager.
2. When prompted, log in with administrator credentials.
3. Click Configure→Zone Admin.
4. On the Alias tab, click New Alias.
5. Provide a name for the Alias, and click OK.
6. In OMEM, reboot the target MX node, and wait for the system POST to complete.
7. On the Web Tools Zone Administration application, click Refresh. Use the drop-down menu to change to the newly created Alias. Select the EMC Clarion device, and click the > to move it to the alias members pane. Select the device with the WWPN associated with your node, and click the > to move it to the alias members pane. Click Save Config, and click Yes.
8. Click the Zone tab. Expand the Aliases folder, select the newly created Alias, and click the > to move it to the Zone Members pane. Click Save Config. To confirm, click Yes.
9. Click Enable Config. To select the current Zone Config, click OK. To confirm, click Yes.

Creating a new host in storage

Note: The following requires at least one reboot after the zone configuration has been saved and enabled, in order for the array to automatically detect the initiators when they attempt to connect to the array.

1. In OMEM, reboot the target MX node, and wait for the system POST to complete.
2. Open a web browser, and connect to the management IP address of the storage array.
3. When prompted, log in with administrative credentials
4. Under Access, click Hosts.
5. Click the + sign to add a new host.
6. Provide a name, and click Next.
7. Click one or more of the automatically discovered initiators associated with your host. Click Next.
8. Review the host configuration, and click Finish.

Migrating one node

1. In OMEM, select Devices→Compute.
2. Check the box beside the target server, and use the Power Control menu to select Power Off (Non-Graceful).
3. Open a browser, and connect to the vSphere vCenter management page.
4. When prompted, provide administrative credentials, and click Login.
5. In vCenter, select the server you want to migrate, and right-click. Select Maintenance Mode→Enter Maintenance mode.
6. When all VMs have moved to other cluster members or the host has entered Maintenance mode, right-click the server, and select Power→Shut Down. Enter Migration as the reason, and click Yes to confirm the operation.
7. With both blade systems down, open the browser window for the array, and click Storage→Block to navigate to the LUN designated as the legacy node BFS LUN. Check the box beside the LUN, and click the pencil icon to edit.
8. Click the Host Access tab. Click the + to add the new target host. Click OK.
9. Check the box beside the old host, and click the trashcan icon to remove it from access. To confirm, click Remove.
10. Click Apply, and click Close.
11. Check the box for the shared data LUN where VMs reside. Click the pencil icon to edit.
12. Click the Host Access tab. Click the + to add the new target host. Click OK.
13. Click Apply, and click Close.
14. Open the browser window containing OMEM, and click Devices→Compute.
15. Check the box next to the target server for the workload you're migrating. Click Power Control→Power On.
16. The server will boot, discover the boot LUN 0, and load the ESXi instance previously in use by the old node.
17. Once the system is showing as responsive in vCenter, right-click on the host, and select Maintenance Mode→Exit Maintenance Mode. The system has been migrated to new hardware and is available for services.

Configuring the Cisco UCS X-Series solution

Initial setup of Fabric Interconnects

1. Connect to FI_A using the included serial cable and Windows jumpbox using PuTTY.
2. Open PuTTY, and set the connection type to Serial.
3. Verify that the serial line is configured as follows (these are PuTTY defaults):
   - Speed: 9600 baud
   - Data bits: 8
   - Stop bits: 1
4. Enter the appropriate serial line in the Serial Line field. In our case this was COM4.
5. Click Open.
6. Press `x`, and press Enter to show the console.
7. When prompted whether to enter setup or restore mode, enter `setup`
8. When prompted to continue, enter `y`
9. When prompted to enforce a strong password, enter `n`
10. Enter the desired admin password.
11. Confirm the desired admin password.
12. When prompted if this Fabric Interconnect is part of a cluster, enter `yes`
13. When prompted for the switch fabric, enter `A`
14. Enter the desired system name.
15. Enter the desired Physical Switch Mgmt0 IP address.
16. Enter the desired Physical Switch Mgmt0 netmask.
17. Enter the desired default gateway.
18. Enter the desired cluster IPv4 address. Note: This item was not specified as part of the included customer pre-configuration checklist.
19. When prompted whether to configure the DNS Server IP address, enter `yes`
20. Enter the desired DNS IP address.
21. When prompted whether to configure the default domain name, enter `yes`
22. Enter the desired default domain name.
23. When prompted to join centralized management environment (UCS Central), enter `n`
24. When prompted to apply and save the configuration, enter `yes`.
25. Wait for the Fabric Interconnect to apply the new configuration and fully reboot back to the console.
26. Once FI A has fully rebooted, connect to the serial port on FI B.
27. When prompted for configuration method, enter `console`
28. When prompted to continue adding this Fabric Interconnect to the cluster of thedetected peer Fabric Interconnect, enter `y`
29. Enter the admin password of FI A specified in steps 10 and 11.
30. Enter the desired Physical Switch Mgmt0 IP address for FI B.
31. When prompted to apply and save the configuration, enter `yes`
32. Wait for FI B to apply the new configuration and fully reboot back to the console.

Upgrading firmware to the minimum required for Intersight management

1. Open a web browser, and enter the cluster IP address specified in the Fabric Interconnect initial configuration.
2. Log into UCS Manager using the username admin and the password specified earlier.
3. In the left side navigation pane, click Equipment.
4. Click the top level Equipment tab.
5. Click the Firmware Management tab.
6. Click Download Firmware.
7. Select Remote File System, and click Browse.
8. Set Protocol to FTP.
9. Specify the IP address of the FTP server.
10. Specify the filename of the infrastructure firmware bundle on the FTP server.
11. Specify the path of the folder containing the firmware bundle on the FTP server.
12. Specify the username and password to use to log onto the remote server.
13. Click OK.

14. Click the Download Tasks tab, and wait for confirmation that the firmware bundle has successfully downloaded to UCS Manager.
15. In the left side navigation pane, click Equipment.
16. Click the top-level Equipment tab.
17. Click the Firmware Management tab.
18. Click the Firmware Auto Install tab.
19. Click Install Infrastructure Firmware.
20. If any warnings are present, check Ignore All.
21. From the Infra Pack drop-down menu, select the infrastructure firmware bundle.
22. Check Upgrade Now.
23. Click OK.
24. Click Pending Activities.
25. Click User Acknowledged Activities.
26. Click the Fabric Interconnects tab.
27. When the option to reboot appears, do so by clicking Reboot Now.
28. Click Yes.
29. Click Yes.

Clearing the Fabric Interconnects and configuring them for Intersight Managed Mode

1. Connect to FI A using the included serial cable and Windows jumpbox using PuTTY.
2. Open PuTTY, and set the Connection Type to Serial.
3. Verify that the serial line is configured as follows (these are PuTTY defaults):
   - Speed: 9600 baud
   - Data bits: 8
   - Stop bits: 1
4. Enter the appropriate serial line in the Serial Line field. In our case, this was COM4.
5. Click Open.
6. Enter `connect local-mgmt`.
7. Log in with the username admin and the password you configured previously.
8. To clear the configuration, enter `erase configuration`
9. Wait for the FI to clear configuration, and reboot.
10. To show the console, press `x`, and press Enter.
11. When prompted for configuration mode, enter `console`
12. When prompted for management mode, enter `intersight`
13. When prompted to continue, enter `y`
14. When prompted to enforce a strong password, enter `n`
15. Enter the desired admin password.
16. Confirm the desired admin password.
17. When prompted if this Fabric Interconnect is part of a cluster, enter `yes`
18. When prompted for the switch fabric, enter `A`
19. Enter the desired system name.
20. Enter the desired Physical Switch Mgmt0 IP address.
21. Enter the desired Physical Switch Mgmt0 netmask.
22. Enter the desired default gateway.
23. Enter the desired cluster IPv4 address. Note: The customer pre-configuration checklist item did not specify this.
24. When prompted whether to configure the DNS Server IP address, enter `yes`
25. Enter the desired DNS IP address.
26. When prompted whether to configure the default domain name, enter `yes`
27. Enter the desired default domain name.
28. When prompted to join centralized management environment (UCS Central), enter `n`
29. When prompted to apply and save the configuration, enter `yes`
30. Wait for the Fabric Interconnect to apply the new configuration and fully reboot back to the console.
31. Once FI A has fully rebooted, connect to the serial port on FI B.
32. When prompted for configuration method, enter `console`
33. When prompted to continue adding this Fabric Interconnect to the cluster of the detected peer Fabric Interconnect, enter `y`
34. Enter the admin password of FI A you specified earlier.
35. Enter the desired Physical Switch Mgmt0 IP address for FI B.
36. When prompted to apply and save the configuration, enter `yes`
37. Wait for FI B to apply the new configuration and fully reboot back to the console.

Enrolling Fabric Interconnects in Intersight

1. Open a web browser, and enter the management IP address of Fabric Interconnect A.
2. Log into the Device Console using the username admin and the password specified during initial configuration.
3. Click the Device Connector tab.
4. Record the Device ID and Claim Code.
5. Browse to Intersight.com, and log in with the appropriate credentials.
6. Click Claim Target.
7. Under Select Target Type, click Cisco UCS Domain (Intersight Managed). Click Start.
8. Copy and paste the Device ID and Claim Code values from step 4 into the corresponding fields.
9. Click Claim.
10. Confirm that the FI has been claimed successfully. Open a web browser, and enter the management IP address of Fabric Interconnect B.
11. Open a web browser, and enter the management IP address of Fabric Interconnect A.
12. Log into the Device Console using the username admin and the password specified during initial configuration.
13. Click the Device Connector tab.
14. Record the Device ID and Claim Code.
15. Browse to Intersight.com, and log in with the appropriate credentials.
16. Click Claim Target.
17. Click Cisco UCS Domain (Intersight Managed) under Select Target Type. Click Start.
18. Copy and paste the Device ID and Claim Code values from step 4 into the corresponding fields.
19. Click Claim.
20. Confirm that the FI has been claimed successfully.

Building and deploying UCS Domain templates and profiles

1. In Intersight, under Infrastructure Service, expand Configure→Profiles.
2. Select UCS Domain Profiles, and click Create UCS Domain Profile.
3. Under General, provide a name and optional description for the profile. Click Next.
4. Under UCS Domain Assignment, select the FI pair you want to assign, and click Next.
5. Under VLAN & VSAN configuration, for Fabric A, click the Select Policy beside VLAN Configuration.
6. In the pop-up side menu, click Create New.
7. For the Create VLAN, add a VLAN name, such as UCS-X_VLAN_Policy, and click Next.
8. Under Policy details, click Add VLANs.
9. Provide a prefix for the VLANs, such as "VLAN". Enter the VLAN numbers you want to add to the policy. We added VLANs 1000, 2000, and 3000. Click Select Multi-cast policy.
10. In the pop-up side menu, click Create New.
11. Under General, provide a name for the multi-cast policy. Click Next.
12. Enable the multicast policies you want for these VLANs. We accepted the defaults (Snooping State and Source IP Proxy State enabled, Querier State disabled). Click Create.
13. Click Add.
14. Click Create.
15. Under VLAN and VSAN configuration, for Fabric A, click Select Policy beside VSAN configuration.
16. In the pop-up side menu, click Create New.
17. Provide a name such as "VSAN", and click Next.
18. Click Add VSAN.
19. Under Add VSAN, provide a name for the VSAN, and select the uplink option. Add a VSAN ID number (required but our switch contained no matching ID), and a FCoE VLAN ID. Click Add.
20. Click Create.
21. Expand Fabric Interconnect B, and beside VLAN configuration, click Select Policy.
22. In the pop-up side menu, select the VLAN policy you just created.
23. Click Select Policy.
24. In the pop-up side menu, select the VSAN policy you just created.
25. Click Next.
26. Under Ports Configuration, click Select Policy besides Ports Configuration under Fabric Interconnect A.
27. Click Create New.
28. Provide a name for the ports configuration policy, and click Next.
29. Click and drag the blue circle along the slider bar to define the number of Fibre Channel ports. We defined all 16 as FC ports. Click Next.
30. Accept the defaults for Breakout Options, and click Next.
31. For Port Roles, select the first 16 ports (designated as FC ports), and click Configure.

32. Under Role, use the drop-down menu to select FC Uplink. Click Save.
33. Set the Admin Speed (we used 16Gbps), and set the VSAN ID to the vSAN you defined in previous steps. Click Save.
34. Select the ports you have configured for uplinks (we selected ports 53 and 54), and click Configure.
35. Under Role, use the drop-down menu to select Ethernet Uplink. Under Ethernet Network Group, click Select Policy.
36. In the pop-up side menu, click Create New.
37. Provide a name for the group, and click Next.
38. Under Allowed VLANs, define the VLANs you want to allow to traverse the uplinks. We used 1000,2000,3000, and left the native VLAN set to 1. Click Create.
39. Under Flow Control, click Select Policy.
40. In the pop-up side menu, click Create New.
41. Provide a name for the policy, and click Next.
42. Accept the default and click Create.
43. Under Link Control, click Select Policy.
44. In the pop-up side menu, click Create New.
45. Provide a name for the policy, and click Next.
46. Accept the defaults and click Create.
47. Click Save.
48. Expand Fabric Interconnect B, and click Select Policy.
49. In the pop-up side menu, select the policy you just created and click Select.
50. Click Next.
51. Under Management, click Select Policy beside NTP.
52. In the pop-up side menu, click Create New.
53. Provide a name for the policy, and click Next.
54. Enter the IP addresses or FQDN you want to use for NTP sources. Click Create.
55. Click Select Policy beside Network Connectivity.
56. In the pop-up side menu, click Create New.
57. Provide a name for the policy, and click Next.
58. Enter the DNS server you want to use and click Create.
59. Under Network, click Select Policy beside System QoS.
60. In the pop-up side menu, click Create New.
61. Provide a name for the policy, and click Next.
62. Beside Best Effort, set MTU to 9216 to enable jumbo frames. Click Create.
63. Next to Switch Control, click Select Policy.
64. In the pop-up side menu, click Create New.
65. Provide a name for the policy, and click Next.
66. Accept the defaults, and click Create.
67. Click Next.
68. Click Deploy.
69. To confirm, click Yes.

Building and deploying UCS Chassis templates and profiles

1. Under Profiles, click UCS Chassis Profiles.
2. Click Create UCS Chassis Profile.
3. Provide a name for the profile, and click Next.
4. Select the chassis you want to assign, and click Next.
5. Under Chassis Configuration, next to IMC Access, click Select Policy.
6. In the pop-up side menu, click Create New.
7. Provide a name for the policy, and click Next.
8. Set the in-band VLAN ID you want to use for management services. Accept the default IPv4 address configuration. Under IP Pool, click Select IP Pool.
9. In the pop-up side menu, click Create New.
10. Provide a name for the pool (such as management), and click Next.
11. Provide the IP address information for the in-band IP addresses you want to use for chassis management. This includes KVM and other remote services. Click Next.
12. Set the toggle to off for Configure IPv6 Pool. Click Create.
13. Click Create.
14. Next to Power, click Select Policy.
15. In the pop-up side menu, click Create New.
16. Provide a name for the policy, and click Next.

17. Accept the defaults, and click Create.
18. Click Next.
19. Click Deploy, and click Yes.

Building and deploying UCS server templates and profiles

Creating a server template

1. In Intersight, expand Configure, and click Templates.
2. Click Create UCS Server Profile Template.
3. Provide a name for the template. Under Target Platform, select UCS Server (FI-attached). Click Next.
4. Under UUID Assignment, under UUID Pool, click Select Pool.
5. In the pop-up widow on the right side, click Create New.
6. Provide a name for the pool, and click Next.
7. Enter a prefix, which must be in hexadecimal format XXXXXXXX-XXXX-XXXX. We used abcdef12-3456-7890. Under UUID Blocks, enter the starting value in hexadecimal format XXXX-XXXXXXXXXXXX. We used abcd-ef1234567890. Provide a size for the UUID block. We used 512. Click Create.
8. Next to Boot Order, click Select Policy.
9. In the pop-up widow on the right side, click Create New.
10. Provide a name for the policy, and click Next.
11. Click the arrow beside Add Boot Device, and select SAN Boot.
12. For the Device Name, use vHBA1. Leave 0 as the target LUN. Use vHBA1 for the Interface Name. Enter a target WWPN of the storage unit containing your boot volume.
13. Click the arrow beside Add Boot Device, and select SAN Boot.
14. For the Device Name, use vHBA2. Leave 0 as the target LUN. Use vHBA2 for the Interface Name. Enter a target WWPN of the storage unit containing your boot volume.
15. Click the arrow beside Add Boot Device, and select Virtual Media.
16. Click the down arrow on the far right of the Virtual Media device to move it to the bottom of the entries. Expand the Virtual Media and give it a device name. Under Sub-Type, select KVM Mapped DVD. Click Create.
17. Under Compute Configuration, click Select Policy beside Virtual Media.
18. In the pop-up menu on the right side, click Create New.
19. Provide a name for the policy, and click Next.
20. Click Create.
21. Click Next.
22. Under Management Configuration, beside IMC Access, click Select Policy.
23. In the pop-up menu on the right side, select the policy you created during the Chassis Profile deployment.
24. Beside IPMI Over LAN, click Select Policy.
25. In the pop-up menu on the right side, click Create New.
26. Provide a name for the policy, and click Next.
27. Accept the default, and click Create.
28. Beside Virtual KVM, click Select Policy.
29. In the pop-up menu on the right, click Create New.
30. Provide a name for the policy, and click Next.
31. Toggle Allow Tunneled KVM to on, and click Create.
32. Click Next.
33. Under Storage Configuration, click Select Policy beside Storage.
34. In the pop-up menu on the right, click Create New.
35. Provide a name for the storage policy, and click Next.
36. Leave all devices disabled (the default) and click Create.
37. Click Next.
38. Under Network Configuration, click Select Policy beside LAN Connectivity.
39. In the pop-up menu on the right, click Create New.
40. Provide a name for the policy, and click Next.
41. Under vNIC Configuration, select Auto vNICs Placement.
42. Click Add vNIC.
43. Provide a name for the vNIC. We used vNIC1. Under MAC Pool, click Select Pool.
44. In the pop-up menu on the right, select Create New.
45. Provide a name for the MAC Pool, and click Next.

46. For the MAC Block's first address, we use the recommended prefix 00:25:B5 and added 00:00:00. For size enter 512. Click Create.
47. For Ethernet Network Group Policy, click Select Policy, and choose the policy you created in the Domain Profile configuration section.
48. For Ethernet Network Control Policy, click Select Policy.
49. In the pop-up menu on the right side, click Create New.
50. Provide a name for the policy, and click Next.
51. Accept the defaults and click Create.
52. For Ethernet QoS, click Select Policy.
53. In the pop-up menu on the right side, click Create New.
54. Provide a name for the policy, and click Next.
55. Accept the defaults, and click Create.
56. For Ethernet Adapter, click Select Policy.
57. In the pop-up menu on the right side, click Create New.
58. Provide a name for the policy, and click Next.
59. Accept the defaults, and click Create.
60. Under Connection, click usNICs. Under usNIC Adapter Policy, click Select Policy.
61. In the pop-up menu on the right side, select the adapter policy you just created. Click Add.
62. On the Policy Details page, click Click Add vNIC.
63. Provide a name for the vNIC. We used vNIC2. Under MAC Pool, click Select Pool.
64. In the pop-up menu on the right, select the policy you created for vNIC1.
65. For Ethernet Network Group Policy, click Select Policy, and choose the policy you created in the Domain Profile configuration section.
66. For Ethernet Network Control Policy, click Select Policy.
67. In the pop-up menu on the right, select the policy you created for vNIC1.
68. For Ethernet QoS, click Select Policy.
69. In the pop-up menu on the right, select the policy you created for vNIC1.
70. For Ethernet Adapter, click Select Policy.
71. In the pop-up menu on the right, select the policy you created for vNIC1.
72. Under Connection click usNICs. Under number of usNICs, enter 1. Under usNIC Adapter Policy, click Select Policy.
73. In the pop-up menu on the right side, select the adapter policy you created for vNIC1. Click Add.
74. Click Create.
75. On the Network Configuration page, click Select Policy beside SAN Connectivity.
76. In the pop-up menu on the right, click Create New.
77. Provide a name for the policy, and click Next.
78. On the Policy details page, click Auto vHBAs Placement. Under WWNN Pool, click Select Pool.
79. In the pop-up menu on the right, click Create New.
80. Provide a name for the pool, and click Next.
81. For WWNN Blocks, keep the recommended prefix of 20:00:00:25:B5 and append 00:00:00. For Size, enter 512. Click Create.
82. On the Policy Details page, click Add vHBA.
83. Provide a name for the vHBA - we used vHBA1. Leave the vHBA type as fc-initiator. Under WWPN Pool, click Select Pool.
84. In the pop-up menu on the right, click Create New.
85. Provide a name for the pool, and click Next.
86. For WWNN Blocks, keep the recommended prefix of 20:00:00:25:B5 and append 11:11:11. For Size, enter 512. Click Create.
87. Under Fibre Channel Network, click Select Policy.
88. In the pop-up menu on the right, click Create New.
89. Provide a name for the policy, and click Next.
90. Enter the VSAN ID you used in during the UCS Domain Profile creation. Click Create.
91. Under Fibre Channel QoS, click Select Policy.
92. In the pop-up menu on the right, click Create New.
93. Provide a name for the policy and click Next.
94. Accept the defaults and click Create.
95. Under Fibre Channel Adapter, click Select Policy.
96. In the pop-up menu on the right, click Create New.
97. Provide a name for the policy, and click Next.
98. Under Fibre Channel Adapter Default Configuration, click Select Default Configuration. Select the built-in policy named Initiator.
99. Click Next.
100. Accept the defaults, and click Create.
101. Click Add.

102. On the Policy Details page, click Add vHBA.
103. Provide a name for the vHBA - we used vHBA2. Leave the vHBA type as fc-initiator. Under WWPN Pool, click Select Pool.
104. In the pop-up menu on the right, select the pool you created for vHBA1.
105. Under Placement, change the Switch ID to B.
106. Under Fibre Channel Network, select the policy you created for vHBA1.
107. Under Fibre Channel QoS, select the policy you created for vHBA1.
108. Under Fibre Channel Adapter, select the policy you created for vHBA1.
109. Click Add.
110. On the Policy Details page, click Create.
111. On the template Network Configuration page, click Next.
112. Click Derive Profiles.
113. Select one or more target servers for profile creation. Click Next.
114. Accept the default name, or edit it as desired, and click Next.
115. Click Derive.

Deploying the server profile

1. On the main Intersight menu, expand Configure→Profiles.
2. Select the newly created profile. On the far right, click the ellipsis, and select Deploy.
3. To confirm, click Deploy.

Performing the migration from UCS to UCS-X

Configuring SAN zoning

Note: This procedure requires the target server to have completed a POST operation, and to have attempted a login to the storage switch. This means at least one reboot after the profile deployment has completed is necessary so the HBA cards can connect and scan for boot LUNs.

1. Open a supported web browser, and enter http://[ip_address_of_storage_switch]/switchExplorer_installed.html. Accept all security prompts to launch the Java-based SAN switch manager.
2. When prompted, log in with administrator credentials.
3. Click Configure→Zone Admin.
4. On the Alias tab, click New Alias.
5. Provide a name for the Alias, and click OK. Click Save Config.
6. In Intersight, reboot the target node and wait for the system POST to complete.
7. On the Web Tools Zone Administration application, click the Refresh button. Use the drop-down menu to change to the newly created Alias. Select the EMC Clarion device and click the > to move it to the alias members pane. Select the device with the WWPN associated with your node, and click the > to move it to the alias members pane. Click Save Config, and click Yes.
8. Click the Zone tab. Expand the Aliases folder and select the newly created Alias and click the > to move it to the Zone Members pane. Click Save Config. Click yes to confirm.
9. Click Enable Config. Click OK to select the current Zone Config. Click Yes to confirm.

Creating a new host in storage

Note: After saving and enabling the zone configuration, the following required at least one reboot for the array to automatically detect the initiators when they attempted to connect to the array.

1. In Intersight, reboot the target UCS node and wait for the system POST to complete.
2. Open a web browser and connect to the management IP address of the storage array.
3. When prompted, log in with administrative credentials
4. Under Access, click Hosts.
5. Click the + sign to add a new host.
6. Provide a name, and click Next.
7. Click one or more of the automatically discovered initiators associated with your host. Click Next.
8. Review the host configuration, and click Finish.

Migrating one node

1. In Intersight, Select Operate → Servers.
2. Check the box beside the target server, and click the ellipses on the far right. Select Power → Power Off to power down the server target for the migrated workload.
3. Open a browser and connect to the vSphere vCenter management page.
4. When prompted, provide administrative credentials and press Login.

5. In vCenter, select the server you want to migrate and right-click. Select Maintenance Mode→Enter Maintenance mode.
6. When all vms have been moved to other cluster members or powered down and the host has entered Maintenance mode, right-click the server and select Power→Shut Down. Enter Migration as the reason and Press Yes to confirm the operation.
7. With both blade systems down, open the browser window for the array, and click Storage→Block to navigate to the LUN designated as the legacy node BFS LUN. Check the box next to the LUN and click the pencil icon to edit.
8. Click the Host Access tab. Click the + to add the new target host. Click OK.
9. Check the box beside the old host, and click the trash can icon to remove it from access. Click Remove to confirm.
10. Click Apply. Click Close.
11. Check the box for the shared data LUN where VMs reside. Click the pencil icon to edit.
12. Click the Host Access tab. Click the + to add the new target host. Click OK.
13. Click Apply. Click Close.
14. Open the browser window containing Intersight, and click Operate→Servers.
15. Check the box beside the target server for the workload you're migrating. Click the ellipsis on the far right, and select Power Control→Power On.
16. Click Power On.
17. The server boots, discovers the boot LUN 0, and loads the ESXi instance the old node was using.
18. Once the system is showing as responsive in vCenter, right-click the host, and select Maintenance Mode→Exit Maintenance Mode. The system has been migrated to new hardware and is available for services.

Read the report at https://facts.pt/f7dDnHy ▶

This project was commissioned by Dell Technologies.

Principled Technologies®

Facts matter.®