



*TOMORROW
starts here.*

Cisco *live!*



UCS C-Series Deployment Options, Best Practice and UCSM Integration

BRKCOM-2640

Steve McQuerry, Manager, Technical Marketing, CCIE #6108

#clmel

Cisco *live!*

Agenda

- Overview and Common Applications
- C-Series Integration Connectivity Options
- Connectivity and Best Practices Recommendations
- Additional I/O consideration
- Discovery of Servers
- Firmware Management
- Network Placement Policies
- Local Drive configuration Considerations
- Case Study




A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and various city buildings are illuminated with lights. The overall scene is a dynamic urban environment.

Overview and Common Applications

Traditional Servers

Hardware Centric

- Servers are specifically purchased for applications
- Hardware dictates the identity of the servers
- Management of servers handled individually or through add-on software



UCS Service Profiles

Workload Centric

- Servers are specifically purchased for additional capacity and provide service elasticity
- Server identities are defined by the service profile
- Management of servers provided through the very system that defines them (UCSM).



UCSM Provides Form Factor Independence

UNIFIED MANAGEMENT
A SINGLE UNIFIED SYSTEM FOR
BLADE AND RACK SERVERS



- A Major Marketing Transformation in Unified Server Management
- Benefits of UCS Manager and Service Profiles for Both Blade and Rack-Optimised Servers
- Add capacity and functionality without complexity

Cisco *live!*

UCSM Provides Form Factor Independence

UNIFIED MANAGEMENT
A SINGLE UNIFIED SYSTEM FOR
BLADE AND RACK SERVERS


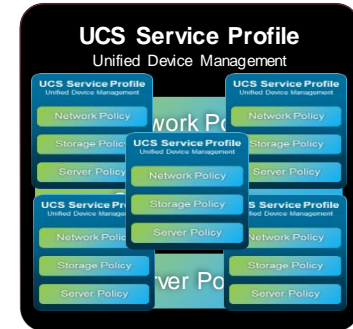


BRKCOM-2640


© 2015 Cisco and/or its affiliates. All rights reserved.

Cisco Public

8



C-Series Rack Optimised
Servers



B-Series
Blade Servers

Cisco *live!*

Common Reasons for Integrating a Rack Mount Server

- Need for other PCIe I/O cards
 - GPU Cards for VDI
 - PCoIP cards for VDI
 - Connecting to Low Latency Network
 - Direct connection to FC SAN
- Need for high number of hard drives
 - Big Data (MPP/NoSQL/Hadoop)
 - Business Analytics
- Need for higher capacity Memory
 - In Memory Database
 - High memory application for Distributed Processing

UCS C-Series Hadoop/Big Data Common Platform Architecture




- Hadoop Common Platform Architecture (CPA)
- Integrated solutions for:
 - Hadoop
 - Massively Parallel Databases
- Integrates with Existing enterprise applications and data repositories
- Rapid Deployment and highly scalable

UCS C-Series VDI Deployments

- Cisco UCS C240 M3 Supports up to 2 NVIDIA Kepler GRID GPUs in an integrated environment
- Support for NVIDIA K1, K2, K10, K20, and K20x
- Supports up to 186 virtual Desktops*

* Based LoginVSI 3.5 Medium work load (Knowledge Worker)



Cisco *live!*


A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and various city buildings are illuminated with lights. The overall scene is a dynamic urban environment.

C-Series Integration Connectivity Options

C-Series Integration Basic Connectivity Overview

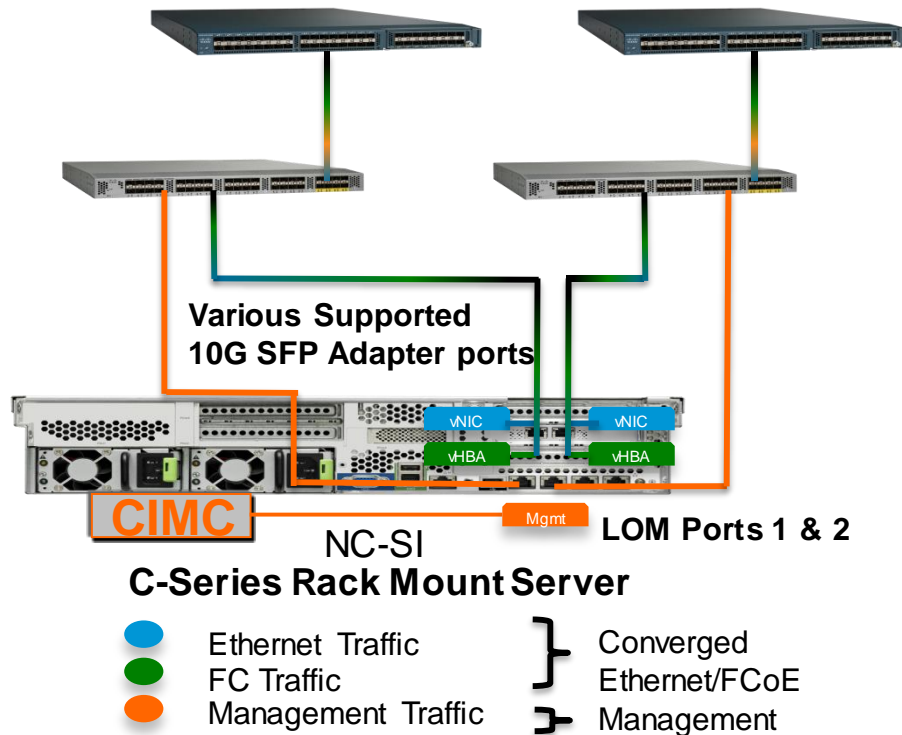
- UCSM integration was introduced in Version 1.4 of UCSM for small scale deployments with separate infrastructure for management.
- Starting with Version 2.0(2) Scalable Mainstream support for integrated servers begins.
- Integration and connectivity evolved with the release of new hardware and software in UCSM versions 2.1 and 2.2.
- Today there are 3 main connectivity modes supported for UCSM integrated C-Series rack Servers.
 - Dual Wire Management
 - Single Connect (a.k.a. Single Wire Management)
 - Direct Connect to the Fabric Interconnect
- Multiple PCIe Slots in the C-Series servers provide a variety of supported connectivity options

Initial Integration Connectivity in UCSM Version 1.4




- Introduction of UCSM Integrated Servers version 1.4 and 2.0(1).
- Limited Scalability as converged infrastructure connected to FI and separate 2248 FEX was needed for management using an additional server port.
- Limited servers supported
 - C200/210
 - C250
 - C460
- Topology deprecated by release 2.0(2)

Dual Wire Management with FEX




- Introduced in Version 2.0(2) and replaced previous topology
- Better Scalability up to 160 C-Series servers in a UCS Domain
- 2232PP FEX provides Converged I/O to the server and Management connectivity.
 - Both must be connected to the same FEX
- Only uplinks require license ports. (up to 8 ports per FEX)
- Uplinks can be port-channelled
- Matches IOM Architecture.

Single Connect with FEX (a.k.a. Single Wire Management)



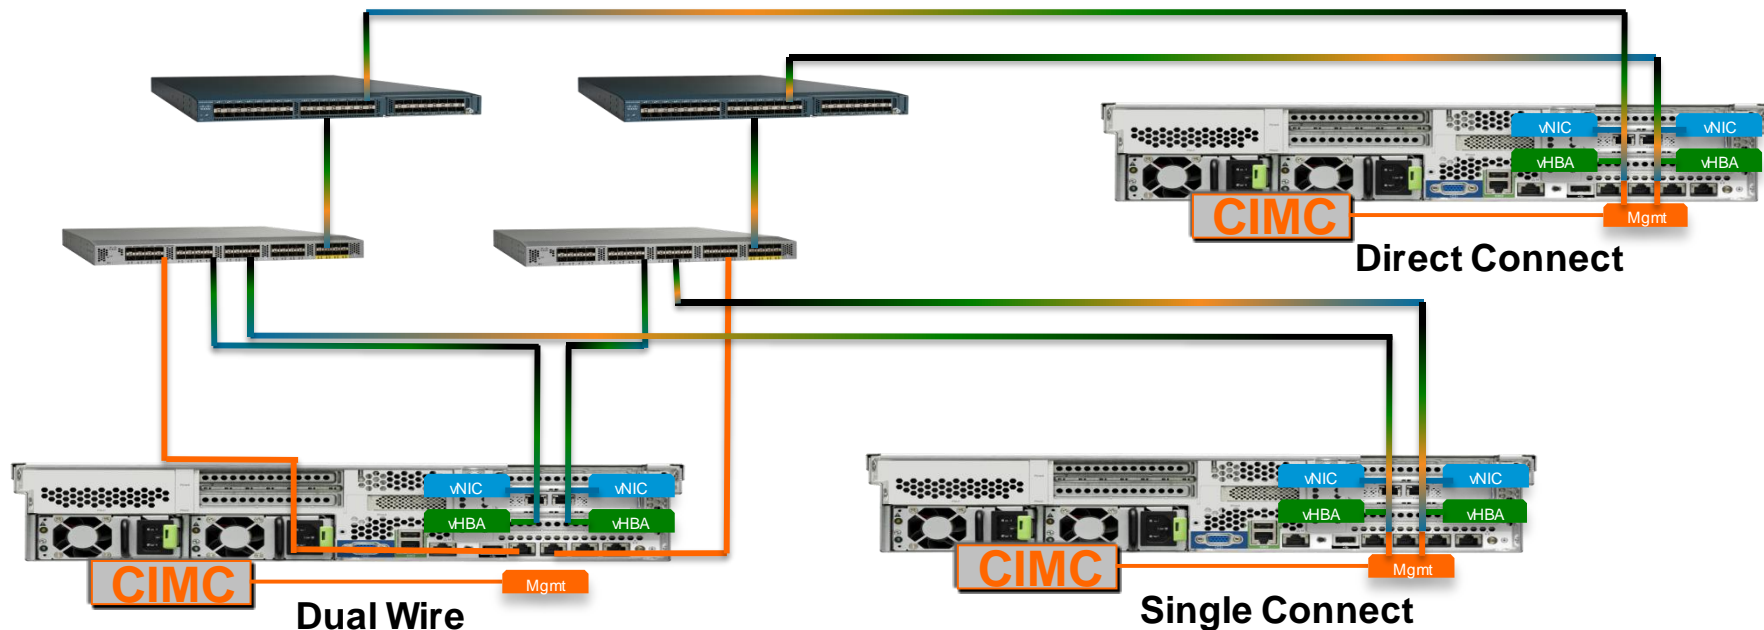
- Introduced in UCSM Version 2.1
- Reduces the need for additional cabling for management interface.
- Requires a VIC 1225 (UCSC-PCIE-CSC-02).
- VIC must be installed in specific slot for standby power.
- Allows up to 32 Servers per FEX reducing infrastructure overhead.

Direct Connect to Fabric Interconnect without FEX




- Introduced in UCSM Version 2.2
- Requires a VIC 1225 (UCSC-PCIE-CSC-02) for Single Connect only.
 - C22 M3, C24 M3, C220 M3, C240, M3, C260 M2, C420 M3, C460 M2
- Reduces the need for additional HW component, but requires a server port license for each server.
- 120 VIFs per VIC installed in the server
- Up to 96 Servers supported with a 6296.
- New licensing model available

Supported Combinations of Standard Connectivity




- Multiple connectivity types are supported within a single UCS Domain provided you are at the appropriate UCSM software level.

Redundant PCIe Adapters Support




- Redundant Adapters are supported with the required PCIe Slots.
- UCSM supports up to 4 adapters.
- Network Placement Policies define which virtual adapters are on which physical card.
- You CAN mix adapter types (e.g. Emulex, Qlogic, Broadcom, Intel, Cisco)

Unsupported Connectivity with Redundant Adapters



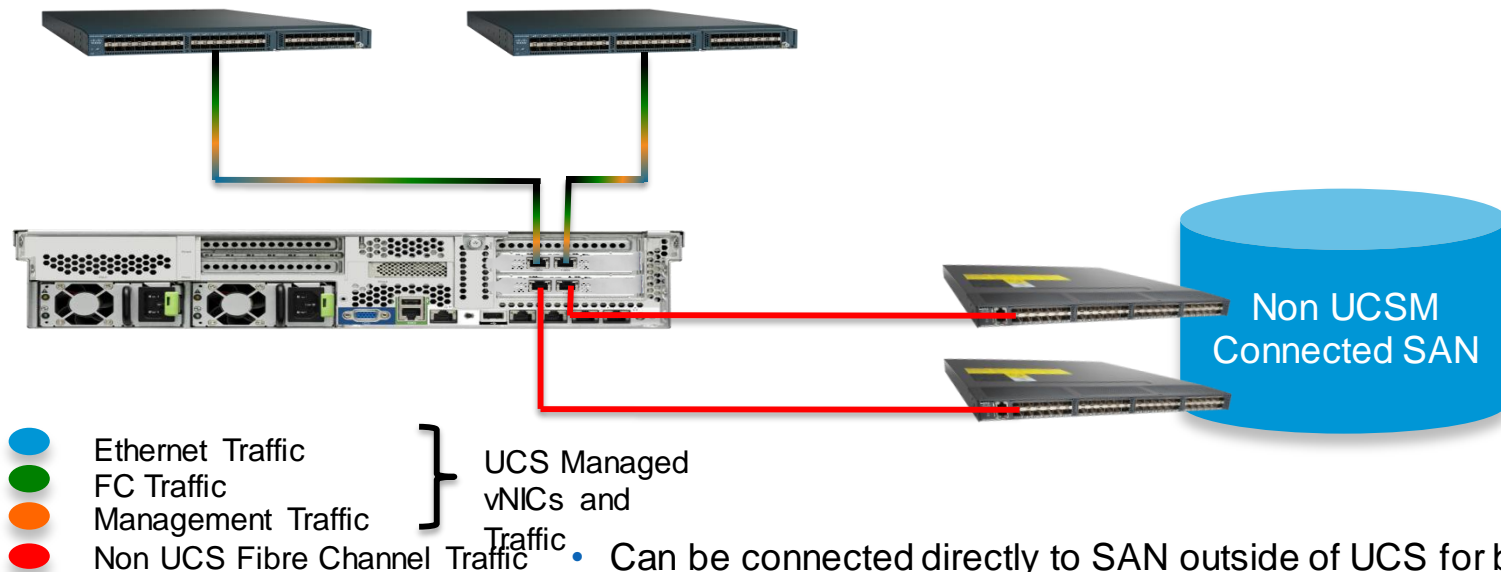
- Redundant adapters **MUST** be cabled so that each card is connected to both sides of the fabric in a cluster.
- Improper cabling will either cause a failure of the discovery or generate errors and create connectivity issues.

Additional Supported Network Connectivity Outside of UCS Fabric



- Useful for connecting to high speed low latency or backup network outside of UCS Fabric.
- Network Configuration must be managed outside UCSM (VIC can NOT be used for external connectivity)

Additional Supported SAN Connectivity Outside of UCS Fabric

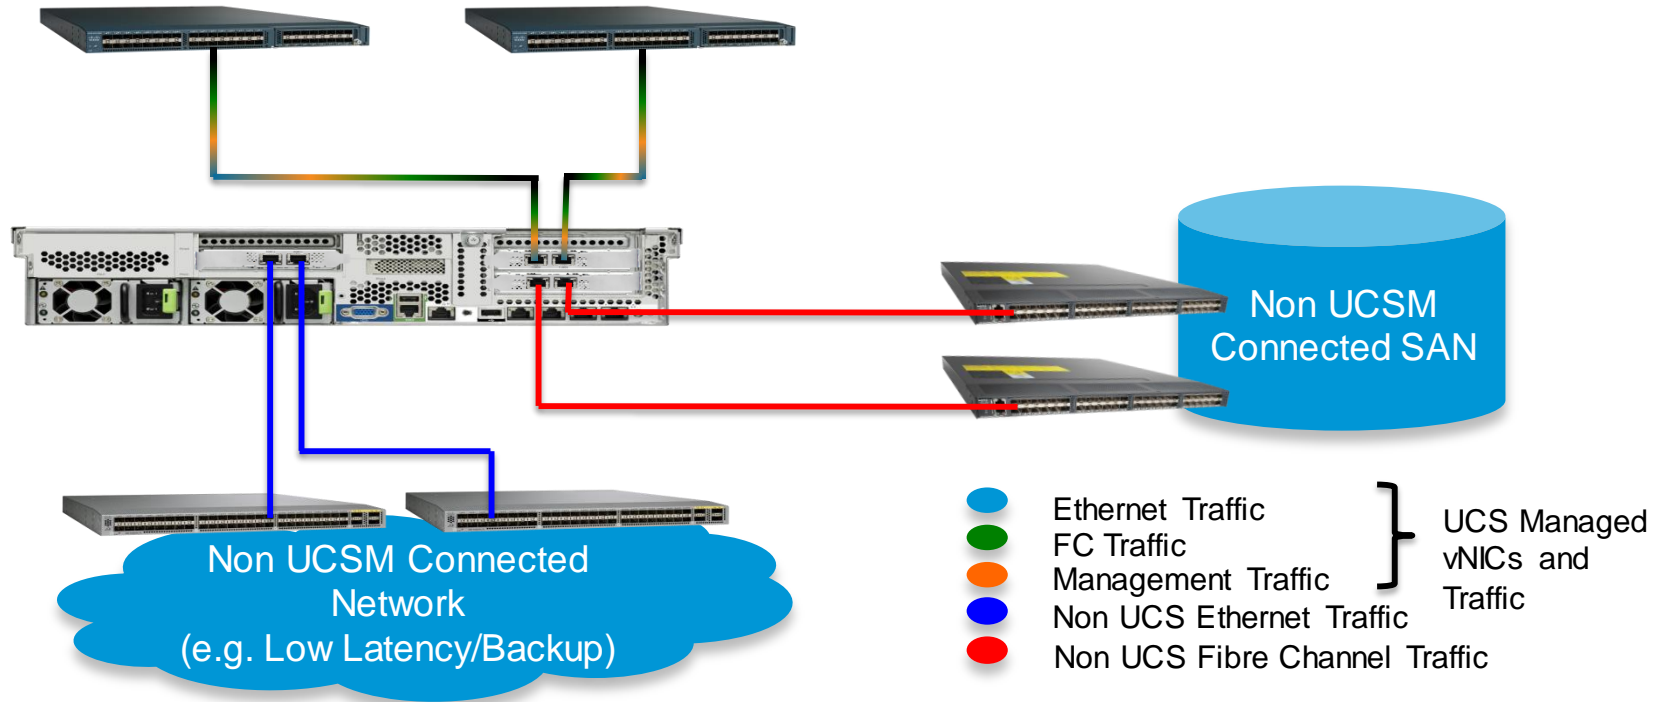


- Can be connected directly to SAN outside of UCS for backup media server or other use.
- SAN configuration must be managed outside of UCSM. (VIC can NOT be used for external connectivity)
- If booting from SAN the appropriate policy must be in place. Prior to 2.2(1) that was local drive.

All VICs Must be UCS Managed

- The VIC interacts directly with the Cisco IMC.
- When a C-Series is integrated in UCS Manager there is no direct way to configure the VIC.
- All configuration of the VIC is through the Service Profile.
- All VICs in an integrated server **MUST BE CONNECTED** to a fabric interconnect.


Combined UCS Fabric with External Network and SAN



A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a modern pedestrian bridge with glass railings spans the street, illuminated with blue lights. Tall city buildings with lit windows and some red flags are visible in the distance under a dark night sky.

Connectivity Best Practices and Recommendations

FEX Uplink Considerations



- When port-channel uplinks from the FEX are connected only within a set of eight ports managed by a single chip, Cisco UCS Manager maximises the number of VIFs used in service profiles deployed on the servers.
- If uplink connections are distributed across ports managed by separate chips, the VIF count is decreased. For example, if you connect seven members of the port channel to ports 1–7, but the eighth member to port 9, this port channel can only support VIFs as though it had one member.

Virtual Interface (VIF) Limitations with FEX


61xx Fabric Interconnect

# Acknowledged links	VIC 1225 Maximum configurable vNIC/vHBA	P81E Maximum configurable vNIC/vHBA
1	Not Supported	Not Supported
2	3	4
4	9	10
8	21	22

62xx Fabric Interconnect

# Acknowledged links	VIC 1225 Maximum configurable vNIC/vHBA	P81E Maximum configurable vNIC/vHBA
1	12	13
2	27	28
4	57	58
8	117	118

- The number of vNIC/vHBAs are limited by the number of uplinks between the FEX and the FI.
- Gen 1 and Gen 2 Fabrics have different VIF limitations.
- The table provides the number of VIFs provided to an adapter based on the number of uplinks.
- A minimum of 2 uplinks are recommended.
- A common issue when moving from a P81E to a VIC 1225 with 2 links on a Gen1 Fabric Interconnect is “not enough resources”



• Static Pinning


- Deterministic Dedicated Traffic Flow
- Best Option for High Throughput
- Failure of Link Results in loss of connectivity (Fabric Failover recommended)

• Port Channel

- Flow Based Load balanced
- May not result in best traffic distribution
- Best option for redundancy & Scaling


Cisco *live!*

Redundancy Considerations



- Fabric Failover provides an alternate path for traffic and is configured and managed through the service profile
- Active/Standby bonding redundancy can be configured through the operating system if desired.

Load Balancing Considerations



- Split Server traffic across both Fabrics, but realise that inter-cluster communication may require an extra hop to the upstream distribution switch.
- Transmit Load Balancing (TLB) and Adaptive Load Balancing (ALB) can be used.
- LACP can not be used because there is no VPC support in the Fabric Interconnect to servers



Additional I/O Consideration

Slot Requirements for Cisco VIC


Server Model	Number of VIC 1225 Supported	PCIe Slots that support VIC 1225	Primary NC-SI Slot (Standby Power) for UCSM integration
UCS C22 M3	1	1	1
UCS C24 M3	1	1	1
UCS C220 M3	1	1	1
UCS C240 M3	2	2 and 5	2
UCS C260 M2	2	1 and 7	7
UCS C420 M3	3	1, 4, and 7	4
UCS C460 M2	2	1 and 2	1

- The Cisco VIC requires an extended slot for communication with the CIMC and NC-SI capabilities.
- The primary slot is preferred, but if the card is placed in the second slot standby power will be supplied there.
- Only one slot will have standby power to provide a management interface in Single Connect.
- Multiple cards are supported for additional Bandwidth.
- The Cisco VIC can **ONLY** be connected through the Fabric Interconnects when integrated with UCSM.

Slot Considerations for GPU

- Cisco C240 Servers can support 2 NVIDIA GPU cards (K1, K2, K10, K20, K20X) in slots 2 and 5.
- GPUs are double height so Slot 1 is blocked when Slot 2 is used.
- Cisco VIC must also reside in slots 2 or 5 which means in a 2GPU deployment the VIC can not be used for integration.
- Consider using other qualified card in Slot 3 for UCSM integration (Qlogic, Emulex, Broadcom, Intel)
- If you are using the mezzanine RAID card with the C240, slot 4 also becomes available.
- If you are using a single GPU consider installing it in Slot 5 for a better thermal signature.

Non UCSM Managed Adapters



- Additional Qualified Adapters may be installed into the C-Series Server and connected to networks external to UCSM.
- These adapters will be inventoried by UCSM but can not be managed as part of the service profile.
- The adapters will be presented to the operating system and configuration and management will be done through traditional vendor tools.



Discovery of Servers

Basics of Discovery

- By default a factory new UCS C-Series Server will discover automatically when plugged into a UCS Fabric using one of the supported connectivity models.
- You must meet the minimum requirements, but in general C-Series servers with CIMC version 1.4(4) or higher can be discovered.
- Plug the server in the way that you want it to be discovered (e.g. Dual Wire or Single Wire).
- Verify that your server is at the minimum levels of required FW to support integration.
- By default management connectivity and server discovery are automatic, you can control this with Global Policies.
- Server and FEX discovery is a RACE condition. First one connected and powered on will be first discovered.
- FEX numbers are the same ID schema as Chassis. If you have 2 chassis then add 2 FEXs the numbers will be 3 and 4 respectively

Discovery Policies

The screenshot shows the 'Equipment' configuration page with the 'Policies' tab selected. The 'Chassis/FEX Discovery Policy' is highlighted with a red box. Below it are the 'Rack Server Discovery Policy' and 'Rack Management Connection Policy'. Other policies visible include 'Power Policy', 'MAC Address Table Aging', 'Global Power Allocation Policy', and 'Firmware Auto Sync Server Policy'.

Chassis/FEX Discovery Policy

Action: Platform Max

Link Grouping Preference: ☒ None ☐ Port Channel

Rack Server Discovery Policy

Action: ☒ Immediate ☐ User Acknowledged

Scrub Policy: <not set>

Rack Management Connection Policy

Action: ☒ Auto Acknowledged ☐ User Acknowledged

Power Policy

Redundancy: ☒ Non Redundant ☐ N+1 ☐ Grid

MAC Address Table Aging

Aging Time: ☐ Never ☒ Mode Default ☐ other

Global Power Allocation Policy

Allocation Method: ☐ Manual Blade Level Cap ☒ Policy Driven Chassis Group Cap

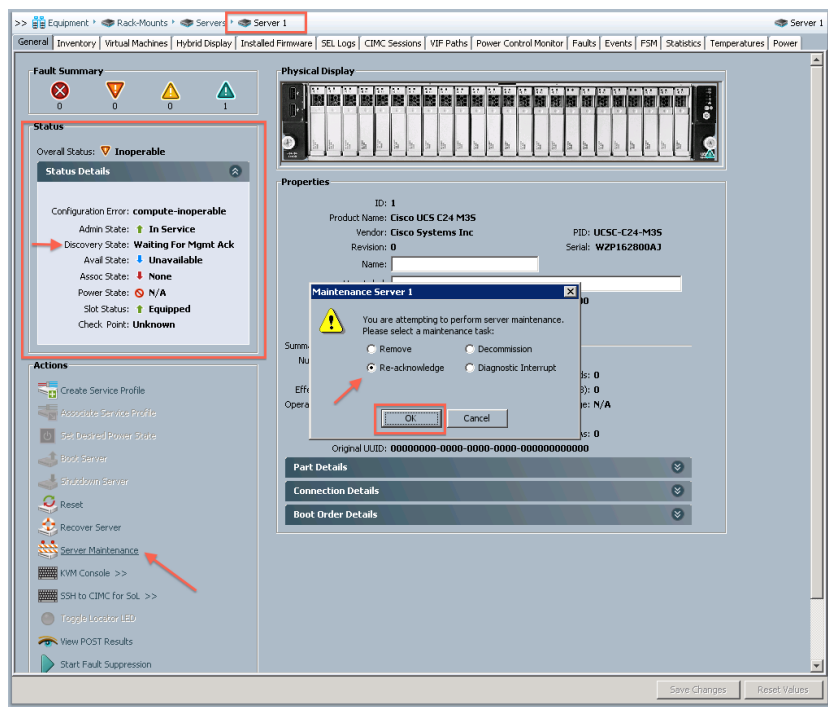
Firmware Auto Sync Server Policy

Sync State: ☒ Auto Acknowledge ☐ User Acknowledge ☐ No Actions

Buttons: Save Changes, Reset Values

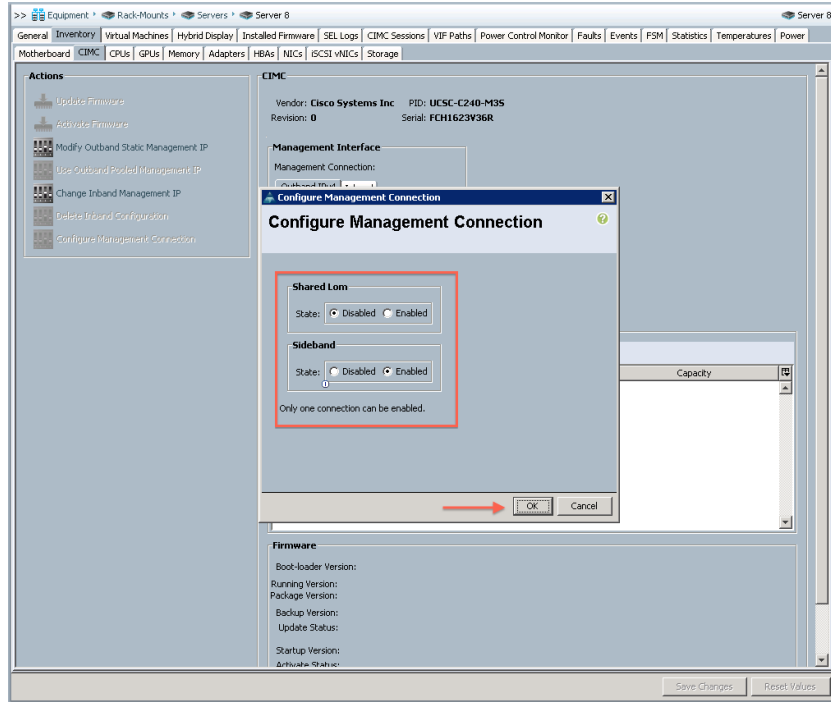
- Global Discovery Policies are located in the Equipment Tab
- 3 Policies apply for C-Series Integration Discovery.
 - 1: FEX Discovery Policy
of links, Port-Channel vs. Pinned
 - 2: Rack Discovery Policy
Immediate discovery and Inventory or wait for user action
 - 3: Rack Management Connectivity Policy
Automatic discovery of Dual or Single Wire or require User input

User Acknowledged Server Discover Policy



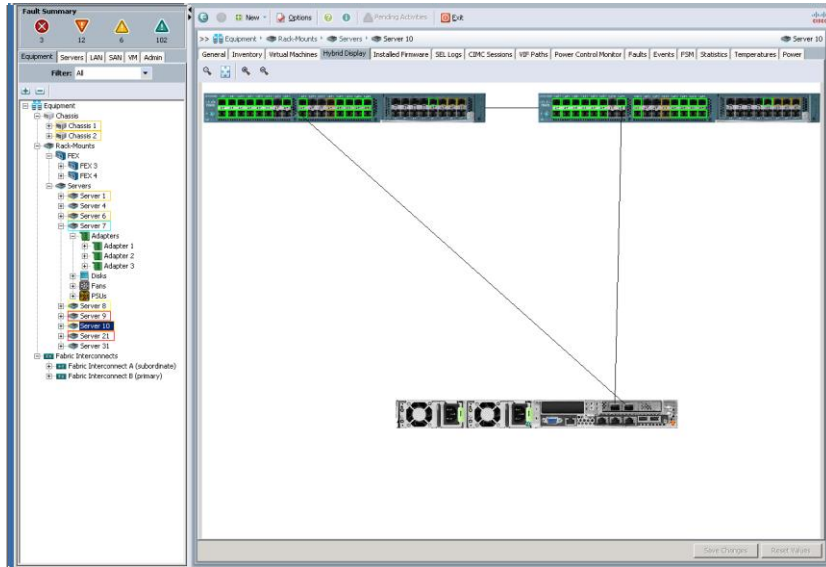
- Server shows up in UCSM, but discovery (inventory and availability) does not occur.
- Overall Status shows as Inoperable until user Re-acknowledges server.
- Server number assigned even though inventory not completed.

User Acknowledged Management Connectivity Policy



- Management connection not established until user provides input.
- Server->Inventory->CIMC
- Select Connection type, only available connection types will show.
- Once management connection type is chosen it can not be changed without a rediscovery (Decom/Recom)

Hybrid Connection Display

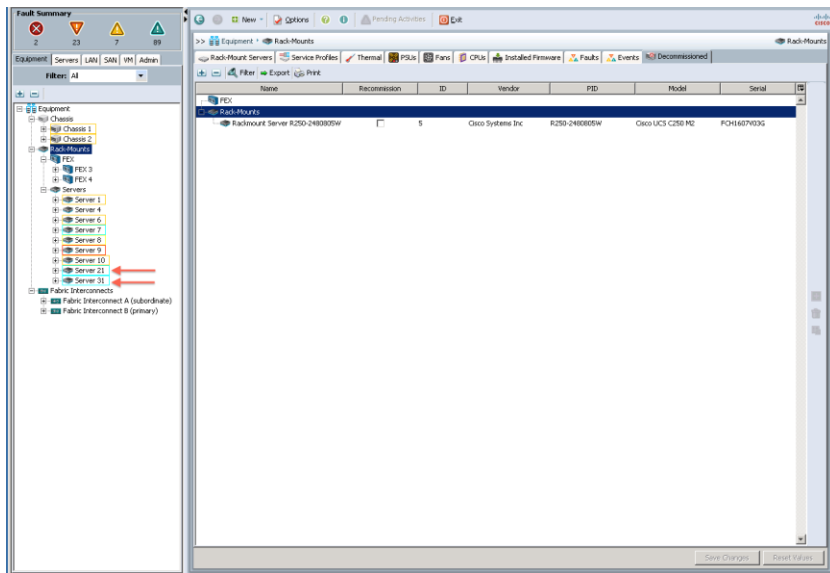


- For UCS C-Series servers you can view the connection path for management and data through the Equipment -> Server->Hybrid Display tab.
- This can be useful for troubleshooting connectivity problems and verifying management connection types.
- Red links indicate that there is not an active VIF on the link. This is normal for Dual-Wire servers that are not active.
- If you reattach an integrated server to different ports than it was discovered on you must decom/recom the server for proper operations.
- Hovering the mouse over the link will provide additional detail.



Server Discovery Video

Renumbering Servers and/or FEX



- Servers and FEX are numbered by when they come on-line. You only need to plug in the power cable and network cables to start discovery.
- You can control this by plugging power in one at a time.
- You can also renumber a FEX or server by decommissioning the device and renumbering prior to decommission.
- This should be done before you put the servers into production as it requires a reboot.

Best Practices with Discovery

- Use Auto-Acknowledge for Rack Server Discovery Policy even if you have a 160 server environment.
- If the server is Single Connect capable, wire for desired connectivity and use Auto-Acknowledge management policy.
 - Note if you wire LOMs and SFPs, Shared-LOM will be the system default mode.
 - If you remove the LOMs after shared-LOM is discovered you will lose management connectivity.
 - When using more than 1 VIC in Single Connect mode the primary card slot will be used for management. If that card has a catastrophic failure, data can continue on the second card, but management through UCSM would be lost.
 - Decom/Recom is Required to establish a different management path between shared-LOM and sideband, or in the event of a VIC failure.
- Enabling both User Acknowledge policies requires 2 re-acknowledgements
- Discover FEXs first and renumber them prior to discovering servers because you will have to take them offline and this could disrupt server discovery.

Troubleshooting Discovery Problems

```
CIMC Configuration Utility  Version 1.7  Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated:      [X]
Shared LOM:     [ ]
Cisco Card:     [ ]
Shared LOM Ext: [ ]
NIC redundancy
None:           [X]
Active-standby:[ ]
Active-active:  [ ]
IPv4 (Basic)
DHCP enabled:   [ ]
CIMC IP:        10.29.140.207
Subnetmask:     255.0.0.0
Gateway:        10.29.140.1
Factory Defaults
CIMC Factory Default:[X] ←
Default User (Basic)
Default password:
Reenter password:
VLAN (Advanced)
VLAN enabled:   [ ]
VLAN ID:        1
Priority:        0
Port Profile
Reset:          [ ]
Name:
Port Properties
Auto Negotiation: [X]
Speed[1000/100 Mbps]: 1000
Duplex mode[half/full]: full
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
```

- Main Causes for Server Discovering Properly
 - CIMC Settings not default. (Most typical problem for servers not fresh out of the box)
 - Discovery Policy or Management connection policy set set to User-Acknowledged
 - Server Not in UCSM Catalog
 - CIMC or VIC Firmware not at appropriate levels (CIMC Minimum for Discover 1.4(4).
 - Improper Connectivity

Troubleshooting Discovery Problems

The screenshot shows the Cisco Prime Network Manager interface for a server named 'Server 10'. The 'FSM Status' is 'Fail'. The current FSM Name is 'Discover', completed at '2014-02-20T20:35:41'. The progress status is 0%. The remote invocation result is 'Intermittent Error' with error code '241' and description 'Disconnected CIMC links'. Below this, a 'Step Sequence' table lists various discovery steps, all of which are marked as 'Skip'.

Order	Name	Description	Status	Timestamp	Try
41	Discover Bmc Unconfig Prnu OS		Skip	1969-12-31T16:00:00	0
42	Discover Hag Disconnect		Skip	1969-12-31T16:00:00	0
43	Discover Serial Debug Discon...		Skip	1969-12-31T16:00:00	0
44	Discover Sol Redirect Disable		Skip	1969-12-31T16:00:00	0
45	Discover Sww Config Port Niv L...		Skip	1969-12-31T16:00:00	0
46	Discover Sww Config Port Niv P...		Skip	1969-12-31T16:00:00	0
47	Discover Nic Inventory Local		Skip	1969-12-31T16:00:00	0
48	Discover Nic Inventory Peer		Skip	1969-12-31T16:00:00	0
49	Discover Config Ww Mode		Skip	1969-12-31T16:00:00	0
50	Discover Bmc Shutdown Disco...		Skip	1969-12-31T16:00:00	0
51	Discover Handle Pooling		Skip	1969-12-31T16:00:00	0

Below the table, there are fields for Name, Status, Description, Order, Try, and Timestamp. At the bottom, there is a 'Scheduled FSM Tasks' section and buttons for 'Save Changes' and 'Reset Values'.

- Improper Cabling
- LOM and Fabric Adapter on different FEX
- LOM connected to FEX with VIC directly connected to Fabric.



Firmware Management

Initial Firmware Policies

- Servers can be discovered as long as they are running an appropriate level of Firmware
 - Version 1.4 and 2.0(1) would only run and discover with CIMC FW version 1.2(2l)
 - Starting with Version 2.0(2) any CIMC FW after the minimum 1.4(4) would allow discovery.
- Only the FW packaged with an appropriate version of UCS Software bundles are supported with UCSM integration.
- After Discovery all servers should be flashed (upgraded/downgraded) to the supported bundle.
- Starting with version 2.1(x) UCSM supports N-1 firmware builds. For example if you are integrating a server that has CIMC code supported in 2.1(1a) into a UCSM 2.1(3) environment the FW will be supported.
- Each version release note describes mixed version support.

Understanding Stand-Alone and UCSM FW Packages

C220 M3 Stand-Alone FW

- C-Series Stand Alone Versions*

- 1.4(4)
- 1.4(5)
- 1.4(6)
- 1.4(7)
- 1.5(1)
- 1.5(2)
- 1.5(3)
- 1.5(4)

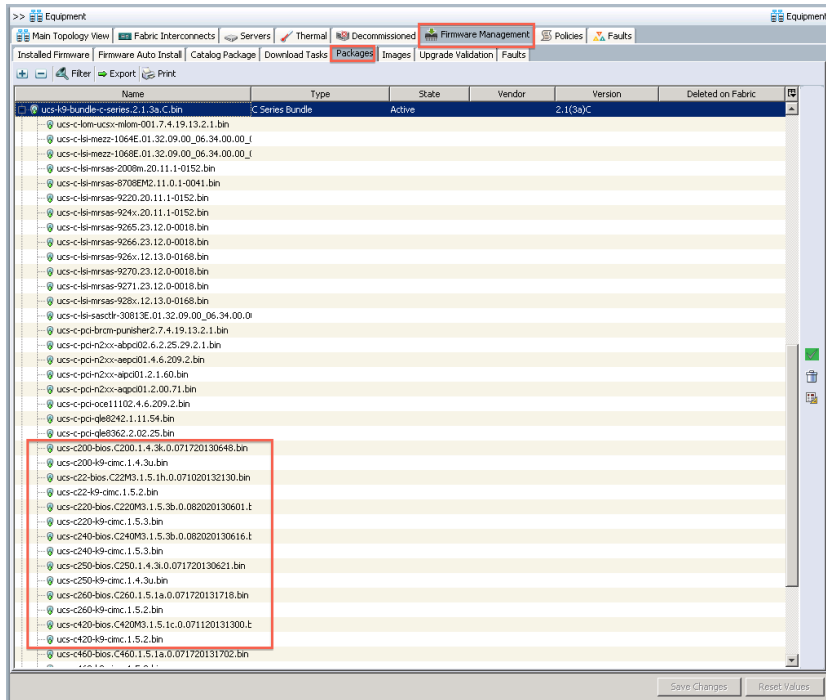
* Does not include patch release detail

C220 M3 UCSM Integrated FW

- UCSM Build Packages*

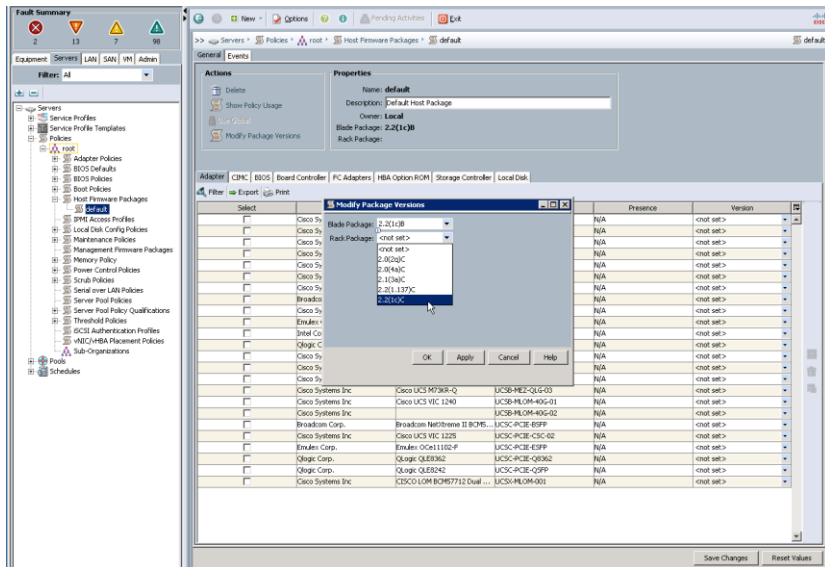
- 2.0(2) Package
 - C-Series FW Build 1.4(4)
- 2.0(3) Package
 - C-Series FW Build 1.4(5)
- 2.0(4) Package
 - C-Series FW Build 1.4(5)
- 2.0(5) Package
 - C-Series FW Build 1.4(5)
- 2.1(1)
 - C-Series FW Build 1.4(7)
- 2.1(2)
 - C-Series FW Build 1.5(1)
- 2.1(3)
 - C-Series FW Build 1.5(3)
- 2.1(4)
 - C-Series FW Build 1.5(4)

Examining UCSM Firmware Bundles



- Each UCSM build has a specific C-Series Package associated with it.
- You can examine the contents of the bundles.
- This can be useful when consulting the Hardware Compatibility Matrix (HCL).
- Not every server will use the same FW build.

Firmware Packages



- Servers DO NOT have to have the same version as the package to be discovered.
- To be supported the firmware must be flashed to a version supported under UCSM.
- Firmware packages are used to upgrade/downgrade integrated servers.
- Best practice when integrating servers is to flash with the Firmware for a given build
- N-1 is supported for C-Series integrated Servers
- Default policy is applied when servers are discovered.


A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a modern pedestrian bridge with blue lighting spans the street. Tall buildings with illuminated windows and storefronts line the street, and several flags are visible on poles to the left.

Network Placement Policies

Placement Policies Overview


- An advantage of UCS C-Series servers is multiple PCIe Adapters.
- By Default UCS will automatically place vNIC/vHBA adapters onto physical adapters.
- A network Placement policy determines which physical adapter a specific vNIC and/or vHBA will be assigned to.
- When using multiple cards and card types it is important to manage the vNICs/vHBAs to achieve the desired network connectivity.
- VCons are abstractions of physical adapters and provide a method for consistent PCIe mappings for a service profile regardless of how many physical cards are installed.

vCon Assignments (Linear)



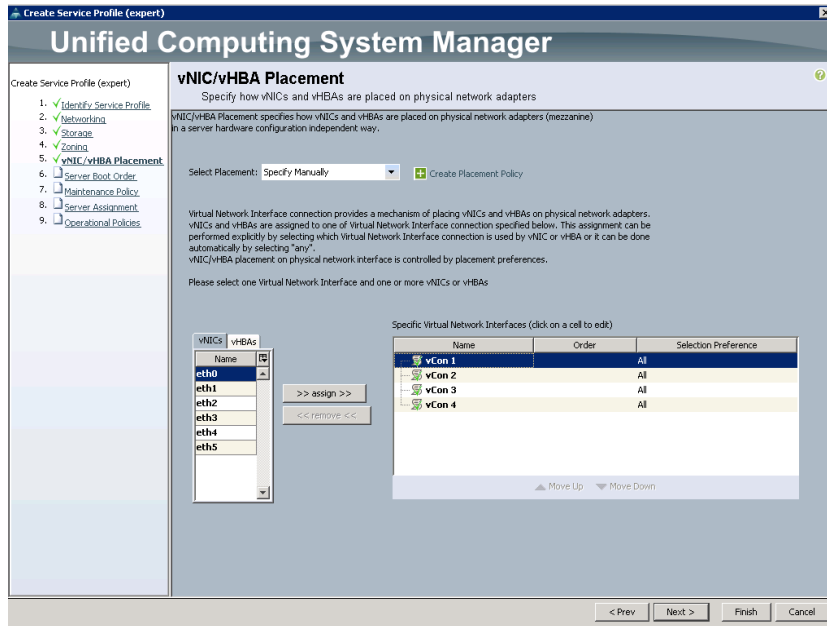
- The Service Profile was designed to have mobility between servers with different resources.
- vCons were designed to provide an abstraction of physical adapters to provide a mechanism for this mobility.
- When building a service profile there are Implicit (system) and Explicit (user) assignments of vNICs and vHBAs to vCONS

vCon Assignments (Round-robin)



- By default the system will place vNICs and vHBAs across 4 vCons and map vCons to adapters.
- Initially this was round robin, but changed to linear in version 2.1.x and above.
- The change allowed for consistent PCIe numbering when the number of adapters in the system changed.

vCon Placement



- Setting the desired order for vNICs and vHBAs allows the system to automatically place them onto vCONs.
- When creating a service profile you can explicitly place the vNICs to provide deterministic connectivity.
- This is accomplished through vNIC/vHBA placement

vCon to Adapter Mapping

Linear Mapping

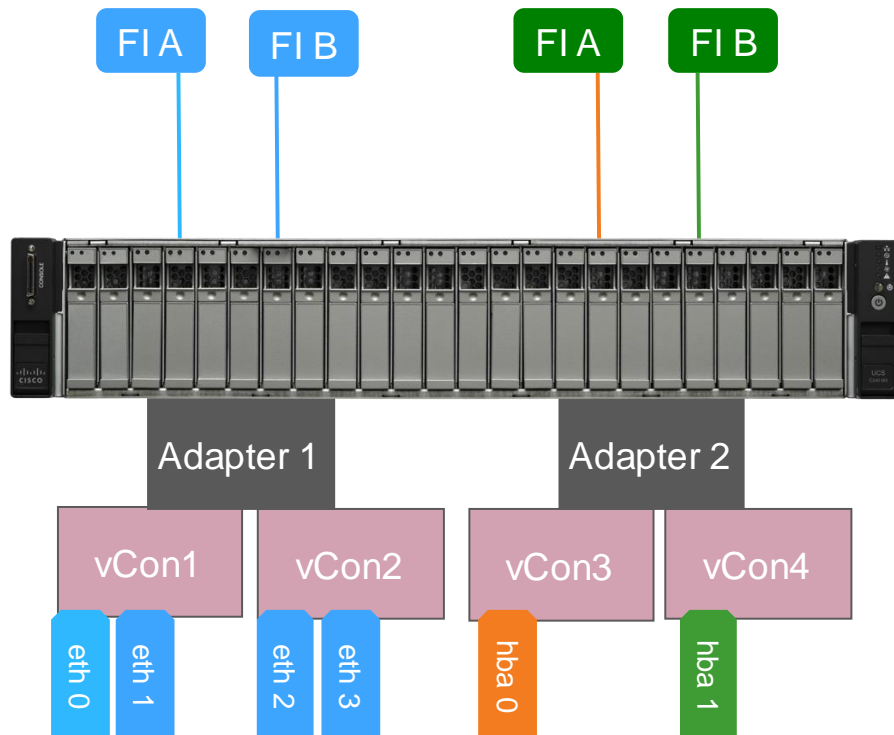
# of Adapters	vCon-1	vCon-2	vCon-3	vCon-4
1	1	1	1	1
2	1	1	2	2
3	1	2	3	3
4	1	2	3	4

Round- Robin Mapping

# of Adapters	vCon-1	vCon-2	vCon-3	vCon-4
1	1	1	1	1
2	1	2	1	2
3	1	2	3	3
4	1	2	3	4

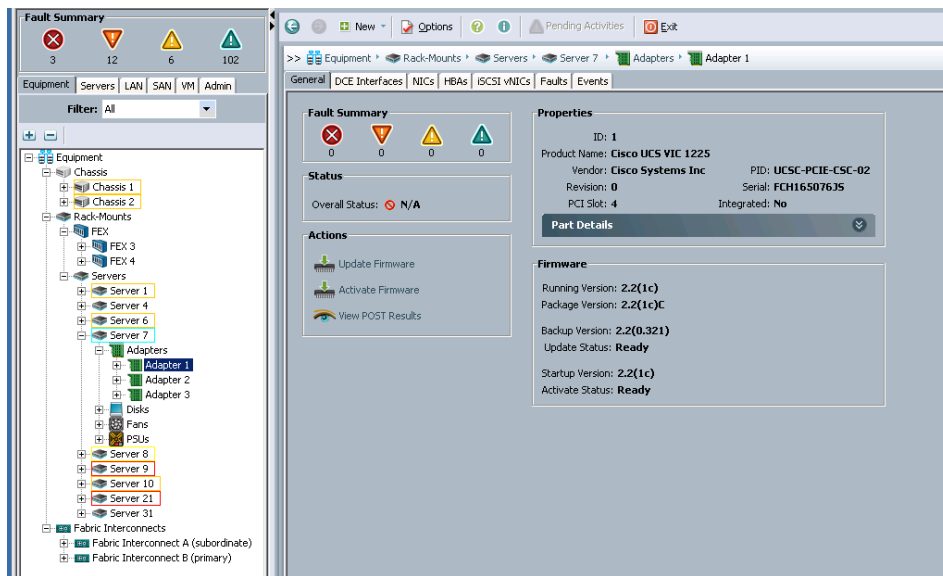
- vCons are automatically mapped to the system adapters. The assignment is dynamic.
- Users can configure a policy that provides round-robin, or linear order.
- As of 2.1.x linear is the default and preferred method.

Reasons for Specifying Placement Policies



- There are 2 main reasons to manually specify vNIC/vHBA placement for C-Series integrated servers.
 - Preventing automatic mapping to adapters that are unused by UCSM (e.g. C460 LOMs or Externally connected adapters)
 - Traffic engineering (e.g. FCoE on one adapter and Ethernet on the other).

Identifying Adapter Inventory in UCSM



- In order to control vNIC/vHBA placement you need to identify the adapter number in UCSM.
- Adapters can be identified through the equipment tab.
- Use the information in this tab along with the knowledge of how vCONs are mapped to adapters by the system achieve the desired placement.

Placement Policy Configuration – Best Practices

- Placement policies are designed to provide consistent PCIe numbering when moving profiles between servers with a different number of adapters.
- Automatic placement will be sufficient unless you have a setup that requires strict control of vNIC/vHBA Placement. Examples Include:
 - Servers with onboard adapters manageable through UCSM that you do not wish to use (e.g. C460).
 - Using UCSM discoverable adapters to connect to external networks (e.g. Broadcom, Intel, etc.)
 - Strict separation of FCoE traffic and Ethernet Traffic.

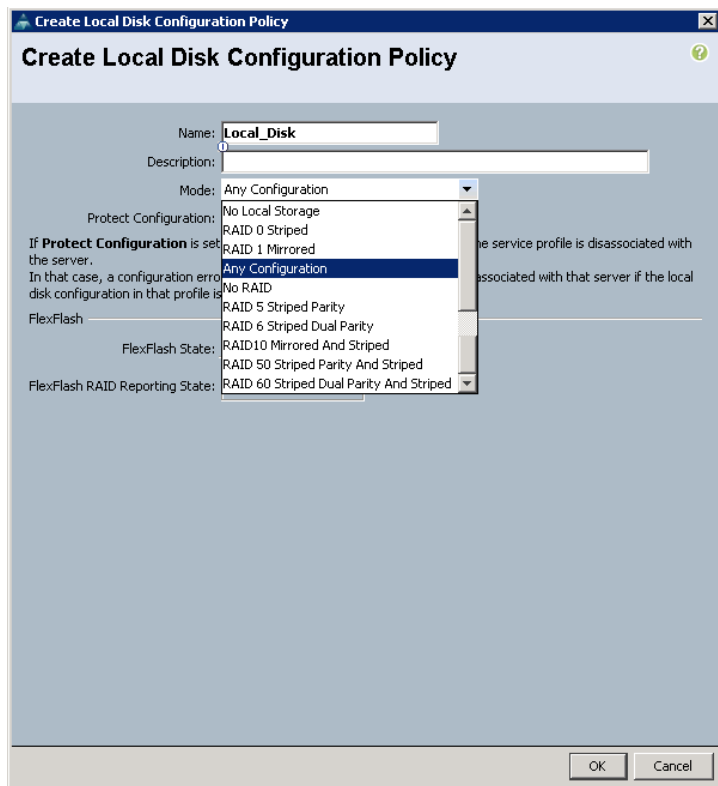
A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and tall buildings with lit windows and signage line the street. The overall scene is a dynamic urban environment.

Local Drive Configuration Considerations

Understanding UCSM Local Disk Policies


- Aside from PCIe slots the other major form factor decision for choosing a C-Series servers is the quantity and flexibility of drive configurations.
- Blade servers support a maximum of 4 drives where as C-Series rack mount servers support up to 24 drives today.
- UCSM provides management of local drive configurations through local disk policies which are part of the service profiles.
- The current implementation of local disk policies is limited in scope and does **not yet** meet some of the requirements for customers. To compensate users can continue to use tools outside of UCSM to customise drive configurations.
- The key to configuring local drives for a C-Series server is an understanding of local disk policies.
- Firmware management of controllers and drives is also achieved through service profiles.

Using Local Disk Policies




- A local disk policy describes how a logical (virtual) drive will be configured on the raid controller using the resources available.
- Depending on the controller capabilities the options for local disk configuration include.
 - RAID 0, 1, 5, 6, 10, 50, 60
 - NO RAID
 - Any Configuration
 - No Local Storage

Using Local Disk Policies



- RAID 0, 1, 5, 6, 10, 50, & 60 will create a single logical drive using all the disks in the system.
- No RAID will create a single drive RAID 0 logical drive using each physical drive in the system.
 - For example a C240 with 24 drives would present 24 logical drives to BIOS using this setting.
 - For the B200 M3 and B22 M3 this configuration will enable JBOD mode and present the drives as pass-through.
 - This is useful for some storage implementations like Hadoop or Ceph.

Using Local Disk Policies



- No local storage will only allow a service profile to be associated with a server that has no local disks in its inventory.
- Any configuration instructs UCSM not to manage the local storage configuration when applying the service profile.
 - This allows the user to set up whatever local drive configuration they wish and prevents UCSM from managing the local drives.
- Note: With Any Configuration UCSM will continue to monitor the drives and report any faults or errors.

Protect Configuration

Create Local Disk Configuration Policy

Name: Local_Disk

Description:

Mode: Any Configuration

Protect Configuration: ☒

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State: ☒ Disable ☐ Enable

FlexFlash RAID Reporting State: ☒ Disable ☐ Enable

OK Cancel

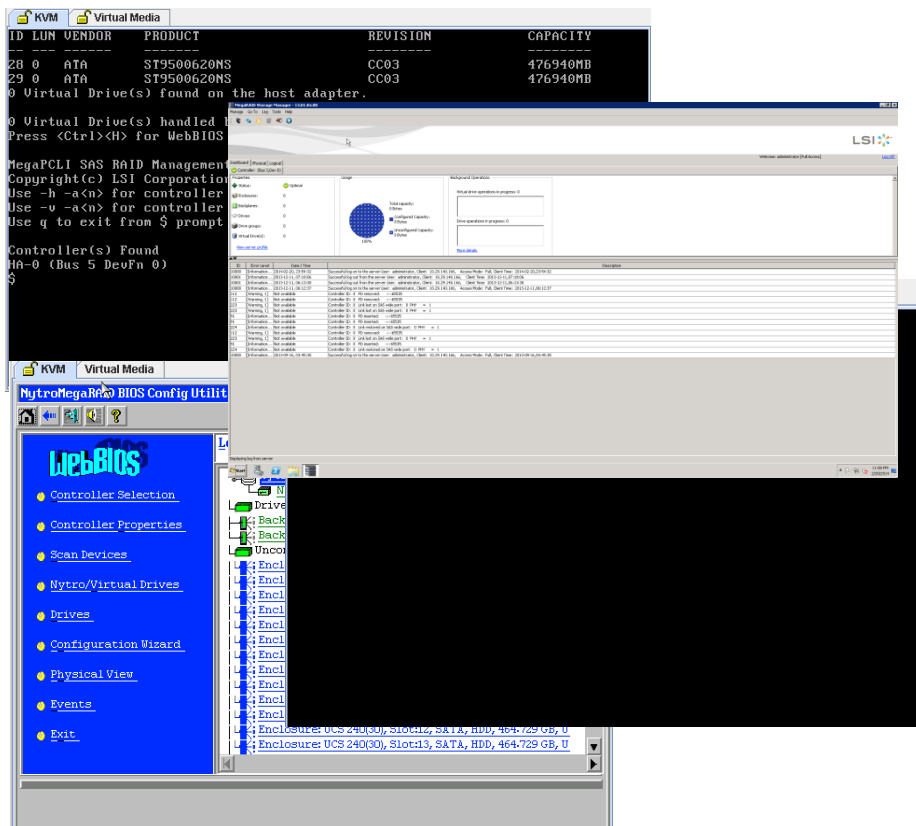
- An EXTREMELY IMPORTANT configuration item in the Local Disk policy is Protect Configuration.
- Protect configuration prevents a service profile with a local disk policy that DOES NOT match the actual configuration on the controller from being applied.
- This will protect the data on an existing volume during Service Profile Mobility.
- This setting should be set to on in almost ALL cases.

Protect Configuration Example



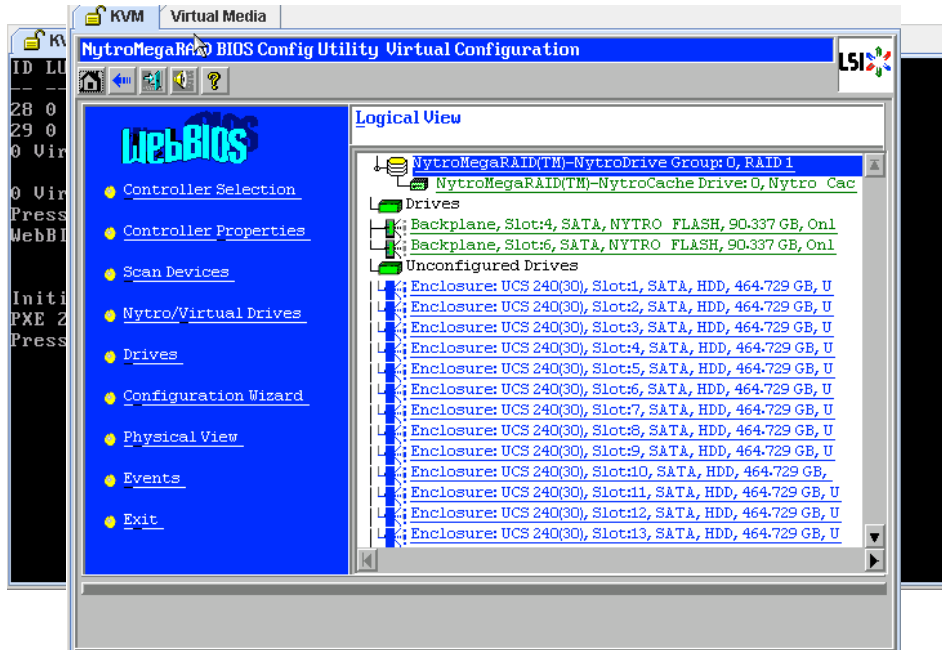
- Protect configuration ONLY applies during a service profile association.
- Let's look at an example.
 - C-Series Server with a local disk policy of Any Configuration.
 - Current Drive configuration is 2 volumes.
 - RAID 1 across the first 2 drives.
 - RAID 0 across the remaining 22 drives.
 - Removing the service profile will not affect the configuration or data (unless a scrub policy is set).
 - Associate a different service profile with a local disk policy of RAID 0 without protect configuration enabled.
 - Existing Drive volumes are destroyed and a new RAID 0 with 24 drives is created.

Local Drive Management Tools Outside of UCSM



- Managing Logical Drive configuration outside of UCSM is fully supported with a local disk policy of any configuration.
- The tools available to achieve this are:
 - LSI Preboot CLI
 - LSI WEBBIOS
 - LSI MegaCLI or StoreCLI
 - LSI MegaRAID Storage Manager
- MegaCLI, StoreCLI, and MegaRAID storage manager require applications to be loaded on the host OS.
- Scripting of drive configuration can be achieved with a Linux image that includes MegaCLI or StoreCLI.

WebBIOS



- Unless you are scripting through a Linux Image WEBBIOS is the most common tool for management outside of UCSM.
- KVM to the console and look for the prompt to enter WEBBIOS and issue the appropriate keystroke (e.g. Ctrl-H).
- Use the utility to configure virtual drives then reboot the system they will be visible in UCSM after the reboot (version 2.2.1 and above).

Multi-Volume Drive Management Best Practices

- **ALWAYS** use protect configuration when moving service profiles between servers unless your intent is to reconfigure the logical drives on the controller thereby destroying all data on the drives.
- If you need any configuration other than a single volume or all the drives to be individual RAID 0 volumes use a local disk policy of any configuration.
- If you do only need a single volume and think you may grow your existing volume or create an additional volume in the future use a local disk policy of any configuration.
- If you have an existing stand-alone C-Series server with logical drives and data on those drives you can still integrate without losing the data if you associate a service provide that has a local disk policy of Any configuration.
- Did I happen to mention that you should **ALMOST ALWAYS** use Any configuration as a local disk policy.*
- **Ask me** about the asterisks in the previous statement.

A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a modern pedestrian bridge with blue lighting spans the street. Tall buildings with illuminated windows and storefronts line the street, and several flags are visible on the left. The overall scene is a dynamic urban environment.

Case Study

Large Scale Integrated C-Series Deployment



- Big Data Implementation
- 700+ Servers in 5 Pods (Initial Install was 3 Pods)
- Mix of C200/C220/C210/C240/C460
- Fusion I/O Drives (PCIe)
- Large Disk footprint
- Short deployment cycle

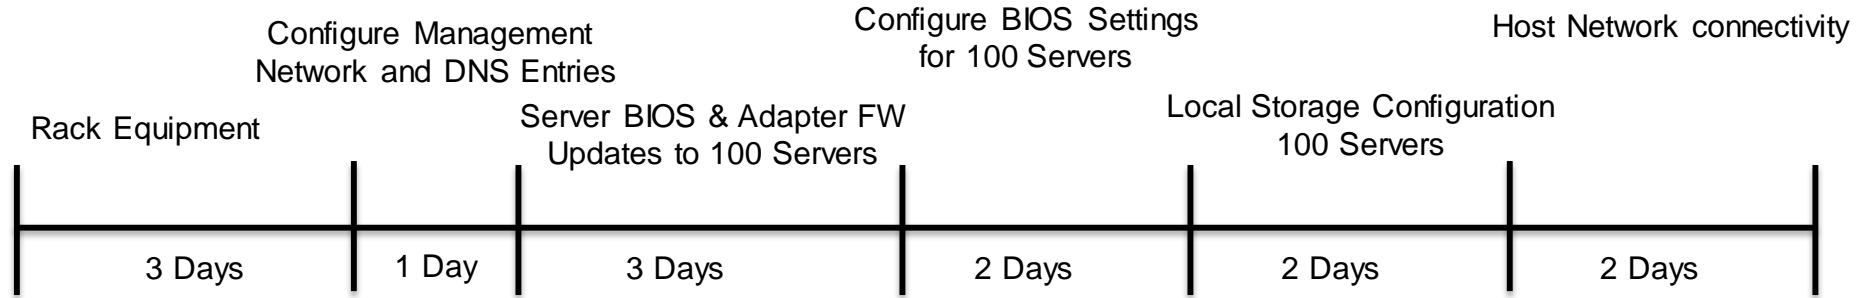
Cisco live!

Typical Server Install Project Tasks

- Rack & Cable Hardware: 2-3 days
- Management Network Configuration: 1-2 days
- DHCP/DNS Configuration: 1-2 days
- BIOS/FW Management: 1-2 hours per server
- BIOS settings: 10-15 minutes per server
- Storage Management (Local): 10-15 minutes per server
- Server Network Configuration: 1-2 days
- OS Install: 15-20 Minutes per server
- Agent Installation: 15 minutes per server
- Monitoring/Troubleshooting: Ongoing

Install Timeline

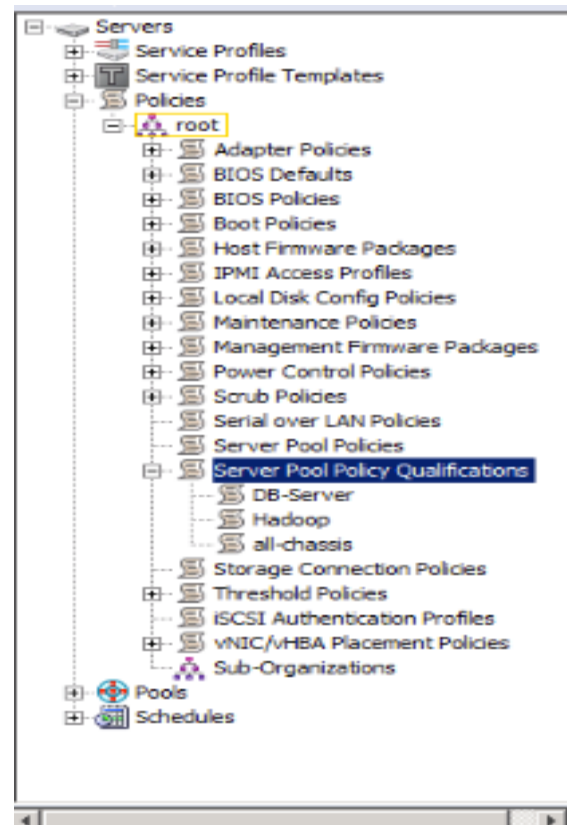
- Typical Timeline to get 100 to a point of OS install (assuming some items can be done in Parallel):



- 13 Business Days (3 WEEKS) – to get to the point of OS install.

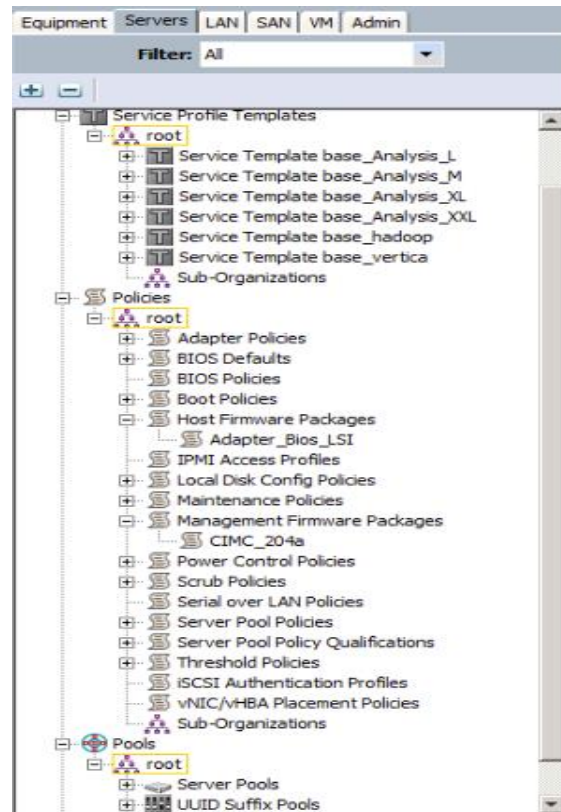
Using Qualification Policies to Pool Servers

- Customer uses a variety of server configurations for different functions
 - C240 10 Drive 96GB Memory
 - C220 8 Drive 96GB Memory
 - C240 2 Drive 128GB Memory
 - C460 10 Drive 256GB Memory
- Each Server has a unique function and uses a different service profile.
- By building qualification policies we were able to identify and add each server to a pool without any manual inspection.
- Service Profile Templates were bound to a pool for SP assignment so that when a server of a particular type was needed and the SP created it picked a physical resource from the pool that matched the qualification.



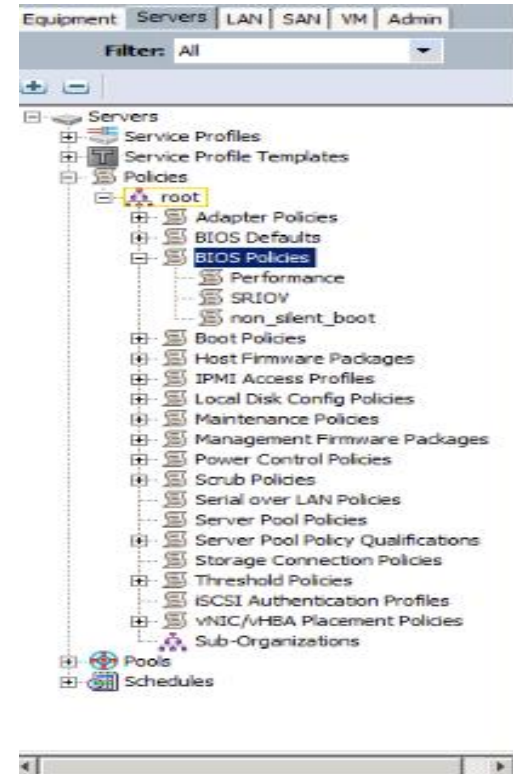
Using Firmware Packages

- When rack servers come from a vendor you almost never know what BIOS and FW will be installed.
- Using Host Firmware Packages provides consistent FW and BIOS loads across all servers in the system.
- Using a Management Firmware package provides a consistent build for the CIMC (management controller) for the server



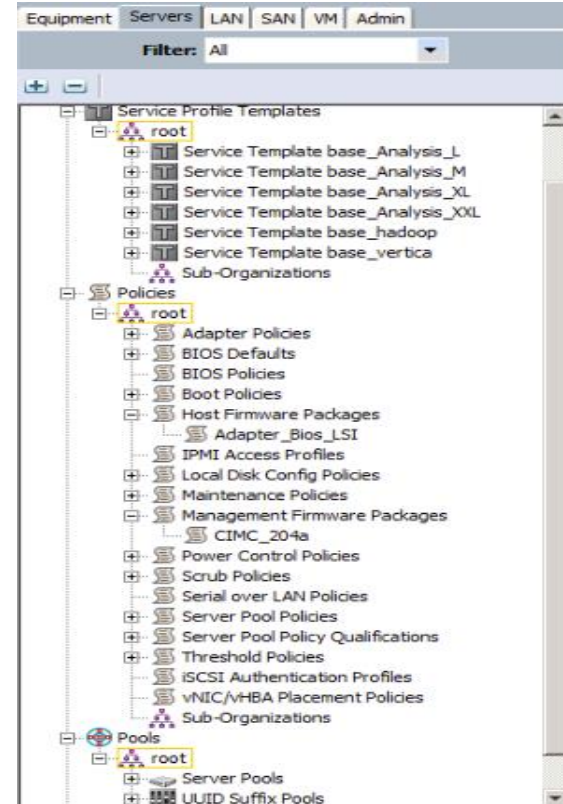
Using BIOS Policies

- Depending on the server type there may be a need to modify BIOS settings for any of the following reasons:
 - Disable any unneeded BIOS settings
 - Improve Memory Performance
 - Control Memory Scrubbing
- By creating a BIOS policy there is no longer a need for anyone to enter the BIOS of each server to change settings.



Using Service Profile Templates

- The final step was to tie it all together using service profile templates.
- A template was created for each server type.
- Policies were used in the template and then when a SP was created it became associated with a piece of HW from the pool.
- Updating templates allow wide spread changes to the system.



Parallel Tasks

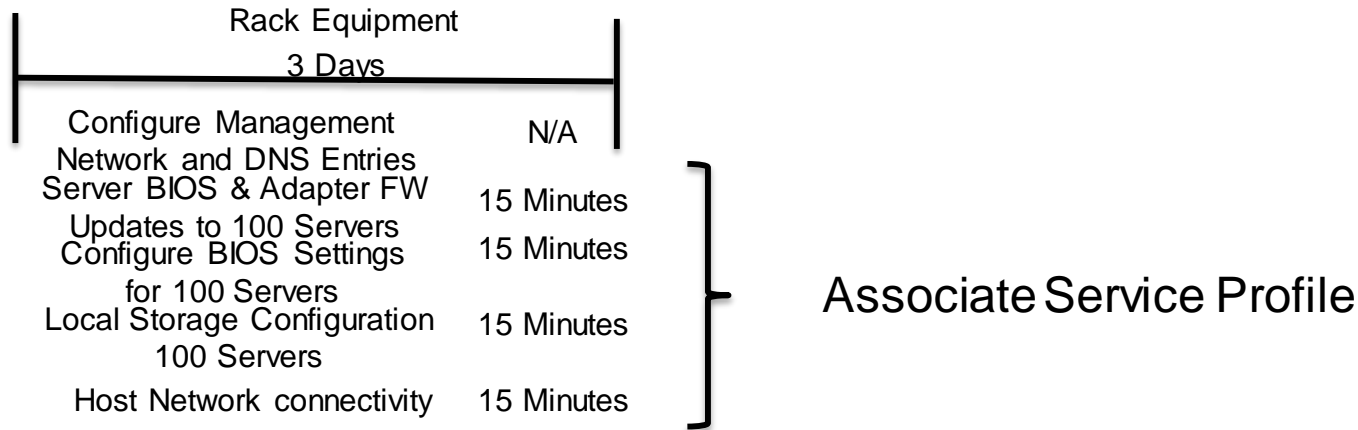
- Once the Fabric Interconnects are racked and initial setup is performed. All of the server configuration tasks can be accomplished without having any hardware installed.
- Discovery can take place at anytime during the process. This greatly reduces the time to deployment

Typical Server Install Project Tasks with UCSM

- Rack & Cable Hardware: 2-3 days
- Management Network Configuration: 5-10 minutes total
- DHCP/DNS Configuration: Not required
- BIOS/FW Management: 10-15 minutes total
- BIOS settings: 10-15 minutes total
- Storage Management (Local): 10-15 minutes total
- Server Network Configuration: 10-15 minutes total
- OS Install: 15-20 Minutes per server
- Agent Installation: Not required
- Monitoring/Troubleshooting: Ongoing

How Long Does It Take With UCSM

- Typical Timeline to get 100 to a point of OS install (SP configuration can be completed before equipment racking is completed):



- 3 Business Days or less – to get to the point of OS install.

Reference Materials

- UCS C-Series Integration Guide

- http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c-series_integration/ucsm2-2/b_C-Series-Integration_UCSM2-2.html

- UCSM GUI Service Profile Configuration Guide

- http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/2-2/b_UCSM_GUI_Configuration_Guide_2_2/b_UCSM_GUI_Configuration_Guide_2_1_chapter_010_0001.html

- LSI MegaRAID SAS Software Users Guide

- https://www.lsi.com/downloads/Public/MegaRAID%20Common%20Files/SASUG_51530-00_RevG.pdf

Participate in the “My Favorite Speaker” Contest

Promote Your Favorite Speaker and You Could be a Winner

- Promote your favorite speaker through Twitter and you could win \$200 of Cisco Press products (@CiscoPress)
- Send a tweet and include
 - Your favorite speaker’s Twitter handle @smcquerry
 - Two hashtags: #CLUS #MyFavoriteSpeaker
- You can submit an entry for more than one of your “favorite” speakers
- Don’t forget to follow @CiscoLive and @CiscoPress
- View the official rules at <http://bit.ly/CLUSwin>

Continue Your Education

- Demos in the Cisco Campus
- Walk-in Self-Paced Labs
- Meet the Expert 1:1 meetings

A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and modern buildings with lit windows and signage line the street. Traffic lights are visible in the distance.

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2015 T-Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site
<http://showcase.genie-connect.com/clmelbourne2015>
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm



Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

Cisco *live!*

Thank you.



CISCO