Server Policies

This chapter includes the following sections:

Ethernet and Fibre Channel Adapter Policies

These policies govern the host-side behavior of the adapter, including how the adapter handles traffic. For example, you can use these policies to change default settings for the following:

- Queues
- Interrupt handling
- Performance enhancement
- RSS hash
- Failover in an cluster configuration with two fabric interconnects

Note For Fibre Channel adapter policies, the values displayed by Cisco UCS Central may not match those displayed by applications such as QLogic SANsurfer. For example, the following values may result in an apparent mismatch between SANsurfer and Cisco UCS Central:

- Max LUNs Per Target—SANsurfer has a maximum of 256 LUNs and does not display more than that number. Cisco UCS Central supports a higher maximum number of LUNs.
- Link Down Timeout—In SANsurfer, you configure the timeout threshold for link down in seconds. In Cisco UCS Central, you configure this value in milliseconds. Therefore, a value of 5500 ms in Cisco UCS Central displays as 5s in SANsurfer.
- Max Data Field Size—SANsurfer has allowed values of 512, 1024, and 2048. Cisco UCS Central allows you to set values of any size. Therefore, a value of 900 in Cisco UCS Central displays as 512 in SANsurfer.

Operating System Specific Adapter Policies

By default, Cisco UCS provides a set of Ethernet adapter policies and Fibre Channel adapter policies. These policies include the recommended settings for each supported server operating system. Operating systems are sensitive to the settings in these policies. Storage vendors typically require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.

Important We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for a Windows OS (instead of using the default Windows adapter policy), you must use the following formulas to calculate values that work with Windows:

Completion Queues = Transmit Queues + Receive Queues
Interrupt Count = (Completion Queues + 2) rounded up to nearest power of 2

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

Completion Queues = 1 + 8 = 9
Interrupt Count = (9 + 2) rounded up to the nearest power of 2 = 16

Creating an Ethernet Adapter Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Adapter Policies** and choose **Create Ethernet Adapter Policy**.

Step 4 In the **Create Ethernet Adapter Policy** dialog box, enter the **Name** and optional description.

Step 5 In the **Resources** area, enter the **Transmit Queues**, **Receive Queues**, and **Completion Queues**, and the **Ring Size** for each queue.

Step 6 In the **Options** area, choose the **Transmit Checksum Offload**, **Receive Checksum Offload**, **TCP Segmentation Offload**, **TCP Large Receive Offload** and **Receive Side Scaling (RSS)**.

Step 7 Enter the **Failback Timeout (Seconds)**, choose the **Interrupt Mode** and **Interrupt Coalescing Type**, and enter the **Interrupt Time (us)**

Step 8 Click **OK**.

Deleting an Ethernet Adapter Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Adapter Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Server BIOS Settings

Cisco UCS provides two methods for making global modifications to the BIOS settings on servers in an Cisco UCS domain. You can create one or more BIOS policies that include a specific grouping of BIOS settings that match the needs of a server or set of servers, or you can use the default BIOS settings for a specific server platform.

Both the BIOS policy and the default BIOS settings for a server platform enable you to fine tune the BIOS settings for a server managed by Cisco UCS Central.

Depending upon the needs of the data center, you can configure BIOS policies for some service profiles and use the BIOS defaults in other service profiles in the same Cisco UCS domain, or you can use only one of them. You can also use Cisco UCS Central to view the actual BIOS settings on a server and determine whether they are meeting current needs.

Note Cisco UCS Central pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Main BIOS Settings

The following table lists the main server BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Reboot on BIOS Settings Change	When the server is rebooted after you change one or more BIOS settings. If you enable this setting, the server is rebooted according to the maintenance policy in the server's service profile. For example, if the maintenance policy requires user acknowledgment, the server is not rebooted and the BIOS changes are not applied until a user acknowledges the pending activity. If you do not enable this setting, the BIOS changes are not applied until the next time the server is rebooted, whether as a result of another server configuration change or a manual reboot.
Quiet Boot	What the BIOS displays during Power On Self-Test (POST). This can be one of the following: • **disabled—**The BIOS displays all messages and Option ROM information during boot. • **enabled—**The BIOS displays the logo screen, but does not display any messages or Option ROM information during boot. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Post Error Pause	What happens when the server encounters a critical error during POST. This can be one of the following: • **disabled—**The BIOS continues to attempt to boot the server. • **enabled—**The BIOS pauses the attempt to boot the server and opens the Error Manager when a critical error occurs during POST. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Resume Ac On Power Loss	How the server behaves when power is restored after an unexpected power loss. This can be one of the following: • **stay-off—**The server remains off until manually powered on. • **last-state—**The server is powered on and the system attempts to restore its last state. • **reset—**The server is powered on and automatically reset. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front Panel Lockout	Whether the power and reset buttons on the front panel are ignored by the server. This can be one of the following: • **disabled—**The power and reset buttons on the front panel are active and can be used to affect the server. • **enabled—**The power and reset buttons are locked out. The server can only be reset or powered on or off from the CIMC GUI. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Processor BIOS Settings

The following table lists the processor BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Turbo Boost	Whether the processor uses Intel Turbo Boost Technology, which allows the processor to automatically increase its frequency if it is running below power, temperature, or voltage specifications. This can be one of the following: • **disabled—**The processor does not increase its frequency automatically. • **enabled—**The processor uses Turbo Boost Technology if required. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Enhanced Intel Speedstep	Whether the processor uses Enhanced Intel SpeedStep Technology, which allows the system to dynamically adjust processor voltage and core frequency. This technology can result in decreased average power consumption and decreased average heat production. This can be one of the following: • **disabled—**The processor never dynamically adjusts its voltage or frequency. • **enabled—**The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Hyper Threading	Whether the processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following: • **disabled—**The processor does not permit hyperthreading. • **enabled—**The processor allows for the parallel execution of multiple threads. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure the operating system supports this feature.

Name	Description
Core Multi Processing	Sets the state of logical processor cores per CPU in a package. If you disable this setting, Intel Hyper Threading technology is also disabled. This can be one of the following: • **all—**Enables multiprocessing on all logical processor cores. • **1** through ***n*—**Specifies the number of logical processor cores per CPU that can run on the server. To disable multiprocessing and have only one logical processor core per CPU running on the server, choose 1. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Execute Disabled Bit	Classifies memory areas on the server to specify where the application code can execute. As a result of this classification, the processor disables code execution if a malicious worm attempts to insert code in the buffer. This setting helps to prevent damage, worm propagation, and certain classes of malicious buffer overflow attacks. This can be one of the following: • **disabled—**The processor does not classify memory areas. • **enabled—**The processor classifies memory areas. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Virtualization Technology (VT)	Whether the processor uses Intel Virtualization Technology, which allows a platform to run multiple operating systems and applications in independent partitions. This can be one of the following: • **disabled—**The processor does not permit virtualization. • **enabled—**The processor allows multiple operating systems in independent partitions. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** If you change this option, you must power cycle the server before the setting takes effect.

Name	Description
Hardware Pre-fetcher	Whether the processor allows the Intel hardware prefetcher to fetch streams of data and instruction from memory into the unified second-level cache when necessary. This can be one of the following: • **Disabled—**The hardware prefetcher is not used. • **Enabled—**The processor uses the hardware prefetcher when cache issues are detected. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** must be set to **Custom** in order to specify this value. For any value other than **Custom**, this option is overridden by the setting in the selected CPU performance profile.
Adjacent Cache Line Pre-fetcher	Whether the processor fetches cache lines in even/odd pairs instead of fetching just the required line. This can be one of the following: • **Disabled—**The processor only fetches the required line. • **Enabled—**The processor fetches both the required line and its paired line. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** must be set to **Custom** in order to specify this value. For any value other than **Custom**, this option is overridden by the setting in the selected CPU performance profile.
DCU Streamer Pre-fetch	Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following: • **Disabled—**The processor does not try to anticipate cache read requirements and only fetches explicitly requested lines. • **Enabled—**The DCU prefetcher analyzes the cache read pattern and prefetches the next line in the cache if it determines that it may be needed. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
DCU IP Pre-fetcher	Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following: • **Disabled—**The processor does not preload any cache data. • **Enabled—**The DCU IP prefetcher preloads the L1 cache with the data it determines to be the most relevant. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Direct Cache Access	Allows processors to increase I/O performance by placing data from I/O devices directly into the processor cache. This setting helps to reduce cache misses. This can be one of the following: • **disabled—**Data from I/O devices is not placed directly into the processor cache. • **enabled—**Data from I/O devices is placed directly into the processor cache. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C State	Whether the system can enter a power savings mode during idle periods. This can be one of the following: • **disabled—**The system remains in a high-performance state even when idle. • **enabled—**The system can reduce power to system components such as the DIMMs and CPUs. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. We recommend that you contact your operating system vendor to make sure your operating system supports this feature.
Processor C1E	Allows the processor to transition to its minimum frequency upon entering C1. This setting does not take effect until after you have rebooted the server. This can be one of the following: • **disabled—**The CPU continues to run at its maximum frequency in the C1 state. • **enabled—**The CPU transitions to its minimum frequency. This option saves the maximum amount of power in the C1 state. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Processor C3 Report	Whether the processor sends the C3 report to the operating system. This can be one of the following: • **disabled—**The processor does not send the C3 report. • **acpi-c2—**The processor sends the C3 report using the advanced configuration and power interface (ACPI) C2 format. • **acpi-c3—**The processor sends the C3 report using the ACPI C3 format. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. On the Cisco UCS B440 Server, the BIOS Setup menu uses enabled and disabled for these options. If you specify acpi-c2 or acpi-c2, the server sets the BIOS value for that option to enabled.
Processor C6 Report	Whether the processor sends the C6 report to the operating system. This can be one of the following: • **disabled—**The processor does not send the C6 report. • **enabled—**The processor sends the C6 report. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C7 Report	Whether the processor sends the C7 report to the operating system. This can be one of the following: • **disabled—**The processor does not send the C7 report. • **enabled—**The processor sends the C7 report. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
CPU Performance	Sets the CPU performance profile for the server. This can be one of the following: • **enterprise—**For M3 servers, all prefetchers and data reuse are enabled. For M1 and M2 servers, data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • **high-throughput—**Data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • **hpc—**All prefetchers are enabled and data reuse is disabled. This setting is also known as high-performance computing.
Max Variable MTRR Setting	Allows you to select the number of mean time to repair (MTRR) variables. This can be one of the following: • **auto-max—**BIOS uses the default value for the processor. • **8—**BIOS uses the number specified for the variable MTRR. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Local X2 APIC	Allows you to set the type of Application Policy Infrastructure Controller (APIC) architecture. This can be one of the following: • **xapic—**Uses the standard xAPIC architecture. • **x2apic—**Uses the enhanced x2APIC architecture to support 32 bit addressability of processors. • **auto—**Automatically uses the xAPIC architecture that is detected. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Power Technology	Enables you to configure the CPU power management settings for the following options: • Enhanced Intel Speedstep Technology • Intel Turbo Boost Technology • Processor Power State C6 Power Technology can be one of the following: • **Disabled—**The server does not perform any CPU power management and any settings for the BIOS parameters mentioned above are ignored. • —The server determines the best settings for the BIOS parameters mentioned above and ignores the individual settings for these parameters. • **Performance—**The server automatically optimizes the performance for the BIOS parameters mentioned above. • **Custom—**The server uses the individual settings for the BIOS parameters mentioned above. You must select this option if you want to change any of these BIOS parameters. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Energy Performance	Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following: • **performance** • **balanced-performance** • **balanced-energy** • **energy-efficient** • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** must be set to **Custom** or the server ignores the setting for this parameter.

Name	Description
Frequency Floor Override	Whether the CPU is allowed to drop below the maximum non-turbo frequency when idle. This can be one of the following: • **Disabled—** The CPU can drop below the maximum non-turbo frequency when idle. This option decreases power consumption but may reduce system performance. • **Enabled—** The CPU cannot drop below the maximum non-turbo frequency when idle. This option improves system performance but may increase power consumption. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
P-STATE Coordination	Allows you to define how BIOS communicates the P-state support model to the operating system. There are 3 models as defined by the Advanced Configuration and Power Interface (ACPI) specification. • **HW_ALL—**The processor hardware is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package). • **SW_ALL—**The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a physical package), and must initiate the transition on all of the logical processors. • **SW_ANY—**The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package), and may initiate the transition on any of the logical processors in the domain. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** must be set to **Custom** or the server ignores the setting for this parameter.

Name	Description
DRAM Clock Throttling	Allows you to tune the system settings between the memory bandwidth and power consumption. This can be one of the following: • **Balanced—** DRAM clock throttling is reduced, providing a balance between performance and power. • **Performance—**DRAM clock throttling is disabled, providing increased memory bandwidth at the cost of additional power. • —DRAM clock throttling is increased to improve energy efficiency. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Channel Interleaving	Whether the CPU divides memory blocks and spreads contiguous portions of data across interleaved channels to enable simultaneous read operations. This can be one of the following: • **Auto—**The CPU determines what interleaving is done. • **1-way—**Some channel interleaving is used. • **2-way** • **3-way** • **4-way—**The maximum amount of channel interleaving is used. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Rank Interleaving	Whether the CPU interleaves physical ranks of memory so that one rank can be accessed while another is being refreshed. This can be one of the following: • **Auto—**The CPU determines what interleaving is done. • **1-way—**Some rank interleaving is used. • **2-way** • **4-way** • **8-way—**The maximum amount of rank interleaving is used. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Demand Scrub	Whether the system corrects single bit memory errors encountered when the CPU or I/O makes a demand read. This can be one of the following: • **Disabled—** Single bit memory errors are not corrected. • **Enabled—** Single bit memory errors are corrected in memory and the corrected data is set in response to the demand read. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Patrol Scrub	Whether the system actively searches for, and corrects, single bit memory errors even in unused portions of the memory on the server. This can be one of the following: • **Disabled—**The system checks for memory ECC errors only when the CPU reads or writes a memory address. • **Enabled—**The system periodically reads and writes memory searching for ECC errors. If any errors are found, the system attempts to fix them. This option may correct single bit errors before they become multi-bit errors, but it may adversely affect performance when the patrol scrub is running. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Altitude	This can be one of the following: • **—** • **—** • **—**

Name	Description
Altitude	The approximate number of meters above sea level at which the physical server is installed. This can be one of the following: • **Auto—**The CPU determines the physical elevation. • —The server is approximately 300 meters above sea level. • —The server is approximately 900 meters above sea level. • —The server is approximately 1500 meters above sea level. • —The server is approximately 3000 meters above sea level. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Package C State Limit **set PackageCStateLimit**	The amount of power available to the server components when they are idle. This can be one of the following: • —The server may enter any available C state. • —The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power. • —When the CPU is idle, the system slightly reduces the power consumption. This option requires less power than C0 and allows the server to return quickly to high performance mode. • —When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • —When the CPU is idle, the system reduces the power consumption further than with the C3 option. This option saves more power than C0, C1, or C3, but there may be performance issues until the server returns to full power. • —When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • —When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves the maximum amount of power but it also requires the longest time for the server to return to high performance mode. • —When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves more power than C7, but it also requires the longest time for the server to return to high performance mode. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Intel Directed I/O BIOS Settings

The following table lists the Intel Directed I/O BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
VT for Directed IO	Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following: • **disabled—**The processor does not use virtualization technology. • **enabled—**The processor uses virtualization technology. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.
Interrupt Remap	Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following: • **disabled—**The processor does not support remapping. • **enabled—**The processor uses VT-d Interrupt Remapping as required. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Coherency Support	Whether the processor supports Intel VT-d Coherency. This can be one of the following: • **disabled—**The processor does not support coherency. • **enabled—**The processor uses VT-d Coherency as required. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
ATS Support	Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following: • **disabled—**The processor does not support ATS. • **enabled—**The processor uses VT-d ATS as required. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Pass Through DMA Support	Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following: • **disabled—**The processor does not support pass-through DMA. • **enabled—**The processor uses VT-d Pass-through DMA as required. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

RAS Memory BIOS Settings

The following table lists the RAS memory BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Memory RAS Config	How the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following: • **maximum performance—**System performance is optimized. • **mirroring—**System reliability is optimized by using half the system memory as backup. • **lockstep—**If the DIMM pairs in the server have an identical type, size, and organization and are populated across the SMI channels, you can enable lockstep mode to minimize memory access latency and provide better performance. Lockstep is enabled by default for B440 servers. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
NUMA	Whether the BIOS supports NUMA. This can be one of the following: • **disabled—**The BIOS does not support NUMA. • **enabled—**The BIOS includes the ACPI tables that are required for NUMA-aware operating systems. If you enable this option, the system must disable Inter-Socket Memory interleaving on some platforms. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Mirroring Mode	Memory mirroring enhances system reliability by keeping two identical data images in memory. This option is only available if you choose the **mirroring** option for **Memory RAS Config**. It can be one of the following: • **inter-socket—**Memory is mirrored between two Integrated Memory Controllers (IMCs) across CPU sockets. • **intra-socket—**One IMC is mirrored with another IMC in the same socket. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Sparing Mode	Sparing optimizes reliability by holding memory in reserve so that it can be used in case other DIMMs fail. This option provides some memory redundancy, but does not provide as much redundancy as mirroring. The available sparing modes depend on the current memory population. This option is only available if you choose **sparing** option for **Memory RAS Config**. It can be one of the following: • **dimm-sparing—**One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM. • **rank-sparing—**A spare rank of DIMMs is held in reserve. If a rank of DIMMs fails, the contents of the failing rank are transferred to the spare rank. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
LV DDR Mode	Whether the system prioritizes low voltage or high frequency memory operations. This can be one of the following: • **power-saving-mode—**The system prioritizes low voltage memory operations over high frequency memory operations. This mode may lower memory frequency in order to keep the voltage low. • **performance-mode—**The system prioritizes high frequency operations over low voltage operations. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
DRAM Refresh Rate	This option controls the refresh interval rate for internal memory.

Serial Port BIOS Settings

The following table lists the serial port BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Serial Port A	Whether serial port A is enabled or disabled. This can be one of the following: • **disabled—**The serial port is disabled. • **enabled—**The serial port is enabled. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

USB BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Make Device Non Bootable	Whether the server can boot from a USB device. This can be one of the following: • **disabled—**The server can boot from a USB device. • **enabled—**The server cannot boot from a USB device. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Legacy USB Support	Whether the system supports legacy USB devices. This can be one of the following: • **disabled—**USB devices are only available to EFI applications. • **enabled—**Legacy USB support is always available. • **auto—**Disables legacy USB support if no USB devices are connected. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB System Idle Power Optimizing Setting	Whether the USB System Idle Power Optimizing setting is used to reduce USB EHCI idle power consumption. Depending upon the value you choose, this setting can have an impact on performance. This can be one of the following: • **high-performance—**The USB System Idle Power Optimizing setting is disabled, because optimal performance is preferred over power savings. Selecting this option can significantly improve performance. We recommend you select this option unless your site has server power restrictions. • **lower-idle-power—**The USB System Idle Power Optimizing setting is enabled, because power savings are preferred over optimal performance. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Front Panel Access Lock	USB front panel lock is configured to enable or disable the front panel access to USB ports. This can be one of the following: • **disabled** • **enabled** • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

PCI Configuration BIOS Settings

The following table lists the PCI configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Max Memory Below 4G	Whether the BIOS maximizes memory usage below 4GB for an operating system without PAE support, depending on the system configuration. This can be one of the following: • **disabled—**Does not maximize memory usage. Choose this option for all operating systems with PAE support. • **enabled—**Maximizes memory usage below 4GB for an operating system without PAE support. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Mapped IO Above 4Gb Config	Whether to enable or disable memory mapped I/O of 64-bit PCI devices to 4GB or greater address space. Legacy option ROMs are not able to access addresses above 4GB. PCI devices that are 64-bit compliant but use a legacy option ROM may not function correctly with this setting enabled. This can be one of the following: • **disabled—**Does not map I/O of 64-bit PCI devices to 4GB or greater address space. • **enabled—**Maps I/O of 64-bit PCI devices to 4GB or greater address space. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Boot Options BIOS Settings

The following table lists the boot options BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
Boot Option Retry	Whether the BIOS retries NON-EFI based boot options without waiting for user input. This can be one of the following: • **disabled—**Waits for user input before retrying NON-EFI based boot options. • **enabled—**Continually retries NON-EFI based boot options without waiting for user input. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel Entry SAS RAID	Whether the Intel SAS Entry RAID Module is enabled. This can be one of the following: • **disabled—**The Intel SAS Entry RAID Module is disabled. • **enabled—**The Intel SAS Entry RAID Module is enabled. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel Entry SAS RAID Module	How the Intel SAS Entry RAID Module is configured. This can be one of the following: • **it-ir-raid—**Configures the RAID module to use Intel IT/IR RAID. • **intel-esrtii—**Configures the RAID module to use Intel Embedded Server RAID Technology II. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Onboard SCU Storage Support	Whether the onboard software RAID controller is available to the server. This can be one of the following: • **disabled—**The software RAID controller is not available. • **enabled—**The software RAID controller is available. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Server Management BIOS Settings

The following tables list the server management BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

General Settings

Name	Description
Assert Nmi on Serr	Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs. This can be one of the following: • **disabled**—The BIOS does not generate an NMI or log an error when a SERR occurs. • **enabled**—The BIOS generates an NMI and logs an error when a SERR occurs. You must enable this setting if you want to enable **Assert Nmi on Perr**. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Assert Nmi on Perr	Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a processor bus parity error (PERR) occurs. This can be one of the following: • **disabled**—The BIOS does not generate an NMI or log an error when a PERR occurs. • **enabled**—The BIOS generates an NMI and logs an error when a PERR occurs. You must enable **Assert Nmi on Serr** to use this setting. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
OS Boot Watchdog Timer	Whether the BIOS programs the watchdog timer with a predefined timeout value. If the operating system does not complete booting before the timer expires, the CIMC resets the system and an error is logged. This can be one of the following: • **disabled**—The watchdog timer is not used to track how long the server takes to boot. • **enabled**—The watchdog timer tracks how long the server takes to boot. If the server does not boot within the predefined length of time, the CIMC resets the system and logs an error. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This feature requires either operating system support or Intel Management software.

Name	Description
OS Boot Watchdog Timer Timeout Policy	What action the system takes if the watchdog timer expires. This can be one of the following: • **power-off—**The server is powered off if the watchdog timer expires during OS boot. • **reset—**The server is reset if the watchdog timer expires during OS boot. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This option is only available if you enable the OS Boot Watchdog Timer.
OS Boot Watchdog Timer Timeout	What timeout value the BIOS uses to configure the watchdog timer. This can be one of the following: • **5-minutes—**The watchdog timer expires 5 minutes after the OS begins to boot. • **10-minutes—**The watchdog timer expires 10 minutes after the OS begins to boot. • **15-minutes—**The watchdog timer expires 15 minutes after the OS begins to boot. • **20-minutes—**The watchdog timer expires 20 minutes after the OS begins to boot. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This option is only available if you enable the OS Boot Watchdog Timer.

Console Redirection Settings

Name	Description
Console Redirection	Allows a serial port to be used for console redirection during POST and BIOS booting. After the BIOS has booted and the operating system is responsible for the server, console redirection is irrelevant and has no effect. This can be one of the following: • **disabled—**No console redirection occurs during POST. • **serial-port-a—**Enables serial port A for console redirection during POST. This option is valid for blade servers and rack-mount servers. • **serial-port-b—**Enables serial port B for console redirection and allows it to perform server management tasks. This option is only valid for rack-mount servers. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** If you enable this option, you also disable the display of the Quiet Boot logo screen during POST.
Flow Control	Whether a handshake protocol is used for flow control. Request to Send / Clear to Send (RTS/CTS) helps to reduce frame collisions that can be introduced by a hidden terminal problem. This can be one of the following: • **none—**No flow control is used. • **rts-cts—**RTS/CTS is used for flow control. • **—**The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This setting must match the setting on the remote terminal application.

Name	Description
BAUD Rate	What BAUD rate is used for the serial port transmission speed. If you disable Console Redirection, this option is not available. This can be one of the following: • **9600**—A 9600 BAUD rate is used. • **19200**—A 19200 BAUD rate is used. • **38400**—A 38400 BAUD rate is used. • **57600**—A 57600 BAUD rate is used. • **115200**—A 115200 BAUD rate is used. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This setting must match the setting on the remote terminal application.
Terminal Type	What type of character formatting is used for console redirection. This can be one of the following: • **pc-ansi**—The PC-ANSI terminal font is used. • **vt100**—A supported vt100 video terminal and its character set are used. • **vt100-plus**—A supported vt100-plus video terminal and its character set are used. • **vt-utf8**—A video terminal with the UTF-8 character set is used. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. **Note** This setting must match the setting on the remote terminal application.
Legacy OS Redirect	Whether redirection from a legacy operating system, such as DOS, is enabled on the serial port. This can be one of the following: • **disabled**—The serial port enabled for console redirection is hidden from the legacy operating system. • **enabled**— The serial port enabled for console redirection is visible to the legacy operating system. • —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

BIOS Policy

The BIOS policy is a policy that automates the configuration of BIOS settings for a server or group of servers. You can create global BIOS policies available to all servers in the root organization, or you can create BIOS policies in sub-organizations that are only available to that hierarchy.

To use a BIOS policy, do the following:

1. Create the BIOS policy in Cisco UCS Central.
2. Assign the BIOS policy to one or more service profiles.
3. Associate the service profile with a server.

During service profile association, Cisco UCS Central modifies the BIOS settings on the server to match the configuration in the BIOS policy. If you do not create and assign a BIOS policy to a service profile, the server uses the default BIOS settings for that server platform.

Default BIOS Settings

Cisco UCS Central includes a set of default BIOS settings for each type of server supported by Cisco UCS. The default BIOS settings are available only in the root organization and are global. Only one set of default BIOS settings can exist for each server platform supported by Cisco UCS. You can modify the default BIOS settings, but you cannot create an additional set of default BIOS settings.

Each set of default BIOS settings are designed for a particular type of supported server and are applied to all servers of that specific type which do not have a BIOS policy included in their service profiles.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Cisco UCS Central applies these server platform-specific BIOS settings as follows:

- The service profile associated with a server does not include a BIOS policy.
- The BIOS policy is configured with the platform-default option for a specific setting.

You can modify the default BIOS settings provided by Cisco UCS Central. However, any changes to the default BIOS settings apply to all servers of that particular type or platform. If you want to modify the BIOS settings for only certain servers, we recommend that you use a BIOS policy.

Creating a BIOS Policy

Cisco UCS Central pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted. We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode and Sparing Mode for RAS Memory, are not supported by all Cisco UCS servers.

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **BIOS Policies** and choose **Create BIOS Policy**.

Step 4 In the **Create BIOS Policy** dialog box, enter the **Name** and optional description.

Note To create a BIOS policy quickly, you can click **Finish** after specifying the name. Cisco UCS Central creates a new BIOS policy with the specified name and all system default values.

Step 5 (Optional) In the **Main** panel, choose the main BIOS settings such as, **Reboot on BIOS Change**, **Quiet Boot**, **Post Error Pause**, **Resume Ac on Power Loss**, and **Front Panel Lockout**, then click **Next**.

Step 6 (Optional) In the **Processor** panel, choose the processor settings, then click **Next**.

Step 7 (Optional) In the **Intel Directed IO** panel, choose the IO settings, then click **Next**.

Step 8 (Optional) In the **RAS Memory** panel, choose the memory settings, then click **Next**.

Step 9 (Optional) In the **Serial Port** panel, choose the **Serial Port A** settings, then click **Next**.

Step 10 (Optional) In the **Processor** panel, choose the processor settings information, then click **Next**.

Step 11 (Optional) In the **USB** panel, choose the USB settings such as, **Make Device Non Bootable**, **Legacy USB Support**, **USB Idle Power Optimizing Setting**, and **USB Front Panel Access Lock**, then click **Next**.

Step 12 (Optional) In the **PCI Configuration** panel, choose the PCI configuration settings such as, **Max Memory Below 4GB** and **Memory Mapped IO Above 4GB Config**, then click **Next**.

Step 13 (Optional) In the **Boot Options** panel, choose the boot settings such as, **Boot Option Retry**, **Intel Entry SAS RAID**, **Intel Entry SAS RAID Module**, and **Onboard SCU Storage Support**, then click **Next**.

Step 14 (Optional) In the **Server Manager** panel, choose the non-maskable interrupt settings and the **OS Boot Watchdog Timer**, specify the **Console Redirection** settings, then click**Finish**.

Modifying a BIOS Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **BIOS Policies**.

Step 4 Click the BIOS policy that you want to modify.

Step 5 In the **Work** pane, click the appropriate tab and then click the desired radio button or make a choice from the drop-down list to modify the BIOS settings.

Step 6 Click **Save**.

Deleting a BIOS Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **BIOS Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

IPMI Access Profile

This policy allows you to determine whether IPMI commands can be sent directly to the server, using the IP address. For example, you can send commands to retrieve sensor data from the CIMC. This policy defines the IPMI access, including a username and password that can be authenticated locally on the server, and whether the access is read-only or read-write.

You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

Creating an IPMI Access Profile

IPMI access profiles require IPMI users. You can create IPMI users at the same time you create the IPMI access profile, or you can add them to an existing IPMI access profile.

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **IPMI Access Profiles** and choose **Create IPMI Access Profile**.

Step 4 In the **Create IPMI Access Profile** dialog box, enter the **Name** and optional description.

Step 5 Click **Create IPMI User** to add IPMI users to the IPMI Access Profile.

Step 6 Click **OK**.

What to Do Next

Include the IPMI profile in a service profile and/or template.

Adding an IPMI User to an IPMI Access Profile

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **IPMI Access Profiles**.

Step 4 Click the IPMI access profile for which you want to add an IPMI user.

Step 5 In the **Work** pane, click the **General Tab**.

Step 6 In the **IPMI Users** area, click **Create IPMI User**.

Step 7 In the **Create IPMI Users** dialog box, enter the **Name** and **Password**, confirm the password, and choose a **Serial over LAN State**.

Step 8 Click **OK**.

Deleting an IPMI Access Profile

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **IPMI Access Profiles**.

Step 4 Right-click the IPMI access profile that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Deleting an IPMI User from an IPMI Access Profile

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **IPMI Access Profiles**.

Step 4 Click the IPMI access profile for which you want to delete an IPMI user.

Step 5 In the **Work** pane, click the **General Tab**.

Step 6 In the **IPMI Users** table, click the IPMI user you want to delete.

Step 7 In the **IPMI Users** toolbar, click **Delete**.

Boot Policy

The boot policy overrides the boot order in the BIOS setup menu, and determines the following:

- Selection of the boot device
- Location from which the server boots
- Order in which boot devices are invoked

For example, you can choose to have associated servers boot from a local device, such as a local disk or CD-ROM (VMedia), or you can select a SAN boot or a LAN (PXE) boot.

You can either create a named boot policy that can be associated with one or more service profiles, or create a boot policy for a specific service profile. A boot policy must be included in a service profile, and that service profile must be associated with a server for it to take effect. If you do not include a boot policy in a service profile, the UCS domain applies the default boot policy.

Note Changes to a boot policy will be propagated to all service profiles created with an updating service profile template that includes that boot policy. Reassociation of the service profile with the server to rewrite the boot order information in the BIOS is automatically triggered.

Boot Order

Cisco UCS Central, release 1.2 enables you to choose one of the following two boot orders for the global boot policies you create in Cisco UCS Central.

- **Standard boot order**: Standard boot order is supported for all Cisco UCS servers, and enables top-level boot order choices. You can add a local device, such as local disk, CD-ROM, or floppy, or you can add SAN, LAN, or iSCSI boot.
- **Enhanced boot order**: Enhanced boot order allows you greater control over the boot devices that you select for your boot policy. Enhanced boot order is supported for all Cisco UCS B-Series M3 Blade Servers and Cisco UCS C-Series M3 Rack Servers at release 2.2(1b) or greater.

Enhanced boot order provides you the following additional second-level boot order choices:

- **Add Local LUN** - Enables boot from local hard disk.
- **Add SD Card** - Enables boot from SD Card.
- **Add Internal USB** - Enables boot from Internal USB.
- **Add External USB** - Enables boot from External USB.
- **Add Local CD/DVD** - Enables boot from local CD/DVD drive.
- **Add Remote CD/DVD** - Enables boot from KVM mapped ISO images.
- **Add Local Floppy** - Enables boot from local floppy drive.
- **Add Remote Floppy** - Enables boot from KVM mapped image files.
- **Add Remote Virtual Drive** - Enables boot from remote virtual drive that is accessible to the server.
- **Add LAN**, SAN or iSCSI Boot - Enables you to select a specific vNIC or vHBA from which to boot.

Local Disk , **CD/DVD ROM** boot are available for backward compatibility.

Note

- If a boot policy with enhanced boot order is applied to Cisco UCS M1 and M2 blade and rack servers, or to Cisco UCS M3 blade and rack servers with a release prior to Release 2.2(1b) installed, the association fails with configuration errors.
- You must enable USB for Virtual Media. If you modify the BIOS settings, that in turn affects the Virtual media. The following USB BIOS default settings are recommended for best performance:
  - **Make Device Non Bootable** - set to disabled
  - **USB Idle Power Optimizing Setting** - set to high-performance

UEFI Boot Mode

Unified Extensible Firmware Interface (UEFI) is a specification that defines a software interface between an operating system and platform firmware. Cisco UCS Manager uses UEFI to replace the BIOS firmware interfaces. This allows the BIOS to run in UEFI mode while still providing legacy support.

You can choose either legacy or UEFI boot mode when you create a boot policy. Legacy boot mode is supported for all Cisco UCS servers. UEFI boot mode is supported only on M3 and M4 servers, and allows you to enable UEFI secure boot mode.

The following limitations apply to the UEFI boot mode:

- UEFI boot mode is only supported on Cisco UCS B-Series M3 and M4 Blade Servers and Cisco UCS C-Series M3 and M4 Rack Servers.
- UEFI boot mode is not supported with the following combinations:
    - Gen-3 Emulex & QLogic adapters on Cisco UCS blade & rack servers integrated with Cisco UCS domain.
    - PXE boot for all adapters on Cisco UCS rack servers integrated with Cisco UCS domain.
    - iSCSI boot for all adapters on Cisco UCS rack servers integrated with Cisco UCS domain.
- You cannot mix UEFI and legacy boot mode on the same server.
- Make sure an UEFI aware operating systems is installed in the device. The server will boot correctly in UEFI mode only if the boot devices configured in the boot policy have UEFI-aware OS installed. If a compatible OS is not present, the boot device is not displayed on the **Actual Boot Order** tab in the **Boot Order Details** area.
- In some corner cases, the UEFI boot may not succeed because the UEFI boot manager entry was not saved correctly in the BIOS NVRAM. You can use the UEFI shell to enter the UEFI boot manager entry manually. This situation could occur in the following situations:
    - If a blade server with UEFI boot mode enabled is disassociated from the service profile, and the blade is manually powered on using the **Equipment** tab or the front panel.
    - If a blade server with UEFI boot mode enabled is disassociated from the service profile, and a direct VIC firmware upgrade is attempted.
    - If a blade or rack server with UEFI boot mode enabled is booted off SAN LUN, and the service profile is migrated.

UEFI Secure Boot

Cisco UCS Central supports UEFI secure boot on Cisco UCS B-Series M3 and M4 Blade Servers. When UEFI secure boot is enabled, all executables, such as boot loaders and adapter drivers, are authenticated by the BIOS before they can be loaded. To be authenticated, the images must be signed by either the Cisco Certificate Authority (CA) or a Microsoft CA.

The following limitations apply to UEFI secure boot:

- UEFI boot mode must be enabled in the boot policy.
- The Cisco UCS Manager software and the BIOS firmware must be at Release 2.2 or greater.
- User-generated encryption keys are not supported.
- UEFI secure boot can only be controlled by Cisco UCS Manager or Cisco UCS Central.
- If you want to downgrade to an earlier version of Cisco UCS Manager, and you have a blade server in secure boot mode, you must disassociate and reassociate the blade server before downgrading. Otherwise, the blade will not be discovered successfully.

Cautions and Guidelines for Downgrading a Boot Policy

You cannot downgrade to an earlier version of Cisco UCS Manager if:

- An associated server has a boot policy with UEFI boot mode enabled.
- An associated server has a boot policy with UEFI secure boot enabled.
- An associated server has a boot policy with enhanced boot order. For example, if an associated server has a boot policy which contains any of the following:
  - SD card
  - Internal USB
  - External USB
- An associated server has a boot policy that includes both SAN and local LUN.

Creating a Boot Policy

You can also create a local boot policy that is restricted to a service profile or service profile template. However, except for iSCSI boot, we recommend that you create a global boot policy that can be included in multiple service profiles or service profile templates.

Note Cisco UCS Central, release 1.2 does not add support to scriptable vMedia.

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Boot Policies** and choose **Create Boot Policy**.

Step 4 In the **Create Boot Policy** dialog box, enter the **Name**and optional description.

Step 5 (Optional) To reboot all servers that use this boot policy after you make changes to the boot policy, check the**Reboot on Boot Order Change** check box.

Important If you apply this boot policy on a server with non VIC adapter, even if the **Reboot on Boot Order Change** check box is not checked, if you add, delete or change the order for SAN devices, when you save the boot policy changes, the server always reboots.

Step 6 (Optional) To enforce that the vNICs, vHBAs, or iSCSI vNICs listed in the **Qualifications** table match the server configuration in the service profile, check the **Enforce vNIC/vHBA/iSCSI Name** check box.

Step 7 To select the **Boot Mode**, click **Legacy** or **UEFI**.

Step 8 In the **Actions** area, configure one or more of the following boot options for the boot policy and set their boot order:

- Local device boot—Click **Add CD/DVD ROM Boot**, **Add Local CD/DVD**, or **Add Local Disk**, **Add Floppy**, or **Add Remote Virtual Drive** to add devices to the boot policy.
- LAN Boot—Click **Add LAN Boot** to boot from a centralized provisioning server.
- SAN Boot—Click **Add SAN Boot** to boot from an operating system image on the SAN.

  If the vHBA points to a bootable SAN image, click **Add SAN Boot Target** to configure it.
- iSCSI vNICs—Click **Add iSCSI Boot** to boot from an iSCSI LUN.

Step 9 (Optional) Click the up and down arrows in the **Qualifications** table to change the boot order.

Step 10 Click **OK**.

What to Do Next

Include the boot policy in a service profile and/or template.

After a server is associated with a service profile that includes this boot policy, you can verify the actual boot order in the Boot Order Details area on the General tab for the server. For more information on boot policy, see Cisco UCS Manager Configuration Guide.

Modifying a Boot Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Boot Policies**.

Step 4 Click the boot policy that you want to modify.

Step 5 In the **Work** pane, click the **General** tab and make the appropriate changes to the boot options and boot order.

Step 6 Click **Save**.

Deleting a Boot Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Boot Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

LAN Boot

You can configure a boot policy to boot one or more servers from a centralized provisioning server on the LAN. A LAN (or PXE) boot is frequently used to install operating systems on a server from that LAN server.

You can add more than one type of boot device to a LAN boot policy. For example, you could add a local disk or virtual media boot as a secondary boot device.

Configuring a LAN Boot for a Boot Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Boot Policies**.

Step 4 Click the boot policy for which you want to configure a LAN boot.

Step 5 In the **Work** pane, on the **General** tab, click **Add LAN Boot**.

Step 6 In the **Add LAN Boot** dialog box, enter the **vNIC** and select primary or secondary from the **Type** drop-down list.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the boot policy.

SAN Boot

You can configure a boot policy to boot one or more servers from an operating system image on the SAN. The boot policy can include a primary and a secondary SAN boot. If the primary boot fails, the server attempts to boot from the secondary.

We recommend that you use a SAN boot, because it offers the most service profile mobility within the system. If you boot from the SAN when you move a service profile from one server to another, the new server boots from the exact same operating system image. Therefore, the new server appears to be the exact same server to the network.

To use a SAN boot, ensure that the following is configured:

- The Cisco UCS domain must be able to communicate with the SAN storage device that hosts the operating system image.
- A boot target LUN on the device where the operating system image is located.

Note SAN boot is not supported on Gen-3 Emulex adapters on Cisco UCS blade & rack servers.

Configuring a SAN Boot for a Boot Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Boot Policies**.

Step 4 Click the boot policy for which you want to configure a SAN boot.

Step 5 In the **Work** pane, on the **General** tab, click **Add SAN Boot**.

Step 6 In the **Add SAN Boot** dialog box, enter the **vHBA** and choose primary or secondary from the **Type** drop-down list.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the boot policy.

Adding a SAN Boot Target

You must have configured a SAN boot for a boot policy before you can add a SAN boot target.

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Boot Policies**.

Step 4 Click the boot policy for which you want add a SAN boot target.

Step 5 In the **Work** pane, on the **General** tab, click **Add SAN Boot Target**.

Step 6 In the **Add SAN Boot Target** dialog box, enter the **Boot Target LUN** and the **Boot Target WWPN**, and select primary or secondary from the **Type** drop-down list.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the boot policy.

iSCSI Boot

iSCSI boot enables a server to boot its operating system from an iSCSI target machine located remotely over a network.

iSCSI boot is supported on the following Cisco UCS hardware:

- Cisco UCS blade servers that have the Cisco UCS M51KR-B Broadcom BCM57711 network adapter and use the default MAC address provided by Broadcom.
- Cisco UCS M81KR Virtual Interface Card
- Cisco UCS VIC-1240 Virtual Interface Card
- Cisco UCS VIC-1280 Virtual Interface Card
- Cisco UCS rack servers that have the Cisco UCS M61KR-B Broadcom BCM57712 network adapter.
- Cisco UCS P81E Virtual Interface Card
- Cisco UCS VIC1225 Virtual Interface Card

There are prerequisites that must be met before you configure iSCSI boot. For a list of these prerequisites, see iSCSI Boot Guidelines and Prerequisites, on page 41.

iSCSI Boot Process

Cisco UCS Central uses the iSCSI vNIC and iSCSI boot information created for the service profile in the association process to program the adapter, located on the server. After the adapter is programmed, the server reboots with the latest service profile values. After the power on self-test (POST), the adapter attempts to initialize using these service profile values. If the adapter can use the values and log in to its specified target, the adapter initializes and posts an iSCSI Boot Firmware Table (iBFT) to the host memory and a valid bootable LUN to the system BIOS. The iBFT that is posted to the host memory contains the initiator and target configuration that is programmed on the primary iSCSI VNIC.

Note Previously, the host would see only one of the boot paths configured, depending on which path completed the LUN discovery first, and would boot from that path. Now, when there are two iSCSI boot vNICs configured, the host will see both of the boot paths. So for multipath configurations, a single IQN needs to be configured on both the boot vNICs If there are different IQNs configured on the boot vNICs on a host, the host will boot with the IQN that is configured on the boot vNIC with the lower PCI order.

The next step, which is the installation of the operating system (OS), requires an OS that is iBFT capable. During installation of the OS, the OS installer scans the host memory for the iBFT table and uses the information in the iBFT to discover the boot device and create an iSCSI path to the target LUN. In some OS's a NIC driver is required to complete this path. If this step is successful, the OS installer finds the iSCSI target LUN on which to install the OS.

Note The iBFT works at the OS installation software level and might not work with HBA mode (also known as TCP offload). Whether iBFT works with HBA mode depends on the OS capabilities during installation. Also, for a server that includes a Cisco UCS M51KR-B Broadcom BCM57711 adapter, the iBFT normally works at a maximum transmission unit (MTU) size of 1500, regardless of the MTU jumbo configuration. If the OS supports HBA mode, you might need to set HBA mode, dual-fabric support, and jumbo MTU size after the iSCSI installation process.

iSCSI Boot Guidelines and Prerequisites

These guidelines and prerequisites must be met before configuring iSCSI boot:

- After the iSCSI boot policies have been created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create iSCSI boot policies.
- To set up iSCSI boot from a Windows 2008 server where the second vNIC (failover vNIC) must boot from an iSCSI LUN, consult Microsoft Knowledge Base Article 976042. Microsoft has a known issue where Windows might fail to boot from an iSCSI drive or cause a bugcheck error if the networking hardware is changed. To work around this issue, follow the resolution recommended by Microsoft.
- The storage array must be licensed for iSCSI boot and the array side LUN masking must be properly configured.
- Two IP addresses must be determined, one for each iSCSI initiator. If possible, the IP addresses should be on the same subnet as the storage array. The IP addresses are assigned statically or dynamically using the Dynamic Host Configuration Protocol (DHCP).
- You cannot configure boot parameters in the Global boot policy. Instead, after configuring boot parameters, you need to include the boot policy in the appropriate service profile.
- The operating system (OS) must be iSCSI Boot Firmware Table (iBFT) compatible.
- For Cisco UCS M51KR-B Broadcom BCM57711 network adapters:
    - Servers that use iSCSI boot must contain the Cisco UCS M51KR-B Broadcom BCM57711 network adapter. For information on installing or replacing an adapter card, see the *Cisco UCS B250 Extended Memory Blade Server Installation and Service Note*. The service note is accessible from

the *Cisco UCS B-Series Servers Documentation Roadmap* at http://www.cisco.com/go/unifiedcomputing/b-series-doc.

    - Set the MAC addresses on the iSCSI device.
    - If you are using the DHCP Vendor ID (Option 43), configure the MAC address of an iSCSI device in /etc/dhcpd.conf.
    - HBA mode (also known as TCP offload) and the boot to target setting are supported. However, only Windows OS supports HBA mode during installation.
    - Before installing the OS, disable the boot to target setting in the iSCSI adapter policy, then after installing the OS, reenable the boot to target setting.

Note Each time you change an adapter policy setting, the adapter reboots to apply the new setting.

    - When installing the OS on the iSCSI target, the iSCSI target must be ordered *before* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the iSCSI target and then the CD.
    - After the server has been iSCSI booted, do not modify the Initiator Name, Target name, LUN, iSCSI device IP, or Netmask/gateway using the Broadcom tool.
    - Do not interrupt the POST (power on self-test) process or the Cisco UCS M51KR-B Broadcom BCM57711 network adapter will fail to initialize.

- For Cisco UCS M81KR Virtual Interface Card and Cisco UCS VIC-1240 Virtual Interface Card:
    - Do not set MAC addresses on the iSCSI device.
    - HBA mode and the boot to target setting are *not* supported.
    - When installing the OS on the iSCSI target, the iSCSI target must be ordered *after* the device where the OS image resides. For example, if you are installing the OS on the iSCSI target from a CD, the boot order should be the CD and then the iSCSI target.
    - If you are using the DHCP Vendor ID (Option 43), the MAC address of the overlay vNIC needs to be configured in /etc/dhcpd.conf.
    - After the server has been iSCSI booted, do not modify the IP details of the overlay vNIC.
- The VMware ESX/ESXi operating system does not support storing a core dump file to an iSCSI boot target LUN. Dump files must be written to a local disk.

Configuring an iSCSI Boot for a Boot Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.

If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Boot Policies**.

Step 4 Click the boot policy for which you want to configure an iSCSI boot.

Step 5 In the **Work** pane, on the **General** tab, click **Add iSCSI Boot**.

Step 6 In the **Add iSCSI Boot** dialog box, enter the **iSCSI vNIC** and choose primary or secondary from the **Type** drop-down list.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the boot policy.

Creating an iSCSI Adapter Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Adapter Policies** and choose **Create iSCSI Adapter Policy**.

Step 4 In the **Create iSCSI Adapter Policy** dialog box, enter the **Name**, optional description, the **Connection Timeout**, **LUN Busy Retry Count**, and **DHCP Timeout**.

Step 5 Choose the **Enable TCP Timestamp**, **HBA Mode**, and **Boot To Target** checkboxes.

Step 6 Click **OK**.

Deleting an iSCSI Adapter Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Adapter Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Creating an iSCSI Authentication Profile

For iSCSI boot, you need to create both an initiator and a target authentication profile

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **iSCSI Authentication Profile** and choose **Create iSCSI Authentication Profile**.

Step 4 In the **Create iSCSI Authentication Profile** dialog box, enter the **Name**, **User ID**, optional description, and **Password**, then confirm the password.

Step 5 Click **OK**.

What to Do Next

Include the authentication profile in a service profile and/or template.

Deleting an iSCSI Authentication Profile

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **iSCSI Authentication Profile**.

Step 4 Right-click the iSCSI authentication profile that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Local Disk Configuration Policy

This policy configures any optional SAS local drives that have been installed on a server through the onboard RAID controller of the local drive. This policy enables you to set a local disk mode for all servers that are associated with a service profile that includes the local disk configuration policy.

The local disk modes include the following:

- **No Local Storage—**For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.

- **RAID 0 Striped—**Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored—**Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration—**For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID—**For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.

  If you choose **No RAID** and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the **No RAID** mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the **Inventory** > **Storage** tab for the server.

  To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the **No RAID** configuration mode.
- **RAID 5 Striped Parity—**Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.
- **RAID 6 Striped Dual Parity—**Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.
- **RAID 10 Mirrored and Striped—**RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates.
- **RAID 50 Striped Parity and Striped —**Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.
- **RAID 60 Striped Dual Parity and Striped —**Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

You must include this policy in a service profile and that service profile must be associated with a server for the policy to take effect.

Guidelines for all Local Disk Configuration Policies

Before you create a local disk configuration policy, consider the following guidelines:

No Mixed HDDs and SSDs

Do not include HDDs and SSDs in a single server or RAID configuration.

Do Not Assign a Service Profile with the Default Local Disk Configuration Policy from a B200 M1 or M2 to a B200 M3

Due to the differences in the RAID/JBOD support provided by the storage controllers of B200 M1 and M2 servers and those of the B200 M3 server, you cannot assign or re-assign a service profile that includes the default local disk configuration policy from a B200M1 or M2 server to a B200 M3 server. The default local disk configuration policy includes those with Any Configuration or JBOD configuration.

JBOD Mode Support

Note Only B200 M1, B200 M2, B200 M3, B250 M1, B250 M2 and B22 M3 blade servers support the JBOD mode for local disks.

Guidelines for Local Disk Configuration Policies Configured for RAID

Configure RAID Settings in Local Disk Configuration Policy for Servers with MegaRAID Storage Controllers

If a blade server or integrated rack-mount server has a MegaRAID controller, you must configure RAID settings for the drives in the Local Disk Configuration policy included in the service profile for that server. You can do this either by configuring the local disk configuration policy in the service profile using one of the defined RAID modes for that server, or you can use the **Any Configuration** mode with the LSI Utilities toolset to create the RAID volumes.

If you do not configure your RAID LUNs before installing the OS, disk discovery failures might occur during the installation and you might see error messages such as “No Device Found.”

Server May Not Boot After RAID1 Cluster Migration if Any Configuration Mode Specified in Service Profile

After RAID1 clusters are migrated, you need to associate a service profile with the server. If the local disk configuration policy in the service profile is configured with **Any Configuration** mode rather than **RAID1**, the RAID LUN remains in "inactive" state during and after association. As a result, the server cannot boot.

To avoid this issue, ensure that the service profile you associate with the server contains the identical local disk configuration policy as the original service profile before the migration and does not include the **Any Configuration** mode.

Do Not Use JBOD Mode on Servers with MegaRAID Storage Controllers

Do not configure or use JBOD mode or JBOD operations on any blade server or integrated rack-mount server with a MegaRAID storage controllers. JBOD mode and operations are not intended for nor are they fully functional on these servers.

Maximum of One RAID Volume and One RAID Controller in Integrated Rack-Mount Servers

A rack-mount server that has been integrated with Cisco UCS Manager and is registered with Cisco UCS Central can have a maximum of one RAID volume irrespective of how many hard drives are present on the server.

All the local hard drives in an integrated rack-mount server must be connected to only one RAID Controller. Integration with Cisco UCS Manager does not support the connection of local hard drives to multiple RAID Controllers in a single rack-mount server. We therefore recommend that you request a single RAID Controller configuration when you order rack-mount servers to be integrated with Cisco UCS Manager.

In addition, do not use third party tools to create multiple RAID LUNs on rack-mount servers. Cisco UCS Manager does not support that configuration.

Maximum of One RAID Volume and One RAID Controller in Blade Servers

A blade server can have a maximum of one RAID volume irrespective of how many drives are present in the server. All the local hard drives must be connected to only one RAID controller. For example, a B200 M3 server has an LSI controller and an Intel Patsburg controller, but only the LSI controller can be used as a RAID controller.

In addition, do not use third party tools to create multiple RAID LUNs on blade servers. Cisco UCS Central does not support that configuration.

Number of Disks Selected in Mirrored RAID Should Not Exceed Two

If the number of disks selected in the Mirrored RAID exceed two, RAID 1 is created as a RAID 10 LUN. This issue can occur with the Cisco UCS B440 M1 and B440 M2 servers.

License Required for Certain RAID Configuration Options on Some Servers

Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Central associates a service profile containing this local disk policy with a server, Cisco UCS Central verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Central displays a configuration error during the service profile association.

For RAID license information for a specific Cisco UCS server, see the *Hardware Installation Guide* for that server.

B420 M3 Server Does Not Support All Configuration Modes

The B420 M3 server does not support the following configuration modes in a local disk configuration policy:

- No RAID
- RAID 6 Striped Dual Parity

In addition, the B420 M3 does not support JBOD modes or operations.

Single-Disk RAID 0 Configurations Not Supported on Some Blade Servers

A single-disk RAID 0 configuration is not supported in the following blade servers:

- Cisco UCS B200 M1
- Cisco UCS B200 M2
- Cisco UCS B250 M1
- Cisco UCS B250 M2

Creating a Local Disk Configuration Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.

If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Local Disk Config Policies** and choose **Create Local Disk Config Policy**.

Step 4 In the **Create Local Disk Config Policy** dialog box, enter the **Name** and other optional details.

Step 5 Click **OK**.

Deleting a Local Disk Configuration Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Local Disk Config Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Power Control Policy

Cisco UCS uses the priority set in the power control policy, along with the blade type and configuration, to calculate the initial power allocation for each blade within a chassis. During normal operation, the active blades within a chassis can borrow power from idle blades within the same chassis. If all blades are active and reach the power cap, service profiles with higher priority power control policies take precedence over service profiles with lower priority power control policies.

Priority is ranked on a scale of 1-10, where 1 indicates the highest priority and 10 indicates lowest priority. The default priority is 5.

For mission-critical application a special priority called no-cap is also available. Setting the priority to no-cap prevents Cisco UCS from leveraging unused power from a particular server. With this setting, the server is allocated the maximum amount of power possible for that type of server.

Note You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

Creating a Power Control Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Power Control Policies** and choose **Create Power Control Policy**.

Step 4 In the **Create Power Control Policy** dialog box, enter the **Name** and optional description, choose whether to use **Power Capping**, and enter the **Power Priority**.

Step 5 Click **OK**.

What to Do Next

Include the policy in a service profile or service profile template.

Deleting a Power Control Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Power Control Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Step 6

Scrub Policy

This policy determines what happens to local data and to the BIOS settings on a server during the discovery process, when the server is reacknowledged, or when the server is disassociated from a service profile.

Note Local disk scrub policies only apply to hard drives that are managed by Cisco UCS Manager and do not apply to other devices such as USB drives.

Depending upon how you configure a scrub policy, the following can occur at those times:

Disk scrub

One of the following occurs to the data on any local drives on disassociation:

- If enabled, destroys all data on any local drives.
- If disabled, preserves all data on any local drives, including local storage configuration.

BIOS Settings Scrub

One of the following occurs to the BIOS settings when a service profile containing the scrub policy is disassociated from a server:

- If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor.
- If disabled, preserves the existing BIOS settings on the server.

FlexFlash Scrub

FlexFlash Scrub enables you to pair new or degraded SD cards, resolve FlexFlash metadata configuration failures, and migrate older SD cards with 4 partitions to single partition SD cards. One of the following occurs to the SD card when a service profile containing the scrub policy is disassociated from a server, or when the server is reacknowledged:

- If enabled, the HV partition on the SD card is formatted using the PNUOS formatting utility. If two SD cards are present, the cards are RAID-1 paired, and the HV partitions in both cards are marked as valid. The card in slot 1 is marked as primary, and the card in slot 2 is marked as secondary.
- If disabled, preserves the existing SD card settings.

Note

- Because the FlexFlash scrub erases the HV partition on the SD cards, we recommend that you take a full backup of the SD card(s) using your preferred host operating system utilities before performing the FlexFlash Scrub.
- To resolve metadata config failures in a service profile, you need to disable FlexFlash in the local disk config policy before you run the FlexFlash scrub, then enable FlexFlash after the server is reacknowledged.
- Disable the scrub policy as soon as the pairing is complete or the metadata failures are resolved.

Creating a Scrub Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.

If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Scrub Policies** and choose **Create Scrub Policy**.

Step 4 In the **Create Scrub Policy** dialog box, enter the **Name** and optional description, and choose whether to use **Disk Scrub** and **BIOS Setting Scrub**.

Step 5 Click **OK**.

Deleting a Scrub Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Scrub Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Serial over LAN Policy

This policy sets the configuration for the serial over LAN connection for all servers associated with service profiles that use the policy. By default, the serial over LAN connection is disabled.

If you implement a serial over LAN policy, we recommend that you also create an IPMI profile.

You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

Creating a Serial over LAN Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Serial over LAN Policies** and choose **Create Serial over LAN Policy**.

Step 4 In the **Create Serial over LAN Policy** dialog box, enter the **Name** and optional description, choose the **Serial over LAN State**, and choose a **Speed** from the drop-down list.

Step 5 Click **OK**.

Deleting a Serial over LAN Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Serial over LAN Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Server Pool Policy

This policy is invoked during the server discovery process. It determines what happens if server pool policy qualifications match a server to the target pool specified in the policy.

If a server qualifies for more than one pool and those pools have server pool policies, the server is added to all those pools.

Creating a Server Pool Policy

Before You Begin

This policy requires that one or more of the following resources already exist in the system:

- A minimum of one server pool
- Server pool policy qualifications, if you choose to have servers automatically added to pools

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.

If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Server Pool Policies** and choose **Create Policy**.

Step 4 In the **Create Policy** dialog box, enter the **Name**, choose a **Target Pool** and **Qualification** from the drop-down lists, and enter an optional description.

Step 5 Click **OK**.

Deleting a Server Pool Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Server Pool Policy Qualifications

This policy qualifies servers based on the inventory of a server conducted during the discovery process. The qualifications are individual rules that you configure in the policy to determine whether a server meets the selection criteria. For example, you can create a rule that specifies the minimum memory capacity for servers in a data center pool.

Qualifications are used in other policies to place servers, not just by the server pool policies. For example, if a server meets the criteria in a qualification policy, it can be added to one or more server pools or have a service profile automatically associated with it.

You can use the server pool policy qualifications to qualify servers according to the following criteria:

- Adapter type
- Chassis location
- Memory type and configuration
- Power group
- CPU cores, type, and configuration
- Storage configuration and capacity
- Server model

Depending upon the implementation, you might need to configure several policies with server pool policy qualifications including the following:

- Autoconfiguration policy
- Chassis discovery policy
- Server discovery policy
- Server inheritance policy
- Server pool policy

Creating Server Pool Policy Qualifications

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **Server Pool Policy Qualifications** and choose **Create Policy Qualification**.

Step 4 In the **Create Policy Qualification** dialog box, enter the **Name** and optional description.

Step 5 In the **Actions** area, configure one or more of the policy qualification options:

- **Create Domain Qualification**
- **Create Adapter Qualification**
- **Create Memory Qualification**
- **Create Processor Qualification**
- **Create Storage Qualification**
- **Create Server PID Qualification**

Step 6 Click **OK**.

Creating a Domain Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, enter the **Name**.

Step 7 In the **Actions** area, configure one or more of the domain qualification options:

- **Create Chassis/Server Qualification**
- **Create Address Qualification**
- **Create Owner Qualification**
- **Create Site Qualification**
- **Create Rack Qualification**

Step 8 Click **OK** to close the dialog box.

Step 9 Click **Save** to save the policy qualification.

Creating an Adapter Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Adapter Qualification**.

Step 6 In the **Create Adapter Qualification** dialog box, choose the **Type** and enter the **PID (RegEx)**.

Step 7 In the **Units** area, enter a number of units or click the **Unspecified** check box.

Step 8 Click **OK** to close the dialog box.

Step 9 Click **Save** to save the policy qualification.

Creating a Memory Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.

If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Memory Qualification**.

Step 6 In the **Create Memory Qualification** dialog box, enter values for **Clock (MHz)**, **Min Cap (MB)**, **Width**, **Speed**, **Latency (ns)**, **Max Cap (MB)**, and **Units**, or leave them unspecified.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the policy qualification.

Creating a Processor Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Processor Qualification**.

Step 6 In the **Create Processor Qualification** dialog box, choose the **Processor Architecture**, then enter values for **Min Number of Cores**, **Max Number of Cores**, **Min Number of Threads**, **Max Number of Threads**, **CPU Speed (MHz)**, **CPU Stepping**, **Min Number of Procs**, and **Max Number of Procs**, or leave them unspecified.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the policy qualification.

Creating a Storage Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Storage Qualification**.

Step 6 In the **Create Storage Qualification** dialog box, choose the **Diskless** state, then enter values for **Number of Blocks**, **Block Size (Bytes)**, **Min Cap (MB)**, **Max Cap (MB)**, **Per Disk Cap (MB)** and **Units**, or leave them unspecified.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the policy qualification.

Creating a Server PID Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Server PID Qualification**.

Step 6 In the **Create Server PID Qualification** dialog box, enter the **PID (RegEx)**.

Step 7 Click **OK** to close the dialog box.

Step 8 Click **Save** to save the policy qualification.

Creating a Chassis/Server Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, click **Create Chassis/Server Qualification**.

Step 7 In the **Create Chassis/Server Qualification** dialog box, enter the **First Chassis Id** and the **Number of Chassis**.

Step 8 Click **Create Server Qualification** to add a service qualification to the **Server Qualifications** table.

Step 9 Click **OK** to close the dialog box.

Step 10 Click **OK** to close the **Domain Qualification** dialog box.

Creating a Server Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, click **Create Chassis/Server Qualification**.

Step 7 In the **Create Chassis/Server Qualification** dialog box, click **Create Server Qualification**.

Step 8 In the **Create Server Qualification** dialog box, enter the **First Slot Id** and **Number of Slots**.

Step 9 Click **OK** to close the dialog box.

Step 10 Click **OK** to close the **Create Domain Qualification** dialog box.

Step 11 Click **OK** to close the **Domain Qualification** dialog box.

Creating an Address Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, click **Create Address Qualification**.

Step 7 In the **Create Address Qualification** dialog box, enter the **Minimum Address** and the **Maximum Address**.

Step 8 Click **OK** to close the dialog box.

Step 9 Click **OK** to close the **Domain Qualification** dialog box.

Creating an Owner Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, click **Create Owner Qualification**.

Step 7 In the **Create Owner Qualification** dialog box, enter the **First Chassis Id** and the **Number of Chassis**.

Step 8 Click **OK** to close the dialog box.

Step 9 Click **OK** to close the **Domain Qualification** dialog box.

Creating a Rack Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, click **Create Rack Qualification**.

Step 7 In the **Create Rack Qualification** dialog box, enter the **First Slot Id** and the **Number of Slots**.

Step 8 Click **OK** to close the dialog box.

Step 9 Click **OK** to close the **Domain Qualification** dialog box.

Creating a Site Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, on the **General** tab, click **Create Domain Qualification**.

Step 6 In the **Create Domain Qualification** dialog box, click **Create Site Qualification**.

Step 7 In the **Create Site Qualification** dialog box, enter the **Name** and the **Regex**.

Step 8 Click **OK** to close the dialog box.

Step 9 Click **OK** to close the **Domain Qualification** dialog box.

Deleting Server Pool Policy Qualifications

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Right-click the policy qualification that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

Deleting a Domain Qualification from a Policy Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 Right-click the qualification that you want to delete and choose **Delete**.

Step 8 Click **Save** to save the policy qualification.

Deleting a Chassis/Server Qualification from a Domain Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 In the **Qualifications** table, expand the domain qualification that you want to modify.

Step 8 Expand **Chassis/Server Qualifications**.

Step 9 Right-click the qualification that you want to delete and choose **Delete**.

Step 10 Click **Save** to save the policy qualification.

Deleting a Server Qualification from a Chassis/Server Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 In the **Qualifications** table, expand the domain qualification that you want to modify.

Step 8 Expand **Chassis Qualifications**.

Step 9 Expand the chassis qualification that you want to modify.

Step 10 Right-click the server qualification that you want to delete and choose **Delete**.

Step 11 Click **Save** to save the policy qualification.

Deleting an Address Qualification from a Domain Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 In the **Qualifications** table, expand the domain qualification that you want to modify.

Step 8 Expand **Address Qualifications**.

Step 9 Right-click the qualification that you want to delete and choose **Delete**.

Step 10 Click **Save** to save the policy qualification.

Deleting an Owner Qualification from a Domain Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 In the **Qualifications** table, expand the domain qualification that you want to modify.

Step 8 Expand **Owner Qualifications**.

Step 9 Right-click the qualification that you want to delete and choose **Delete**.

Step 10 Click **Save** to save the policy qualification.

Deleting a Rack Qualification from a Domain Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 In the **Qualifications** table, expand the domain qualification that you want to modify.

Step 8 Expand **Rack Qualifications**.

Step 9 Right-click the qualification that you want to delete and choose **Delete**.

Step 10 Click **Save** to save the policy qualification.

Deleting a Site Qualification from a Domain Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Domain Qualifications**.

Step 7 In the **Qualifications** table, expand the domain qualification that you want to modify.

Step 8 Expand **Site Qualifications**.

Step 9 Right-click the qualification that you want to delete and choose **Delete**.

Step 10 Click **Save** to save the policy qualification.

Deleting an Adapter Qualification from a Policy Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Expand **Adapter Qualifications**.

Step 7 Right-click the qualification that you want to delete and choose **Delete**.

Step 8 Click **Save** to save the policy qualification.

Deleting a Memory Qualification from a Policy Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Right-click the qualification that you want to delete and choose **Delete**.

Step 7 Click **Save** to save the policy qualification.

Deleting a Processor Qualification from a Policy Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Right-click the qualification that you want to delete and choose **Delete**.

Step 7 Click **Save** to save the policy qualification.

Deleting a Storage Qualification from a Policy Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Right-click the qualification that you want to delete and choose **Delete**.

Step 7 Click **Save** to save the policy qualification.

Deleting a Server Qualification from a Policy Qualification

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **Server Pool Policy Qualifications**.

Step 4 Click the policy qualification that you want to modify.

Step 5 In the **Work** pane, click the **General** tab.

Step 6 Right-click the qualification that you want to delete and choose **Delete**.

Step 7 Click **Save** to save the policy qualification.

vNIC/vHBA Placement Policies

vNIC/vHBA placement policies are used to determine the following:

- How the virtual network interface connections (vCons) are mapped to the physical adapters on a server.
- What types of vNICs or vHBAs can be assigned to each vCon.

Each vNIC/vHBA placement policy contains four vCons that are virtual representations of the physical adapters. When a vNIC/vHBA placement policy is assigned to a service profile, and the service profile is associated with a server, the vCons in the vNIC/vHBA placement policy are assigned to the physical adapters and the vNICs and vHBAs are assigned to those vCons.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the type of server and the selected virtual slot mapping scheme, which can be **Round Robin** or **Linear Ordered**. For details about the available mapping schemes, see vCon to Adapter Placement, on page 68.

After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the **Selection Preference** for each vCon. This can be one of the following:

- —All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.
- —vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- —Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- —Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
- —Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.

Note An SRIOV usNIC that is explicitly assigned to a vCon set to will remain assigned to that vCon.

If you do not include a vNIC/vHBA placement policy in the service profile, Cisco UCS Central defaults to the **Round Robin** vCon mapping scheme and the **All** vNIC/vHBA selection preference, distributing the vNICs and vHBAs between the adapters based on the capabilities and relative capacities of each adapter.

Creating a vNIC/vHBA Placement Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Right-click **vNIC/vHBA Placement Policies** and choose **Create Placement Policy**.

Step 4 In the **Create Placement Policy** dialog box, enter the **Name** and other optional details.

Step 5 Click **OK**.

Deleting a vNIC/vHBA Placement Policy

Procedure

Step 1 On the menu bar, click **Servers**.

Step 2 In the **Navigation** Pane, expand **Servers** > **Policies** > **root**.
If you want to create or access a policy in a sub-organization, expand **Sub-Organizations** > ***Organization_Name***.

Step 3 Expand **vNIC/vHBA Placement Policies**.

Step 4 Right-click the policy that you want to delete and choose **Delete**.

Step 5 If Cisco UCS Central GUI displays a confirmation dialog box, click **Yes**.

vCon to Adapter Placement

Cisco UCS maps every vCon in a service profile to a physical adapter on the server. How that mapping occurs and how the vCons are assigned to a specific adapter in a server depends on the following:

- The type of server. N20-B6620-2 and N20-B6625-2 blade servers with two adapter cards use a different mapping scheme than other supported rack or blade servers.
- The number of adapters in the server.
- The setting of the virtual slot mapping scheme in the vNIC/vHBA placement policy, if applicable.

You must consider this placement when you configure the vNIC/vHBA selection preference to assign vNICs and vHBAs to vCons.

Note vCon to adapter placement is not dependent upon the PCIE slot number of the adapter. The adapter numbers used for the purpose of vCon placement are not the PCIE slot numbers of the adapters, but the ID assigned to them during server discovery.

vCon to Adapter Placement for N20-B6620-2 and N20-B6625-2 Blade Servers

In N20-B6620-2 and N20-B6625-2 blade servers, the two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme:

- —Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default.
- —Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2.

vCon to Adapter Placement for All Other Supported Servers

For all other servers supported by Cisco UCS in addition to the N20-B6620-2 and N20-B6625-2 blade servers, the vCon assignment depends on the number of adapters in the server and the virtual slot mapping scheme.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme: Round Robin or Linear Ordered.

Table 1: vCon to Adapter Placement Using the Round - Robin Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter2	Adapter1	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter2
4	Adapter1	Adapter2	Adapter3	Adapter4

Round Robin is the default mapping scheme.

Table 2: vCon to Adapter Placement Using the Linear Ordered Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter1	Adapter2	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter3
4	Adapter1	Adapter2	Adapter3	Adapter4

Note If you are using a vCon policy with two adapters in the Cisco UCS B440 M2 Blade Server, be aware of the following mapping.

- vCon 2 to adapter 1 maps first
- vCon 1 to adapter 2 maps second ZXA Q

vNIC/vHBA to vCon Assignment

Cisco UCS Central provides two options for assigning vNICs and vHBAs to vCons through the vNIC/vHBA placement policy: explicit assignment and implicit assignment.

Explicit Assignment of vNICs and vHBAs

With explicit assignment, you specify the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned. Use this assignment option when you need to determine how the vNICs and vHBAs are distributed between the adapters on a server.

To configure a vCon and the associated vNICs and vHBAs for explicit assignment, do the following:

- Set the vCon configuration to any of the available options. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server. If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon.
- Assign the vNICs and vHBAs to a vCon. You can make this assignment through the virtual host interface placement properties of the vNIC or vHBA or in the service profile associated with the server.

If you attempt to assign a vNIC or vHBA to a vCon that is not configured for that type of vNIC or vHBA, Cisco UCS Central displays a message advising you of the configuration error.

During service profile association, Cisco UCS Central validates the configured placement of the vNICs and vHBAs against the number and capabilities of the physical adapters in the server before assigning the vNICs and vHBAs according to the configuration in the policy. Load distribution is based upon the explicit assignments to the vCons and adapters configured in this policy.

If the adapters do not support the assignment of one or more vNICs or vHBAs, Cisco UCS Central raises a fault against the service profile.

Implicit Assignment of vNICs and vHBAs

With implicit assignment, Cisco UCS Central determines the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned according to the capability of the adapters and their relative capacity. Use this assignment option if the adapter to which a vNIC or vHBA is assigned is not important to your system configuration.

To configure a vCon for implicit assignment, do the following:

- Set the vCon configuration to **All**, **Exclude Dynamic**, or **Exclude Unassigned**. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server.
- Do not set the vCon configuration to **Assigned Only**. Implicit assignment cannot be performed with this setting.
- Do not assign any vNICs or vHBAs to a vCon.

During service profile association, Cisco UCS Central verifies the number and capabilities of the physical adapters in the server and assigns the vNICs and vHBAs accordingly. Load distribution is based upon the capabilities of the adapters, and placement of the vNICs and vHBAs is performed according to the actual order determined by the system. For example, if one adapter can accommodate more vNICs than another, that adapter is assigned more vNICs.

If the adapters cannot support the number of vNICs and vHBAs configured for that server, Cisco UCS Central raises a fault against the service profile.

Implicit Assignment of vNICs in a Dual Adapter Environment

When you use implicit vNIC assignment for a dual slot server with an adapter card in each slot, Cisco UCS Central typically assigns the vNICs/vHBAs as follows:

- If the server has the same adapter in both slots, Cisco UCS Central assigns half the vNICs and half the vHBAs to each adapter.
- If the server has one non-VIC adapter and one VIC adapter, Cisco UCS Central assigns two vNICs and two vHBAs to the non-VIC adapter and the remaining vNICs and vHBAs to the VIC adapter.
- If the server has two different VIC adapters, Cisco UCS Central assigns the vNICs and vHBAs proportionally, based on the relative capabilities of the two adapters.

The following examples show how Cisco UCS Central would typically assign the vNICs and vHBAs with different combinations of supported adapter cards:

- If you want to configure four vNICs and the server contains two Cisco UCS M51KR-B Broadcom BCM57711 adapters (with two vNICs each), Cisco UCS Central assigns two vNICs to each adapter.
- If you want to configure 50 vNICs and the server contains a Cisco UCS CNA M72KR-E adapter (2 vNICs) and a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs), Cisco UCS Central assigns two vNICs to the Cisco UCS CNA M72KR-E adapter and 48 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter.
- If you want to configure 150 vNICs and the server contains a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs) and a Cisco UCS VIC-1240 Virtual Interface Card adapter (256 vNICs), Cisco UCS Central assigns 50 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter and 100 vNICs to the Cisco UCS VIC-1240 Virtual Interface Card adapter.

Note Exceptions to this implicit assignment occur if you configure the vNICs for fabric failover and if you configure dynamic vNICs for the server.

For a configuration that includes vNIC fabric failover where one adapter does not support vNIC failover, Cisco UCS Central implicitly assigns all vNICs that have fabric failover enabled to the adapter that supports them. If the configuration includes only vNICs that are configured for fabric failover, no vNICs are implicitly assigned to the adapter that does not support them. If some vNICs are configured for fabric failover and some are not, Cisco UCS Central assigns all failover vNICs to the adapter that supports them and a minimum of one nonfailover vNIC to the adapter that does not support them, according to the ratio above.

For a configuration that includes dynamic vNICs, the same implicit assignment would occur. Cisco UCS Central assigns all dynamic vNICs to the adapter that supports them. However, with a combination of dynamic vNICs and static vNICs, at least one static vNIC is assigned to the adapter that does not support dynamic vNICs.