



Configuring Server Policies

- [Server Policies](#), on page 2
- [Creating a Policy](#), on page 8
- [Supported UCS Server Policies](#), on page 9
- [Creating a Certificate Management Policy](#), on page 12
- [Creating an Adapter Configuration Policy](#), on page 14
- [Creating a LAN Connectivity Policy](#), on page 17
- [Creating an Ethernet Adapter Policy](#), on page 25
- [Creating an Ethernet QoS Policy](#), on page 33
- [Creating an Ethernet Network Policy](#), on page 35
- [Creating an Ethernet Network Group Policy](#), on page 40
- [Creating an Ethernet Network Control Policy](#), on page 41
- [Creating a SAN Connectivity Policy](#), on page 43
- [Creating a Fibre Channel Adapter Policy](#), on page 49
- [Creating a Fibre Channel Network Policy](#), on page 52
- [Creating a Fibre Channel QoS Policy](#), on page 53
- [Create FC Zone Policy](#), on page 55
- [Creating a Firmware Policy](#), on page 56
- [Creating a BIOS Policy](#), on page 57
- [Creating a Boot Order Policy](#), on page 71
- [Configuring an iSCSI Boot Policy](#), on page 81
- [Creating an iSCSI Adapter Policy](#), on page 84
- [Creating an iSCSI Static Target Policy](#), on page 85
- [Creating a Device Connector Policy](#), on page 86
- [Creating a Drive Security Policy](#), on page 87
- [Creating a Disk Group Policy](#), on page 88
- [Creating an IMC Access Policy](#), on page 90
- [Creating an IPMI Over LAN Policy](#), on page 92
- [Creating an LDAP Policy](#), on page 94
- [Creating a Local User Policy](#), on page 98
- [Creating an NTP Policy](#), on page 101
- [Creating an SD Card Policy](#), on page 102
- [Create a Serial Over LAN Policy](#), on page 104
- [Create SSH Policy](#), on page 106

- [Creating a Virtual KVM Policy, on page 107](#)
- [Creating a Virtual Media Policy, on page 108](#)
- [Creating a Network Connectivity Policy, on page 112](#)
- [Creating a SMTP Policy, on page 114](#)
- [Creating an SNMP Policy, on page 115](#)
- [Creating a Storage Policy, on page 118](#)
- [Creating a Syslog Policy, on page 129](#)
- [Creating a Power Policy for Server, on page 130](#)
- [Creating a Thermal Policy for Server, on page 132](#)
- [Creating vNIC or vHBA Templates, on page 133](#)

Server Policies

Policies in Cisco Intersight provide different configurations for UCS servers, including BIOS settings, firmware versions, disk group creation, Simple Mail Transfer Protocol (SMTP), Intelligent Platform Management Interface (IPMI) settings, and more. A policy that is once configured can be assigned to any number of servers to provide a configuration baseline. Policies in Cisco Intersight are native to the application and are not directly imported from the UCS Systems. Policy-based configuration with Server Profiles is a Cisco Intersight Essentials functionality.

The Server Policy creation wizard in Cisco Intersight has two pages:

- **General**—The general page allows you to select the organization and enter a name for your policy. Optionally, include a short description and tag information to help identify the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
- **Policy Details**—The policy details page has properties that are applicable to standalone UCS servers, FI-attached UCS servers, or both. You can view these properties separately for **All Platforms**, **UCS Servers (Standalone)**, and **UCS Servers (FI-Attached)** by clicking on these options.

Server Policies can be imported as part of importing configuration details (server profiles and policies) of a Cisco C-Series Standalone server from Cisco IMC. For more information, see [Importing a Server Profile](#).

The following list describes the server policies that you can configure in Cisco Intersight.

- **Adapter Configuration Policy**—Configures the Ethernet and Fibre-Channel settings for the VIC adapter.
- **BIOS Policy**—Automates the configuration of BIOS settings on the managed devices. You can create one or more BIOS policies which contain a specific grouping of BIOS settings. If you do not specify a BIOS policy for a server, the BIOS settings remain as they are. If a BIOS policy is specified, the values that are specified in the policy replace any previously configured values on a server (including bare metal server configuration settings). To apply the BIOS policy settings, you must reboot the server.
- **Boot Order Policy**—Configures the linear ordering of devices and enables you to change the boot order and boot mode. You can also add multiple devices under various device types, rearrange the boot order, and set parameters for each boot device type.

The inventory view enables you to view the actual boot order configured on a server. The boot order displays the details that include device name, device type, configuration details such as Boot Mode (Legacy or UEFI), and Secure Boot Mode (Enabled or Disabled).



Note A device configured in the server profile of Boot Order Policy may not appear in the actual boot order, if the server BIOS does not detect the device during server boot.

Intersight provides a One-Time Boot (OTB) option to set a boot device that temporarily overrides the Boot Order Policy and the existing boot order. To set a One-Time Boot Device, select **Power Cycle** or **Power On** from the **Servers Table view** or from the **Server Details** page and toggle ON the **Set One Time Boot Device** Option. This operation attempts to boot from the One Time Boot device as part of the power cycle or power on action. After power cycle or power on, OTB configuration will be cleared to enable the next reboot to follow the default Boot Order.



-
- Note**
- The OTB option is available for servers that have been configured with a Boot Order Policy that is associated with a server profile. For a successful OTB configuration, you must deploy a server profile with a Boot Order Policy in Intersight in advance.
 - Any out-of-band- boot order change will not reflect on the Intersight UI for OTB device configuration.
-

In the case of **PXE Boot** configuration, importing the server policy will not create the PXE device under boot policy if either the MAC address or both the slot and port are not present for a given PXE device under the Boot policy on the server. However, if both slot and port are present, boot order is set to **ANY** for the bootable interface on a given slot on the server. For non-VIC adapters you can configure PXE Boot with the MAC address, or both the slot and port, or slot only.

In the case of **SAN Boot** device configuration in the legacy mode, provide the boot target Logical Unit Number (LUN), device slot ID, interface name, and target WWPN. For **SAN Boot** device configuration in the Unified Extensible Firmware Interface (UEFI) mode, provide the bootloader name, description, and path in addition to the fields listed in the legacy mode.

In the case of **iSCSI Boot** provide the target interface details, authentication mechanism, and initiator IP source.

- In the case of **Non-Volatile Memory Express (NVMe) Boot**, configure the NVMe drive as bootable in the UEFI mode. During the server profile deployment, this NVMe configuration setting enables selecting the BIOS in a defined order.
- **Certificate Management Policy**—Allows you to specify the certificate details for an external certificate and attach the policy to servers. Cisco Intersight currently supports the following certificates:
 - Root CA certificates
 - IMC certificates
- **Disk Group Policy**—Disk Group Policy is now a part of Storage Policy.
- **Device Connector Policy**—Lets you choose the **Configuration from Intersight only** option to control configuration changes allowed from Cisco IMC. The **Configuration from Intersight only** option is

enabled by default. You will observe the following changes when you deploy the Device Connector policy in Intersight:

- Validation tasks will fail:
 - If Intersight Read-only mode is enabled in the claimed device.
 - If the firmware version of the Cisco UCS Standalone C-Series Servers is lower than 4.0(1).
- If Intersight Read-only mode is enabled, firmware upgrades will be successful only when performed from Intersight. Firmware upgrade performed locally from Cisco IMC will fail.
- IPMI over LAN privileges will be reset to **read-only** level if **Configuration from Intersight only** is enabled through the Device Connector policy, or if the same configuration is enabled in the Device Connector in Cisco IMC.



Attention The Device Connector Policy will not be imported as part of the Server Profile Import.

- **Ethernet Adapter Policy**—Governs the host-side behavior of the adapter, including how the adapter handles traffic. For each VIC Virtual Ethernet Interface, you can configure various features such as VXLAN, NVGRE, ARFS, Interrupt settings, and TCP Offload settings.

This policy includes the recommended default configurations for the supported server operating systems. The policy supports 16 default configurations. During the policy creation, you can select and import a default configuration.



Note You cannot modify the default configurations. However, the policy that has the imported default configuration can be modified.

- **Ethernet Network Policy**—Allows to define the port to carry single VLAN(Access) or multiple VLANs(Trunk) traffic. You can configure the Default VLAN and QinQ VLAN settings for vNICs. You can specify the VLAN to be associated with an Ethernet packet if no tag is found.
- **Ethernet Network Control Policy**—Configures the network control settings for the appliance ports, appliance port channels, or vNICs.
- **Ethernet Network Group Policy**—Configures the VLAN settings that include Native VLAN and QinQ VLAN for appliance ports, appliance port channels, or vNICs.
- **Ethernet QoS Policy**—Assigns a system class to the outgoing traffic for a vNIC. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.
- **Fibre Channel Adapter Policy**—Governs the host-side behavior of the adapter, including how the adapter handles traffic. You can enable FCP Error Recovery, change the default settings of Queues, and Interrupt handling for performance enhancement.

This policy includes the recommended default configurations for the supported server operating systems. The policy supports nine default configurations. During the policy creation, you can select and import a default configuration.



Note You cannot modify the default configurations. However, the policy that has the imported default configuration can be modified.

- **Fibre Channel Network Policy**—Governs the VSAN configuration for the virtual interfaces.
- **Fibre Channel QoS Policy**—Assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.
- **IPMI over LAN Policy**—Defines the protocols for interfacing with a service processor that is embedded in a server platform. The Intelligent Platform Management Interface (IPMI) enables an operating system to obtain information about the system health and control system hardware and directs the Cisco IMC to perform the required actions. You can create an IPMI Over LAN policy to manage the IPMI messages through Cisco Intersight. You can assign these user roles to an IPMI user per session:
 - **admin**—IPMI users can perform all available actions. If you select this option, IPMI users with the "Administrator" user role can create admin, user, and read-only sessions on this server.
 - **read-only**—Can view information but cannot make any changes. IPMI users with the "Administrator", "Operator", or "User" user roles can only create read-only IPMI sessions, regardless of their other IPMI privileges.
 - **user**—IPMI users can perform some functions but cannot perform administrative tasks. If you select this option, IPMI users with the "Administrator" or "Operator" user role can create user and read-only sessions on this server.



Important The encryption key to use for IPMI Communication. The key must have an even number of hexadecimal characters and not exceeding 40 characters. You can use "00" to disable the encryption key use. If the encryption key specified is less than 40 characters, then the IPMI commands must add zeroes to the encryption key to achieve a length of 40 characters.

- **LAN Connectivity Policy**—Determines the connections and the network communication resources between the server and the LAN on the network. You must create the Ethernet Adapter, Ethernet QoS, and Ethernet Network policies as part of the LAN connectivity policy. For IMM servers, use a MAC pool, or static MAC addresses, to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For more information about creating Network Policies, see [Creating Network Policies](#).
- **LDAP Policy**—Specifies the LDAP configuration settings and preferences for an endpoint. The endpoints support LDAP to store and maintain directory information in a network. The LDAP policy determines configuration settings for LDAP Servers, DNS parameters including options to obtain a domain name used for the DNS SRV request, Binding methods, Search parameters, and Group Authorization preferences. Through an LDAP policy, you can also create multiple LDAP groups and add them to the LDAP server database.
- **Local User Policy**—Automates the configuration of local user preferences. You can create one or more Local User policies which contain a list of local users that need to be configured.

- **Persistent Memory Policy**—Persistent Memory Modules (PMem Modules) are non-volatile memory modules that bring together the low latency of memory and the persistence of storage. PMem Modules provide faster access to data and retain across power cycles, based on the mode. Intersight supports the configuration of Intel® Optane™ PMem Module modules on the UCS M5 servers that are based on the Second Generation Intel® Xeon® Scalable processors. Intel® Optane™ PMem Modules can be used only with the Second-Generation Intel® Xeon® Scalable processors. The Persistent Memory Policy allows the configuration of security, Goals, and Namespaces of Persistent Memory Modules:

- **Security**—Used to configure the secure passphrase for all the persistent memory modules.
- **Goal**—Used to configure volatile memory and regions in all the PMem Modules connected to all the sockets of the server. Intersight supports only the creation and modification of a Goal as part of the Persistent Memory policy. Some data loss occurs when a Goal is modified during the creation or modification of a Persistent Memory Policy. For information on the data loss, see the Data Loss during Persistent Memory Policy Configuration and Deployment table in [Resources](#).
- **Namespaces**—Used to partition a region mapped to a specific socket or a PMem Module on a socket. Intersight supports only the creation and deletion of Namespaces as part of the Persistent Memory Policy. Modifying a Namespace is not supported. Some data loss occurs when a Namespace is created or deleted during the creation of a Persistent Memory policy. For information on the data loss, see the Data Loss during Persistent Memory Policy Configuration and Deployment table in [Resources](#).

It is important to consider the memory performance guidelines and population rules of the Persistent Memory Modules before they are installed or replaced, and the policy is deployed. The population guidelines for the PMem Modules can be divided into the following categories, based on the number of CPU sockets:

- Dual CPU for UCS [C220 M6](#), [C240 M6](#), and [B200 M6](#) servers
- Dual CPU for UCS [C220 M5](#), [C240 M5](#), and [B200 M5](#) servers
- Quad CPU for UCS [C480 M5](#) and [B480 M5](#) servers
- Dual CPU for UCS [S3260 M5](#) servers

For more information about creating a Persistent Memory policy, exceptions to the policy, and other caveats regarding the policy, see Persistent Memory Policy in [Resources](#).

- **SAN Connectivity Policy**—Determines the network storage resources and the connections between the server and the SAN on the network. This policy enables you to configure vHBAs that the servers use to communicate with the Storage Area Network. You can use WWNN and WWPNN address pools, or static WWNN and WWPNN addresses to add vHBAs and to configure them. You must create the Fibre Channel Adapter, Fibre Channel QoS, and Fibre Channel Network policies as part of the SAN connectivity policy. For more information about creating Network policies, see [Creating Network Policies](#).
- **SD Card Policy**—Configures the Cisco FlexFlash and FlexUtil Secure Digital (SD) cards for the Cisco UCS C-Series Standalone M4 and M5 servers. This policy specifies details of virtual drives on the SD cards. You can configure the SD cards in the Operating System Only, Utility Only, or Operating System + Utility modes.

When two cards are present in the Cisco FlexFlash controller and Operating System is chosen in the SD card policy, the configured OS partition is mirrored. If only single card is available in the Cisco FlexFlash controller, the configured OS partition is non-RAID. The utility partitions are always set as non-RAID.



- Note**
1. This policy is currently not supported on Cisco UCS M6 servers.
 2. You can enable up to two utility virtual drives on Cisco UCS M5 servers, and any number of supported utility virtual drives on Cisco UCS M4 servers.
 3. Diagnostics is supported only for Cisco UCS M5 servers.
 4. User Partition drives can be renamed only on Cisco UCS M4 servers.
 5. FlexFlash configuration is not supported on Cisco UCS C460 M4 servers.
 6. For the Operating System+Utility mode, the Cisco UCS M4 servers require two FlexFlash cards, and the Cisco UCS M5 servers require at least 1 FlexFlash + 1 FlexUtil card.

- **SMTP Policy**—Sets the state of the SMTP client in the managed device. You can specify the preferred settings for outgoing communication and select the fault severity level to report and the mail recipients.
- **SOL Policy**—Enables the input and output of the serial port of a managed system to be redirected over IP. You can create one or more Serial over LAN policies which contain a specific grouping of Serial over LAN attributes that match the needs of a server or a set of servers.
- **SSH Policy**—Enables an SSH client to make a secure, encrypted connection. You can create one or more SSH policies that contain a specific grouping of SSH properties for a server or a set of servers.
- **Simple Network Management Protocol (SNMP) Policy**—Configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the server are removed but not replaced.
- **Storage Policy**—A Storage policy allows you to create drive groups, virtual drives, configure the storage capacity of a virtual drive, and configure the M.2 RAID controllers.
- **Syslog Policy**—Defines the logging level (minimum severity) to report for a log file collected from an endpoint, the target destination to store the Syslog messages, and the Hostname/IP Address, port information, and communication protocol for the Remote Logging Server(s).
- **Virtual Media Policy**—Enables you to install an Operating System on the server using the KVM console and virtual media, mount files to the host from a remote file share, and enable virtual media encryption. You can create one or more Virtual Media policies, which can contain virtual media mappings for different OS images, and configure up to two virtual media mappings, one for ISO files (through CDD), and the other for IMG files (through HDD).

For more information about the various mount options for the Virtual Media volumes, see [Virtual Media Mount options](#).
- **Virtual KVM Policy**—Enables specific grouping of virtual KVM properties. This policy allows you specify the number of allowed concurrent KVM sessions, port information, and video encryption options.
- **IMC Access Policy**—Enables to manage and configure your network through mapping of IP pools to the server profile. This policy allows you to configure a VLAN and associate it with an IP address through the IP pool address.

In-Band IP address, Out-of-Band IP address, or both In-Band and Out-of-Band IP addresses can be configured using IMC Access Policy and are supported on the following:

- Drive Security, SNMP, Syslog, and vMedia policies
- vKVM, IPMI, SOL, and vMedia policies using vKVM client
- **Power Policy**—Enables the management of power for FI-attached servers and chassis. This policy allows you to set the power profiling the power priority of the server, and the power restore state of the system. For more information, see [Creating a Power Policy for Server](#)
- **NTP Policy**—Allows you to enable the NTP service on an Intersight Managed Cisco IMC (Standalone) server. The NTP service synchronizes the time with an NTP server. You must enable and configure the NTP service by specifying the IP address or DNS of a minimum of one to a maximum of four NTP servers.

NTP policy also allows you to configure the timezone on Cisco IMC (Standalone) server. When you enable the NTP service and select Timezone, Cisco Intersight configures the NTP details and Timezone on the endpoint.
- **FC Zone Policy**—Allows you to set up access control between hosts and storage devices. You can create a Single Initiator Single Target, or Single Initiator Multiple Target Zone on a VSAN with the scope FC Storage, and attach the Zone policy to the SAN Connectivity policy using the vHBA.



Note You can configure zones only when the Fabric Interconnect is in FC switching mode

Configuration drift is not supported for the FC Zone policy

Creating a Policy

In Cisco Intersight, you can create a UCS Server or UCS Domain policy by using the policy wizard. To create and configure a new policy, do the following:

-
- Step 1** Log in to Cisco Intersight with your Cisco ID and select admin role.
 - Step 2** From the **Service Selector** drop-down list, select **Infrastructure Service**.
 - Step 3** Navigate to **Configure > Policies**, and then click **Create Policy**.
 - Step 4** Select **UCS Server > <A UCS server policy>**.
 - Step 5** Click **Start** to begin configuring the policy.
 - Step 6** On the **General** page, enter the **Name** of the policy. Optionally, enter a **Description** and **Tags**.
 - Step 7** On the **Policy Details** page, configure policy properties.

Some policy properties may be applicable to specific target platforms—Standalone UCS servers, FI-attached UCS servers, or both. You can view these properties separately for **All Platforms**, **UCS Servers (Standalone)**, and **UCS Servers (FI-Attached)** by clicking on these options. The properties that are applicable only to Standalone servers or FI-Attached servers are indicated by an icon alongside the property.

Step 8 Click Create.

Supported UCS Server Policies

The following table provides a list of UCS server policies and the managed devices on which they are supported. All the server policies listed in this table are available with a Cisco Intersight Essentials license.

UCS Server Policy	Supported Servers										
	Cisco UCS C-Series							Cisco UCS B-Series		Cisco UCS X-Series	
	Standalone				IMM			IMM		IMM	
	M4	M5	M6	M7	M5	M6	M7	M5	M6	M6	M7
Certificate Management Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Device Connector Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—
IPMI Over LAN Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LDAP Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—
Local User Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NTP Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—
Network Connectivity Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—
Persistent Memory Policy	—	Yes	Yes	Yes	—	—	—	—	—	—	—
Power Policy	—	—	—	—	—	—	—	Yes	Yes	Yes	Yes

UCS Server Policy	Supported Servers											
	Cisco UCS C-Series							Cisco UCS B-Series		Cisco UCS X-Series		
	Standalone				IMM			IMM		IMM		
	M4	M5	M6	M7	M5	M6	M7	M5	M6	M6	M7	
SD Card Policy	Yes	Yes	—	—	Yes	—	—	Yes	—	—	—	
SMTP Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	
SNMP Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
SSH Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—	
Serial Over LAN (SoL) Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Syslog Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Virtual KVM Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
BIOS Token Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Virtual Media Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
LAN Connectivity Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
SAN Connectivity Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Boot Order Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	

UCS Server Policy	Supported Servers										
	Cisco UCS C-Series							Cisco UCS B-Series		Cisco UCS X-Series	
	Standalone				IMM			IMM		IMM	
	M4	M5	M6	M7	M5	M6	M7	M5	M6	M6	M7
Adapter Configuration Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—
Drive Security Policy	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Storage Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IMC Access Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Adapter Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Network Policy	Yes	Yes	Yes	Yes	—	—	—	—	—	—	—
Ethernet QoS Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Network Control Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ethernet Network Group Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FC Zone Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fibre Channel Adapter Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

UCS Server Policy	Supported Servers											
	Cisco UCS C-Series							Cisco UCS B-Series		Cisco UCS X-Series		
	Standalone				IMM			IMM		IMM		
	M4	M5	M6	M7	M5	M6	M7	M5	M6	M6	M7	
Fibre Channel Network Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Fibre Channel QoS Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
iSCSI Boot Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
iSCSI Adapter Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
iSCSI Static Target Policy	—	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Firmware Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Thermal Policy	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	

Creating a Certificate Management Policy

In Intersight Managed Mode, the Certificate Management policy allows you to specify the certificate details for an external certificate and attach the policy to servers. Cisco Intersight currently supports the following certificates:

- **Root CA certificates:** A Root CA certificate is necessary for HTTPS boot authentication. You can deploy a maximum of 10 Root CA certificates using the Certificate Management Policy. For a successful boot, at least one valid and unexpired Root CA certificate is required. For more information, see [Creating a Boot Order Policy](#).



Note In Intersight Managed Mode servers, removing a server profile will delete the Root CA certificates from the CIMC.

However, for C-Series servers in Standalone mode, the Root CA certificates are not automatically removed. You must manually delete them from CIMC or perform a factory reset on the server. Additionally, when you export the profile of a C-Series server in Standalone mode, the certificate management policy will not be included.

• **IMC certificates:** This option is available only for Intersight Managed Mode servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Certificate Management**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, add the certificate that you want to provide, and configure the following parameters:

Property	Essential Information
Root CA	<ul style="list-style-type: none"> • Certificate Name—Enter the name of the certificate. • Certificate—Enter the certificate details.
IMC	<ul style="list-style-type: none"> • Certificate—Enter the certificate details. • Private Key—Enter the private key details for the certificate.

7. Click **Create**.

Creating an Adapter Configuration Policy

An Adapter Configuration Policy configures the Ethernet and Fibre-Channel settings for the Virtual Interface Card (VIC) adapter.



Note This policy, if attached to a server profile that is assigned to an Intersight Managed Fabric Attached server, will be ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Adapter Configuration**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, click **Add VIC Adapter Configuration** and configure the following parameters:

Property	Essential Information
Add VIC Adapter Configuration	
PCI Slot	The PCI slot in which the adapter is installed. The range is from 1 to 15 and MLOM.

Property	Essential Information
<p>LLDP</p>	<p>The LLDP protocol status on the adapter interface.</p> <p>If checked, then Link Layer Discovery Protocol (LLDP) enables all the Data Center Bridging Capability Exchange protocol (DCBX) functionality, which includes FCoE, priority based flow control.</p> <p>Note LLDP is available only on some UCS C-Series servers.</p> <p>We recommend that you do not disable LLDP option, as it disables all the DCBX functionality.</p>
<p>FIP</p>	<p>The FIP protocol status on the adapter interface.</p> <p>If checked, then FCoE Initialization Protocol (FIP) mode is enabled. FIP mode ensures that the adapter is compatible with current FCoE standards.</p> <p>Note We recommend that you use FIP option only when explicitly directed to do so by a technical support representative.</p>
<p>Port Channel</p>	<p>The port channel status on the adapter interface.</p> <p>When Port Channel is enabled, two vNICs and two vHBAs are available for use on the adapter card. When disabled, four vNICs and four vHBAs are available for use on the adapter card. Disabling port channel reboots the server.</p> <p>Note Port Channel is supported only for Cisco VIC 1455/1457 adapters.</p>

Property	Essential Information
Enable Physical NIC Mode	<p>When Physical NIC Mode is enabled, uplink ports of the VIC are set to pass-through mode. This allows the host to transmit packets without any modification. VIC ASIC does not rewrite the VLAN tag of the packets based on the VLAN and CoS settings for the vNIC.</p> <p>Note</p> <ul style="list-style-type: none"> • Enabling Physical NIC Mode reboots the server. • Physical NIC Mode supports UCS VIC 1400 Series and VIC 15000 Series adapters. • The minimum supported Cisco Server firmware version 4.2(2a) and later and Adapter firmware version 5.2(2a). • This feature is not supported for Cisco Intersight Managed FI Attached servers. • Only default vNICs will be added if the Physical NIC mode is enabled. • This option cannot be enabled on an adapter that has: <ul style="list-style-type: none"> • Port Channel mode enabled • VNTAG mode enabled • LLDP enabled • FIP mode enabled • Cisco IMC Management Enabled value set to Yes <p>When Physical NIC Mode is enabled, the following message is displayed in a pop-up window:</p> <p>After physical nic-mode mode switch, vNIC configurations will be lost and new default vNICs will be created.</p> <p>Click Ok.</p>

Property	Essential Information
DCE Interface	<p>The Forward Error Correction (FEC) mode setting for the DCE interfaces of the adapter.</p> <p>Note FEC mode setting is supported only for Cisco VIC 14xx adapters. FEC mode 'cl74' is unsupported for Cisco VIC 1495/1497. This setting will be ignored for unsupported adapters and for unavailable DCE interfaces</p>

7. Click **Add**.
8. Click **Create**.

Creating a LAN Connectivity Policy

A LAN Connectivity Policy determines the connections and the network communication resources between the server and the LAN on the network. You can specify MAC address pools, or static MAC addresses, to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network.

Prerequisites

Choose the following sub-policies or pool as per your requirement to create the LAN Connectivity policy

- **Ethernet Network Policy**—Specify if the port should carry single VLAN (Access) or multiple VLANs (Trunk) traffic. You can specify the VLAN to be associated with an Ethernet packet if no tag is found.
- **Ethernet QoS Policy**—Configure the maximum size for a Fibre Channel frame payload that the virtual interface supports, limit the data rate on the virtual interface, associate a Class of Service to the traffic on the virtual interface.
- **Ethernet Adapter Policy**—Configure features like VXLAN, NVGRE, ARFS, Interrupt settings, RoCE, and TCP Offload settings to govern the host side behavior of the adapter.
- **IQN Pool**—You can configure the Prefix and Suffix for the IQN block, the first suffix number in the block and the number of identifiers the block can hold .

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **LAN Connectivity**, and then click **Start**.
5. On the **General** page, enter the following information:
 - **Name** of your policy.
 - **Target Platform** for which the policy is applicable. This can be **Standalone** servers or **FI Attached** servers.

A LAN Connectivity Policy created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a LAN Connectivity Policy created for FI Attached servers cannot be deployed on Standalone servers.

- **Description** to help identify the policy.
- **Tag** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following:

- For an FI-attached server, turn the **Enable Azure Stack Host QoS** button ON, to successfully deploy the Azure Stack QoS capability on the adapter with RDMA enabled.

Enabled—Enabling AzureStack-Host QoS on an adapter allows the user to carve out traffic classes for RDMA traffic and ensure a desired portion of the bandwidth is allocated to it.

Disabled—Disables the Azure Stack Host QoS feature on the adapter.

- Specify whether no IQN, an IQN pool, or a unique IQN identifier is to be associated with the policy by selecting **None**, **Pool**, or **Static**.
 - **None**—If you select this option, you do not have to specify any IQN details.
 - **Pool**—If you select this option, select the IQN pool that you want to associate with the LAN Connectivity policy.
 - **Static**—If you select this option, enter a static IQN for use as initiator identifiers by iSCSI vNICs in a Fabric Interconnect domain.
- Select the placement option for each vNIC—**Manual** or **Auto**
 - **Manual vNIC Placement**—If you select this option, you must manually specify the placement for each vNIC. You can also use the **Graphic vNICs Editor** to create and specify the placement for each vNIC manually by adding vNICs and slots, and defining the connection between them.



Note

- For manual placement, **PCI Link** is not supported on UCS VIC 1400 Series adapters.
- If a LAN Connectivity Policy has both Simple and Advanced placements, ensure the number provided in PCI Order is appropriate to prevent Server Profile deployment failure.
- **Auto vNIC Placement**—If you select this option, vNIC placement will be done automatically during profile deployment. This option is available only for Cisco Intersight Managed FI Attached servers.



- Note**
- Cisco UCS VIC 1300 Series adapters auto-upgrade is supported on B-Series server with Cisco Server firmware version 4.2(2e) and above.
 - Discovery of a C-Series server will not get triggered if the server with Cisco UCS VIC 1300 Series adapters has a Cisco Server firmware version lower than 4.2(2g). Upgrade the Cisco Server firmware to 4.2(2g) to enable server discovery.

7. To set up a vNIC without using a template, click **Add vNIC** and configure the following parameters:

Property	Essential Information
Add vNIC	Ensure that you configure eth0 and eth1 interfaces for each VIC adapter you configure. You can add additional vNICs depending on your network requirements.
Name	vNIC name.
Pin Group Name	<p>Name of the pin group that contains the specific port/port channels. All traffic from the vNIC is pinned to the specified uplink Ethernet ports or port channels.</p> <p>Note The pin group can be defined while creating a Port policy.</p> <p>If you do not assign a pin group to a vNIC, an uplink Ethernet port or port channel for traffic is chosen from that server interface dynamically. This choice is not permanent. A different uplink Ethernet port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.</p>
MAC Address Pool	Click Select Pool and choose a MAC address pool for MAC address assignment.
Static	Click Static and enter a static MAC address for MAC address assignment. This option is available only for Cisco Intersight Managed FI Attached servers.
Placement	Placement Settings for the virtual interface.

Property	Essential Information
<p>Simple</p> <p>When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vNICs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0.</p>	
<p>Switch ID</p>	<p>Refers to the Fabric Interconnect that carries the vNIC traffic.</p>
<p>PCI Order</p>	<p>The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter.</p> <p>Note You cannot change the PCI order of two vNICs without deleting and recreating the vNICs.</p>
<p>Advanced</p>	
<p>Automatic Slot ID Assignment</p>	<p>When enabled, slot ID is determined automatically by the system.</p>
<p>Slot ID</p>	<p>When automatic slot ID assignment is disabled, the slot ID needs to be entered manually.</p> <p>Supported values are (1-15) and MLOM.</p>
<p>PCI link</p> <p>The PCI link used as transport for the virtual interface.</p> <p>PCI Link is only applicable for select Cisco UCS VIC 1300 Series models (UCSC-PCIE-C40Q-03, UCSB-MLOM-40G-03, UCSB-VIC-M83-8P) that support two PCI links. The value, if specified, for any other VIC model will be ignored.</p> <p>Note The host device order can get impacted when using both the PCI links and while adding or removing vNICs.</p>	

Property	Essential Information
<p>Automatic PCI link Assignment</p>	<p>When enabled, PCI link is determined automatically by the system.</p> <p>Note</p> <ul style="list-style-type: none"> • If Automatic assignment is enabled for both Slot ID and PCI link, then the behavior is same as Simple placement. All the vNICs are placed on the same PCI link (link 0). • If Automatic Slot ID assignment is disabled but automatic PCI link assignment is enabled, then you need to provide the slot ID and the vNIC will be placed on PCI link 0.
<p>Load Balanced</p>	<p>When Automatic PCI link assignment is disabled and Load Balanced is enabled, the system uniformly distributes the interfaces across the PCI Links.</p> <ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to specify the PCI order to load balance the vNICs. • If both automatic PCI link assignment and automatic Slot ID are disabled, you need to specify the slot and the PCI order to load balance the vNICs. <p>Note You cannot change the PCI link mode of two vNICs from Load Balanced mode to Custom mode without deleting and recreating the vNICs.</p>
<p>Custom</p>	<ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to provide the value of the PCI order, PCI link, and Switch ID. • If both automatic PCI link assignment and automatic Slot ID assignment are disabled, you need to provide the values of the Slot ID, PCI order and the PCI link. <p>Note You cannot change the PCI link mode of two vNICs from Custom mode to Load Balanced mode without deleting and recreating the vNICs.</p>

Property	Essential Information
Consistent Device Naming (CDN) Consistent Device Naming configuration for the virtual NIC.	
Source	Whether the source of the CDN name is the name of the vNIC instance or a user-defined name.
Failover Enabling failover ensures that traffic automatically fails over from one uplink to another in case of an uplink failure.	
Ethernet Network Policy	Select or create an Ethernet Network policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on Standalone servers.
Ethernet Network Group Policy	Select or create an Ethernet Network Group policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Ethernet Network Control Policy	Select or create an Ethernet Network Control policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Ethernet QoS Policy	Select or create an Ethernet QoS policy.
Ethernet Adapter Policy	Select or create an Ethernet Adapter policy.
iSCSI Boot Policy	Select or create an iSCSI Boot policy. Note This sub-policy is applicable only for the LAN Connectivity Policy on FI-attached servers.
Connection: Disabled/usNIC/VMQ/SR-IOV	
Disabled	Does not configure a connection policy.
usNIC User Space NIC Settings that enable low-latency and higher throughput by bypassing the kernel layer when sending/receiving packets.	
Number of usNICs	Number of usNIC interfaces to be created.
usNIC Adapter Policy	Select the Ethernet Adapter policy to be associated with the usNICs.

Property	Essential Information
Class of Service	Class of service to be used for traffic on the usNIC.
VMQ Virtual Machine Queue Settings for the virtual interface that allow efficient transfer of network traffic to the guest operating system.	
Enable Multi Queue Support	Whether Virtual Machine Multi-Queue (VMMQ) is enabled in the policy. With VMMQ, multiple queues are allocated to a single VM.
Number of Sub vNICs	Number of sub vNICs that are available for Multi Queue.
Enable RoCE Settings	Whether Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) is enabled over this virtual interface.
Memory Regions	The number of memory regions per adapter. Enter an integer between 1 and 524288. It is recommended that this number be an integer power of 2.
Queue Pairs	The number of queue pairs per adapter. Enter an integer between 1 and 8192. It is recommended that this number be an integer power of 2.
Resource Groups	The number of resource groups per adapter. Enter an integer between 1 and 128. It is recommended that this number be an integer power of 2 greater than or equal to the number of CPU cores on the system for optimum performance.
Version	Version of the RDMA protocol Version 1 is a link layer protocol. It allows communication between any two hosts in the same Ethernet broadcast domain. RoCEv2 is an internet layer protocol. RoCEv2 packets can be routed. This is possible because RoCEv2 packets now include an IP and UDP header.
SR-IOV Single Root Input/Output Virtualization (SR-IOV) allows multiple VMs running a variety of Linux guest operating systems to share a single PCIe network adapter within a host server. SR-IOV allows a VM to move data directly to and from the vNIC, bypassing the hypervisor for increased network throughput and lower server CPU overhead.	

Property	Essential Information
Number of VFs	Number of VFs to create. Enter a value between 1 and 64. Default value is 64.
Receive Queue Count Per VF	Number of Receive Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 4.
Transmit Queue Count Per VF	Number of Transmit Queue resources to configure for each VF. Enter a value between 1 to 8. Default value is 1.
Completion Queue Count Per VF	Number of Completion Queue resources to configure for each VF. Enter a value between 1 to 16. Default value is 5.
Interrupt Count Per VF	Number of Interrupt count to configure for each VF. Enter a value between 1 to 16. Default value is 8.

- To derive vNIC for FI-attached servers using a vNIC template, choose **vNIC from Template** from the **Add** drop-down list. For more information on creating vNIC templates, see *Creating vNIC or vHBA Templates*.



Note

- When deriving a vNIC from a template, the vNIC configuration is auto-populated from the template configuration. You can edit or delete parameters, which are enabled for configuration override through the vNIC template. For parameters that are not enabled for override, you can only view the configurations using the Eye icon.
- The parameters that have been overridden are indicated using an Overridden label. In the case of override-enabled parameters, the changes applied in the template are not reflected in the derived vNIC.
- Only those parameters can be modified in the derived vNIC instance which are not included in the template.
- If you attempt to derive a vNIC from a template while profile deployment is in progress, the task will be retried until the profile deployment is completed. You can find these details in the Requests tab.

- Click **Create**.

Configuration Feature Matrix for Supported Adapters in IMM

The following table shows the features supported by various adapters in Intersight Managed Mode.

Feature	Cisco UCS 1300 Series Adapter	Cisco UCS 1400/14000 Series Adapter	Cisco UCS 15000 Series Adapter
usNIC	Yes	Yes	Yes
VMQ	Yes	Yes	Yes
VMMQ	No	Yes	Yes

Feature	Cisco UCS 1300 Series Adapter	Cisco UCS 1400/14000 Series Adapter	Cisco UCS 15000 Series Adapter
SR-IOV	No	Yes	Yes
NetQueue	Yes	Yes	Yes
RoCEv1	Yes	No	No
RoCEv2	No	Yes	Yes
Geneve Offload	No	Yes	Yes
AzureQoS	No	Yes	Yes
RSS	Yes	Yes	Yes
RSSv2	No	No	Yes
NVGRE	Yes	Yes	Yes
ARFS	Yes	Yes	Yes
VIC QinQ Tunneling	No	Yes	Yes
VXLAN	Yes	Yes	Yes
Advance Filter	Yes	Yes	Yes
Interrupt Scaling/Group Interrupt	Yes	Yes	Yes
Host Port Configuration	Yes	No	No
vHBA Type	Yes	Yes	Yes
16K Ring Size	No	No	Yes
Precision Time Protocol	No	No	Yes
FC MQ	Yes	Yes	Yes
FC NVMe	Yes	Yes	Yes
ENS	No	Yes	Yes

Creating an Ethernet Adapter Policy

An Ethernet adapter policy governs the host-side behavior of the adapter, including how the adapter handles traffic. For each VIC Virtual Ethernet Interface, you can configure various features like Virtual Extensible LAN (VXLAN), Network Virtualization using Generic Routing Encapsulation (NVGRE), Accelerated Receive Flow Steering (ARFS), Interrupt settings, and TCP Offload settings.

The Ethernet Adapter policy include the recommended settings for the virtual Ethernet interface, for each supported server operating system. Operating systems are sensitive to the settings in these policies. In general, the storage vendors require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.

GENEVE Offload

Cisco Intersight now supports Generic Network Virtualization Encapsulation (GENEVE) Offload on the ESXi platform, which allows essentially any information to be encoded in a packet and passed between tunnel endpoints. GENEVE provides the overlay capability to create isolated, multi-tenant broadcast domains across data center fabrics on 1400 Series adapters. Using the GENEVE protocol allows you to create logical networks that span physical network boundaries.

GENEVE offload is present in all Ethernet adapter policies and is disabled by default. It is the recommended setting if using VMWare ESXi GENEVE.

For more information on how to implement GENEVE offload end-to-end configuration, see [Cisco UCS Manager Network Management](#) documentation.

Cisco recommends configuring the following values in the Ethernet adapter policy when GENEVE offload is enabled:

- Transmit Queues :1
- TX Ring Size: 4096
- Receive Queues: 8
- RX Ring Size: 4096
- Completion Queues : 16
- Interrupts : 32

The following features are not supported when GENEVE offload is enabled on any interface:

- Azure Stack QoS
- RoCEv2 - you cannot have GENEVE enabled on one vNIC and RoCEv2 enabled on another.
- Advanced Filters
- VIC QinQ Tunneling

Support for usNIC and VIC QinQ Tunneling features on interfaces:



Note

- usNIC or VMQ is not compatible with GENEVE Offload on the same interface only for 1400 Series adapters.
- usNIC or VMQ is compatible with GENEVE Offload on different interfaces for 1400 Series adapters.
- usNIC and VMQ is compatible with GENEVE Offload on both the same and different interfaces for 1500 Series adapters.



Note On switching from GENEVE offload feature to Azure Stack QoS feature or vice versa, please do the following:

1. Disable the current feature
2. Reboot the server
3. Enable the required feature

Other limitations with GENEVE offload include:

- External outer IPV6 is NOT supported with GENEVE offload.
- GENEVE offload is supported with ESX 7.0 (NSX-T 3.0) and ESX 6.7U3(NSX-T 2.5).
- GENEVE offload is supported only with Cisco UCS VIC 1400/14000 and 15000 Series adapters. It is not supported on Cisco UCS VIC 1300 Series adapters or Cisco UCS VIC 1200 Series adapters.
- Cisco UCS VIC 1400/14000 and 15000 Series adapters.
- Minimum server firmware version for UCS C-Series Standalone: 4.1(2a)
- Minimum adapter firmware version: 5.1(2f)
- Cisco recommends that you remove the GENEVE offload configuration before downgrading to any non-supported release.

For details on supported features matrix with GENEVE offload, refer the table below.

Table 1: GENEVE Offload Supported Features Matrix for 1400 Series Adapters

	KVM VM - FEX	VXLAN	NVGRE	RoCEv2	usNIC	Netflow	Advanced Filters	VMQ/ VMMQ/ netqueue	arfs	Azure QoS
GENEVE offload enabled on the interface vnic1 and feature is enabled on vnic1	No	Yes	Yes	No	No	No	No	No	No	No
GENEVE offload that is enabled on the interface vnic1 and feature is enabled on vnic2	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	No



Note 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.

Table 2: GENEVE Offload Supported Features Matrix for 15000 Series Adapters

	VXLAN	NVGRE	RoCEv2	usNIC	Netflow	Advanced Filters	VMQ/VMMQ/netqueue	arfs	quad port per adapter	physical nic node per adapter
GENEVE offload enabled on the same interface (vnic1) and feature is enabled on vnic1	Yes	Yes	No	Yes	Yes	Yes	Yes	No	Yes	Yes
GENEVE offload enabled on different interface (vnic1) and feature is enabled on vnic2	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Adapter**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Set Tags	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.
Ethernet Adapter Default Configuration	
Select a default configuration	Click to view and import a default Cisco provided configuration. The policy currently supports up to 16 default configurations.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Virtual Extensible LAN	Enables the Virtual Extensible LAN protocol on the virtual Ethernet interface.

Property	Essential Information
<p>Enable Network Virtualization using Generic Routing Encapsulation</p>	<p>Enables Network Virtualization using Generic Routing Encapsulation on the virtual Ethernet interface.</p> <p>Note The Transmit checksum offload and TSO must be enabled for the NVGRE offloading to be effective.</p>
<p>Enable Accelerated Receive Flow Steering</p>	<p>Enables Accelerated Receive Flow Steering (ARFS) on the virtual Ethernet interface. ARFS is hardware-assisted receive flow steering that can increase CPU data cache hit rate by steering kernel level processing of packets to the CPU where the application thread consuming the packet is running.</p>
<p>Enable Advanced Filter</p>	<p>Enables advanced filtering on the virtual Ethernet interface.</p>
<p>Enable Interrupt Scaling</p>	<p>Enables Interrupt Scaling of resources on the virtual Ethernet interface.</p>
<p>Enable Geneve Offload</p>	<p>Enables GENEVE overlay hardware offloads.</p>
<p>RoCE Settings</p> <p>Intersight supports RDMA over Converged Ethernet (RoCE) for Microsoft SMB Direct. It sends additional configuration information to the adapter while creating or modifying an Ethernet adapter policy.</p>	
<p>Enable RDMA over converged Ethernet</p>	<p>Enables RDMA over Converged Ethernet (RoCE) on the virtual Ethernet interface.</p> <p>RoCE allows direct memory access over an Ethernet network. RoCE is a link layer protocol, and hence, it allows communication between any two hosts in the same Ethernet broadcast domain. RoCE delivers superior performance compared to traditional network socket implementations because of lower latency, lower CPU utilization, and higher utilization of network bandwidth.</p>
<p>Queue Pairs</p>	<p>The number of queue pairs per adapter.</p> <p>Enter an integer between 0 and 8192. It is recommended that this number be an integer power of 2.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>

Property	Essential Information
Memory Regions	<p>The number of memory regions per adapter.</p> <p>Enter an integer between 0 and 524288. It is recommended that this number be an integer power of 2.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Resource Groups	<p>The number of resource groups per adapter. It is recommended that this number be an integer power of 2 greater than or equal to the number of CPU cores on the system for optimum performance.</p> <p>Enter an integer between 0 and 128.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Version	<p>Version of the RDMA protocol</p> <p>Version 1 is a link layer protocol. It allows communication between any two hosts in the same Ethernet broadcast domain.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Interrupt Settings	
Interrupts	<p>Enter the number of interrupt resources to allocate. Typically this value is equal to the number of completion queue resources.</p> <p>Enter an integer between 1 and 1024.</p>
Interrupt Mode	<p>Select the preferred driver interrupt mode that include:</p> <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—Message Signaled Interrupts (MSI) only • INTx—PCI INTx interrupts

Property	Essential Information
Interrupt Timer, us	<p>The time to wait between interrupts or the idle period that must be encountered before an interrupt is sent. To turn off interrupt coalescing, enter 0 (zero) in this field.</p> <p>Enter an integer between 0 and 65535.</p>
Interrupt Coalescing Type	<p>Select the Interrupt Coalescing Type:</p> <ul style="list-style-type: none"> • Min - The system waits for the time specified in the Coalescing Time field before sending another interrupt event. • Idle - The system does not send an interrupt until there is a period of no activity lasting as least the time specified in the Coalescing Time field.
<p>Receive Receive Queue resource settings.</p>	
Receive Queue Count	<p>The number of queue resources to allocate.</p> <p>Enter an integer between 1 and 1000.</p>
Receive Ring Size	<p>The number of descriptors in each queue.</p> <p>Enter an integer between 64 and 4096.</p>
<p>Transmit Transmit Queue resource settings</p>	
Transmit Queue Count	<p>The number of queue resources to allocate.</p> <p>Enter an integer between 1 and 1000.</p>
Transmit Ring Size	<p>The number of descriptors in each queue.</p> <p>Enter an integer between 64 and 4096.</p>
<p>Completion Completion Queue resources settings</p>	
Completion Queue Count	<p>The number of completion queue resources to allocate. In general, the number of completion queue resources to allocate is equal to the number of transmit queue resources plus the number of receive queue resources.</p> <p>Enter an integer between 1 and 2000.</p>

Property	Essential Information
Completion Ring Size	<p>The number of descriptors in each queue.</p> <p>Enter an integer between 1 and 256.</p> <p>Note This property is displayed only when Enable RDMA over converged Ethernet is enabled.</p>
Uplink Failback Timeout (seconds)	<p>Uplink Failback Timeout in seconds when uplink failover is enabled for a vNIC. After a vNIC has started using its secondary interface, this setting controls how long the primary interface must be available before the system resumes using the primary interface for the vNIC.</p> <p>Enter an integer between 0 and 600.</p>
<p>TCP Offload</p> <p>The TCP offload settings decide whether to offload the TCP related network functions from the CPU to the network hardware or not. These options help reduce the CPU overhead and increase the network throughput.</p>	
Enable Tx Checksum Offload	Enables the CPU to send all packets to the hardware so that the checksum can be calculated.
Enable Rx Checksum Offload	Enables the CPU to send all packet checksums to the hardware for validation.
Enable Large Send Offload	Enables the CPU to send large packets to the hardware for segmentation.
Enable Large Receive Offload	Enables the CPU to reassemble the segmented packets in hardware before sending them to the CPU.
<p>Receive Side Scaling: Receive Side Scaling (RSS)/Receive Side Scaling Version 2 (RSSv2) supports multiple cores to process the incoming data traffic.</p> <p>RSSv2 is supported on Windows 2019 OS and later versions and it requires Windows NENIC driver. With RSS enabled Windows NENIC driver and Cisco UCS VIC adapter, you can configure multiple hardware receive queues on the Physical Function(PF). With VMMQ enabled on the VIC, you can configure multiple hardware receive queues per Virtual Machine(VM).</p> <p>Before using the RSSv2 functionality, ensure the NENIC driver supports RSSv2. In general, a NENIC driver supports 4 queues. With RSSv2, the NENIC driver has no upper limit on the number of hardware queues for PF or VM.</p>	

Property	Essential Information
Enable Receive Side Scaling	<p>Enables receive side scaling and allows the incoming traffic to be spread across multiple CPU cores. This property supports both RSS and RSSv2.</p> <p>By default, RSS is enabled. RSSv2 is compatible with RSS. Based on the NENIC driver support on RSS or RSSv2, this property is supported accordingly.</p> <p>Note RSSv2 is supported on the following:</p> <ul style="list-style-type: none"> • Cisco UCS VIC 15000 Series adapters • Cisco UCS M6 and M7 servers
Enable IPv4 Hash	Enables the IPv4 address for traffic distribution.
Enable IPv6 Extension Hash	Enables the IPv6 extensions for traffic distribution.
Enable IPv6 Hash	Enables the IPv6 address for traffic distribution.
Enable TCP and IPv4 Hash	Enables both the IPv4 address and TCP port number for traffic distribution.
Enable TCP and IPv6 Extensions Hash	Enables both the IPv6 extensions and TCP port number for traffic distribution.
Enable TCP and IPv6 Hash	Enables both the IPv6 address and TCP port number for traffic distribution.
Enable UDP and IPv4 Hash	Enables both the IPv4 address and UDP port number for traffic distribution.
Enable UDP and IPv6 Hash	Enables both the IPv6 address and UDP port number for traffic distribution.

7. Click **Create**.

Creating an Ethernet QoS Policy

An Ethernet Quality of Service (QoS) policy assigns a system class to the outgoing traffic for a vNIC. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.

4. Select **Ethernet QoS**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
MTU, Bytes	The Maximum Transmission Unit (MTU) or packet size that the virtual interface accepts. The valid range is between 1500 and 9000. The default value is 1500.
Rate Limit, Mbps	The value in Mbps (0-100000) to use for limiting the data rate on the virtual interface. Setting this to zero will turn rate limiting off.
Class of Service	The Class of Service to be associated to the traffic on the virtual interface. The valid range is between 0 and 6. The default value is 3. Note This property is supported only on Standalone servers.
Burst	The burst traffic allowed on the vNIC in bytes. The valid range is between 1024 and 1000000. The default value is 1024. Note This property is supported only on FI-attached servers.

Property	Essential Information
Priority	<p>Select the priority matching the System QoS defined in the domain profile that include:</p> <ul style="list-style-type: none"> • Best-effort • Fibre Channel (FC) • Platinum • Gold • Silver • Bronze <p>Note</p> <ul style="list-style-type: none"> • The Best-effort system class is enabled by default. • This property is supported only on FI-attached servers.
Enable Trust Host CoS	Select to enable the usage of the Class of Service to be associated to the traffic on the virtual interface.

7. Click **Create**.

Creating an Ethernet Network Policy

An Ethernet Network policy sets the rules for the port to handle network traffic. This policy determines whether the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic.

This policy also supports VIC QinQ Tunneling. A QinQ (802.1Qin802.1Q) tunnel allows segregation and isolation of different VLANs within a network. To configure QinQ VLAN, you can specify the desired VLAN ID as part of the VLAN settings for the specific port, port channel, or vNIC. This enables the transmission of multiple VLANs over a single VLAN trunk.



Important This policy is supported only on C-Series Standalone servers.

An Ethernet Network policy determines if the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic. You can specify the VLAN to be associated with an Ethernet packet if no tag is found.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
VLAN Mode	

Property	Essential Information
	<p>Assign traffic flow to the VLAN to determine if the port can carry single VLAN (Access) or multiple VLANs (Trunk) traffic.</p> <ul style="list-style-type: none"> <p>Access Mode—Traffic is received and sent in native formats with no VLAN tagging. Anything arriving on an access port is assumed to belong to the VLAN assigned to the port.</p> <p>You can configure a port in access mode and specify the VLAN to carry the traffic for that interface. If you do not configure the VLAN for a port in access mode, or an access port, the interface carries the traffic for the default VLAN, which is VLAN 1. You can change the access port membership in a VLAN by configuring the VLAN. You must create the VLAN before you can assign it as an access VLAN for an access port. If you change the access VLAN on an access port to a VLAN that is not yet created, the UCS Manager shuts down that access port.</p> <p>If an access port receives a packet with an 802.1Q tag in the header other than the access VLAN value, that port drops the packet without learning its MAC source address. If you assign an access VLAN that is also a primary VLAN for a private VLAN, all access ports with that access VLAN receives all the broadcast traffic for the primary VLAN in the private VLAN mode.</p> <p>Trunk Mode—Trunk ports allow multiple VLANs to transport between switches over that trunk link. A trunk port can carry untagged packets simultaneously with the 802.1Q tagged packets. When you assign a default port VLAN ID to the trunk port, all untagged traffic travels on the default port VLAN ID for the trunk port, and all untagged traffic is assumed to belong to this VLAN. This VLAN is referred to as the native VLAN ID for a trunk port. The native VLAN ID is the VLAN that carries untagged traffic on trunk ports.</p> <p>The trunk port sends an egressing packet with a VLAN that is equal to the default port VLAN ID as untagged; all the other egressing packets are tagged by the trunk port. If you do not configure a native VLAN ID, the trunk port uses the default VLAN.</p>

Property	Essential Information
	This property is applicable only to Standalone servers, and not to FI Attached servers. For FI Attached mode, VLAN Mode is configured as Trunk .
Access Mode	
Enable QinQ Tunneling	Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.
Default VLAN	Refers to the VLAN ID assigned to the traffic on the virtual interface by default. The range for the Default VLAN ID is from 0 to 4094.
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. This supported VLAN ID range is from 2 to 4093, allowing you to effectively manage and segregate the network traffic.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>
Trunk Mode	
Enable QinQ Tunneling	Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.
Default VLAN	Refers to the VLAN ID assigned to the traffic on the virtual interface by default. The range for the Default VLAN ID is from 0 to 4094.
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. This supported VLAN ID range is from 2 to 4093, allowing you to effectively manage and segregate the network traffic.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>

7. Click **Create**.

Creating an Ethernet Network Group Policy

An Ethernet Network Group policy enables you to manage settings for VLANs on a UCS Server. These settings include defining which VLANs are allowed, designating a Native VLAN, and specifying a QinQ VLAN.

This policy also supports VIC QinQ Tunneling. A QinQ (802.1Qin802.1Q) tunnel allows segregation and isolation of different VLANs within a network. To configure QinQ VLAN, you can specify the desired VLAN ID as part of the VLAN settings for the specific port, port channel, or vNIC. This enables the transmission of multiple VLANs over a single VLAN trunk.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network Group**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
VLAN Settings	
Native VLANs	<p>This property allows you to specify the native VLAN ID for the virtual interface or its corresponding vethernet in a range of 1-4093.</p> <ul style="list-style-type: none"> • If the native VLAN is not already part of the allowed VLANs, it will be automatically added to the list of allowed VLANs. • If QinQ Tunneling is enabled, the native VLAN and Allowed VLAN properties are combined.
Enable QinQ Tunneling	Slide to enable VIC QinQ (802.1Qin802.1Q) Tunneling.

Property	Essential Information
Allowed VLANs	<p>Refers to the VLANs that are permitted for the virtual interface. You can specify the allowed VLANs by providing a list of comma-separated VLAN IDs and VLAN ID ranges.</p> <p>For example, you can enter VLAN IDs 10, 20, 30-40 to allow VLANs 10, 20, and a range from 30 to 40.</p> <p>Note This property is displayed only when <i>Enable QinQ Tunneling</i> slider is disabled.</p>
QinQ VLAN	<p>This property enables the configuration of QinQ Tunneling, that facilitates the encapsulation of multiple VLANs within a single VLAN. The supported VLAN IDs range from 2 to 4093 that allows you to effectively manage and segregate the network traffic.</p> <p>Note This property is available only when <i>Enable QinQ Tunneling</i> slider is enabled.</p>



Note To make the server an Isolated host or a Community host, specify the ID of an Isolated VLAN or a Community VLAN in both Allowed VLANs and Native VLAN

7. Click **Create**.

Creating an Ethernet Network Control Policy

Ethernet Network Control policies configure the network control settings for the UCS Domain. This policy is applicable only for the Appliance Ports defined in a Port Policy and for the vNICs defined in a LAN Connectivity Policy, on an FI-Attached UCS Servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Ethernet Network Control**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.

Property	Essential Information
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable CDP	Enables the Cisco Discovery Protocol (CDP) on an interface.
MAC Register Mode	<p>Determines the MAC addresses to be registered with the switch. This can be:</p> <ul style="list-style-type: none"> • Only Native VLAN—MAC addresses are only added to the native VLAN. This option is the default, and it maximizes the port+VLAN count. • All Host VLANs—MAC addresses are added to all VLANs with which they are associated. Select this option if your VLANs are configured to use trunking but are not running in Promiscuous mode.
Action on Uplink Fail	<p>Determines how the interface behaves if no uplink port is available when the switch is in end-host mode.</p> <ul style="list-style-type: none"> • Link Down—Changes the operational state of a vNIC to down when uplink connectivity is lost on the switch, and enables fabric failover for vNICs. This is the default option. • Warning—Maintains server-to-server connectivity even when no uplink port is available, and disables fabric failover when uplink connectivity is lost on the switch.

Property	Essential Information
MAC Security Forge	Determines whether forged MAC addresses are allowed or denied when packets are sent from the server to the switch. This can be: <ul style="list-style-type: none"> • Allow— All server packets are accepted by the switch, regardless of the MAC address associated with the packets. This is the default option. • Deny— After the first packet has been sent to the switch, all other packets must use the same MAC address or they will be silently rejected by the switch. In effect, this option enables port security for the associated vNIC.
LLDP	Determines whether interfaces can transmit or receive LLDP packets. <ul style="list-style-type: none"> • To enable or disable the transmission of LLDP packets on an interface, click Enable Transmit. • To enable or disable the receipt of LLDP packets on an interface, click Enable Receive.

7. Click **Create**.

Creating a SAN Connectivity Policy

A Storage Area Network (SAN) connectivity policy determines the network storage resources and the connections between the server and the storage device on the network. This policy enables you to specify WWPN address pools, or a static WWPN address to add a vHBA. Similarly, you can specify a WWNN pool, or a static WWNN address to configure vHBAs that the servers use to communicate with the SAN.

Prerequisites

The following sub-policies are required to create the SAN Connectivity policy:

- **Fibre Channel Network Policy**—Configure the VSAN ID on the virtual interfaces.
- **Fibre Channel QoS Policy**—Limit the data rate on the virtual interface, configure the maximum size for a Fibre Channel frame payload bytes that the virtual interface supports, associate a Class of Service to the traffic on the virtual interface.
- **Fibre Channel Adapter Policy**—Govern the host side behavior of the adapter. You can enable FCP Error Recovery, change the default settings of Queues, and change Interrupt handling for performance enhancement.
- **Fibre Channel Zone Policy**—Specify direct access storage path configurations in the FC Zone policy, to set up access control between hosts and storage devices. You can create a Single Initiator Single Target, or Single Initiator Multiple Target zone on a VSAN with FC Storage scope.

- **WWNN Pool**—A World Wide Name (WWN) pool that contains only WW node names for use by the Fibre Channel vHBAs in a Cisco UCS Domain. You can also assign a static WWNN to a Fibre Channel vHBA in a Cisco UCS Domain.
 - **WWPN Pool**—A World Wide Name (WWN) pool that contains only WW port names for use by the Fibre Channel vHBAs in a Cisco UCS Domain. You can also assign a static WWPN to a Fibre Channel vHBA in a Cisco UCS Domain.
1. Log in to Cisco Intersight with your Cisco ID and select admin role.
 2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
 3. Navigate to **Configure > Policies**, and then click **Create Policy**.
 4. Select **SAN Connectivity**, and then click **Start**.
 5. On the **General** page, enter the following information:
 - **Name** of your policy.
 - **Target Platform** for which the policy is applicable. This can be **Standalone** servers or **FI Attached** servers.
A SAN Connectivity Policy created for Standalone servers cannot be deployed on FI Attached servers. Similarly, a SAN Connectivity Policy created for FI Attached servers cannot be deployed on Standalone servers.
 - **Description** to help identify the policy.
 - **Tag** for the policy. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
 6. On the **Policy Details** page, configure the following:
 - Select the placement option—**Manual** or **Auto**
 - **Manual vHBAs Placement**—If you select this option, you must manually specify the PCI slot and PCI order for each vHBA. You can also use the **Graphic vHBAs Editor** to create and specify the placement for each vHBA manually by adding vHBAs and slots, and defining the connection between them.

**Note**

- For manual placement, **PCI Link** is not supported on UCS VIC 1400 Series adapters
 - If a SAN Connectivity Policy has both Simple and Advanced placements, ensure the number provided in PCI Order is appropriate to prevent Server Profile deployment failure.
-
- **Auto vHBAs Placement**—If you select this option, vHBA placement will be done automatically during profile deployment. This option is available only for Cisco Intersight Managed FI Attached servers.
- Create or select a **WWNN Address Pool**, or select **Static** and enter a WWNN address. The Static option is available only for Cisco Intersight Managed FI Attached servers.

7. To set up a vHBA without using a template, click **Add vHBA** and configure the following parameters:

Property	Essential Information
Add vHBA	
Name	Name of the virtual Fibre Channel interface.
vHBA Type	<p>Type of vHBA configuration for SAN Connectivity Policy.</p> <ul style="list-style-type: none"> • fc-initiator—The type of Fibre Channel zoning to be configured for the vHBA is of the initiator type. • fc-target—The type of Fibre Channel zoning to be configured for the vHBA is of the target type. • fc-nvme-initiator—The vHBA type is initiator and applies the NVMe interface to Fibre Channel. • fc-nvme-target—The vHBA type is target and applies the NVMe interface to Fibre Channel. <p>The NVM Express (NVMe) interface allows host software to communicate with a non-volatile memory subsystem. It is optimized for Enterprise non-volatile storage, which is typically attached as a register level interface to the PCI Express (PCIe) interface.</p> <p>Note</p> <ul style="list-style-type: none"> • This configuration is supported only on Cisco VIC 1400 series and higher series of adapters. • 1300 series adapters support only fc-initiator, and fc-nvme-initiator. • Prior to connection, association with adapter should be fine. • After connection with adapter, check vhba_type in the vnic.cfg file. <p>For fc-nvme-initiator type, vhba_type should read the name.</p> <p>For fc-initiator type, vhba_type should not be present.</p>

Property	Essential Information
<p>Pin Group Name</p>	<p>Name of the pin group that contains the specific port/port channels. All traffic from the vHBA is pinned to the specified FC/FCoE uplink ports or port channels.</p> <p>Note The pin group can be defined while creating a Port policy.</p> <p>If you do not assign a pin group to a vHBA, an uplink FC/FCoE uplink port or port channel for traffic is chosen from that server interface dynamically. This choice is not permanent. A different FC/FCoE uplink port or port channel may be used for traffic from that server interface after an interface flap or a server reboot.</p>
<p>WWPN Address Pool</p>	<p>Click Select Pool and choose a WWPN address pool.</p>
<p>Static</p>	<p>Click Static and enter a static WWPN address. This option is available only for Cisco Intersight Managed FI Attached servers.</p>
<p>Placement</p> <p>Placement Settings for the virtual interface.</p>	
<p>Simple</p> <p>When you select Simple Placement, the Slot ID and PCI Link are determined automatically by the system. vHBAs are deployed on the first VIC. The slot ID determines the first VIC. Slot ID numbering begins with MLOM, and thereafter it keeps incrementing by 1, starting from 1. The PCI link is always set to 0.</p>	
<p>Switch ID</p>	<p>Refers to the Fabric Interconnect that carries the vHBA traffic.</p>
<p>PCI Order</p>	<p>The order in which the virtual interface is brought up. The order assigned to an interface should be unique and in sequence starting with "0" for all the Ethernet and Fibre-Channel interfaces on each PCI link on a VIC adapter. The maximum value of PCI order is limited by the number of virtual interfaces (Ethernet and Fibre-Channel) on each PCI link on a VIC adapter.</p> <p>Note You cannot change the PCI order of two vHBAs without deleting and recreating the vHBAs.</p>

Property	Essential Information
Advanced	
Automatic Slot ID Assignment	When enabled, slot ID is determined automatically by the system.
Slot ID	When automatic slot ID assignment is disabled, the slot ID needs to be entered manually. Supported values are (1-15) and MLOM.
<p>PCI link</p> <p>The PCI link used as transport for the virtual interface.</p> <p>PCI Link is only applicable for select Cisco UCS VIC 1300 Series models (UCSC-PCIE-C40Q-03, UCSB-MLOM-40G-03, UCSB-VIC-M83-8P) that support two PCI links. The value, if specified, for any other VIC model will be ignored.</p> <p>Note The host device order can get impacted when using both the PCI links.</p>	
Automatic PCI link Assignment	<p>When enabled, PCI link is determined automatically by the system.</p> <p>Note</p> <ul style="list-style-type: none"> • If Automatic assignment is enabled for both Slot ID and PCI link, then the behavior is same as Simple placement. All the vHBAs are placed on the same PCI link (link 0). • If Automatic Slot ID assignment is disabled but automatic PCI link assignment is enabled, then you need to provide the slot ID and the vHBA will be placed on PCI link 0.

Property	Essential Information
Load Balanced	<p>When Automatic PCI link assignment is disabled and Load Balanced is enabled, the system uniformly distributes the interfaces across the PCI Links.</p> <ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you can specify the PCI order to load balance the vHBAs. • If both automatic PCI link assignment and automatic Slot ID are disabled, you can specify the slot and the PCI order to load balance the vHBAs. <p>Note You cannot change the PCI link mode of two vHBAs from Load Balanced mode to Custom mode without deleting and recreating the vHBAs.</p>
Custom	<ul style="list-style-type: none"> • If automatic PCI link assignment is disabled and automatic Slot ID is enabled, you need to provide the value of the PCI order, PCI link, and Switch ID. • If both automatic PCI link assignment and automatic Slot ID assignment are disabled, you need to provide the values of the Slot ID, PCI order, and the PCI link. <p>Note You cannot change the PCI link mode of two vHBAs from Custom mode to Load Balanced mode without deleting and recreating the vHBAs.</p>
Persistent LUN Bindings	
Enable Persistent LUN Bindings	Enables retention of LUN ID associations in memory until they are manually cleared.
Fibre Channel Network	Select or create a Fibre Channel Network policy.
Fibre Channel QoS	Select or create a Fibre Channel QoS policy.
Fibre Channel Adapter	Select or create a Fibre Channel Adapter policy.
FC Zone	Select or create the FC Zone policy to be attached.

- To derive vHBA for FI-attached servers using a vHBA template, choose **vHBA from Template** from the **Add** drop-down list. For more information on creating vHBA templates, see *Creating vNIC or vHBA Templates*.

**Note**

- When deriving a vHBA from a template, the vHBA configuration is auto-populated from the template configuration. You can edit or delete parameters, which are enabled for configuration override through the vHBA template. For parameters that are not enabled for override, you can only view the configurations using the **Eye** icon.
- The parameters that have been overridden are indicated using an **Overridden** label. In the case of override-enabled parameters, the changes applied in the template are not reflected in the derived vHBA.
- Only those parameters can be modified in the derived vHBA instance which are not included in the template.
- If you attempt to derive a vHBA from a template while profile deployment is in progress, the task will be retried until the profile deployment is completed. You can find these details in the **Requests** tab.

9. Click **Create**.

Creating a Fibre Channel Adapter Policy

A Fibre Channel adapter policy governs the host-side behavior of the adapter, including how the adapter handles traffic. You can enable FCP Error Recovery, change the default settings of Queues, and Interrupt handling for performance enhancement.

**Note**

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.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Fibre Channel Adapter**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.
Fibre Channel Adapter Default Configuration	

Property	Essential Information
Select a default configuration	Click to view and import a default configuration. The policy currently supports nine (9) default configurations.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Error Recovery	
FCP Error Recovery	Enables the use of FCP Sequence Level Error Recovery protocol (FC-TAPE) on the virtual interface.
Port Down Timeout, ms	The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. Enter an integer between 0 and 240000.
I/O Retry Timeout, Seconds	The number of seconds the adapter waits before aborting the pending command and resending the same I/O request. Enter an integer between 1 and 59.
Link Down Timeout, ms	The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost. Enter an integer between 0 and 240000.
Port Down IO Retry, ms	The number of times an IO request to a port is returned because the port is busy before the system decides the port is unavailable. Enter an integer between 0 and 255.
Error Detection	
Error Detection Timeout	Error Detection Timeout, also referred to as EDTOV, is the number of milliseconds to wait before the system assumes that an error has occurred. Enter an integer between 1000 and 10000.
Resource Allocation	

Property	Essential Information
Resource Allocation Timeout	The number of milliseconds to wait before the system assumes that a resource cannot be properly allocated. Enter an integer between 5000 and 100000.
Flogi	
Flogi Retries	The number of times that the system tries to log in to the fabric after the first failure.
Flogi Timeout, ms	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1000 and 255000.
Plogi	
Plogi Retries	The number of times that the system tries to log into a port after the first failure. Enter an integer between 0 and 255.
Plogi Timeout, ms	The number of milliseconds that the system waits before it tries to log in again. Enter an integer between 1000 and 255000
Interrupt	
Mode	Select the preferred driver interrupt mode: <ul style="list-style-type: none"> • MSIx—Message Signaled Interrupts (MSI) with the optional extension. This is the recommended option. • MSI—Message Signaled Interrupts (MSI) only • INTx—PCI INTx interrupts
IO Throttle	
I/O Throttle Count	The number of I/O operations that can be pending in the vHBA at one time. Enter an integer between 1 and 1024.
LUN	

Property	Essential Information
Maximum LUNs Per Target	<p>The maximum number of LUNs that the driver will export. This is usually an operating system platform limitation.</p> <p>Enter an integer between 1 and 1024.</p> <p>For fc-initiator vHBA type, enter an integer between 1 and 4096.</p> <p>Note The fc-initiator vHBA maximum LUN configuration requires the minimum server firmware version 4.2(3d). For more information on the supported firmware for adapters, see Supported Hardware.</p>
LUN Queue Depth	<p>The number of commands that the HBA can send and receive in a single transmission per LUN.</p> <p>Enter an integer between 1 and 254.</p>
Receive	
Receive Ring Size	<p>The number of descriptors in each queue.</p> <p>Enter an integer between 64 and 2048.</p>
Transmit	
Transmit Ring Size	<p>The number of descriptors in each queue.</p> <p>Enter an integer between 64 and 2048.</p>
SCSI I/O	
SCSI I/O Queues	<p>The number of SCSI I/O queue resources the system should allocate.</p> <p>Enter an integer between 1 and 245.</p>
SCSI I/O Ring Size	<p>The number of descriptors in each SCSI I/O queue.</p> <p>Enter an integer between 64 and 512.</p>

7. Click **Create**.

Creating a Fibre Channel Network Policy

A Fibre Channel Network policy governs the Virtual Storage Area Network (VSAN) configuration for the virtual interfaces.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.

2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Fibre Channel Network**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Default VLAN	Default VLAN of the virtual interface in Standalone Rack server. Setting the value to 0 is equivalent to None and will not associate any default VLAN to the traffic on the virtual interface. Valid values are 0 to 4094.
VSAN ID	Default VSAN ID of the virtual interface. Setting the ID to 0 will not associate any default VSAN to the traffic on the virtual interface.

7. Click **Create**.

Creating a Fibre Channel QoS Policy

The Fibre Channel QoS policy assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. For certain adapters, you can also specify additional controls like burst and rate on the outgoing traffic.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Fibre Channel QoS**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.

Property	Essential Information
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Rate Limit, Mbps	Used for limiting the data rate on the virtual interface. The valid range is between 0 and 100000. The default value is Zero.
Maximum Data Field Size, Bytes	The maximum size of the Fibre Channel frame payload bytes that the virtual interface supports. The valid range is between 256 and 2112. The default value is 2112.
Class of Service	The Class of Service to be associated to the traffic on the virtual interface. The valid range is between 0 and 6. The default value is 3. Note <ul style="list-style-type: none"> • FCoE traffic has a reserved QoS system class that should not be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is remarked to 0. • This property is supported only on Standalone servers.
Burst	The burst traffic allowed on the vNIC in bytes. The valid range is between 1024 and 1000000. The default value is 1024. Note This property is supported only on FI-attached servers.
Priority	The priority matching the System QoS defined in the domain profile. The Fibre Channel (FC) is enabled by default. Note This property is supported only on FI-attached servers.

7. Click **Create**.

Create FC Zone Policy

This policy allows you to set up access control between hosts and storage devices.

Certain points to be noted when creating the FC Zone policy:

- Deploying a storage VSAN using a domain profile, for the first time, clears all the unmanaged zones from the Fabric Interconnect.
 - SAN boot targets with a storage VSAN have a zone entry in the Fabric Interconnect.
 - A one-time SAN boot with a storage VSAN has a zone entry in the Fabric Interconnect.
 - Editing the FC Zone policy causes the server profile status to be changed to Pending Changes.
 - When the Fabric Interconnect is rebooted, there is a replay of zones in the configuration.
 - Detection of configuration drift is not supported for FC Zone policy.
1. Log in to Cisco Intersight with your Cisco ID and select admin role.
 2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
 3. Navigate to **Configure > Policies**, and then click **Create Policy**.
 4. Select **FC Zone**, and then click **Start**.
 5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
FC Target Zoning Type	Type of FC Zoning. FC Zoning can be of type: <ul style="list-style-type: none"> • Single Initiator Single Target • Single Initiator Multiple Target • None <p>Note If you select FC Zoning Type as None, you cannot add targets nor view the table of added FC Zone sets.</p>
Add Target	Click to add target details of the FC Zone policy.
Name	Name of the FC Zone policy.
WWPN	WWPN that is a member of the FC Zone.
Switch ID	Unique identifier of the Fabric object. The Switch ID can be A or B.
VSAN ID	Unique identifier of the VSAN on which the FC Zone is to be created. Valid values for the VSAN ID are 1 to 4093. <p>Note The VSAN ID scope should be Storage in the VSAN policy specified for the domain.</p>

7. Click **Create**.

Creating a Firmware Policy

This policy allows you to see the firmware present in your systems, as against the firmware baseline. Firmware policy also enables you to bring the firmware of your systems in line with the desired version and thereby enables the drive to compliance.

1. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

- On the **Policy Details** page, onfigure the following parameters:

Property	Essential Information
Advanced Mode	Enable Advanced Mode to exclude components during firmware upgrade.
Exclude Drives	Enable Advanced Mode and select the Exclude Drives checkbox to exclude drives from the firmware upgrade.
Exclude Storage Controllers	Enable Advanced Mode and select the Exclude Storage Controllers checkbox to exclude storage controllers from the firmware upgrade.
Server Model	Select the server family for the firmware upgrade. Click + to add more server models. Note You can select a maximum of six server models.
Firmware Version	Select the bundle version to which the server is to be upgraded.

- Click **Create**.

Creating a BIOS Policy

A BIOS policy automates the configuration of BIOS settings on servers. You can create one or more BIOS policies that contain a specific grouping of BIOS settings, matching the needs of a server or a set of servers. If you do not specify a BIOS policy for a server, the BIOS settings will default to set of values for a brand new baremetal server or to a set of values previously configured using Cisco IMC. If a BIOS policy is specified, its values replace any previously configured values on the server.

All BIOS tokens are not applicable to all servers. If unsupported tokens are pushed to a server, those tokens are ignored.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **BIOS**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description

Property	Essential Information
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following BIOS policy options:

Property	Essential Information
LOM and PCIe Slots	
ACS Control GPU-<i>n</i> <i>n</i> = 1-8	Access Control Services (ACS) allow the processor to enable or disable peer-to-peer communication between multiple devices for GPUs.
ACS Control Slot <i>n</i> <i>n</i> = 11-14	Access Control Services (ACS) allow the processor to enable or disable peer-to-peer communication between multiple devices for Control Slot <i>n</i> .
CDN Support for LOM	Whether the Ethernet Networking Identifier naming convention is according to Consistent Device Naming (CDN) or the traditional way of naming conventions.
LOM Port <i>n</i> OptionROM <i>n</i> = 0-3	Whether Option ROM is available on the LOM port <i>n</i>
All Onboard LOM Ports	Whether all onboard LOM ports are enabled or disabled
All PCIe Slots OptionROM	Whether Option ROM is available on all PCIe slots
PCI ROM CLP	PCI ROM Command Line Protocol (CLP) controls the execution of different Option ROMs such as PxE and iSCSI that are present in the card.
PCIe Slot:<i>n</i> Link Speed <i>n</i> = 1-12	This option allows you to restrict the maximum speed of an adapter card installed in PCIe slot <i>n</i> .
Slot <i>n</i> state <i>n</i> = 1-12	The state of the adapter card installed in PCIe slot <i>n</i> .
PCIe Slot:FLOM Link Speed	This option allows you to restrict the maximum speed of an adapter card installed in PCIe FLOM slot.
PCIe Slot:Front Nvmen Link Speed <i>n</i> = 1-2	This option allows you to restrict the maximum speed of an NVMe card installed in the front PCIe slot <i>n</i> .

Property	Essential Information
PCIe Slot:Frontn Link Speed $n= 1-2$	This option allows you to restrict the maximum speed of an adapter card installed in the front PCIe slot n .
GPUn OptionROM $n= 1-8$	Whether the Option ROM is enabled on GPU slot n .
PCIe Slot:HBA Link Speed	This option allows you to restrict the maximum speed of an adapter card installed in PCIe HBA slot.
PCIe Slot:HBA OptionROM	Whether the Option ROM is enabled on the HBA slot.
PCIe LOM:n Link $n= 1-2$	Whether Option ROM is available on the LOM port.
Slot Mezz state	State of the Mezzanine card slot.
PCIe Slot:MLOM Link Speed	This option allows you to restrict the maximum speed of an MLOM adapter card installed in a PCIe slot.
PCIe Slot MLOM OptionROM	Whether the Option ROM is enabled on the MLOM slot.
MRAID Link Speed	This option allows you to restrict the maximum speed of MRAID.
PCIe Slot MRAID OptionROM	Whether Option ROM is available on the MRAID port.
PCIe Slot Nn OptionROM $n= 1-24$	Whether the Option ROM is enabled on the PCIe slot.
RAID Link Speed	This option allows you to restrict the maximum speed of MRAID.
PCIe Slot RAID OptionROM	Whether the Option ROM is enabled on the RAID slot.
PCIe Slot:Rear Nvmen Link Speed $n= 1-2$	This option allows you to restrict the maximum speed of an NVMe card installed in the rear PCIe slot n .
PCIe Slot:Rear NVME n OptionRom $n= 1-8$	Whether the Option ROM is enabled on the rear NVMe slot n .
PCIe Slot:Risern Link Speed $n= 1-2$	This option allows you to restrict the maximum speed of Riser card n installed in the PCIe slot.

Property	Essential Information
PCIe Slot:Riser1 Slotn Link Speed $n= 1-3$	This option allows you to restrict the maximum speed of slot n on Riser card1 installed in the PCIe slot.
PCIe Slot:Riser2 Slotn Link Speed $n= 4-6$	This option allows you to restrict the maximum speed of slot n on Riser card2 installed in the PCIe slot.
PCIe Slot:SAS OptionROM	Whether the Option ROM is enabled on the SAS slot.
PCIe Slot:FrontPcin Link Speed $n= 1-2$	This option allows you to restrict the maximum speed of the front PCIe n .
Processor	
X2APIC Opt-Out Flag	Prevents the OS from enabling extended xAPIC (x2APIC) mode when the OS is not working with x2APIC.
Adjacent Cache Line Prefetcher	Whether the processor fetches cache lines in even/odd pairs instead of fetching just the required line.
Altitude	The approximate number of meters above sea level at which the physical server is installed.
Autonomous Core C-state	When the Operating System requests CPU core C1 state, system hardware automatically changes the request to core C6 state.
CPU Autonomous Cstate	Enables CPU Autonomous C-State, which converts the HALT instructions to the MWAIT instructions.
Boot Performance Mode	Allows the user to select the BIOS performance state that is set before the operating system handoff.
Downcore control	Allows AMD processors to disable cores and, thus, select how many cores to enable.
Channel Interleaving	Whether the CPU divides memory blocks and spreads contiguous portions of data across interleaved channels to enable simultaneous read operations.
Closed Loop Therm Throt	Allows for the support of Closed-Loop Thermal Throttling, which improves reliability and reduces CPU power consumption through the automatic voltage control while the CPUs are in the idle state.
Processor CMCI	Enables CMCI generation.

Property	Essential Information
Config TDP	Allows you to configure the Thermal Design Power (TDP) settings for the system. TDP is the maximum amount of power allowed for running applications without triggering an overheating event.
Core MultiProcessing	Sets the state of logical processor cores per CPU in a package. If you disable this setting, Intel Hyper Threading technology is also disabled.
Energy Performance	Allows you to determine whether system performance or energy efficiency is more important on this server.
Frequency Floor Override	Whether the CPU is allowed to drop below the maximum non-turbo frequency when idle.
CPU Performance	Sets the CPU performance profile for the server.
Power Technology	Enables you to configure the CPU power management settings.
Demand Scrub	Whether the system corrects single bit memory errors encountered when the CPU or I/O makes a demand read.
Direct Cache Access Support	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.
DRAM Clock Throttling	Allows you to tune the system settings between the memory bandwidth and power consumption.
Energy Efficient Turbo	Allows the processor to switch to a minimum performance state when it is idle.
Energy Performance Tuning	Determines if the BIOS or Operating System can turn on the energy performance bias tuning.
Enhanced Intel Speedstep(R) Technology	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.
EPP Profile	Determines the processor Enhanced Performance Profile.
Local X2 Apic	Allows you to set the type of Application Policy Infrastructure Controller (APIC) architecture.

Property	Essential Information
Hardware Prefetcher	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.
CPU Hardware Power Management	Enables processor Hardware Power Management (HWPM).
IMC Interleaving	This BIOS option controls the interleaving between the Integrated Memory Controllers (IMCs).
Intel HyperThreading Tech	Whether the processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor.
Intel Speed Select	Allows improved CPU performance by using Intel Speed Select technology to tune the CPU to run at one of three operating profiles, based on number of logical processor cores, frequency, and TDP thread setting, to improve performance over the basic Platform Default setting. These profiles correspond to High, Medium, and Low Core settings
Intel Turbo Boost Tech	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.
Intel(R) VT	Whether the processor uses Intel Virtualization Technology, which allows a platform to run multiple operating systems and applications in independent partitions.
IIO Error Enable	Allows you to generate the IIO-related errors.
DCU IP Prefetcher	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.
KTI Prefetch	KTI prefetch is a mechanism to get the memory read started early on a DDR bus.
LLC Prefetch	Whether the processor uses the LLC Prefetch mechanism to fetch the data into the LLC.
Memory Interleaving	Whether the CPU interleaves the physical memory so that the memory can be accessed while another is being refreshed.

Property	Essential Information
Package C State Limit	The amount of power available to the server components when they are idle.
Patrol Scrub	Whether the system actively searches for, and corrects, single bit memory errors even in unused portions of the memory on the server.
Patrol Scrub Interval	Controls the time interval between each patrol scrub memory access. A lower interval scrubs the memory more often but requires more memory bandwidth. Select a value between 5 and 23. The default value is 8. This option is used only if Patrol Scrub is enabled.
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.
Processor C3 Report	Whether the BIOS sends the C3 reports to the operating system. When the OS receives the report, it can transition the processor into the lower C3 power state to decrease energy usage while maintaining optimal processor performance.
Processor C6 Report	Whether the BIOS sends the C6 reports to the operating system. When the OS receives the report, it can transition the processor into the lower C6 power state to decrease energy usage while maintaining optimal processor performance.
CPU C State	Whether the system can enter a power savings mode during idle periods.
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.
Power Performance Tuning	Determines if the BIOS or Operating System can turn on the energy performance bias tuning.
Rank Interleaving	Whether the CPU interleaves physical ranks of memory so that one rank can be accessed while another is being refreshed.
Single PCTL	Facilitates single PCTL support for better processor power management.

Property	Essential Information
SMT Mode	Whether the processor uses AMD Simultaneous MultiThreading Technology, which allows multithreaded software applications to execute threads in parallel within each processor.
Sub Numa Clustering	Whether the CPU supports sub NUMA clustering, in which the tag directory and the memory channel are always in the same region.
DCU Streamer Prefetch	Whether the processor uses the DCU Streamer Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache.
SVM Mode	Whether the processor uses AMD Secure Virtual Machine Technology.
Workload Configuration	This feature allows for workload optimization.
XPT Prefetch	Whether XPT prefetch is used to enable a read request sent to the last level cache to issue a copy of that request to the memory controller prefetcher.
USB	
All USB Devices	Whether all physical and virtual USB devices are enabled or disabled.
Legacy USB Support	Whether the system supports legacy USB devices.
Make Device Non Bootable	Whether the server can boot from a USB device.
xHCI Mode	Whether xHCI mode is enabled or disabled.
Port 60/64 Emulation	Whether the system supports 60h/64h emulation for complete USB keyboard legacy support.
USB Port Front	Whether the front panel USB devices are enabled or disabled.
USB Port Internal	Whether the internal USB devices are enabled or disabled.
USB Port KVM	Whether the KVM ports are enabled or disabled.
USB Port Rear	Whether the rear panel USB devices are enabled or disabled.
USB Port SD Card	Whether the SD card drives are enabled or disabled.
USB Port VMedia	Whether the virtual media devices are enabled or disabled.

Property	Essential Information
XHCI Legacy Support	Whether the legacy xHCI mode is enabled or disabled.
Property	
ASPM Support	Allows you to set the level of ASPM (Active Power State Management) support in the BIOS.
IOH Resource Allocation	Enables you to distribute 64KB of 16-bit IO resources between IOH0 and IOH1 as per system requirement.
Memory mapped IO above 4GB	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.
MMCFG BASE	Sets the low base address for PCIe adapters within 4GB.
Onboard 10Gbit LOM	Whether 10Gbit LOM is enabled or disabled on the server.
Onboard Gbit LOM	Whether Gbit LOM is enabled or disabled on the server.
NVMe SSD Hot-Plug Support	Allows you to replace an NVMe SSD without powering down the server.
SR-IOV Support	Whether SR-IOV (Single Root I/O Virtualization) is enabled or disabled on the server.
VGA Priority	Allows you to set the priority for VGA graphics devices if multiple VGA devices are found in the system.
Server Management	
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.
Assert NMI on SERR	Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs.
Baud rate	What Baud rate is used for the serial port transmission speed. If you disable Console Redirection, this option is not available.

Property	Essential Information
Consistent Device Naming	Whether the Ethernet Network naming convention is according to Consistent Device Naming (CDN) or the traditional way of naming conventions.
Adaptive Memory Training	The BIOS saves the memory training results (optimized timing/voltage values) along with CPU/memory configuration information and reuses them on subsequent reboots to save boot time. The saved memory training results are used only if the reboot happens within 24 hours of the last save operation.
BIOS Techlog Level	The BIOS Tech log output to be controlled at more a granular level. This reduces the number of BIOS Tech log messages that are redundant, or of little use.
OptionROM Launch Optimization	The Option ROM launch is controlled at the PCI Slot level, and is enabled by default. In configurations that consist of a large number of network controllers and storage HBAs having Option ROMs, all the Option ROMs may get launched if the PCI Slot Option ROM Control is enabled for all. However, only a subset of controllers may be used in the boot process. When this token is enabled, Option ROMs are launched only for those controllers that are present in boot policy.
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.
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.
FRB-2 Timer	Whether the FRB-2 timer is used to recover the system if it hangs during POST.
Legacy OS redirection	Whether redirection from a legacy operating system, such as DOS, is enabled on the serial port.

Property	Essential Information
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. Note The OS Boot Watchdog Timer value must not exceed 5 minutes.
OS Boot Watchdog Timer Policy	What action the system takes if the watchdog timer expires.
OS Boot Watchdog Timer Timeout	What timeout value the BIOS uses to configure the watchdog timer.
Out-of-Band Mgmt Port	Used for Windows Special Administration Control (SAC). This option allows you to configure the COM port 0 that can be used for Windows Emergency Management services. ACPI SPCR table is reported based on this setup option.
Putty KeyPad	Allows you to change the action of the PuTTY function keys and the top row of the numeric keypad.
Redirection After BIOS POST	Whether BIOS console redirection should be active after BIOS POST is complete and control given to the OS bootloader.
Terminal Type	What type of character formatting is used for console redirection.
Boot Order Rules	How the server changes the boot order list defined when there are no devices of a particular device type available or when the user defines a different boot order using the server's BIOS Setup Utility.
Memory	
BME DMA Mitigation	Allows you to disable the PCI BME bit to mitigate the threat from an unauthorized external DMA.
IOMMU	Input Output Memory Management Unit (IOMMU) allows AMD processors to map virtual addresses to physical addresses.
Bank Group Swap	Determines how physical addresses are assigned to applications.
Chipselect Interleaving	Whether memory blocks across the DRAM chip selects for node 0 are interleaved.

Property	Essential Information
Memory interleaving	Whether the CPU interleaves the physical memory so that the memory can be accessed while another is being refreshed. This controls fabric level memory interleaving. Channel, die and socket have requirements based on memory populations and will be ignored if the memory does not support the selected option.
Memory interleaving size	Determines the size of the memory blocks to be interleaved. It also determines the starting address of the interleave (bit 8,9,10 or 11).
DCPMM Firmware Downgrade	Whether DCPMM firmware downgrade is enabled.
SMEE	Whether the processor uses the Secure Memory Encryption Enable (SMEE) function, which provides memory encryption support.
Boot Options	
Number of Retries	Number of attempts to boot.
Cool Down Time (sec)	The time to wait (in seconds) before the next boot attempt.
Boot option retry	Whether the BIOS retries NON-EFI based boot options without waiting for user input.
IPV6 PXE Support	Enables or disables IPV6 support for PXE.
Onboard SCU Storage Support	Whether the onboard software RAID controller is available to the server.
Onboard SCU Storage SW Stack	Whether the onboard software stack is available to the server.
Power ON Password	This token requires that you set a BIOS password before using the F2 BIOS configuration. If enabled, password needs to be validated before you access BIOS functions such as IO configuration, BIOS set up, and booting to an operating system using BIOS.
P-SATA mode	This options allows you to select the P-SATA mode.
SATA mode	This options allows you to select the SATA mode.
VMD Enablement	Whether NVMe SSDs that are connected to the PCIe bus can be hot swapped. It also standardizes the LED status light on these drives. LED status lights can be optionally programmed to display specific Failure indicator patterns.

Property	Essential Information
Power and Performance	
Core Performance Boost	Whether the AMD processor increases its frequency on some cores when it is idle or not being used much.
Global C-state Control	Whether the AMD processors control IO-based C-state generation and DF C-states
L1 Stream HW Prefetcher	Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L1 cache when necessary.
L2 Stream HW Prefetcher	Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L2 cache when necessary.
Determinism Slider	Allows AMD processors to determine how to operate - Performance or Power.
cTDP Control	Allows you to set customized value for Thermal Design Power (TDP).
RAS Memory	
CKE Low Policy	Controls the DIMM power savings mode policy.
DRAM Refresh Rate	The refresh interval rate for internal memory.
Low Voltage DDR Mode	Whether the system prioritizes low voltage or high frequency memory operations.
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.
NUMA optimized	Whether the BIOS supports NUMA.
Select Memory RAS configuration	How the memory reliability, availability, and serviceability (RAS) is configured for the server.

Property	Essential Information
Sparing Mode	<p>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.</p> <p>This option is only available if you choose sparing option for Memory RAS Config.</p>
Intel Directed IO	
Intel VT for directed IO	Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d).
Intel(R) VT-d Coherency Support	Whether the processor supports Intel VT-d Coherency.
Intel(R) VT-d Interrupt Remapping	Whether the processor supports Intel VT-d Interrupt Remapping.
Intel(R) VT-d PassThrough DMA support	Whether the processor supports Intel VT-d Pass-through DMA.
Intel VTD ATS support	Whether the processor supports Intel VT-d Address Translation Services (ATS).
Main	
POST Error Pause	What happens when the server encounters a critical error during POST.
QPI	
QPI Link Frequency Select	The Intel QuickPath Interconnect (QPI) link frequency, in megatransfers per second (MT/s).
QPI Snoop Mode	The Intel QuickPath Interconnect (QPI) snoop mode.
Serial Port	
Serial A Enable	Whether serial port A is enabled or disabled.
Trusted Platform	
Trusted Platform Module State	Determines whether the TPM has been initialized and attached to the Operating System.
Intel Trusted Execution Technology Support	Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business server. This option allows you to control the TXT support for the system.

Property	Essential Information
DMA Control Opt-In Flag	Enabling this token enables Windows 2022 Kernel DMA Protection feature. The OS treats this as a hint that the IOMMU should be enabled to prevent DMA attacks from possible malicious devices.
Security Device Support	Enables or disables BIOS support for the security device.

7. Click **Create**.

Creating a Boot Order Policy

The Boot Order policy configures the linear ordering of devices and enables you to change the boot order and boot mode. You can also add multiple devices under various device types, rearrange the boot order, and set parameters for each boot device type.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Boot Order**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
BootMode	<p>The type of boot mode that is enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • Legacy—Uses the Master Boot Record (MBR) partitioning scheme. Select if the system is not UEFI-enabled. • UEFI—Uses the GUID Partition Table (GPT). Select Unified Extensible Firmware Interface (UEFI) if the system is UEFI-enabled. <p>Note The Legacy boot mode is currently not supported on Cisco UCS C225 M6, C245 M6, C220 M7, and C240 M7 servers.</p>
Enable Secure Boot Mode	<p>This option is available only when UEFI Boot Mode is enabled.</p> <p>Secure boot mode enforces that a device boots using the software that is trusted by the Original Equipment Manufacturer (OEM).</p>

Property	Essential Information
Add Boot Device	

Property	Essential Information
	<p>Select to add and configure a boot device. The configuration options vary with boot device types. The supported boot devices and its configuration options for UCS standalone and FI-attached servers are listed below:</p> <ul style="list-style-type: none"> • FlexMMC Boot <ul style="list-style-type: none"> Note <ul style="list-style-type: none"> • FlexMMC boot is supported only with UEFI Boot Mode for C-series standalone servers. • Secure Boot option is supported for FlexMMC. <p>For more information on the firmware requirements for FlexMMC Boot, see Firmware Requirements for FlexMMC Boot Option.</p> <p>Configuration options:</p> <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Sub-Type—The sub-type for the selected device <ul style="list-style-type: none"> • None • FlexMMC Mapped DVD • FlexMMC Mapped HDD <ul style="list-style-type: none"> • HTTP Boot <ul style="list-style-type: none"> Note HTTP/HTTPS boot is supported only with UEFI Boot Mode for both IMM servers and C-series standalone servers. For more information on the firmware requirements for HTTP Boot, see Firmware Requirements for HTTP Boot Option. <p>Configuration options:</p> <ul style="list-style-type: none"> • Device Name—Name of the boot device. • IP Type—The IP address family type to use during the HTTP boot process. • IP Config Type—The IP config type to

Property	Essential Information
	use during the HTTP Boot process.

Property	Essential Information
	<ul style="list-style-type: none"> • DHCP <ul style="list-style-type: none"> • [Optional] URI—The boot resource location in URI format. <p>Note If you do not enter a URI, ensure that DHCP is configured with client extensions.</p> • Interface Name (Only for UCS Server (FI-Attached))—The name of the underlying vNIC that will be used by the HTTP boot device. You can select a vNIC that was configured using the LAN Connectivity Policy. For more information, see the LAN Connectivity Policy section. • Static <p><i>When IP Config Type is Static and IP Type is IPv4:</i></p> <ul style="list-style-type: none"> • DNS IP—The IP address of DNS server. • Gateway IP—The IP address of default gateway. • Static IP—IPv4 or IPv6 static Internet Protocol address. • Network Mask—Network mask of the IPv4 address. • URI—The boot resource location in URI format. • Interface Name—The name of the underlying vNIC that will be used by the HTTP boot device. You can select a vNIC that was configured using the LAN Connectivity Policy. <p><i>When IP Config Type is Static and IP Type is IPv6:</i></p> <ul style="list-style-type: none"> • DNS IP—The IP address of

Property	Essential Information
	<p>DNS server.</p> <ul style="list-style-type: none"> • Gateway IP—The IP address of default gateway. • Static IP—IPv4 or IPv6 static Internet Protocol address. • Prefix Length—A prefix length which masks the IP address and divides the IP address into network address and host address. • URI—The boot resource location in URI format. • Interface Name—The name of the underlying vNIC that will be used by the HTTP boot device. You can select a vNIC that was configured using the LAN Connectivity Policy. <ul style="list-style-type: none"> • Protocol—The protocol used for HTTP Boot. <p>To use the HTTPS protocol, you must have a valid Root CA Certificate for authentication. You can deploy Root CA certificates using the Certificate Management Policy. For more information, see the <i>Creating a Certificate Management Policy</i> section.</p> <p>Note Certificate Management Policy does not support addition, deletion, and modification of a single certificate. Even if one of the certificates is added, deleted or modified in policy, the Server Profile will need to be redeployed or Server Action must be performed, for certificate changes to take effect.</p>

Property	Essential Information
	<ul style="list-style-type: none"> • Interface Source (Only for C-series standalone servers)—Lists the supported Interface Source for HTTP device. • Interface Name (Only for VIC Adapters) <ul style="list-style-type: none"> • Slot—The slot ID of the adapter on which the underlying virtual ethernet interface is present. • Interface Name—The name of the underlying virtual ethernet interface used by the HTTP boot device. • Port (Only for VIC Adapters) <ul style="list-style-type: none"> • Slot—The slot ID of the adapter on which the underlying virtual ethernet interface is present. • Port—The Port ID of the adapter on which the underlying virtual ethernet interface is present. If no port is specified, the default value is -1. Supported values are 0 to 255. • MAC Address <ul style="list-style-type: none"> • Slot—The slot ID of the adapter on which the underlying virtual ethernet interface is present. • MAC—The MAC address of the underlying virtual ethernet interface used by the HTTP boot device. • iSCSI Boot <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Slot—The slot id of the boot device. • Port—The port id of the boot device.

Property	Essential Information
	<ul style="list-style-type: none"> • Local CDD <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Local Disk <ul style="list-style-type: none"> Note This device allows the host to use the virtual drive as a bootable device. • Device Name—Name of the boot device. • Slot—The slot id of the boot device. • NVMe <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Bootloader Name—Name of the bootloader image. • Bootloader Description—Description of the bootloader. • Bootloader Path—Path to the boatloader image. Note The NVMe device can be configured only on UEFI mode. • PCH Storage <ul style="list-style-type: none"> • Device Name—Name of the boot device. • LUN—The Logical Unit Number (LUN) of the boot device (0-255). Note Only UEFI boot mode is supported with software RAID configuration. • PXE Boot <ul style="list-style-type: none"> • Device Name—Name of the boot device. • IP Type—The IP address family type to use during the PXE boot process. • Slot—The slot ID of the adapter on which the virtual ethernet interface is present. • Interface Name/Port/ MAC Address—The name or address of the underlying virtual ethernet interface used by the PXE boot device.

Property	Essential Information
	<ul style="list-style-type: none"> • SAN Boot <ul style="list-style-type: none"> • Device Name—Name of the boot device. • LUN—The Logical Unit Number (LUN) of the boot device (0-255). • Slot—The slot id of the boot device. This field is applicable only for Standalone servers. • Interface Name—The name of the underlying vHBA interface. • Target WWPN—The WWPN Address of the underlying fibre channel interface • Bootloader Name — The name of the bootloader image. This field is available only in UEFI Mode. • Bootloader Description— The details of the bootloader image. This field is available only in UEFI Mode. • Bootloader Path— The path of the bootloader image. This field is available only in UEFI Mode. • SD Card <ul style="list-style-type: none"> • Device Name—Name of the boot device. • LUN—The Logical Unit Number (LUN) of the boot device (0-255). • Sub-Type— The sub-type for the selected device: <ul style="list-style-type: none"> • FlexUtil • FlexFlash • SDCard • UEFI Shell <ul style="list-style-type: none"> • Device Name—Name of the boot device.

Property	Essential Information
	<ul style="list-style-type: none"> • USB <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Sub-Type— The sub-type for the selected device: <ul style="list-style-type: none"> • CD • FDD • HDD • Virtual Media <ul style="list-style-type: none"> • Device Name—Name of the boot device. • Sub-Type— The sub-type for the selected device: <ul style="list-style-type: none"> • None <ul style="list-style-type: none"> Note This option is not supported on UCS FI-attached servers. • CIMC Mapped DVD • CIMC Mapped HDD • KVM Mapped DVD • KVM Mapped HDD • KVM Mapped FDD <p>Note The device name of the boot devices can be any string that adheres to the following constraints. It should start and end with an alphanumeric character. It can have underscores and hyphens. It cannot be more than 30 characters.</p>

7. Click **Create**.

Configuring an iSCSI Boot Policy

iSCSI boot support allows you to initialize the Operating System on FI-attached blade and rack servers from a remote disk across a Storage Area Network. The remote disk, known as the target, is accessed using TCP/IP and iSCSI boot firmware.

Prerequisites

The following are required to configure the iSCSI boot device:

- **iSCSI Static Target Policy**—When you select **Static** as the mode for configuring the iSCSI boot policy, you can use the iSCSI Static Target policy to specify the primary target details. You can also specify the details of a secondary target, if required.
 - **iSCSI Adapter Policy**—Using this policy you can specify the TCP and DHCP Connection Timeout and the retry count when the logical unit number of the boot device is busy.
 - **Creating an IQN Pool**—Using this policy you can specify the TCP and DHCP Connection Timeout and the retry count when the logical unit number of the boot device is busy.
1. Log in to Cisco Intersight with your Cisco ID and select admin role.
 2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
 3. Navigate to **Configure > Policies**, and then click **Create Policy**.
 4. Select **iSCSI Boot**, and then click **Start**.
 5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Target Interface	Target interface can be Auto or Static.
DHCP Vendor ID/IQN	If you select Auto for the target interface, specify the Initiator name, or the DHCP vendor ID. The vendor ID can be up to 32 alphanumeric characters.
Static	If the target interface is Static specify the following parameters.
Primary Target	Select the Primary Target policy. iSCSI target is the remote disk in the storage area network from which the operating system is initialized. This policy specifies the Target Name, the IP Address of the target, the Port, and the LUN ID.

Property	Essential Information
Secondary Target	Select the Secondary Target policy. Secondary Target is optional
Adapter Policy	Select the Adapter Policy for the iSCSI boot device. The Adapter Policy specifies the TCP and DHCP Timeouts, and the Retry Count if the LUN ID is busy.
<p>Authentication</p> <p>You can select CHAP or Mutual CHAP as the authentication method and specify the parameters. If you have selected CHAP, specify the CHAP authentication parameters for iSCSI Target. Mutual CHAP is a two-way DHCP mechanism and is more secure.</p>	
CHAP	<p>For CHAP authentication, enter:</p> <ul style="list-style-type: none"> • Username: The user Id of the Initiator/Target Interface. Enter between 1 and 128 characters, spaces, or special characters. • Password: Password of Initiator or Target Interface. Enter between 12 and 16 characters, including special characters except spaces, tabs, line breaks. • Password Confirmation: Re-enter the password that you entered. Both the password and password confirmation have to match.
Mutual CHAP	<p>Mutual CHAP is a two-way CHAP mechanism. For Mutual CHAP authentication, enter:</p> <ul style="list-style-type: none"> • Username: The user Id of the Initiator or Target Interface. Enter between 1 and 128 characters, spaces, or special characters. • Password: Password of Initiator or Target Interface. Enter between 12 and 16 characters, including special characters except spaces, tabs, line breaks. • Password Confirmation: Re-enter the password that you entered. Both the password and password confirmation have to match.

Property	Essential Information
Initiator IP Source	<p>Select the method that determines the Initiator IP Source. The methods to determine the Initiator IP Source are:</p> <ul style="list-style-type: none"> • Pool: You can select an IP pool • Auto: The IP is automatically determined • Static: You can specify a static IP address as the Initiator IP. Select Static and specify: <ul style="list-style-type: none"> • IP Address: Enter the Static IP address provided for iSCSI Initiator. • Subnet Mask: Enter the 32-bit number that masks an IP address and divides the IP address into network address and host address.. • Default Gateway: Enter the IP address of the default IPv4 gateway. • Primary DNS: Enter the IP address of the primary Domain Name System server. • Secondary DNS: Enter the IP address of the secondary Domain Name System server.

7. Click **Create**.

Creating an iSCSI Adapter Policy

The iSCSI Adapter policy allows you to configure values for TCP Connection Timeout, DHCP Timeout, and the Retry Count if the specified LUN ID is busy.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **iSCSI Adapter**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.

Property	Essential Information
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
TCP Connection Timeout	Enter the number of seconds after which the TCP connection times out.
DHCP Timeout	Enter the number of seconds after which the DHCP times out.
LUN Busy Retry Count	Enter the number of times connection is to be attempted when the LUN ID is busy.

7. Click **Create**.

Creating an iSCSI Static Target Policy

The iSCSI Static Target policy allows you to specify the name, IP address, port, and logical unit number of the primary target for iSCSI boot. You can optionally specify these details for a secondary target as well.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **iSCSI Static Target**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Target Name	Enter the name of the target.

Property	Essential Information
IP Address	Enter the target IP address.
Port	Enter the port number of the target.
LUN ID	Enter the ID of the boot logical unit number.

- Click **Create**.

Creating a Device Connector Policy

Device Connector Policy lets you choose the **Configuration from Intersight only** option to control configuration changes allowed from Cisco IMC. The **Configuration from Intersight only** option is enabled by default. You will observe the following changes when you deploy the Device Connector policy in Intersight:

- Validation tasks will fail:
 - If Intersight Read-only mode is enabled in the claimed device.
 - If the firmware version of the Cisco UCS Standalone C-Series Servers is lower than 4.0(1).
 - If Intersight Read-only mode is enabled, firmware upgrades will be successful only when performed from Intersight. Firmware upgrade performed locally from Cisco IMC will fail.
 - IPMI over LAN privileges will be reset to read-only level if Configuration from Intersight only is enabled through the Device Connector policy, or if the same configuration is enabled in the Device Connector in Cisco IMC.
- Log in to Cisco Intersight with your Cisco ID and select admin role.
 - From the **Service Selector** drop-down list, select **Infrastructure Service**.
 - Navigate to **Configure > Policies**, and then click **Create Policy**.
 - Select **Device Connector**, and then click **Start**.
 - On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- In the **Policy Details** page, enable or disable **Configuration from Intersight only**. This option is enabled by default.
- Click **Create**.

Creating a Drive Security Policy

In Intersight Managed Mode, the Drive Security Policy allows you to specify the KMIP server details and attach the policy to the server profile.

1. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

2. On the **Policy Details** page:
 - a. Use the toggle button to enable the primary KMIP server.
 - b. Configure the following parameters:

Property	Essential Information
Hostname/IP Address	Enter the IP address of the KMIP server that you want to use.
Port	Enter the port number for the KMIP server. The default port is 5696.
Timeout	Enter the time that will be allowed to elapse within which the KMIP client should connect. The recommended timeout interval is up to 65 seconds.

- c. [Optional] To configure a fallback KMIP server, add the details of an additional KMIP server under the **Secondary KMIP Server**.
- d. In the **Server Public Root CA Certificate** field, copy-paste the root certificate from the KMIP server.
- e. [Optional] If your KMIP server supports authentication, click the **Enable Authentication** option for additional security and enter your username and password.



Note You can use authentication only if the KMIP server supports it.

3. Click **Create**.

The newly created policy is displayed in the table view on the **Policy Details** page.

Creating a Disk Group Policy

The Disk Group policy defines how a disk group (a group of physical disks that are used for creating virtual drives) is created and configured, and specifies the RAID level to be used for the disk group. With this policy, you can select the physical disks that have to be part of a disk group. When a Disk Group policy is associated with multiple virtual drives in a Storage policy, the virtual drives share the same disk group space.



Note This policy is not applicable for virtual drives for a Cisco Boot Optimised M.2 RAID Controller.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Disk Group**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Virtual Drive Configuration	

Property	Essential Information
<p>RAID Level</p>	<p>Set the Redundant Array of Inexpensive Disks (RAID) level to ensure availability and redundancy of data, and I/O performance.</p> <p>The supported RAID levels for the disk group are:</p> <ul style="list-style-type: none"> • RAID0—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. • RAID1—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. • RAID5—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. • RAID6—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored. • RAID10—This RAID uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10. • RAID50—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID60—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.
<p>Local Disk Configuration - Disk Group (Span 0)</p>	
<p>Drive Number</p>	<p>Specify the drive number for the disk group associated with the RAID controller.</p>
<p>Dedicated Hot Spares</p>	

Property	Essential Information
Dedicated Hot Spares	Select Enable to use a hot spare drive in the case of disk failure in the disk group.
Drive Number	Specify the identified drive number to act as a dedicated hot spare for the disk group.
Set Disks in JBOD state to Unconfigured good	Select to allow users to convert any disks in JBOD to be un-configured good disks so that they can be used in the RAID group.



Attention All virtual drives in a disk group should be managed by using the same disk group policy.

7. Click **Create**.

Creating an IMC Access Policy

The IMC Access policy allows you to configure your network and associate an IP address from an IP Pool with a server. In-Band IP address, Out-Of-Band IP address, or both In-Band and Out-Of-Band IP addresses can be configured using IMC Access Policy and is supported on Drive Security, SNMP, Syslog, and vMedia policies.



Note The Out-of-Band IP address support for SNMP policy is available only for the Fabric Interconnects running on Infrastructure Firmware 4.3(2.230129) or later versions.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **IMC Access**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information	
In-Band Configuration	Enable, to have the server management services made available using the uplink port.	
	VLAN ID	Enter the VLAN ID to be used for server access over the inband network. The field value can be between 4 and 4093.
	IPv4 address configuration	Select to determine the type of network for this policy. Note You can select only IPv4 address configuration or both IPv4 and IPv6 configurations.
	IPv6 address Configuration	Select to determine the type of network for this policy. Note You can select only IPv6 address configuration or both IPv4 and IPv6 configurations.
	IP Pool	
Select IP Pool	Click to view the list of IP Pools available and select an IP pool for In-Band configuration. Note Ensure that the default gateway specified in the IP Pool used for IMC Access Policy has connectivity to Cisco IMC. For more information, see the <i>Creating an IP Pool</i> section.	

Property	Essential Information
Out-Of-Band Configuration	Enable, to have the server management services made available using the management port.
	IP Pool
	Select IP Pool Click to view the list of IP Pools available and select an IP pool for the Out-Of-Band configuration. Note Only IPv4 addresses are supported for Out-Of-Band configuration.

Creating an IPMI Over LAN Policy

The IPMI over LAN policy defines the protocols for interfacing with a service processor that is embedded in a server platform. The Intelligent Platform Management Interface (IPMI) enables an operating system to obtain information about the system health and control system hardware and directs the Cisco IMC to perform the required actions. You can create an IPMI Over LAN policy to manage the IPMI messages through Cisco Intersight.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **IPMI Over LAN**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable IPMI Over LAN	The state of the IPMI Over LAN service on the endpoint.

Property	Essential Information
<p>Privilege Level</p>	<p>You can assign these privileges to the IPMI sessions on the server:</p> <ul style="list-style-type: none"> • admin—You can create admin, user, and read-only sessions on servers with the "Administrator" user role. • read-only—You can only create read-only IPMI sessions on servers with the "Read-only" user role. • user—You can create user and read-only sessions, but not admin sessions on servers with the "User" role. <p>Note</p> <ul style="list-style-type: none"> • This configuration is supported only on Cisco UCS C-Series Standalone and C-Series Intersight Managed Mode Servers. • The value of the Privilege field must match exactly the role assigned to the user attempting to log in. For example, if this field is set to read-only and a user with the admin role attempts to log in through IPMI, that login attempt will fail.
<p>Encryption Key</p>	<p>The encryption key to use for IPMI Communication. The key must have an even number of hexadecimal characters and not exceeding 40 characters. You can use "00" to disable the encryption key use. If the encryption key specified is less than 40 characters, then the IPMI commands must add zeroes to the encryption key to achieve a length of 40 characters.</p> <p>Note</p> <p>This encryption key configuration is supported only on Cisco UCS C-Series Standalone and C-Series Intersight Managed Mode servers. To support this configuration on Intersight Managed Mode servers, a minimum firmware version 4.2(3a) is required.</p>

7. Click **Create**.

Creating an LDAP Policy

Lightweight Directory Access Protocol (LDAP) stores and maintains directory information in a network. When LDAP is enabled in the Cisco IMC, user authentication and role authorization is performed by the LDAP server for user accounts not found in the local user database. You can enable and configure LDAP, and configure LDAP servers and LDAP groups.



Note This policy, if attached to a server profile that is assigned to an Intersight Managed FI-attached UCS server, will be ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **LDAP**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable LDAP	The state of the LDAP service on the endpoint.
Base Settings	
Base DN	Base Distinguished Name. This field describes where to load users and groups from. It must be in the dc=domain,dc=com format for Active Directory servers.
Domain	The IPv4 domain that all users must be in. This field is required unless you specify at least one Global Catalog server address.

Property	Essential Information
<p>Timeout</p>	<p>The number of seconds that Intersight waits until the LDAP search operation times out.</p> <p>If the search operation times out, Intersight tries to connect to the next server listed on this tab, if one is available.</p> <p>Note The value you specify for this field could impact the overall time.</p>
<p>Enable Encryption</p>	<p>If enabled, the server encrypts all information it sends to the LDAP server.</p>
<p>Binding Parameters</p>	
<p>Bind Method</p>	<p>It can be one of the following:</p> <p>Anonymous—requires NULL username and password. If this option is selected and the LDAP server is configured for Anonymous logins, then the user can gain access.</p> <p>Configured Credentials—requires a known set of credentials to be specified for the initial bind process. If the initial bind process succeeds, then the distinguished name (DN) of the user name is queried and re-used for the re-binding process. If the re-binding process fails, then the user is denied access.</p> <p>Login Credentials—requires the user credentials. If the bind process fails, the user is denied access. By default, the Login Credentials option is selected.</p>
<p>Bind DN</p>	<p>The distinguished name (DN) of the user. This field is editable only if you have selected Configured Credentials option as the binding method.</p>
<p>Bind Password</p>	<p>The password of the user. This field is editable only if you have selected Configured Credentials option as the binding method.</p>
<p>Search Parameters</p>	
<p>Filter</p>	<p>This field must match the configured attribute in the schema on the LDAP server.</p> <p>By default, this field displays sAMAccountName.</p>
<p>Group Attribute</p>	<p>This field must match the configured attribute in the schema on the LDAP server.</p> <p>By default, this field displays memberOf.</p>

Property	Essential Information
<p>Attribute</p>	<p>An LDAP attribute that contains the role and locale information for the user. This property is always a name-value pair. The system queries the user record for the value that matches this attribute name.</p> <p>The LDAP attribute can use an existing LDAP attribute that is mapped to the Cisco IMC user roles and locales, or can modify the schema such that a new LDAP attribute can be created. For example, CiscoAvPair.</p> <p>Note If you do not specify this property, the user cannot login. Although the object is located on the LDAP server, it should be an exact match of the attribute that is specified in this field.</p>
<p>Group Authorization</p>	
<p>Group Authorization</p>	<p>If enabled, user authentication is also done on the group level for LDAP users that are not found in the local user database.</p>
<p>Nested Group Search Depth</p>	<p>Parameter to search for an LDAP group nested within another defined group in an LDAP group map. The parameter defines the depth of a nested group search.</p>
<p>Configure LDAP Servers</p>	
<p>Enable DNS</p>	<p>If enabled, you can use DNS to configure access to the LDAP servers.</p>
<p>Source</p>	<p>Specifies how to obtain the domain name used for the DNS SRV request. It can be one of the following:</p> <ul style="list-style-type: none"> • Extracted—specifies using domain name extracted-domain from the login ID • Configured—specifies using the configured-search domain. • Configured-Extracted—specifies using the domain name extracted from the login ID than the configured-search domain.
<p>Server</p>	<p>The IP address or host name of the LDAP server.</p>
<p>Port</p>	<p>The LDAP server port numbers.</p>

Property	Essential Information
User Search Precedence	The order of search between the local user database and LDAP user database. This can be one of the following: <ul style="list-style-type: none"> • Local User Database (Default setting) • LDAP User Database
Add New LDAP Group	
Name	The name of the group in the LDAP server database that is authorized to access the server.
Domain	The LDAP server domain the group must reside in.
Role	The role assigned to all users in this LDAP server group. This can be one of the following: <ul style="list-style-type: none"> • read-only—A user with this role can view information but cannot make any changes. • user—A user with this role can perform the following tasks: <ul style="list-style-type: none"> • View all information • Manage the power control options such as power on, power cycle, and power off • Launch the KVM console and virtual media • Clear all logs • Toggle the locator LED • Set time zone • Ping • admin—A user with this role can perform all actions available through the GUI, CLI, and IPMI.
Port	The LDAP server port numbers.
User Search Precedence	The order of search between the local user database and LDAP user database. This can be one of the following: <ul style="list-style-type: none"> • Local User Database (Default setting) • LDAP User Database

- Click **Create**.

Creating a Local User Policy

The Local User policy automates the configuration of local user preferences. You can create one or more Local User policies which contain a list of local users that need to be configured.



Note By default, IPMI support is enabled for all users

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Local User**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Password Properties	Password properties apply only to Rack servers and not to Blade Servers.
Enforce Strong Password	Enables strong password policy.
Change Password	Enables changing the existing password.
Enable Password Expiry	Enables password expiry on the endpoint. Note Password expiry once set by the admin is applicable for all users that are subsequently created. The valid Password Expiry Duration must be greater than the Notification Period and the Grace Period . If otherwise, you will see an User Password Expiry Policy configuration error .

Property	Essential Information
Password Expiry Duration	The time period that you can set for the existing password to expire (from the time you set a new password or modify an existing one). The range is between 1 to 3650 days.
Notification Period	Notifies the time by when the password expires. Enter a value between 0 to 15 days. Entering 0 disables this field.
Grace Period	Time period till when the existing password can still be used, after it expires. Enter a value between 0 to 5 days. Entering 0 disables this field.
Password History	The number of occurrences when a password was entered. When this is enabled, you cannot repeat a password. Enter a value between 0 to 5. Entering 0 disables this field.
Always Send User Password	When enabled, the user password is always sent to the endpoint device. When not enabled, the user password is sent to the endpoint device for new users and when the password is changed for existing users.
Add New User	
Enable	Enables the user account on the endpoint.
New User	Enables new user configuration.
Username	The username for the user. Enter between 1 and 16 characters.

Property	Essential Information
Role	<p>The role associated with the user on the endpoint.</p> <ul style="list-style-type: none">• read-only—A user with this role can view information but cannot make any changes.• user—The user role type is supported only in racks. A user with this role can perform the following tasks:<ul style="list-style-type: none">• View all information• Manage the power control options such as power on, power cycle, and power off• Launch the KVM console and virtual media• Clear all logs• Ping• admin—A user with this role can perform all actions available through the GUI, CLI, and IPMI.

Property	Essential Information
<p>Password</p>	<p>The password for this user name. When you move the mouse over the help icon beside the field, the following guidelines to set the password are displayed:</p> <ul style="list-style-type: none"> • The password must have a minimum of 8 and a maximum of 20 characters. This is an Intersight platform limitation. • The password must not contain the User Name. • The password must contain characters from three of the following four categories: <ul style="list-style-type: none"> • English uppercase characters (A through Z). • English lowercase characters (a through z). • Base 10 digits (0 through 9). • Non-alphabetic characters (!, @, #, \$, %, ^, &, *, -, _ , =, "). <p>These rules are meant to define a strong password for the user, for security reasons. However, if you want to set a password of your choice ignoring these guidelines, click the Disable Strong Password button on the Local Users tab. While setting a password when the strong password option is disabled, you can use between 1- 20 characters.</p> <p>Note You can change the password of a Local User policy by editing the policy. However, the Change Password option is disabled once the policy is deployed.</p>
<p>Password Confirmation</p>	<p>The password repeated for confirmation purposes.</p>

7. Click **Create**.

Creating an NTP Policy

The NTP policy enables the NTP service to configure a UCS system that is managed by Cisco Intersight to synchronize the time with an NTP server. You must enable and configure the NTP service by specifying the IP/DNS address of at least one server or a maximum of four servers that function as NTP servers. When you enable the NTP service, Cisco Intersight configures the NTP details on the endpoint.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **NTP**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable NTP	Enables NTP policy configuration.
NTP Servers	A collection of NTP Server IP addresses or hostnames.
Time Zone	A collection of time zones from which you can select a time zone for the endpoint. This property is applicable to switches and to Cisco IMC (standalone) servers.

When a hostname is used for NTP configuration, DNS server information must be configured in the Network Connectivity policy.

7. Click **Create**.

Creating an SD Card Policy

The SD Card policy in Cisco Intersight configures the Cisco FlexFlash and FlexUtil Secure Digital (SD) cards for the Cisco UCS C-Series Standalone M4, M5 servers, and Cisco UCS C-Series M5 servers in a Cisco Intersight-Managed Fabric Interconnect Domain. This policy specifies details of virtual drives on the SD cards. You can configure the SD cards in the Operating System Only, Utility Only, or Operating System + Utility modes.

When two cards are present in the Cisco FlexFlash controller and Operating System is chosen in the SD card policy, the configured OS partition is mirrored. If only single card is available in the Cisco FlexFlash controller, the configured OS partition is non-RAID. The utility partitions are always set as non-RAID.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.

2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SD Card**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Operating System Only	
Operating System	Enables the Operating System partition.
Operating System Partition Name	The name for the Operating System partition.
Utility Only	
Diagnostics	Enables the Operating System health diagnostics utility.
Drivers	Enables virtual driver utility.
Host Upgrade Utility	Enables Host Upgrade Utility (HUU).
Server Configuration Utility	Enables Server Configuration Utility (SCU).
User Partition	Enables user partition.
User Partition Name	The user partition name.
Operating System + Utility	
Diagnostics	Enables the operating system health diagnostics utility.
Drivers	Enables virtual driver utility.
Host Upgrade Utility	Enables Host Upgrade Utility (HUU).
Server Configuration Utility	Enables Server Configuration Utility (SCU).
User Partition	Enables user partition.

Property	Essential Information
User Partition Name	The user partition name.
Operating System Partition	Enables the Operating System partition.
Operating System Partition Name	The name for the Operating System partition.

- Click **Create**.

Exceptions

- **SD Card Policy is not supported on M6 servers.**
- SD Card Policy is not imported with a Server Profile when the SD Cards are not present in the server.
- Diagnostics is applicable for M5 Series only.
- For the Operating System+Utility mode the M5 servers require at least 1 FlexFlash + 1 FlexUtil card.

Create a Serial Over LAN Policy

The Serial Over LAN policy enables the input and output of the serial port of a managed system to be redirected over IP. You can create one or more Serial over LAN policies which contain a specific grouping of Serial over LAN attributes that match the needs of a server or a set of servers.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Serial Over LAN**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Serial Over LAN	The state of Serial Over LAN service on the endpoint.

Property	Essential Information
COM Port	<p>The serial port through which the system routes Serial Over LAN communication.</p> <ul style="list-style-type: none"> • com0—SoL communication is routed through COM port 0, an externally accessible serial port that supports either a physical RJ45 connection to an external device or a virtual SoL connection to a network device. <p>If you select this option, the system enables SoL and disables the RJ45 connection, which means that the server can no longer support an external serial device.</p> <ul style="list-style-type: none"> • com1—SoL communication is routed through COM port 1, an internal port accessible only through SoL. <p>If you select this option, you can use SoL on COM port 1 and the physical RJ45 connection on COM port 0.</p> <p>Note</p> <ul style="list-style-type: none"> • This is applicable to Cisco UCS C-Series Standalone M4, M5, and M6 servers only. • Serial Port is available only on some Cisco UCS C-Series servers. If it is unavailable, the server uses COM port 0 by default. Changing the Com Port setting disconnects any existing SoL sessions.
Baud Rate	<p>The Baud Rate used for Serial Over LAN communication. The rate can be:</p> <ul style="list-style-type: none"> • 9600 bps • 19.2 kbps • 38.4 kbps • 57.6 kbps • 115.2 kbps <p>Note The baud rate must match the baud rate configured in the server serial console.</p>

Property	Essential Information
SSH Port	<p>The SSH port used to access Serial Over LAN directly. Enables bypassing Cisco IMC shell to provide direct access to Serial Over LAN.</p> <p>The valid range is 1024 to 65535. The default value is 2400.</p> <p>Note</p> <ul style="list-style-type: none"> • This is applicable to Cisco UCS C-Series Standalone M4, M5 and M6 servers only. • Changing the SSH Port setting disconnects any existing SSH sessions.

7. Click **Create**.

Create SSH Policy

The SSH policy enables an SSH client to make a secure, encrypted connection. You can create one or more SSH policies that contain a specific grouping of SSH properties for a server or a set of servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SSH**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SSH Policy	Enables SSH.
SSH Port	The port used for secure shell access.

Property	Essential Information
SSH Timeout (seconds)	The number of seconds to wait before the system considers a SSH request to have timed out. Enter an integer between 60 and 10,800. The default is 1,800 seconds.

- Click **Create**.

Creating a Virtual KVM Policy

The KVM console is an interface that emulates a direct keyboard, video, and mouse (KVM) connection to the server. It allows you to control the server from a remote location and to map physical locations to virtual drives that can be accessed by the server during this KVM session.

Enables specific grouping of virtual KVM properties. This policy lets you specify the number of allowed concurrent KVM sessions, port information, and video encryption options.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Virtual KVM**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Virtual KVM	The state of the vKVM service on the endpoint.
Max Sessions	The maximum number of concurrent KVM sessions allowed.
Remote Port	The port for remote KVM communication. The port range is from 1024 to 49151. The default is 2068.

Property	Essential Information
Enable Video Encryption	<p>Enables encryption on all video information sent through KVM. The Video Encryption is enabled by default.</p> <p>Note For firmware versions 4.2(1a) or higher, this encryption parameter is deprecated and disabling the encryption will further result in validation failure during the server profile deployment.</p>
Enable Local Server Video	<p>Enables KVM session displays on any monitor attached to the server.</p> <p>Note This is applicable to Cisco UCS C-Series Standalone M4, M5, and M6 servers only.</p>
Allow Tunneled vKVM	<p>Enable to allow tunneled vKVM on the endpoint.</p> <p>Note Applies only to Device Connectors that support Tunneled vKVM.</p>

7. Click **Create**.

Exceptions

- The virtual media viewer is accessed through the KVM. If you disable the KVM console, Cisco IMC also disables access to all virtual media devices attached to the host.
- After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Creating a Virtual Media Policy

The Virtual Media policy enables you to install an operating system on the server using the KVM console and virtual media, mount files to the host from a remote file share, and enable virtual media encryption. You can create one or more virtual media policies, which could contain virtual media mappings for different OS images, and configure up to two virtual media mappings, one for ISO files through CDD and the other for IMG files through HDD.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Virtual Media**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format. For example, Org: IT or Site: APJ.
Description (optional)	Enter a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable Virtual Media	Select this option to enable the virtual media policy. This property is enabled by default.
Enable Virtual Media Encryption	Select this option to enable encryption of the virtual media communications. This property is enabled by default. Note For firmware versions 4.2(1a) or higher, this encryption parameter is deprecated and disabling the encryption will further result in validation failure during the server profile deployment.
Enable Low Power USB	Select this option to enable the appearance of virtual drives on the boot selection menu after mapping the image and rebooting the host. This property is enabled by default.
Add Virtual Media	
Virtual Media Type	Select the remote virtual media type: <ul style="list-style-type: none"> • CDD • HDD
NFS/CIFS/HTTP/HTTPS	
The properties below vary depending on the tab that is selected.	
Name	The identity of the image for virtual media mapping.

Property	Essential Information
File Location	<p>Provide the remote file location path: Host Name or IP address/file path/file name</p> <ul style="list-style-type: none"> • IP Address—The IP address or the hostname of the remote server. • File Path—The path to the location of the image on the remote server. • File Name—The name of the remote file in .iso or .img format. <p>The remote file location path for virtual media mapping, the options include:</p> <ul style="list-style-type: none"> • HDD Virtual Media: hostname or IP address /filePath/fileName.img • CDD Virtual Media: hostname or IP address /filePath/fileName.iso • HDD Virtual media for HTTP: http://server-hostname-or-ip/filePath/fileName.img • CDD Virtual media for HTTP: http://server-hostname-or-ip/filePath/fileName.iso • HDD Virtual media for HTTPS: https://server-hostname-or-ip/filePath/fileName.img • CDD Virtual media for HTTPS: https://server-hostname-or-ip/filePath/fileName.iso
Username	The username to log in to the remote server. This field is displayed on selecting CIFS, HTTP, or HTTPS.
Password	The password associated with the username. This field is displayed on selecting CIFS, HTTP, or HTTPS.

Property	Essential Information
Mount Options	<p>The mount options for the virtual media mapping. The field can be left blank or filled in a comma separated list using the following options:</p> <ul style="list-style-type: none"> • For NFS, supported options are ro, rw, noexec, noexec, soft, port=VALUE, timeo=VALUE, retry=VALUE. • For CIFS, supported options are soft, nounix, noserverino, guest, ver=3.0, or ver=2.0. <p>Note If the firmware version is 4.1 or higher, and the CIFS version is lower than 3.0, the mount option field must be entered with the version value (vers=VALUE). For example, vers=2.0.</p> <ul style="list-style-type: none"> • For HTTP and HTTPS, the only supported option is noauto.
Authentication Protocol	<p>Select the authentication protocol when CIFS is used for communication with the remote server. This field is displayed on selecting CIFS.</p> <ul style="list-style-type: none"> • None—No authentication is used • ntlm—NT LAN Manager (NTLM) security protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • ntlmi—NTLMI security protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • ntlmv2—NTLMv2 security protocol. Use this option only with Samba Linux. • ntlmv2i—NTLMv2i security protocol. Use this option only with Samba Linux. • ntlmssp—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • ntlmsspi—NT LAN Manager Security Support Provider (NTLMSSPI) protocol. Use this option only when you enable Digital Signing in the CIFS Windows server.
Add	Click Add to confirm adding the virtual media.

- Click **Create**.

Exceptions

- When an answer file is embedded in the OS ISO, it fails to boot from vMedia when the bootmode is set to UEFI, and the OS installation fails on Cisco UCS C-Series Standalone M4 servers.
- vMedia mapping of the OS image for HTTPS based share fails to mount.

Creating a Network Connectivity Policy

The Network Connectivity policy enables you to configure and assign IPv4 and IPv6 addresses.

Dynamic DNS

Dynamic DNS (DDNS) is used to add or update the resource records on the DNS server. When you enable the DDNS option, the DDNS service records the current hostname, Domain name, and the management IP address and updates the resource records in the DNS server.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Network Connectivity**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following properties:

Common Properties

Property	Essential Information
Enable Dynamic DNS	Enables Dynamic DNS. This property is not applicable to Fabric Interconnects.
Dynamic DNS Update Domain	Specify the dynamic DNS Domain. The Domain can be either a main Domain or a sub-Domain. This property is not applicable to Fabric Interconnects.

IPv4 Properties

Property	Essential Information
Obtain IPv4 DNS Server Addresses from DHCP	<p>Whether the IPv4 addresses are obtained from Dynamic Host Configuration Protocol (DHCP) or from a specifically configured set of DNS servers.</p> <ul style="list-style-type: none"> • Enabled—Intersight uses DHCP • Disabled—Intersight uses a configured set of IPv4 DNS servers. <p>This property is not applicable to Fabric Interconnects.</p>
Preferred IPv4 DNS Server	The IP address of the primary DNS server. This property is displayed only when Obtain IPv4 DNS Server Addresses from DHCP is disabled.
Alternate IPv4 DNS Server	The IP address of the secondary DNS server. This property is displayed only when Obtain IPv4 DNS Server Addresses from DHCP is disabled.

Property	Essential Information
Enable IPv6	Whether IPv6 is enabled. You can configure IPv6 properties only if this property is enabled.

IPv6 Properties

Property	Essential Information
Obtain IPv6 DNS Server Addresses from DHCP	<p>Whether the IPv6 addresses are obtained from Dynamic Host Configuration Protocol (DHCP) or from a specifically configured set of DNS servers.</p> <ul style="list-style-type: none"> • Enabled—Intersight uses DHCP • Disabled—Intersight uses a configured set of IPv6 DNS servers. <p>This property is not applicable to Fabric Interconnects.</p>
Preferred IPv6 DNS Server	The IP address of the primary DNS server. This property is displayed only when Obtain IPv6 DNS Server Addresses from DHCP is disabled.
Alternate IPv6 DNS Server	The IP address of the secondary DNS server. This property is displayed only when Obtain IPv6 DNS Server Addresses from DHCP is disabled.

7. Click **Create**.

Creating a SMTP Policy

Simple Mail Transfer Protocol (SMTP) sends server faults as email alerts to the configured SMTP server.

Sets the state of the SMTP client in the managed device. You can specify the preferred settings for outgoing communication and select the fault severity level to report and the mail recipients.



Note This policy, if attached to a server profile that is assigned to an Intersight Managed FI-attached UCS server, will be ignored.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SMTP**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. In the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SMTP	Enables or disables the SMTP policy.
SMTP Server Address	The IP address or host name of the SMTP server.
SMTP Port	The port number used by the SMTP server for outgoing SMTP communication. The range is from 1 to 65535. The default is 25.
Minimum Severity	The minimum fault severity level to receive email notifications. Email notifications are sent for all faults whose severity is equal to or greater than the chosen level.
SMTP Alert Sender Address	The sender IP address or hostname of all the SMTP mail alerts.

Property	Essential Information
Mail Alert Recipients	A list of email addresses that will receive notifications for faults.

7. Click **Create**.

Creating an SNMP Policy

The SNMP policy configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. This policy supports SNMP versions such as SNMPv1, SNMPv2(includes v2c), and SNMPv3. Any existing SNMP Users or SNMP Traps configured previously on the managed devices are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the server are removed.

Using the SNMP Policy you can enable or disable SNMP, specify the access and community strings, and provide the SNMP user details that is used to retrieve data.



Note The Out-of-Band IP address support for SNMP policy is available only for the Fabric Interconnects running on Infrastructure Firmware 4.3(2.230129) or later versions.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **SNMP**, and then click **Start**.
5. In the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the organization.
Name	Enter a name for your policy.
Tag (optional)	Enter a tag in the key value format.
Description (optional)	Enter a short description.

6. In the **Policy Details** page, configure the following parameters:

Property	Essential Information
Enable SNMP	Displays the state of the SNMP Policy on the endpoint. Enable this option for the endpoint to send SNMP traps to the designated host.
SNMP Port	The port on which Cisco IMC SNMP agent runs.

Property	Essential Information
Access Community String	<p>Enter the SNMPv1, SNMPv2 community string or the SNMPv3 username. This field allows maximum of 18 characters.</p> <p>Note If the field is empty, it indicates that the SNMPv1 and SNMPv2c users are disabled.</p>
SNMP Community Access	<p>The controls access to the information in the inventory tables. Applicable only for SNMPv1 and SNMPv2c users.</p> <p>Note This property is supported only in UCS Standalone C-Series M4, M5, and M6 servers.</p>
Trap Community String	<p>Enter the SNMP community group name used for sending SNMP trap to other devices.</p> <p>Note This field is applicable only for SNMPv2c trap host or destination.</p>
System Contact	<p>The contact person responsible for the SNMP implementation. Enter a string up to 64 characters, such as an email address or a name and telephone number.</p> <p>Note This property is supported only in UCS Standalone C-Series M4, M5, and M6 servers.</p>
System Location	<p>The location of host on which the SNMP agent (server) runs.</p> <p>Note This property is supported only in UCS Standalone C-Series M4, M5, and M6 servers.</p>
SNMP Engine Input ID	<p>The user-defined unique identification of the static engine.</p> <p>Note This property is supported only in UCS Standalone C-Series M4, M5, and M6 servers.</p>
SNMP Users	
Name	Enter the SNMP username. This field must have a minimum of 1 and a maximum of 31 characters.

Property	Essential Information
Security Level	Select the security mechanism for communication between the agent and the manager that include: <ul style="list-style-type: none"> • AuthPriv • AuthNoPriv
Auth Type	Select SHA as the authorization protocol for authenticating the user. Note The MD5 authorization protocol is not supported.
Auth Password	Enter the authorization password for the user.
Auth Password Confirmation	Enter the authorization password confirmation for the user.
Privacy Type	Select AES as the privacy protocol for the user. Note The DES privacy type is deprecated to meet security standards.
Privacy Password	Enter the privacy password for the user.
Privacy Password Confirmation	Enter the privacy password confirmation for the user.
SNMP Trap Destinations	
Enable	Enable this option to use the SNMP policy.
SNMP Version	Select v2 or v3 as the SNMP version for the trap.
User	Select the SNMP user for the trap. You can define maximum of 15 trap users. Note This field is applicable only to SNMPv3.
Trap Type	Select the trap type to receive a notification when a trap is received at the destination: <ul style="list-style-type: none"> • Trap • Inform
Destination Address	Provide the address to which the SNMP trap information can be sent. You are allowed to define maximum of 15 trap destinations.

Property	Essential Information
Port	Enter the port number for the server to communicate with trap destination. The range is from 1 to 65535. The default is 162.

- Click **Create**.

Creating a Storage Policy

The Storage policy allows you to create drive groups, virtual drives, configure the storage capacity of a virtual drive, and configure the M.2 RAID controllers.

- Log in to Cisco Intersight with your Cisco ID and select admin role.
- From the **Service Selector** drop-down list, select **Infrastructure Service**.
- Navigate to **Configure > Policies**, and then click **Create Policy**.
- Select **Storage**, and then click **Start**.
- On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

- On the **Policy Details** page, configure the following parameters:

Property	Essential Information
General Configuration	
Use JBOD Drives for Virtual Drive creation	Enable this option to use disks in JBOD state for creating virtual drives.
Unused Disks State	Select the state to which unused disks in this policy are to be moved. The state can be any one of UnconfiguredGood , or JBOD . Selecting No Change leaves the state unchanged.

Property	Essential Information
<p>Default Drive Mode</p>	<p>Select the default disk state that should be set on supported storage controller for newly inserted drives or on reboot. The state can be any one of UnconfiguredGood, JBOD, or RAID0.</p> <p>Unused Disks State should be No Change if Default Drive Mode is set to JBOD or RAID0.</p> <p>Note The default drive mode is supported only on M6 servers and for the following storage controllers.</p> <ul style="list-style-type: none"> • UCSC-RAID-M6T • UCSC-RAID-M6HD • UCSC-RAID-M6SD • UCSX-X10C-RAIDF <p>Configuration Limitations:</p> <ul style="list-style-type: none"> • When Default Drive State is JBOD or RAID0, then Unused Disks State should be No Change. • Use JBOD for VD creation cannot be enabled if Default Drive Mode is JBOD. • When Default Drive State is UnconfiguredGood, the drive state does not change on reboot. <p>Refer the table Default Drive Mode Scenarios for different Default Drive Mode scenarios.</p>
<p>Secure JBOD Disk Slots</p>	<p>Specify the JBOD drive slots that you want to encrypt. You may enter a comma or hyphen separated number range. For example: 1, 3 or 4-6, 8.</p>
<p>M.2 RAID Configuration</p>	<p>Enable this option to specify the Virtual Drive Name and Slot of the M.2 RAID controller for virtual drive creation.</p> <p>The disk slots used by the M.2 controller are automatically added.</p>

Property	Essential Information
Virtual Drive Name	<p>This field comes pre-filled with a default name. You can change it to your preferred name. A suffix will be added to your preferred name based on the selected controller slot.</p> <p>The name must be between 1 and 15 characters in length and can include letters, numbers, and the special characters hyphen (-), underscore (_), colon (:), and period (.).</p>
Slot of the M.2 RAID Controller for Virtual Drive Creation	<p>Select the slot of the M.2 RAID controller for virtual drive creation. The slots that can be selected are:</p> <ul style="list-style-type: none"> • MSTOR-RAID-1 — Select this option if there is only one M.2 RAID controller slot, or if there are two slots for the M.2 RAID controller and the virtual drive has to be created on the controller in the first slot. • MSTOR-RAID-2 — Select this option if there are two slots for the M.2 RAID controller and the virtual drive has to be created on the controller in the second slot. • MSTOR-RAID-1,MSTOR-RAID-2 — Select this option to create virtual drives on controllers in either or both slots.
Drive Group Configuration	<p>Enable to add RAID drive groups that can be used to create virtual drives. You can also specify the Global Hot Spares information.</p> <p>This configuration is not applicable for M.2 RAID controllers.</p>
Global Hot Spares	<p>Specify the disks that are to be used as hot spares, globally for all the RAID groups.</p> <p>The allowed value is a number range separated by a comma or a hyphen.</p>
Add Drive Group	Click to add a drive group.
Drive Group Name	<p>Enter the name of the drive group.</p> <p>The name must be between 1 and 15 characters in length and can include letters, numbers, and the special characters hyphen (-), underscore (_), colon (:), and period (.).</p>

Property	Essential Information
<p>RAID Level</p>	<p>The RAID level of a disk group describes how the data is organized on the disk group for the purpose of ensuring availability, redundancy of data, and I/O performance. The levels are:</p> <ul style="list-style-type: none"> • RAID0—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. • RAID1—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. • RAID5—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. • RAID6—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored. • RAID10—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10. • RAID50—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID60—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.
<p>Secure Drive Group</p>	<p>Enable this option to configure encryption for drives that are part of the Virtual Drive.</p>

Property	Essential Information
Number of Spans	Number of span groups to be created for the RAID group. RAID levels with no nesting have a single span. Note Number of spans appears only when a RAID level with spans is selected.
Drive Selection	
Drive Array Span 0	Enter the drive array span. RAID levels RAID0, RAID1, RAID5, and RAID6 that do not have spans have only one disk group. RAID levels with spans have multiple disk groups with each disk group representing a span. RAID levels without spans have one span group and RAID levels with spans have two to eight span groups. Note If you have selected a RAID level without spans, then the field Drive Array Span 0 alone appears. If you have selected a RAID level with spans, you would have had to specify the number of spans. In this scenario, as many Drive Array Span fields as there are spans appear for you to specify the details.
Dedicated Hot Spares	Specify the collection of drives to be used as hot spares for this drive group. The allowed value is a number range separated by a comma or a hyphen.
Add	Click Add to add the drive group.
Add Virtual Drive	
Drive Groups	Select the drive groups on which the virtual drive is to be created.
Number of Copies	Enter the number of copies of the virtual drive that is to be created. You can create a maximum of 10 copies.
Virtual Drive Configuration	

Property	Essential Information
Virtual Drive Name	Enter the name of the virtual drive. The name can be 1 to 15 characters long and can contain alphanumeric characters, and special characters '-' (hyphen), '_' (underscore), ':' (colon), and '.' (period).
Size (MiB)	Virtual drive size in MebiBytes. Size is mandatory except when the Expand to Available option is enabled.
Secured	Set this to enable encryption for the virtual drive. Note This option is not supported for UCS-M2-NVRAID (M.2 NVMe controller) as there are no SED drives that are supported on this controller.
RAID Type	Select the RAID type.
Expand to Available	Enable for the virtual drive to use all the space available in the disk group. When this flag is enabled, the size property is ignored.
Set as Boot Drive	Select to use this virtual drive as a boot drive. Note For standalone racks, you cannot set a drive, with a native block size of 4K, as the boot drive.
Strip Size	Select the strip size required. Allowed values are 64KiB, 128KiB, 256KiB, 512KiB, 1 MiB.
Access Policy	Select the type of access the host has to this virtual drive: <ul style="list-style-type: none"> • Read Write—Enables host to perform read-write on the virtual drive • Read Only—Host can only read from the virtual drive. • Blocked—Host can neither read nor write to the virtual drive.
Read Policy	Select the read ahead mode for this virtual drive: <ul style="list-style-type: none"> • Always Read Ahead • No Read Ahead

Property	Essential Information
Write Policy	Select the mode to be used to write to this virtual drive: <ul style="list-style-type: none"> • Write Through—Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back Good BBU—With this policy, write caching remains Write Back even if the battery backup unit is in good condition. • Always Write Back—Data is stored in the cache, and is only written to the physical drives when space in the cache is needed.
Disk Cache	Select the disk cache policy for this virtual drive. The values are: <ul style="list-style-type: none"> • Unchanged • Enabled • Disabled
Add	Click Add to add the virtual drive.
Single Drive RAID Configuration	Enable to create RAID0 virtual drives on each physical drive.
Drive Slots	Specify the set of drive slots where RAID0 virtual drives are to be created. <p>Note Single drive RAID allows you to add slots only where disks are planned to be inserted in future.</p>
Strip Size	Select the strip size required. Allowed values are 64KiB, 128KiB, 256KiB, 512KiB, 1MiB.
Access Policy	Select the type of access the host has to this virtual drive: <ul style="list-style-type: none"> • Read Write—Enables host to perform read-write on the virtual drive • Read Only—Host can only read from the virtual drive. • Blocked—Host can neither read nor write to the virtual drive.

Property	Essential Information
Read Policy	Select the read ahead mode for this virtual drive: <ul style="list-style-type: none">• Always Read Ahead• No Read Ahead
Write Policy	Select the mode to be used to write to this virtual drive: <ul style="list-style-type: none">• Write Through—Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache.• Write Back Good BBU—With this policy, write caching remains Write Back even if the battery backup unit is in good condition.• Always Write Back—Data is stored in the cache, and is only written to the physical drives when space in the cache is needed.
Disk Cache	Select the disk cache policy for this virtual drive. The values are: <ul style="list-style-type: none">• Unchanged• Enabled• Disabled

Property	Essential Information
<p>Hybrid Slot Configuration</p>	<p>Select the following modes for server that supports Hybrid Drive Slots configuration:</p> <ul style="list-style-type: none"> • Direct Attached NVMe Slots—NVMe drives specified in the slot range will be moved to direct attached mode. • RAID Attached NVMe Slots—NVMe drives specified in the slot range will be moved to RAID attached mode. <p>Note</p> <ul style="list-style-type: none"> • NVMe Hybrid slots are supported only for UCSC-C240-M7 and UCSC-C220-M7 servers in Standalone mode and Intersight Managed Mode. • Hybrid slots support is available for Slots 1–4 and Slots 101–104. • If an endpoint has Trimode 24G SAS RAID controller with PID UCSC-RAID-HP and Micron 7450 4GC cache Drive then the RAID attached NVMe slots can be used to create RAID configuration. • Combination of U.2 and U.3 drive PIDs are not recommended in the hybrid slots.

7. Click **Create**.



Note The Delete Virtual Drives option is not available in Storage Policy. Use the Storage Controllers page to delete virtual drives



Note Decommissioning or recommissioning operation will not delete the RAIDs or data on the disks.

The following table shows the behavior of Default Drive State in different scenario.

Table 3: Default Drive Mode Scenarios

Default Drive State	Host Reboot/ Host Boot	Hotplug	User Action (Service Profile deployment with Default Drive State)
UnconfiguredGood (OFF)	<ul style="list-style-type: none"> • All UnconfiguredGood drives remain UnconfiguredGood. • All previously converted JBOD continue to be JBOD. 	<ul style="list-style-type: none"> • Inserted drive remains UnconfiguredGood • JBOD from a different server remains UnconfiguredGood on this controller. 	<ul style="list-style-type: none"> • Setting UnconfiguredGood has no impact on the existing configuration. • Any JBOD device will remain as JBOD across controller boot. • Any UnconfiguredGood will remain UnconfiguredGood across controller boot.
JBOD	All unconfigured drives (non-user configured) are converted to JBOD.	Newly inserted unconfigured drive is converted to JBOD.	<ul style="list-style-type: none"> • All unconfigured drives (non-user configured drives) on the controller will be converted to JBOD. • User created UnconfiguredGood drive will remain UnconfiguredGood.

Default Drive State	Host Reboot/ Host Boot	Hotplug	User Action (Service Profile deployment with Default Drive State)
RAID0(RAID0 WriteBack)	All unconfigured drives will be converted to RAID0 WriteBack (WB). Note Unconfigured drives are the drives whose state remains unchanged by any user action.	Newly inserted unconfigured drive is converted to RAID0 WB.	<ul style="list-style-type: none"> All unconfigured drives (non-user created UnconfiguredGood) on the controller will be converted to RAID0 WriteBack (WB). User created UnconfiguredGood will remain UnconfiguredGood across controller reboot. Any RAID0 WriteBack device will remain as RAID0 WB across controller boot/reboot.



Note The Virtual Drives created by the system due to default drive state being **RAID0** will have **Server Profile Derived** as **No**.

The following table shows sample use cases for different Default Drive State scenarios.

Table 4: Various Drive Mode Use Cases

Use Case Scenario	Default Drive State
Using the server for JBOD Only (for example: Hyper converged, Hadoop data node and so on)	JBOD
Using the server for RAID volume (for example: SAP HANA database)	UnconfiguredGood
Using the server for Mixed JBOD and RAID volume	UnconfiguredGood
Using the server for per drive ROWB (for example: Hadoop data node)	RAID0 WriteBack

Creating a Syslog Policy

The Syslog policy defines the logging level (minimum severity) to report for a log file collected from an endpoint, the target destination to store the Syslog messages, and the Hostname/IP Address, port information, and communication protocol for the Remote Logging Server(s).

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Syslog**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Description (Optional)	Provide a short description
Add Tag (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Local Logging	
Minimum Severity to Report	Select the lowest severity level to report in the remote log. The severity levels are: <ul style="list-style-type: none"> • 0 Emergency • 1 Alert • 2 Critical • 3 Error • 4 Warning • 5 Notice • 6 Informational • 7 Debug
Remote Logging - Syslog Server 1 and Syslog Server 2	

Property	Essential Information
Enable	Select this option to enable or disable the Syslog policy. Note When the Syslog Policy is created with Syslog Server 1 disabled and Syslog Server 2 enabled, it is observed that the Syslog server 1 always gets enabled first in the end point server.
Hostname/IP Address	Enter the hostname or IP address of the Syslog server to store the Cisco IMC log. You can set an IPv4 or IPv6 address or a domain name as the remote system address.
Port	Enter the destination port number of the Syslog server between 1 and 65535. The default port number is 514.
Protocol	Select the transport layer protocol for transmission of log messages to the syslog server. The options are: <ul style="list-style-type: none"> • TCP • UDP
Minimum Severity To Report	Select the lowest severity level to report in the remote log. The severity levels are: <ul style="list-style-type: none"> • 0 Emergency • 1 Alert • 2 Critical • 3 Error • 4 Warning • 5 Notice • 6 Informational • 7 Debug

7. Click **Create**.

Creating a Power Policy for Server

This policy enables configuration of power redundancy, power profiling, and power restore for servers.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Power**, and then click **Start**.
5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.
Description (Optional)	Provide a short description

6. On the **Policy Details** page, navigate to **All Platforms** tab.
7. Configure the following parameters:

Property	Essential Information
Power Profiling	<p>Enables/disables the power profiling of the system</p> <p>Enabled—When enabled, it allows the CIMC to run power profiling utility during BIOS boot to determine the power needs of the server.</p> <p>Disabled—When disabled, power profiling is not run.</p> <p>Note This property is supported only on Cisco UCS X-Series servers.</p>
Power Priority	<p>Each server is assigned a power priority, which can be High, Medium, or Low. The power budgeted for the server depends on the power priority of the server. A server with higher priority gets a higher power budget. The default power priority of a server is Low.</p> <p>Note This property is supported on the following:</p> <ul style="list-style-type: none"> • Servers in the Cisco-UCSX-9508 chassis with the minimum Cisco IMC firmware version of 4.2(1e). • Servers in the Cisco-UCSB-5108 chassis with the minimum Cisco IMC firmware version of 4.3(2a).

Property	Essential Information
<p>Power Restore</p> <p>Allows the user to configure the power restore state of the server on the CIMC. In the absence of IMM connectivity, the CIMC will use this policy to recover the host power after a power loss event.</p> <p>Note This property is supported only on:</p> <ul style="list-style-type: none"> • Cisco UCS X-Series IMM servers in Cisco-UCSX-9508 chassis with the minimum Cisco IMC firmware version of 4.2(1e). • Cisco UCS B-Series IMM servers in Cisco-UCSB-5108 chassis with the minimum Cisco IMC firmware version of 4.2(1d). 	
Last State	Sets the host power to whatever state it was in before the power loss event.
Always On	Always power on the host after a power loss event.
Always Off	Always keep the host power off after a power loss event.

8. Click **Create**.

Creating a Thermal Policy for Server

This policy enables controlling the speed of the chassis fan.

1. Log in to Cisco Intersight with your Cisco ID and select admin role.
2. From the **Service Selector** drop-down list, select **Infrastructure Service**.
3. Navigate to **Configure > Policies**, and then click **Create Policy**.
4. Select **Thermal**, and then click **Start**.



Note Thermal Policy is not supported for Cisco UCS Standalone M4 servers, Cisco UCS B-Series servers, and Cisco UCS X-Series servers.

5. On the **General** page, configure the following parameters:

Property	Essential Information
Organization	Select the Organization.
Name	Enter a name for your policy.
Set Tags (Optional)	Enter a tag in the key:value format. For example, Org: IT or Site: APJ.

Property	Essential Information
Description (Optional)	Provide a short description.

6. On the **Policy Details** page, configure the following parameters:

Property	Essential Information
Fan Control Mode controls the fan speed of the chassis.	
Balanced	The fans run faster when needed based on the heat generated by the server. When possible, the fans return to the minimum required speed.
Low Power	The fans run at slightly lower minimum speeds than the Balanced mode, to consume less power when possible.
High Power	The fans are kept at higher speed to emphasize performance over power consumption. Note This mode is supported for all Cisco UCS C-Series servers.
Maximum Power	The fan is always kept at the maximum speed. This option provides the most cooling and consumes most power. Note This mode is supported for all Cisco UCS C-Series servers.
Acoustic	The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Note This mode is supported for all Cisco UCS C-Series servers.

7. Click **Create**.

Creating vNIC or vHBA Templates

Creating vNIC or vHBA Templates

A vNIC or vHBA template consists of common configurations that you can reuse across multiple vNICs or vHBAs, used in various Server Profiles. This approach simplifies network configuration across multiple servers. You can create vNICs or vHBAs from the template using the **Derive** operation while creating the policy. Additionally, you can attach an existing vNIC or vHBA to a template to utilize the configurations set in the template. These templates can be created with or without override options. The override option allows the configuration of the derived vNIC or vHBA to override the template configuration.

To create a vNIC or vHBA template:

1. Log in to Cisco Intersight.
2. Navigate to the **Templates** tab, and click one of the following:
 - To create a vNIC template, click **Create vNIC Template**.
 - To create a vHBA template, click **Create vHBA Template**.
3. On the **General** page:
 - a. Choose an Organization for the template from the list. This field supports the capability of configuration sharing across Organizations.
 - b. Enter a name for the template.
 - c. Enter a Tag for the template. Tags must be in the key:value format. For example, Org: IT or Site: APJ.
 - d. Enter a description to help identify the template.
 - e. Click **Next**.
4. On the **Configuration** page:
 - a. If you want to allow the configuration of the derived vNIC or vHBA to override the template configuration, select the **Allow Override** checkbox.



Note Parameters that can be overridden are indicated by an **Override Allowed** label.

- b. Configure the template properties as required.
- c. Click **Create** to create the template.



-
- Note**
- You cannot modify a template to remove mandatory configurations when there are active derived vNICs or vHBAs with override enabled.
 - If there are active overridden properties, you cannot disable Override option in the template. It is required to detach the derived vNIC or vHBA from the template usage page and then attempt to disable the Override option in the template.
-

Deriving a vNIC from a template is done as part of creating the LAN Connectivity Policy. Similarly, deriving a vHBA from a template is done as part of creating the SAN Connectivity Policy. For more information, see *Creating a LAN Connectivity Policy* and *Creating a SAN Connectivity Policy*.

Tutorial 1: Working with a vNIC template with active derived vNICs, when Override is enabled

Consider the following scenario after a vNIC has been derived from a template and attached to a LAN Connectivity policy:

1. Do the following in the vNIC template:
 - a. Enable the **Override** option.

- b. Modify **Failover** to **Enabled**. Note that the **Override** option is not available for **Failover**; hence, this property change will be propagated to the vNICs derived from this template.
 - c. Retain **MAC Pool** as is with no change. Note that the **Override** option is available for **MAC Pool**.
 - d. Create an **Ethernet Adapter** policy and attach it to the template. Note that the **Override** option is available for **Ethernet Adapter**. The **Ethernet Adapter** policy is indicated with the **Overridden** label, which means that the configuration of the derived vNICs will override the configuration propagated from the template.
2. Review the modifications in the LAN Connectivity policy that utilizes the vNIC derived from the template:
 - a. **Failover** is propagated from the template and modified to **Enabled**.
 - b. **MAC Pool** is propagated from the template. Even though **Override** is allowed for MAC Pool, it does not display the **Overridden** label since the configuration remains consistent with the template. If you modify **MAC Pool** now, as the **Override** option is available for this property, the new configuration is applied to the vNIC instance and the **Overridden** label is displayed.
 - c. Create or attach a different **Ethernet Adapter** policy to the derived vNIC of the LAN connectivity policy. The **Ethernet Adapter** policy is marked with the **Overridden** label. The **Ethernet Adapter** policy is not propagated from the template to the derived vNIC in the LAN Connectivity policy.

Working with a vHBA template with active derived vHBA, when Override is disabled

Consider the following scenario after a vHBA has been derived from a template and attached to a SAN Connectivity policy:

1. Maintain the **Override** option as **Disabled** in the vHBA template.
2. Review the SAN Connectivity policy that utilizes the vHBA derived from the template:
 - Since override is not enabled, the parameters inherited from the template can only be viewed and not be modified.
 - Only those parameters can be modified that are excluded from the template. For example, under **Placement**, **Switch ID** is part of the template, it can only be viewed. The other parameters, which are not included in the template, can be modified.
3. To make modifications to the parameters included in the template, you must enable **Override** in the template, and then retry overriding the configuration in the SAN Connectivity policy.

