



Equipment Policies

- [Chassis/FEX Discovery Policy, on page 1](#)
- [Chassis Connectivity Policy, on page 8](#)
- [Rack Server Discovery Policy, on page 9](#)
- [Aging Time for the MAC Address Table, on page 10](#)
- [HA Version Holder Replacement, on page 11](#)

Chassis/FEX Discovery Policy

The chassis/FEX discovery policy determines how the system reacts when you add a new chassis or FEX. Cisco UCS Manager uses the settings in the chassis/FEX discovery policy to determine the minimum threshold for the number of links between the chassis or FEX and the fabric interconnect and whether to group links from the IOM to the fabric interconnect in a fabric port channel.

In a Cisco UCS Mini setup, chassis discovery policy is supported only on the extended chassis.

Chassis Links

If you have a Cisco UCS domain with some of the chassis' wired with one link, some with two links, some with four links, and some with eight links, Cisco recommends configuring the chassis/FEX discovery policy for the minimum number links in the domain so that Cisco UCS Manager can discover all chassis.



Tip To establish the highest available chassis connectivity in a Cisco UCS domain where Fabric Interconnect is connected to different types of IO Modules supporting different max number of uplinks, select platform max value. Setting the platform max ensures that Cisco UCS Manager discovers the chassis including the connections and servers only when the maximum supported IOM uplinks are connected per IO Module.

After the initial discovery of a chassis, if chassis/FEX discovery policy changes are done, acknowledge IO Modules rather than the entire Chassis to avoid disruption. The discovery policy changes can include increasing the number of links between Fabric Interconnect and IO Module, or changes to the Link Grouping preference.

Make sure that you monitor for faults before and after the IO Module acknowledgement to ensure that the connectivity is restored before proceeding to the other IO Module for the chassis.

Cisco UCS Manager cannot discover any chassis that is wired for fewer links than are configured in the chassis/FEX discovery policy. For example, if the chassis/FEX discovery policy is configured for four links,

Cisco UCS Manager cannot discover any chassis that is wired for one link or two links. Re-acknowledgement of the chassis resolves this issue.

The following table provides an overview of how the chassis/FEX discovery policy works in a multi-chassis Cisco UCS domain:

Table 1: Chassis/FEX Discovery Policy and Chassis Links

Number of Links Wired for the Chassis	1-Link Discovery Policy	2-Link Discovery Policy	4-Link Discovery Policy	8-Link Discovery Policy	Platform-Max Discovery Policy
1 link between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.
2 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 link.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.

Number of Links Wired for the Chassis	1-Link Discovery Policy	2-Link Discovery Policy	4-Link Discovery Policy	8-Link Discovery Policy	Platform-Max Discovery Policy
4 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 link.	Chassis connections and servers cannot be discovered by Cisco UCS Manager and are not added to the Cisco UCS domain.	If the IOM has 4 links, the chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. If the IOM has 8 links, the chassis is not fully discovered by Cisco UCS Manager.
8 links between IOM and fabric interconnects	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 1 link. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 2 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 4 links. After initial discovery, reacknowledge the chassis and Cisco UCS Manager recognizes and uses the additional links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.	Chassis is discovered by Cisco UCS Manager and added to the Cisco UCS domain as a chassis wired with 8 links.

Link Grouping

For hardware configurations that support fabric port channels, link grouping determines whether all of the links from the IOM to the fabric interconnect are grouped in to a fabric port channel during chassis discovery. If the link grouping preference is set to **Port Channel**, all of the links from the IOM to the fabric interconnect are grouped in a fabric port channel. If set to **None**, links from the IOM are pinned to the fabric interconnect.

After a fabric port channel is created through Cisco UCS Manager, you can add or remove links by changing the link group preference and re-acknowledging the chassis, or by enabling or disabling the chassis from the port channel.



Note The link grouping preference only takes effect if both sides of the links between an IOM or FEX and the fabric interconnect support fabric port channels. If one side of the links does not support fabric port channels, this preference is ignored and the links are not grouped in a port channel.

Multicast Hardware Hash

In a port channel, by default, ingress multicast traffic on any port in the fabric interconnect (FI) selects a particular link between the IOM and the fabric interconnect to egress the traffic. To reduce potential issues with the bandwidth, and to provide effective load balancing of the ingress multicast traffic, hardware hashing is used for multicast traffic. When multicast hardware hashing is enabled, all links between the IOM and the fabric interconnect in a port channel can be used for multicast traffic.

Pinning

Pinning in Cisco UCS is only relevant to uplink ports. If you configure **Link Grouping Preference** as **None** during chassis discovery, the IOM forwards traffic from a specific server to the fabric interconnect through its uplink ports by using static route pinning.

The following table showcases how pinning is done between an IOM and the fabric interconnect based on the number of active fabric links between the IOM and the fabric interconnect.

Table 2: Pinning on an IOM

Number of Active Fabric Links	Server slot pinned to fabric link
1-Link	All the HIF ports are pinned to the active link
2-Link	1,3,5,7 to link-1 2,4,6,8 to link-2
4-Link	1,5 to link-1 2,6 to link-2 3,7 to link-3 4,8 to link-4

Number of Active Fabric Links	Server slot pinned to fabric link
8-Link (Applies only to 2208XP)	1 to link-1 2 to link-2 3 to link-3 4 to link-4 5 to link-5 6 to link-6 7 to link-7 8 to link-8

Only 1,2,4 and 8 links are supported. 3,5,6, and 7 links are not valid configurations.

Port-Channeling

While pinning traffic from a specific server to an uplink port provides you with greater control over the unified fabric and ensures optimal utilization of uplink port bandwidth, it could also mean excessive traffic over certain circuits. This issue can be overcome by using port channeling. Port channeling groups all links between the IOM and the fabric interconnect into one port channel. The port channel uses a load balancing algorithm to decide the link over which to send traffic. This results in optimal traffic management.

Cisco UCS supports port-channeling only through the Link Aggregation Control Protocol (LACP). For hardware configurations that support fabric port channels, link grouping determines whether all of the links from the IOM to the fabric interconnect are grouped into a fabric port channel during chassis discovery. If the **Link Grouping Preference** is set to **Port Channel**, all of the links from the IOM to the fabric interconnect are grouped in a fabric port channel. If this parameter is set to **None**, links from the IOM to the fabric interconnect are not grouped in a fabric port channel.

Once a fabric port channel is created, links can be added or removed by changing the link group preference and reacknowledging the chassis, or by enabling or disabling the chassis from the port channel.

Configuring the Chassis/FEX Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode. Note The chassis/FEX discovery policy can be accessed only from the root organization.
Step 2	UCS-A /org # scope chassis-disc-policy	Enters organization chassis/FEX discovery policy mode.

	Command or Action	Purpose
Step 3	UCS-A /org/chassis-disc-policy # set action { 1-link 2-link 4-link 8-link platform-max }	Specifies the minimum threshold for the number of links between the chassis or FEX and the fabric interconnect.
Step 4	(Optional) UCS-A /org/chassis-disc-policy # set descr <i>description</i>	Provides a description for the chassis/FEX discovery policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/chassis-disc-policy # set link-aggregation-pref { none port-channel }	Specifies whether the links from the IOMs or FEXes to the fabric interconnects are grouped in a port channel. Note The link grouping preference only takes effect if both sides of the links between an IOM or FEX and the fabric interconnect support fabric port channels. If one side of the links does not support fabric port channels, this preference is ignored and the links are not grouped in a port channel.
Step 6	UCS-A /org/chassis-disc-policy # set multicast-hw-hash { disabled enabled }	Specifies whether the all the links between the IOM and the fabric interconnect in a port channel can be used for multicast traffic. <ul style="list-style-type: none"> • disabled—Only one link between the IOM and the fabric interconnect is used for multicast traffic • enabled—All links between the IOM and the fabric interconnect can be used for multicast traffic
Step 7	(Optional) UCS-A /org/chassis-disc-policy # set qualifier <i>qualifier</i>	Uses the specified server pool policy qualifications to associate this policy with a server pool.
Step 8	UCS-A /org/chassis-disc-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example scopes to the default chassis/FEX discovery policy, sets it to discover chassis with four links to a fabric interconnect, provides a description for the policy, specifies the server pool policy qualifications that will be used to qualify the chassis, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope chassis-disc-policy
UCS-A /org/chassis-disc-policy* # set action 4-link
UCS-A /org/chassis-disc-policy* # set descr "This is an example chassis/FEX discovery
policy."
UCS-A /org/chassis-disc-policy* # set qualifier ExampleQual
UCS-A /org/chassis-disc-policy* # commit-buffer
UCS-A /org/chassis-disc-policy #
```

The following example scopes to the default chassis/FEX discovery policy, sets it to discover chassis with eight links to a fabric interconnect, provides a description for the policy, sets the link grouping preference to port channel, specifies the server pool policy qualifications that will be used to qualify the chassis, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope chassis-disc-policy
UCS-A /org/chassis-disc-policy* # set action 8-link
UCS-A /org/chassis-disc-policy* # set descr "This is an example chassis/FEX discovery
policy."
UCS-A /org/chassis-disc-policy* # set link-aggregation-pref port-channel
UCS-A /org/chassis-disc-policy* # set qualifier ExampleQual
UCS-A /org/chassis-disc-policy* # commit-buffer
UCS-A /org/chassis-disc-policy #
```

The following example scopes to the default chassis/FEX discovery policy, sets it to discover chassis with four links to a fabric interconnect, provides a description for the policy, sets the link grouping preference to port channel, enables multicast hardware hashing, specifies the server pool policy qualifications that will be used to qualify the chassis, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope chassis-disc-policy
UCS-A /org/chassis-disc-policy* # set action 4-link
UCS-A /org/chassis-disc-policy* # set descr "This is an example chassis/FEX discovery
policy."
UCS-A /org/chassis-disc-policy* # set link-aggregation-pref port-channel
UCS-A /org/chassis-disc-policy* # set multicast-hw-hash enabled
UCS-A /org/chassis-disc-policy* # set qualifier ExampleQual
UCS-A /org/chassis-disc-policy* # commit-buffer
UCS-A /org/chassis-disc-policy #
```

What to do next

To customize fabric port channel connectivity for a specific chassis, configure the chassis connectivity policy.

Chassis Connectivity Policy

The chassis connectivity policy determines the whether a specific chassis is included in a fabric port channel after chassis discovery. This policy is helpful for users who want to configure one or more chassis differently from what is specified in the global chassis discovery policy. The chassis connectivity policy also allows for different connectivity modes per fabric interconnect, further expanding the level of control offered with regards to chassis connectivity.

By default, the chassis connectivity policy is set to global. This means that connectivity control is configured when the chassis is newly discovered, using the settings configured in the chassis discovery policy. Once the chassis is discovered, the chassis connectivity policy controls whether the connectivity control is set to none or port channel.



Important

The 40G backplane setting is not applicable for 22xx IOMs.

The chassis connectivity policy is created by Cisco UCS Manager only when the hardware configuration supports fabric port channels.

In a Cisco UCS Mini setup, the creation of a chassis connectivity policy is supported only on the extended chassis.

Configuring a Chassis Connectivity Policy

Changing the connectivity mode for a chassis might result in decreased VIF namespace.



Caution

Changing the connectivity mode for a chassis results in chassis re-acknowledgement. Traffic might be disrupted during this time.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope chassis-conn-policy <i>chassis-num</i> [a b]	Enters chassis connection policy organization mode for the specified chassis and fabric.
Step 3	UCS-A /org/chassis-conn-policy # set link-aggregation-pref { global none port-channel }	Specifies whether the links from the IOMs or FEXes to the fabric interconnects are grouped in a port channel. <ul style="list-style-type: none"> • None—No links are grouped in a port channel

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Port Channel—All links from an IOM to a fabric interconnect are grouped in a port channel. • Global—The chassis inherits this configuration from the chassis discovery policy. This is the default value.
Step 4	UCS-A /org/chassis-conn-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to change the fabric port channel connectivity for two chassis. Chassis 6, fabric A is changed to port channel and chassis 12, fabric B is changed to discrete links:

```
UCS-A# scope org /
UCS-A /org # scope chassis-conn-policy 6 a
UCS-A /org/chassis-conn-policy # set link-aggregation-pref port-channel
UCS-A /org/chassis-conn-policy* # up
UCS-A /org* # scope chassis-conn-policy 12 b
UCS-A /org/chassis-conn-policy* # set link-aggregation-pref none
UCS-A /org/chassis-conn-policy* # commit-buffer
UCS-A /org/chassis-conn-policy #
```

Rack Server Discovery Policy

The rack server discovery policy determines how the system reacts when you add a new rack-mount server. Cisco UCS Manager uses the settings in the rack server discovery policy to determine whether any data on the hard disks are scrubbed and whether server discovery occurs immediately or needs to wait for explicit user acknowledgement.

Cisco UCS Manager cannot discover any rack-mount server that has not been correctly cabled and connected to the fabric interconnects. For information about how to integrate a supported Cisco UCS rack-mount server with Cisco UCS Manager, see the appropriate [rack-mount server integration guide](#).

Configuring the Rack Server Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode. Note The rack server discovery policy can be accessed only from the root organization.

	Command or Action	Purpose
Step 2	UCS-A /org # scope rackserver-disc-policy	Enters organization rack server discovery policy mode.
Step 3	UCS-A /org/rackserver-disc-policy # set action { immediate user-acknowledged }	Specifies the way the system reacts when you add a new rack server.
Step 4	(Optional) UCS-A /org/rackserver-disc-policy # set descr <i>description</i>	Provides a description for the rack server discovery policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/rackserver-disc-policy # set scrub-policy <i>scrub-pol-name</i>	Specifies the scrub policy that should run on a newly discovered rack server.
Step 6	UCS-A /org/rackserver-disc-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example scopes to the default rack server discovery policy, sets it to immediately discover new rack servers, provides a description for the policy, specifies a scrub policy called scrubpoll, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope rackserver-disc-policy
UCS-A /org/rackserver-disc-policy* # set action immediate
UCS-A /org/rackserver-disc-policy* # set descr "This is an example rackserver discovery
policy."
UCS-A /org/rackserver-disc-policy* # set scrub-policy scrubpoll
UCS-A /org/rackserver-disc-policy* # commit-buffer
UCS-A /org/rackserver-disc-policy #
```

Aging Time for the MAC Address Table

To efficiently switch packets between ports, the fabric interconnect maintains a MAC address table. It dynamically builds the MAC address table by using the MAC source address from the packets received and the associated port on which the packets were learned. The fabric interconnect uses an aging mechanism, defined by a configurable aging timer, to determine how long an entry remains in the MAC address table. If an address remains inactive for a specified number of seconds, it is removed from the MAC address table.

You can configure the amount of time (age) that a MAC address entry (MAC address and associated port) remains in the MAC address table.

Configuring the Aging Time for the MAC Address Table

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope eth-uplink	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # set mac-aging { <i>dd hh mm ss</i> mode-default never }	Specifies the aging time for the MAC address table. Use the mode-default keyword to set the aging time to a default value dependent on the configured Ethernet switching mode. Use the never keyword to never remove MAC addresses from the table regardless of how long they have been idle.
Step 3	UCS-A /eth-uplink # commit-buffer	Commits the transaction to the system configuration.

Example

The following example sets the aging time for the MAC address table to one day and 12 hours and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set mac-aging 01 12 00 00
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```

HA Version Holder Replacement

In releases earlier than Cisco UCS Manager Release 3.1(2), version holders are selected on a first come first serve basis. As chassis and rack servers are discovered, they can become version holders if they meet the requirements, and if the number of version holders has not reached the maximum permitted number. After a device is marked as a version holder, it persists as a version holder until it is decommissioned or removed. For example, if the connection status between the device and one or both fabric interconnects goes down, the device will not be removed as version holder.

In some situations, the shared storage devices that are selected as high availability (HA) version holders become unreachable for an extended period of time. Cisco UCS Manager Release 3.1(2) introduces the ability to specify new preferred HA version holders corresponding to the devices that are functioning correctly. When you trigger a reelection of version holders, these new preferred HA devices are selected first.

Guidelines for Preferred HA Version Holder Replacement

Consider the following guidelines when replacing HA version holders:

- Both fabric interconnects must be up for device reelection to be triggered.
- Cisco UCS Mini does not support preferred HA version holder replacement.

- A preferred version holder can be any device that is currently supported for shared storage.
- You can specify up to five preferred version holder devices. However, only three devices will be selected for active HA access.
- When you trigger shared storage device reelection, it removes all currently active devices and selects a new set of active devices. This set of devices may include previously active devices. Devices that are specified as preferred version holders are selected first as active devices.
- You can trigger reelection of shared storage devices at any time. However, the device will be selected as a version holder only in the following scenarios:
 - When the connection path is both fabric interconnect A and B for UCS B Series blade chassis
 - When the connection status is both fabric interconnect A and B for UCS C Series racks
- For a device to be selected as a version holder, the following requirements must be met:
 - There must be less than three devices selected for active HA access.
 - Chassis removal must not be in progress.
 - A chassis that has been removed from the system must not be used as a version holder.
 - The connection path must be both fabric interconnect A and B.
- Replacement of HA version holders can be done only through Cisco UCS Manager CLI.

Creating a Preferred Version Holder

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # create preferred-ha-device <i>device-serial</i>	Creates the specified preferred HA device.
Step 3	UCS-A /system/ preferred-ha-device # commit-buffer	Commits the transaction to the system configuration.
Step 4	UCS-A /system/ preferred-ha-device* # exit	Enters system mode.
Step 5	UCS-A /system # show preferred-ha-devices	Displays the list of preferred HA version holders and whether they are active or not.

Example

This example shows how to create a preferred version holder:

```
UCS-A# scope system
UCS-A /system # create preferred-ha-device FCH1606V02F
UCS-A /system/ preferred-ha-device* # commit-buffer
UCS-A /system/ preferred-ha-device # exit
```

```
UCS-A /system # show preferred-ha-devices
```

```
Preferred Version Holder:
  Chassis Serial Active
-----
FCH1606V02F   Yes
FOX1636H6R3   Yes
FOX1636H6R4   No
```

What to do next

Trigger a reelection of version holders.

Deleting a Preferred Version Holder

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # delete preferred-ha-device <i>device-serial</i>	Deletes the specified preferred HA device.
Step 3	UCS-A /system/ preferred-ha-device* # commit-buffer	Commits the transaction to the system configuration.
Step 4	UCS-A /system/ preferred-ha-device # exit	Enters system mode.
Step 5	UCS-A /system # show preferred-ha-devices	Displays the list of preferred HA version holders and whether they are active or not.

Example

This example shows how to delete a preferred version holder:

```
UCS-A# scope system
UCS-A /system # delete preferred-ha-device FCH1606V02F
UCS-A /system/ preferred-ha-device* # commit-buffer
UCS-A /system/ preferred-ha-device # exit
UCS-A /system # show preferred-ha-devices
```

```
Preferred Version Holder:
  Chassis Serial Active
-----
FOX1636H6R3   Yes
FOX1636H6R4   No
```

Triggering the Reelection of Version Holders

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # re-elect-ha-devices	Triggers reelection of version holders for HA devices.

Example

This example shows how to trigger the reelection of version holders:

```
UCS-A# scope system
UCS-A /system # re-elect-ha-devices
```

Displaying Operational Version Holders

You can use this command to display all operational version holders, including preferred version holders.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # show operational-ha-devices	Displays the list of all currently operational HA version holders.

Example

This example shows how to display all currently operational version holders:

```
UCS-A# scope system
UCS-A /system # show operational-ha-devices
```

```
Current Version Holder:
  Serial
  -----
  FOX1636H6R5
```