Configuring System-Related Policies

This chapter includes the following sections:

Configuring the Chassis Discovery Policy

Chassis Discovery Policy

This discovery policy determines how the system reacts when you add a new chassis. If you create a chassis discovery policy, Cisco UCS Manager configures the chassis for the number of links between the chassis and the fabric interconnect specified in the policy.

Configuring the Chassis Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# **scope org /**	Enters the root organization mode. **Note** The chassis discovery policy can only be accessed from the root organization.
Step 2	UCS-A /org # **scope chassis-disc-policy**	Enters organization chassis discovery policy mode.
Step 3	UCS-A /org/chassis-disc-policy # **set action** {**1-link** | **2-link** | **4-link**}	Specifies the number of links to the fabric interconnect that the chassis must have before it can be discovered.
Step 4	UCS-A /org/chassis-disc-policy # **set descr** *description*	(Optional) Provides a description for the chassis discovery policy.

	Command or Action	Purpose
		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 qualifier** *qualifier*	(Optional) Uses the specified server pool policy qualifications to associates this policy with a server pool.
Step 6	UCS-A /org/chassis-disc-policy # **commit-buffer**	Commits the transaction to the system configuration.

The following example scopes to the default chassis discovery policy, sets it to discover chassis with four links to a fabric interconnect, provides a description for the policy, and 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 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 #
```

Configuring the Power Policy

Power Policy

The power policy is a global policy that specifies the redundancy for power supplies in all chassis in the Cisco UCS instance. This policy is also known as the PSU policy.

For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Configuring the Power Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# **scope org** *org-name*	Enters organization mode for the specified organization. To enter the root organization mode, type / as the *org-name*.
Step 2	UCS-A /org # **scope psu-policy**	Enters PSU policy mode.
Step 3	UCS-A /org/psu-policy # **set redundancy** {**grid** | **n-plus-1** | **non-redund**}	Specifies one of the following redundancy types:

	Command or Action	Purpose
		• **grid**—Provides power redundancy when two power sources are used to power the chassis. If one power source fails, the surviving power supplies on the other power circuit continue to provide power to the chassis. • **n-plus-1**—Balances the power load for the chassis across the number of power supplies needed to satisfy non-redundancy plus one additional power supply for redundancy. If any additional power supplies are installed, they are recognized and powered off. • **non-redund**—Balances the power load for the chassis evenly across all installed power supplies. For more information about power redundancy, see the *Cisco UCS 5108 Server Chassis Installation Guide*.
Step 4	UCS-A /org/psu-policy # **commit-buffer**	Commits the transaction to the system configuration.

The following example configures the power policy to use grid redundancy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope psu-policy
UCS-A /org/psu-policy # set redundancy grid
UCS-A /org/psu-policy* # commit-buffer
UCS-A /org/psu-policy #
```

Configuring the Aging Time for the MAC Address Table

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.

	Command or Action	Purpose
Step 2	UCS-A /eth-uplink # **set mac-aging** {*seconds* | **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. For end-host mode, the default aging time is 7,200 seconds; for switch mode, the default aging time is 300 seconds. 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.

The following example sets the aging time for the MAC address table to 10,000 seconds and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # set mac-aging 10000
UCS-A /eth-uplink* # commit-buffer
UCS-A /eth-uplink #
```