



Administering the Switch

This chapter describes how to perform various administrative tasks on the Catalyst 6000 family switches.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 6000 Family Command Reference* publication.

This chapter consists of these sections:

- [Setting the System Name and System Prompt, page 20-1](#)
- [Setting the System Contact and Location, page 20-3](#)
- [Setting the System Clock, page 20-4](#)
- [Creating a Login Banner, page 20-4](#)
- [Defining Command Aliases, page 20-5](#)
- [Defining IP Aliases, page 20-6](#)
- [Configuring Static Routes, page 20-7](#)
- [Configuring Permanent and Static ARP Entries, page 20-8](#)
- [Scheduling a System Reset, page 20-9](#)
- [Power Management, page 20-11](#)
- [Environmental Monitoring, page 20-16](#)
- [Displaying System Status Information for Technical Support, page 20-17](#)

Setting the System Name and System Prompt

The system name on the switch is a user-configurable string used to identify the device. The default configuration has no system name configured.

If you do not manually configure a system name, the system name is obtained through the Domain Name System (DNS) if you configure the switch as follows:

- Assign the sc0 interface an IP address that is mapped to the switch name on the DNS server
- Enable DNS on the switch
- Specify at least one valid DNS server on the switch

If the DNS lookup is successful, the DNS host name of the switch is configured as the system name of the switch and is saved in NVRAM (the domain name is removed).

If you have not configured a system prompt, the first 20 characters of the system name are used as the system prompt (a greater-than symbol [>] is appended). The prompt is updated whenever the system name changes, unless you manually configure the prompt using the **set prompt** command.

The switch performs a DNS lookup for the system name whenever one of the following occurs:

- The switch is initialized (power on or reset)
- You configure the IP address on the sc0 interface using the command-line interface (CLI) or Simple Network Management Protocol (SNMP)
- You configure a route using the **set ip route** command
- You clear the system name using the **set system name** command
- You enable DNS or specify DNS servers

If the system name is user configured, no DNS lookup is performed.

Setting the Static System Name and Prompt

These sections describe how to set the static system name and prompt:

- [Setting the Static System Name, page 20-2](#)
- [Setting the Static System Prompt, page 20-3](#)
- [Clearing the System Name, page 20-3](#)

Setting the Static System Name

To set a static system name, perform this task in privileged mode:

Task	Command
Set the static system name.	set system name <i>name_string</i>



Note

When you set the system name, the system name is used as the system prompt. You can override the prompt string with the **set prompt** command.

This example shows how to configure the system name on the switch:

```
Console> (enable) set system name Catalyst 6000
System name set.
Catalyst 6000> (enable)
```

Setting the Static System Prompt

To set the static system prompt, perform this task in privileged mode:

Task	Command
Set the static system prompt.	set prompt <i>prompt_string</i>

This example shows how to set the static system prompt on the switch:

```
Console> (enable) set prompt Catalyst6509>
Catalyst6509> (enable)
```

Clearing the System Name

To clear the system name, perform this task in privileged mode:

Task	Command
Clear the system name.	set system name

This example shows how to clear the system name:

```
Console> (enable) set system name
System name cleared.
Console> (enable)
```

Setting the System Contact and Location

You can set the system contact and location to help you with resource management tasks.

To set the system contact and location, perform this task in privileged mode:

	Task	Command
Step 1	Set the system contact.	set system contact [<i>contact_string</i>]
Step 2	Set the system location.	set system location [<i>location_string</i>]
Step 3	Verify the global system information.	show system

This example shows how to set the system contact and location and verify the configuration:

```
Catalyst 6000> (enable) set system contact sysadmin@corp.com
System contact set.
Catalyst 6000> (enable) set system location Sunnyvale CA
System location set.
Catalyst 6000> (enable) show system
PS1-Status PS2-Status Fan-Status Temp-Alarm Sys-Status Uptime d,h:m:s Logout
-----
ok          none          ok           off          ok           0,04:04:07    20 min
```

```

PS1-Type   PS2-Type   Modem   Baud   Traffic Peak Peak-Time
-----
other      none       disable 9600   0%     0% Tue Jun 23 1998, 16:51:36

System Name           System Location           System Contact
-----
Catalyst 6000         Sunnyvale CA              sysadmin@corp.com
Catalyst 6000> (enable)

```

Setting the System Clock



Note

You can configure the switch to obtain the time and date using the Network Time Protocol (NTP). For information on configuring NTP, see [Chapter 31, “Configuring NTP.”](#)

To set the system clock, perform this task in privileged mode:

	Task	Command
Step 1	Set the system clock.	set time [<i>day_of_week</i>] [<i>mm/dd/yy</i>] [<i>hh:mm:ss</i>]
Step 2	Display the current date and time.	show time

This example shows how to set the system clock and display the current date and time:

```

Console> (enable) set time Mon 06/15/98 12:30:00
Mon Jun 15 1998, 12:30:00
Console> (enable) show time
Mon Jun 15 1998, 12:30:02
Console> (enable)

```

Creating a Login Banner

You can create a single or multiline message banner that appears on the screen when someone logs in to the switch. The first character following the **motd** keyword is used to delimit the beginning and end of the banner text. Characters following the ending delimiter are discarded. After entering the ending delimiter, press **Return**. The banner must be fewer than 3070 characters.

These sections describe how to configure and clear a login banner:

- [Configuring a Login Banner, page 20-5](#)
- [Clearing the Login Banner, page 20-5](#)

Configuring a Login Banner

To configure a login banner, perform this task in privileged mode:

	Task	Command
Step 1	Enter the message of the day.	set banner motd <i>c message_of_the_day c</i>
Step 2	Display the login banner by logging out and logging back into the switch.	

This example shows how to configure the login banner on the switch using the # symbol as the beginning and ending delimiter:

```
Console> (enable) set banner motd #
Welcome to the Catalyst 6000 Switch!
Unauthorized access prohibited.
Contact sysadmin@corp.com for access.
#
MOTD banner set
Console> (enable)
```

Clearing the Login Banner

To clear the login banner, perform this task in privileged mode:

	Task	Command
	Clear the message of the day.	set banner motd <i>cc</i>

This example shows how to clear the login banner:

```
Console> (enable) set banner motd ##
MOTD banner cleared
Console> (enable)
```

Defining Command Aliases

You can use the **set alias** command to define command aliases (shorthand versions of commands) for frequently used or long and complex commands. Command aliases can save you time and can help prevent typing errors when you are configuring or monitoring the switch.

The *name* argument defines the command alias. The *command* and *parameter* arguments define the command to enter when the command alias is entered at the command line.

To define a command alias on the switch, perform this task in privileged mode:

	Task	Command
Step 1	Define a command alias on the switch.	set alias <i>name command [parameter] [parameter]</i>
Step 2	Verify the currently defined command aliases.	show alias [<i>name</i>]

This example shows how to define two command aliases, **sm8** and **sp8**. **sm8** issues the **show module 8** command, and **sp8** issues the **show port 8** command. This example also shows how to verify the currently defined command aliases and what happens when you enter the command aliases at the command line:

```

Console> (enable) set alias sm8 show module 8
Command alias added.
Console> (enable) set alias sp8 show port 8
Command alias added.
Console> (enable) show alias
sm8          show module 8
sp8          show port 8
Console> (enable) sm8
Mod Module-Name          Ports Module-Type          Model      Serial-Num Status
-----
8                          2      DS3 Dual PHY ATM        WS-X5166   007243262 ok

Mod MAC-Address(es)
-----
8  00-60-2f-45-26-2f          2.0    1.3    51.1 (103)
Console> (enable) sp8
Port Name                Status Vlan    Level Duplex Speed Type
-----
8/1                       notconnect trunk   normal full   45 DS3 ATM
8/2                       notconnect trunk   normal full   45 DS3 ATM

Port  ifIndex
-----
8/1   285
8/2   286

Use 'session' command to see ATM counters.

Last-Time-Cleared
-----
Thu Sep 10 1998, 16:56:08
Console> (enable)

```

Defining IP Aliases

You can use the **set ip alias** command to define textual aliases for IP addresses. IP aliases can make it easier to refer to other network devices when using **ping**, **telnet**, and other commands, even when DNS is not enabled.

The *name* argument defines the IP alias. The *ip_addr* argument defines the IP address to which the name refers.

To define an IP alias on the switch, perform this task in privileged mode:

	Task	Command
Step 1	Define an IP alias on the switch.	set ip alias name ip_addr
Step 2	Verify the currently defined IP aliases.	show ip alias [name]

This example shows how to define two IP aliases, **sparc** and **cat6509**. **sparc** refers to IP address 172.20.52.3, and **cat6509** refers to IP address 172.20.52.71. This example also shows how to verify the currently defined IP aliases and what happens when you use the IP aliases with the **ping** command:

```

Console> (enable) set ip alias sparc 172.20.52.3
IP alias added.
Console> (enable) set ip alias cat6509 172.20.52.71
IP alias added.
Console> (enable) show ip alias
default          0.0.0.0
sparc            172.20.52.3
cat6509         172.20.52.71
Console> (enable) ping sparc
sparc is alive
Console> (enable) ping cat6509
cat6509 is alive
Console> (enable)

```

Configuring Static Routes



Note

For information on configuring a default gateway (default route), see the [“Configuring Default Gateways” section on page 3-6](#).

In some situations, you might need to add a static routing table entry for one or more destination networks. Static route entries consist of the destination IP network address, the IP address of the next hop router, and the metric (hop count) for the route.

The destination IP network address can be variably subnetted to support Classless Interdomain Routing (CIDR). You can specify the subnet mask (*netmask*) for a destination network using the number of subnet bits or using the subnet mask in dotted decimal format. If no subnet mask is specified, the default (classful) mask is used.

The switch forwards IP traffic generated by the switch using the longest address match in the IP routing table. The switch does not use the IP routing table to forward traffic from connected devices, only IP traffic generated by the switch itself (for example, Telnet, TFTP, and ping).

To configure a static route, perform this task in privileged mode:

	Task	Command
Step 1	Configure a static route to the remote network.	set ip route <i>destination</i> [/ <i>netmask</i>] <i>gateway</i> [<i>metric</i>]
Step 2	Verify that the static route appears correctly in the IP routing table.	show ip route

This example shows how to configure a static route on the switch and how to verify that the route is configured properly in the routing table:

```

Console> (enable) set ip route 172.16.16.0/20 172.20.52.127
Route added.
Console> (enable) show ip route
Fragmentation  Redirect  Unreachable
-----
enabled        enabled   enabled

```

```

The primary gateway: 172.20.52.121
Destination      Gateway          RouteMask        Flags    Use      Interface
-----
172.16.16.0     172.20.52.127  0xfffff000      UG       0        sc0
default         172.20.52.121  0x0             UG       0        sc0
172.20.52.120  172.20.52.124  0xfffffffff8    U        1        sc0
default         default         0xff000000      UH       0        sl0
Console> (enable)

```

Configuring Permanent and Static ARP Entries

To enable your Catalyst LAN switch to communicate with devices that do not respond to Address Resolution Protocol (ARP) requests, you can configure a static or permanent ARP entry that maps the IP addresses of those devices to their MAC addresses. You can configure an ARP entry so that it does not age out by configuring it as either static or permanent. When you configure a static ARP entry using the **set arp static** command, the entry is removed from the ARP cache after a system reset. When you configure a permanent ARP by using the **set arp permanent** command, the ARP entry is retained even after a system reset.

Because most hosts support dynamic resolution, you usually do not need to specify static or permanent ARP cache entries. When a device does not respond to ARP requests, you can configure an ARP entry to be statically or permanently entered into the ARP cache so that those devices can still be reached.

To configure a static or permanent ARP entry, perform this task in privileged mode:

	Task	Command
Step 1	Configure a static or permanent ARP entry.	set arp [dynamic permanent static] {ip_addr hw_addr}
Step 2	(Optional) Specify the ARP aging time.	set arp agingtime seconds
Step 3	Verify the ARP configuration.	show arp

This example shows how to define a static ARP entry:

```

Console> (enable) set arp static 20.1.1.1 00-80-1c-93-80-40
Static ARP entry added as
20.1.1.1 at 00-80-1c-93-80-40 on vlan 1
Console> (enable)

```

This example shows how to define a permanent ARP entry:

```

Console> (enable) set arp permanent 10.1.1.1 00-80-1c-93-80-60
Permanent ARP entry added as
10.1.1.1 at 00-80-1c-93-80-60 on vlan 1
Console> (enable)

```

This example shows how to set the ARP aging time:

```

Console> (enable) set arp agingtime 300
ARP aging time set to 300 seconds.
Console> (enable)

```

This example shows how to display the ARP cache:

```
Console> (enable) show arp
ARP Aging time = 300 sec
+ - Permanent Arp Entries
* - Static Arp Entries
+ 10.1.1.1 at 00-80-1c-93-80-60 on vlan 1
172.20.52.1 at 00-60-5c-86-5b-28 port 8/1 on vlan 1
* 20.1.1.1 at 00-80-1c-93-80-40 port 8/1 on vlan 1
Console> (enable)
```

To clear ARP entries, perform this task in privileged mode:

	Task	Command
Step 1	Clear a dynamic, static, or permanent ARP entry.	clear arp [dynamic permanent static] {ip_addr hw_addr}
Step 2	Verify the ARP configuration.	show arp

This example shows how to clear all permanent ARP entries and verify the configuration:

```
Console> (enable) clear arp permanent
Permanent ARP entries cleared.
Console> (enable)
Console> (enable) show arp
ARP Aging time = 300 sec
+ - Permanent Arp Entries
* - Static Arp Entries
172.20.52.1 at 00-60-5c-86-5b-28 port 8/1 on vlan 1
* 20.1.1.1 at 00-80-1c-93-80-40 port 8/1 on vlan 1
Console> (enable)
```

Scheduling a System Reset

These sections describe how to schedule a system reset:

- [Scheduling a Reset at a Specific Time, page 20-10](#)
- [Scheduling a Reset Within a Specified Amount of Time, page 20-10](#)

You can use the **schedule reset** command to schedule a system to reset at a future time. This feature allows you to upgrade software during business hours and schedule the system upgrade after business hours to avoid a major impact on users.

You can also use the schedule reset feature when trying out new features on a switch. To avoid misconfiguration or the possibility of losing network connectivity to the device, you can set up the startup configuration feature and schedule a reset to occur in 30 minutes. You can then change the configuration, and if connectivity is lost, the system will reset in 30 minutes and return to the previous configuration.

Scheduling a Reset at a Specific Time

You can specify an absolute time and date at which the reset should take place with the **reset at** command. Entering the month and day argument with this command is optional. If you do not specify the month and day, the reset will take place on the current day if the time specified is later than the current time. If the time scheduled for reset is earlier than the current time, the reset will take place on the following day.



Note The maximum scheduled reset time is 24 days.

To schedule a reset at a specific time, perform this task in privileged mode:

	Task	Command
Step 1	Schedule the reset time at a specific time.	reset [mindown] at {hh:mm} [mm/dd] [reason]
Step 2	Verify the scheduled reset.	show reset



Note The minimum downtime argument is valid only if the system has a standby supervisor engine.

This example shows how to schedule a reset at a specific time:

```
Console> (enable) reset at 20:00
Reset scheduled at 20:00:00, Wed Aug 18 1999.
Proceed with scheduled reset? (y/n) [n]? y
Reset scheduled for 20:00:00, Wed Aug 18 1999 (in 0 day 5 hours 40 minutes).
Console> (enable)
```

This example shows how to schedule a reset at a specific time and include a reason for the reset:

```
Console> (enable) reset at 23:00 8/18 Software upgrade to 5.3(1).
Reset scheduled at 23:00:00, Wed Aug 18 1999.
Reset reason: Software upgrade to 5.3(1).
Proceed with scheduled reset? (y/n) [n]? y
Reset scheduled for 23:00:00, Wed Aug 18 1999 (in 0 day 8 hours 39 minutes).
Console> (enable)
```

This example shows how to schedule a reset with a minimum downtime:

```
Console> (enable) reset mindown at 23:00 8/18 Software upgrade to 5.3(1).
Reset scheduled at 23:00:00, Wed Aug 18 1999.
Reset reason: Software upgrade to 5.3(1).
Proceed with scheduled reset? (y/n) [n]? y
Reset mindown scheduled for 23:00:00, Wed Aug 18 1999 (in 0 day 8 hours 39 minutes).
Console> (enable)
```

Scheduling a Reset Within a Specified Amount of Time

You can schedule a reset within a specified time with the **reset in** command. For instance, if the current system time is 9:00 a.m. and reset is scheduled in one hour, the scheduled reset will take place at 10:00 a.m. If you or NTP advances the system clock to 10:00 a.m., the reset will take place at 11:00 a.m. If the clock is advanced ahead of the scheduled reset time, the reset will take place 5 minutes after the current time.

To schedule a reset within a specified time, perform this task in privileged mode:

	Task	Command
Step 1	Schedule the reset time within a specific amount of time.	reset [mindown] in [hh] {mm} [reason]
Step 2	Verify the scheduled reset.	show reset



Note

The minimum downtime argument is valid only if the system has a standby supervisor engine.

This example shows how to schedule a reset in a specified time:

```
Console> (enable) reset in 5:20 Configuration update
Reset scheduled in 5 hours 20 minutes.
Reset reason: Configuration update
Proceed with scheduled reset? (y/n) [n]? y
Reset scheduled for 19:56:01, Wed Aug 18 1999 (in 5 hours 20 minutes).
Reset reason: Configuration update
Console> (enable)
```

Power Management

This section describes power management in the Catalyst 6000 family switches and includes the following information:

- [Enabling or Disabling Power Redundancy, page 20-11](#)
- [Using the CLI to Power Modules Up or Down, page 20-13](#)
- [Determining System Power Requirements, page 20-14](#)



Note

In systems with redundant power supplies, both power supplies must be of the same wattage. The Catalyst 6000 family switches allow you to mix AC-input and DC-input power supplies in the same chassis. For detailed information on supported power supply configurations for each chassis, refer to the *Catalyst 6000 Family Installation Guide*.

Catalyst 6000 family modules have different power requirements and, depending upon the wattage of the power supply, certain switch configurations might require more power than a single power supply can provide. Although the power management feature allows you to power all installed modules with two power supplies, redundancy is not supported in this configuration. Redundant and nonredundant power configurations are discussed in the following sections.

Enabling or Disabling Power Redundancy

Enter the **set power redundancy enable | disable** command to enable or disable redundancy (redundancy is enabled by default). With redundancy enabled and two power supplies of equal wattage installed, the total power drawn from both supplies is at no time greater than the capability of one supply. If one supply malfunctions, the other supply can take over the entire system load. When you install and

turn on two power supplies of equal wattage, each concurrently provides approximately half of the required power to the system. Load sharing and redundancy are enabled automatically; no software configuration is required.

With redundancy enabled, if you power up the system with two power supplies of unequal wattage, both power supplies come online but a syslog message displays that the lower wattage power supply will be disabled. If the active power supply fails, the lower wattage power supply that was disabled comes online and, if necessary, modules are powered down to accommodate the lower wattage power supply.

In a nonredundant configuration, the power available to the system is the combined power capability of both power supplies. The system powers up as many modules as the combined capacity allows. However, if one supply should fail and there is not enough power for all previously powered up modules, the system powers down some modules. These modules are marked as *power-deny* in the **show module** Status field.

You can change the configuration of the power supplies to redundant or nonredundant at any time. If you switch from a redundant to a nonredundant configuration, both power supplies are enabled (even a power supply that was disabled because it was of a lower wattage than the other power supply). If you change from a nonredundant to a redundant configuration, both power supplies are initially enabled, and if they are of the same wattage, remain enabled. If they are of different wattage, a syslog message displays and the lower wattage supply is disabled.

Table 20-1 describes how the system responds to changes in the power supply configuration.

Table 20-1 Effects of Power Supply Configuration Changes

Configuration Change	Effect
Redundant to nonredundant	<ul style="list-style-type: none"> System log and syslog messages are generated. System power is increased to the combined power capability of both supplies. The modules marked as <i>power-deny</i> in the show module Status field are brought up if there is sufficient power.
Nonredundant to redundant	<ul style="list-style-type: none"> System log and syslog messages are generated. System power is the power capability of the larger wattage supply. If there is not enough power for all previously powered-up modules, some modules are powered down and marked as <i>power-deny</i> in the show module Status field.
Equal wattage power supply is inserted with redundancy enabled	<ul style="list-style-type: none"> System log and syslog messages are generated. System power equals the power capability of one supply. No change in the module status because the power capability is unchanged.
Equal wattage power supply is inserted with redundancy disabled	<ul style="list-style-type: none"> System log and syslog messages are generated. System power is the combined power capability of both supplies. The modules marked as <i>power-deny</i> in the show module Status field are brought up if there is sufficient power.
Higher wattage power supply is inserted with redundancy enabled	<ul style="list-style-type: none"> System log and syslog messages are generated. The system disables the lower wattage power supply; the higher wattage supply powers the system.

Table 20-1 Effects of Power Supply Configuration Changes (continued)

Configuration Change	Effect
Lower wattage power supply is inserted with redundancy enabled	<ul style="list-style-type: none"> System log and syslog messages are generated. The system disables the lower wattage power supply; the higher wattage supply powers the system.
Higher or lower wattage power supply is inserted with redundancy disabled	<ul style="list-style-type: none"> System log and syslog messages are generated. System power is increased to the combined power capability of both supplies. The modules marked as <i>power-deny</i> in the show module Status field are brought up if there is sufficient power.
Power supply is removed with redundancy enabled	<ul style="list-style-type: none"> System log and syslog messages are generated. If the power supplies are of equal wattage, there is no change in the module status because the power capability is unchanged. <p>If the power supplies are of unequal wattage and the lower wattage supply is removed, there is no change in the module status.</p> <p>If the power supplies are of unequal wattage and the higher wattage supply is removed, and if there is not enough power for all previously powered-up modules, some modules are powered down and marked as <i>power-deny</i> in the show module Status field.</p>
Power supply is removed with redundancy disabled	<ul style="list-style-type: none"> System log and syslog messages are generated. System power is decreased to the power capability of one supply. If there is not enough power for all previously powered-up modules, some modules are powered down and marked as <i>power-deny</i> in the show module Status field.
System is booted with power supplies of different wattage installed and redundancy enabled	<ul style="list-style-type: none"> System log and syslog messages are generated. The lower wattage supply is disabled.
System is booted with power supplies of equal or different wattage installed and redundancy disabled	<ul style="list-style-type: none"> System log and syslog messages are generated. System power equals the combined power capability of both supplies. The system powers up as many modules as the combined capacity allows.

Using the CLI to Power Modules Up or Down

You can power down a properly working module from the command-line interface (CLI) by entering the **set module power down mod** command. The module is marked as *power-down* in the **show module** Status field. Enter the **set module power up mod** command to check if adequate power is available in the system to turn the power on for a module that was previously powered down. If not enough power is available, the module status changes from *power-down* to *power-deny*.

Determining System Power Requirements

This section describes how to determine the system power requirements for 6-, 9-, and 13-slot chassis. See [Table 20-2](#) to determine the exact power requirements for your configuration.



Note

Enter the **show environment power** command to display current system power usage.

Table 20-2 Module Power Requirements

Module	Power Requirement
Supervisor Engine 1:	
WS-X6K-SUP1A-2GE	1.70A
WS-X6K-SUP1-2GE	1.70A
Supervisor Engine 1 with PFC:	
WS-X6K-SUP1A-PFC	2.50A
Supervisor Engine 1 with PFC and MSFC:	
WS-X6K-SUP1A-MSFC	3.30A
Supervisor Engine 1 with PFC and MSFC2:	
WS-X6K-S1A-MSFC2	2.90A
Supervisor Engine 2 with PFC2:	
WS-X6K-S2-PFC2	3.06A
Supervisor Engine 2 with PFC2 and MSFC2:	
WS-X6K-S2-MSFC2	3.46A
MSFC2 (spare):	
WS-F6K-MSFC2=	0.40A
Multilayer Switching Module:	
WS-X6302-MSM	5.20A
24-Port 10BASE-FL:	
WS-X6024-10FL-MT	1.52A
Switch Fabric Modules:	
WS-C6500-SFM	2.79A
WS-X6500-SFM2	3.09A
24-Port 100FX:	
WS-X6224-100FX-MT	1.90A
WS-X6324-100FX-SM	1.52A
WS-X6324-100FX-MM	1.52A

Table 20-2 Module Power Requirements (continued)

Module	Power Requirement
48-Port 10/100TX:	
WS-X6248-RJ-45	2.69A
WS-X6248-TEL	2.69A
WS-X6248A-TEL	2.69A
WS-X6348-RJ-45	2.39A
WS-X6548-RJ-45	2.90A
WS-X6648-PWR	2.39A
8-Port Gigabit Ethernet:	
WS-X6408-GBIC	2.00A
WS-X6408A-GBIC	2.00A
16-Port Gigabit Ethernet:	
WS-X6416-GBIC	2.81A
WS-X6416-GE-MT	2.50A
WS-X6316-GE-TX	5.15A
WS-X6516-GE-TX	3.45A
1-Port OC-12 ATM:	
WS-X6101-OC12-MMF	2.10A
WS-X6101-OC12-SMF	2.10A
WAN module:	
WS-X6182-2PA (FlexWAN)	2.38A ¹
Optical Services Modules:	
OSM-2OC12-POS-MM, -SI, -SL	3.36A
OSM-4OC12-POS-MM, -SI, -SL	4.78A
OSM-8OC3-POS-MM, -SI, -SL	3.57A
OSM-16OC3-POS-MM, -SI, -SL	5.09A
OSM-10C48-POS-SS, -SI, -SL	4.25A
OSM-4GE-WAN (GBIC)	3.59
Server load balancing:	
WS-X6066-SLB-APG	3.00A
8-Port T1/E1 PSTN Interface:	
WS-X6608-T1	1.98A
WS-X6608-E1	1.98A
24-Port FXS Analog Interface:	
WS-X6624-FXS	1.54A
Cisco IP Phone 7960 (when plugged into the WS-X6348-RJ-45 and WS-X6648-PWR modules)	0.167A (default) 0.120A (after bootup, initialization)

The total power available with the 4000W power supply is 95.70A.. The total power available with the 2500W power supply is 55.50A. The total power available with the 1300W power supply is 27.46A. The total power available with the 1000W power supply is 21.40A.

1. Based on the base FlexWAN module power draw plus a worst-case 15W per port adapter, plus margin.

Environmental Monitoring

Environmental monitoring of chassis components provides early warning indications of possible component failure to ensure safe and reliable system operation and avoid network interruptions. This section describes how to monitor these critical system components, enabling you to identify and rapidly correct hardware-related problems in your system.

The following sections describe the environmental monitors:

- [Environmental Monitoring Using CLI Commands, page 20-16](#)
- [LED Indications, page 20-16](#)

Environmental Monitoring Using CLI Commands

Enter the **show test** [*mod*] command to display the errors reported from the diagnostic tests. If you do not specify a module number, test statistics are given for the general system and for the module in slot 1. If there are no errors, PASS is displayed in the Line Card Status field.

Enter the **show environment** [**temperature** | **all** | **power**] command to display system status information. Keyword descriptions follow:

- **temperature**—(Optional) Displays temperature information.
- **all**—(Optional) Displays environmental status (for example, power supply, fan status, and temperature information) and information about the power available to the system.
- **power**—(Optional) Displays environmental power information.

LED Indications

There are two alarm types, major and minor. Major alarms indicate a critical problem that could lead to the system being shut down. Minor alarms are for informational purposes only, giving you notice of a problem that could turn critical if corrective action is not taken.

When the system has an alarm (major or minor), indicating an overtemperature condition, the alarm is not canceled or any action taken (such as a module reset or shutdown) for 5 minutes. If the temperature falls 5°C (41°F) below the alarm threshold during this period, the alarm is canceled.

[Table 20-3](#) lists the environmental indicators for the supervisor engine and switching modules.

**Note**

For additional information on LED indications, refer to the *Catalyst 6000 Family Module Installation Guide*.

Table 20-3 Environmental Monitoring for Supervisor Engine and Switching Modules

Component	Alarm Type	LED Indication	Action
Supervisor engine temperature sensor exceeds major threshold ¹	Major	STATUS ² LED red ³	syslog message and SNMP trap generated. If redundancy, system switches to redundant supervisor engine and the active supervisor engine shuts down. If there is no redundancy and the overtemperature condition is not corrected, the system shuts down after 5 minutes.
Supervisor engine temperature sensor exceeds minor threshold	Minor	STATUS LED orange	syslog message and SNMP trap generated. Monitor the condition.
Redundant supervisor engine temperature sensor exceeds major or minor threshold	Major	STATUS LED red	syslog message and SNMP trap generated. If major alarm and the overtemperature condition is not corrected, the system shuts down after 5 minutes.
	Minor	STATUS LED orange	If minor alarm, monitor the condition.
Switching module temperature sensor exceeds major threshold	Major	STATUS LED red	syslog message and SNMP trap generated. Power down the module ⁴ .
Switching module temperature sensor exceeds minor threshold	Minor	STATUS LED orange	syslog message and SNMP trap generated. Monitor the condition.

1. Temperature sensors monitor key supervisor engine components including daughter cards.
2. A STATUS LED is located on the supervisor engine front panel and all module front panels.
3. The STATUS LED is red on the failed supervisor engine. If there is no redundant supervisor, the SYSTEM LED is red also.
4. See the “Power Management” section on page 20-11 for instructions.

Displaying System Status Information for Technical Support

These sections describe how to display system status information for technical support:

- [Generating a System Status Report, page 20-18](#)
- [Using System Dump Files, page 20-18](#)

Generating a System Status Report

Using a single command, you can generate a report that contains status information about your switch. The information generated is useful when reporting a problem to Cisco Technical Assistance Center (TAC). This command is a combination of several **show system status** commands. You can upload the output of the command to a TFTP server, where you can send it to TAC.

You can use keywords to limit the output to certain areas, such as specific modules, VLANs, ports, and so forth. If you do not specify any keywords, a report for the entire system is generated.

To generate a report and upload the report to a TFTP server, perform this task in privileged mode:

Task	Command
Generate a system status report that you can send to TAC.	write tech-support { <i>host</i> } { <i>filename</i> } [module <i>mod</i>] [port <i>mod/port</i>] [vlan <i>vlan</i>] [memory] [config]

This example shows a report sent to host 172.20.32.10 to a filename you supply. No keywords are specified, so the complete status of the switch will be included in the report.

```
Console> (enable) write tech-support 172.20.32.10 tech.txt
Upload tech-report to tech.txt on 172.20.32.10 (y/n) [n]? y
/
Finished network upload. (67784 bytes)
Console> (enable)
```

Using System Dump Files

The core dump and the stack dump features generate reports that contain status information about your switch. Send images captured by the core dump or the stack dump to the Cisco TAC for analysis.

Enabling and Disabling the Core Dump

A core dump produces a comprehensive report of images when your system fails due to a software error. This report contains system memory content, including text, code, and stack segments. The core image is produced in Cisco core file format and is stored in the file system. By examining the core dump file, TAC can analyze the error condition of a terminated process.

Enter the **set system core-dump** command to enable or disable the core dump feature. If the switch has a redundant supervisor engine, the standby supervisor engine takes over automatically before the core dump occurs. The previously active supervisor engine resets itself after the core dump is complete.

To enable or disable the core dump feature, perform this task in privileged mode:

Task	Command
Enable or disable the core dump feature.	set system core-dump { enable disable }

This example shows how to enable the core dump feature:

```
Console> (enable) set system core-dump enable
(1) In the event of a system crash, this feature will
    cause a core file to be written out.
(2) Core file generation may take up to 20 minutes.
(3) Selected core file is slot0:crash.hz
(4) Please make sure the above device has been installed,
    and ready to use
Core-dump enabled
Console> (enable)
```

This example shows how to disable the core dump feature:

```
Console> (enable) set system core-dump disable
Core-dump disabled
Console> (enable)
```

The size of the file system depends on the size of your memory card. An error process will generate a core image that is proportional to the size of the system DRAM. Make sure that you have enough memory available to store the core dump file.

Specifying the Core Image Filename

Enter the **set system core-file** command to specify the core image filename. The default filename is “slot0:crash.hz.” This command automatically checks the validity of the device name that you input.

To specify the core image filename, perform this task in privileged mode:

Task	Command
Specify the core image filename.	set system core-file { <i>device:filename</i> }

This example shows how to specify the core image filename:

```
Console> (enable) set system core-file slot0:core.hz
System core-file set.
Console> (enable)
```

Displaying the Stack Dump

A stack dump provides only the images related to a particular process that has caused the system to fail. This image stack is displayed on the console and is also saved in the log area. The stack dump is automatic and becomes available when you enter the **show log** command after you reboot your system.

To display log information, perform this task in normal mode:

Task	Command
Display the stack dump.	show log

The following is an example of an image stack that may display after you enter the **show log** command:

```

Breakpoint Exception occurred.
Software version = 6.2(0.83)
Process ID #52, Name
= Console
      EPC: 807523F4
Stack content:
sp+00: 00000000 80A75698 00000005 00000005
sp+10: BE000A00 00000000 83F84150 801194B8
sp+20: 80A75698 80A74BC8 80C8DBDC 000006E8
sp+30: 8006AF30 8006AE98 82040664 00000630
sp+40: 801AC744 801AC734 80A32488 80A32484
sp+50: 80A3249C 00000000 00000002 000009E4
sp+60: 8204067B 82040670 8011812C 81CAFC98
sp+70: 8011814C 82040670 8011812C 81CAFC98
sp+80: 00000002 000009E4 80110160 80110088
sp+90: 82040670 80A71EB4 81F1E9F8 00000004
sp+A0: 00000000 81F25EAC 81FF5750 00000000
sp+B0: 00000000 00000000 81F1E314 800840BC
sp+C0: 0000000B 80084EB0 00000001 8073A358
sp+D0: 00000003 0000000D 00000000 0000000A
sp+E0: 00000020 00000000 800831B4 0000001A
sp+F0: 00000000 00000000 00000000 000D84F0
Register content:
      Status: 3401FC23      Cause: 00000024
AT: 81640000
      V0: 00000007      V1: 00000007
      A0: 00000000      A1: 80A756A6
      A2: 00000011      A3: BE000BD0
      T0: BFFFFFFE      T1: 80000000
      T2: 00000000      T3: 00000001
      T4: 00000000      T5: 00000007
      T6: 00000000      T7: 00000000
      S0: 00000001      S1: 00000032
      S2: 81F1E9F8      S3: 80A74BC8
      S4: 80C8DBDC      S5: 000006E8
      S6: 00000000      S7: 00000000
      T8: F0D09E3A      T9: 82940828
      K0: 3041C001      K1: 80C73038
      GP: 811F39C0      SP: 83F84010
      S8: 83F84010      RA: 807523F4
      HIGH: 00000001    LOW: D5555559
      BADVADDR: 7DFF7FFF  ERR EPC: 58982466
GDB: Breakpoint Exception
GDB: The system has trapped into the debugger.
GDB: It will hang until examined with gdb.

```