



CHAPTER 6

Clustering Switches

This chapter provides the concepts and procedures to create and manage Catalyst 2960 and 2960-S switch clusters. Unless otherwise noted, the term *switch* refers to a standalone switch and to a switch stack.

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

You can create and manage switch clusters by using Cisco Network Assistant (hereafter known as Network Assistant), the command-line interface (CLI), or SNMP. For complete procedures, see the online help. For the CLI cluster commands, see the switch command reference.

**Note**

Network Assistant supports switch clusters, but we recommend that you instead group switches into *communities*. Network Assistant has a Cluster Conversion Wizard to help you convert a cluster to a *community*. For more information about Network Assistant, including introductory information on managing switch clusters and converting a switch cluster to a community, see *Getting Started with Cisco Network Assistant*, available on Cisco.com.

This chapter focuses on Catalyst 2960 and 2960-S switch clusters. It also includes guidelines and limitations for clusters mixed with other cluster-capable Catalyst switches, but it does not provide complete descriptions of the cluster features for these other switches. For complete cluster information for a specific Catalyst platform, refer to the software configuration guide for that switch.

This chapter consists of these sections:

- [Understanding Switch Clusters, page 6-2](#)
- [Planning a Switch Cluster, page 6-5](#)
- [Using the CLI to Manage Switch Clusters, page 6-17](#)
- [Using SNMP to Manage Switch Clusters, page 6-18](#)

**Note**

We do not recommend using the **ip http access-class** global configuration command to limit access to specific hosts or networks. Access should be controlled through the cluster command switch or by applying access control lists (ACLs) on interfaces that are configured with IP address. For more information on ACLs, see [Chapter 33, “Configuring Network Security with ACLs.”](#)

Understanding Switch Clusters

A *switch cluster* is a set of up to 16 connected, cluster-capable Catalyst switches that are managed as a single entity. The switches in the cluster use the switch clustering technology so that you can configure and troubleshoot a group of different Catalyst desktop switch platforms through a single IP address.

In a switch cluster, 1 switch must be the *cluster command switch* and up to 15 other switches can be *cluster member switches*. The total number of switches in a cluster cannot exceed 16 switches. The cluster command switch is the single point of access used to configure, manage, and monitor the cluster member switches. Cluster members can belong to only one cluster at a time.



Note

A switch cluster is different from a *switch stack*. A switch stack is a set of Catalyst 2960-S switches connected through their stack ports. For more information about how switch stacks differ from switch clusters, see the “[Switch Clusters and Switch Stacks](#)” section on page 6-15. Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

The benefits of clustering switches include:

- Management of Catalyst switches regardless of their interconnection media and their physical locations. The switches can be in the same location, or they can be distributed across a Layer 2 or Layer 3 (if your cluster is using a Catalyst 3550, Catalyst 3560, or Catalyst 3750 switch as a Layer 3 router between the Layer 2 switches in the cluster) network.

Cluster members are connected to the cluster command switch according to the connectivity guidelines described in the “[Automatic Discovery of Cluster Candidates and Members](#)” section on page 6-5. This section includes management VLAN considerations for the Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

- Command-switch redundancy if a cluster command switch fails. One or more switches can be designated as *standby cluster command switches* to avoid loss of contact with cluster members. A *cluster standby group* is a group of standby cluster command switches.
- Management of a variety of Catalyst switches through a single IP address. This conserves on IP addresses, especially if you have a limited number of them. All communication with the switch cluster is through the cluster command switch IP address.

[Table 6-1](#) lists the Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and the required software versions.

Table 6-1 Switch Software and Cluster Capability

Switch	Cisco IOS Release	Cluster Capability
Catalyst 3750-E or Catalyst 3560-E	12.2(35)SE2 or later	Member or command switch
Catalyst 3750	12.1(11)AX or later	Member or command switch
Catalyst 3560	12.1(19)EA1b or later	Member or command switch
Catalyst 3550	12.1(4)EA1 or later	Member or command switch
Catalyst 2975	12.2(46)EX or later	Member or command switch
Catalyst 2970	12.1(11)AX or later	Member or command switch
Catalyst 2960-S	12.2(53)SE or later	Member or command switch

Table 6-1 *Switch Software and Cluster Capability (continued)*

Switch	Cisco IOS Release	Cluster Capability
Catalyst 2960	12.2(25)FX or later	Member or command switch
Catalyst 2955	12.1(12c)EA1 or later	Member or command switch
Catalyst 2950	12.0(5.2)WC(1) or later	Member or command switch
Catalyst 2950 LRE	12.1(11)JY or later	Member or command switch
Catalyst 2940	12.1(13)AY or later	Member or command switch
Catalyst 3500 XL	12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (8-MB switches)	12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (4-MB switches)	11.2(8.5)SA6 (recommended)	Member switch only
Catalyst 1900 and 2820	9.00(-A or -EN) or later	Member switch only

Cluster Command Switch Characteristics

A cluster command switch must meet these requirements:

- It is running Cisco IOS Release 12.2(25)FX or later for a Catalyst 2960 switch, or Cisco IOS Release 12.2(53)SE or later for a Catalyst 2960-S switch.
- It has an IP address.
- It has Cisco Discovery Protocol (CDP) version 2 enabled (the default).
- It is not a command or cluster member switch of another cluster.
- It is connected to the standby cluster command switches through the management VLAN and to the cluster member switches through a common VLAN.

Standby Cluster Command Switch Characteristics

A standby cluster command switch must meet these requirements:

- It is running Cisco IOS 12.2(25)FX or later for a Catalyst 2960 switch, or Cisco IOS Release 12.2(53)SE or later for a Catalyst 2960-S switch.
- It has an IP address.
- It has CDP version 2 enabled.
- It is connected to the command switch and to other standby command switches through its management VLAN.
- It is connected to all other cluster member switches (except the cluster command and standby command switches) through a common VLAN.
- It is redundantly connected to the cluster so that connectivity to cluster member switches is maintained.
- It is not a command or member switch of another cluster.

**Note**

Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2960 switch, the standby cluster command switches must also be Catalyst 2960 switches. If the cluster command switch is a Catalyst 2960-S switch, the standby cluster command switches must also be Catalyst 2960-S switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

Candidate Switch and Cluster Member Switch Characteristics

Candidate switches are cluster-capable switches and switch stacks that have not yet been added to a cluster. Cluster member switches are switches and switch stacks that have actually been added to a switch cluster. Although not required, a candidate or cluster member switch can have its own IP address and password (for related considerations, see the “[IP Addresses](#)” section on page 6-13 and “[Passwords](#)” section on page 6-14).

**Note**

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

To join a cluster, a candidate switch must meet these requirements:

- It is running cluster-capable software.
- It has CDP version 2 enabled.
- It is not a command or cluster member switch of another cluster.
- If a cluster standby group exists, it is connected to every standby cluster command switch through at least one common VLAN. The VLAN to each standby cluster command switch can be different.
- It is connected to the cluster command switch through at least one common VLAN.

**Note**

Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL candidate and cluster member switches must be connected through their management VLAN to the cluster command switch and standby cluster command switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

This requirement does not apply if you have a Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switch. Candidate and cluster member switches can connect through any VLAN in common with the cluster command switch.

Planning a Switch Cluster

Anticipating conflicts and compatibility issues is a high priority when you manage several switches through a cluster. This section describes these guidelines, requirements, and caveats that you should understand before you create the cluster:

- [Automatic Discovery of Cluster Candidates and Members, page 6-5](#)
- [HSRP and Standby Cluster Command Switches, page 6-10](#)
- [IP Addresses, page 6-13](#)
- [Hostnames, page 6-14](#)
- [Passwords, page 6-14](#)
- [SNMP Community Strings, page 6-14](#)
- [Switch Clusters and Switch Stacks, page 6-15](#)
- [TACACS+ and RADIUS, page 6-16](#)
- [LRE Profiles, page 6-16](#)

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and for the required software versions and browser and Java plug-in configurations.

Automatic Discovery of Cluster Candidates and Members

The cluster command switch uses Cisco Discovery Protocol (CDP) to discover cluster member switches, candidate switches, neighboring switch clusters, and edge devices across multiple VLANs and in star or cascaded topologies.

**Note**

Do not disable CDP on the cluster command switch, on cluster members, or on any cluster-capable switches that you might want a cluster command switch to discover. For more information about CDP, see [Chapter 26, “Configuring CDP.”](#)

Following these connectivity guidelines ensures automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices:

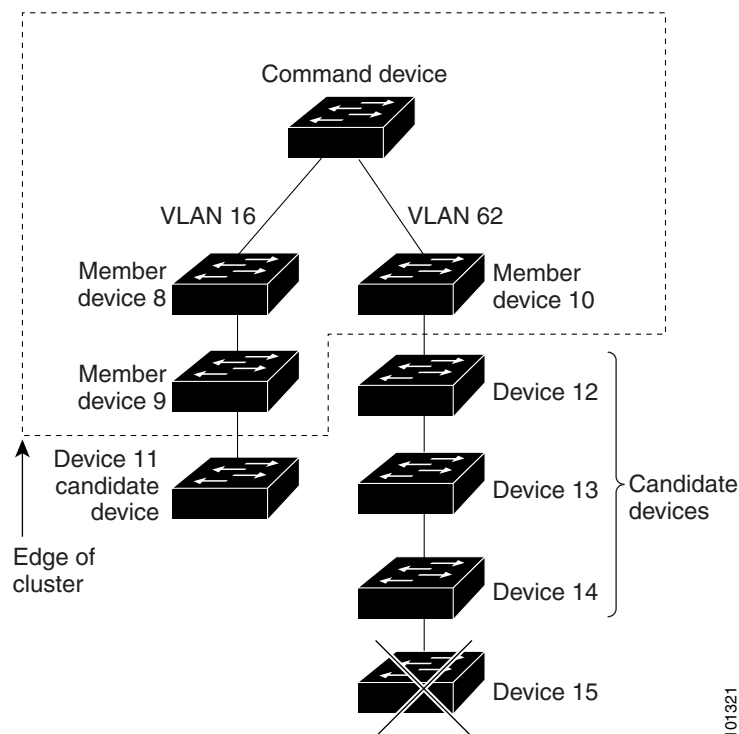
- [Discovery Through CDP Hops, page 6-6](#)
- [Discovery Through Non-CDP-Capable and Noncluster-Capable Devices, page 6-7](#)
- [Discovery Through Different VLANs, page 6-7](#)
- [Discovery Through Different Management VLANs, page 6-8](#)
- [Discovery of Newly Installed Switches, page 6-9](#)

Discovery Through CDP Hops

By using CDP, a cluster command switch can discover switches up to seven CDP hops away (the default is three hops) from the edge of the cluster. The edge of the cluster is where the last cluster member switches are connected to the cluster and to candidate switches. For example, cluster member switches 9 and 10 in [Figure 6-1](#) are at the edge of the cluster.

In [Figure 6-1](#), the cluster command switch has ports assigned to VLANs 16 and 62. The CDP hop count is three. The cluster command switch discovers switches 11, 12, 13, and 14 because they are within three hops from the edge of the cluster. It does not discover switch 15 because it is four hops from the edge of the cluster.

Figure 6-1 Discovery Through CDP Hops




Discovery Through Non-CDP-Capable and Noncluster-Capable Devices

If a cluster command switch is connected to a *non-CDP-capable third-party hub* (such as a non-Cisco hub), it can discover cluster-enabled devices connected to that third-party hub. However, if the cluster command switch is connected to a *noncluster-capable Cisco device*, it cannot discover a cluster-enabled device connected beyond the noncluster-capable Cisco device.

Figure 6-2 shows that the cluster command switch discovers the switch that is connected to a third-party hub. However, the cluster command switch does not discover the switch that is connected to a Catalyst 5000 switch.

Figure 6-2 Discovery Through Non-CDP-Capable and Noncluster-Capable Devices



Discovery Through Different VLANs


If the cluster command switch is a Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 switch, the cluster can have cluster member switches in different VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. The cluster command switch in Figure 6-3 has ports assigned to VLANs 9, 16, and 62 and therefore discovers the switches in those VLANs. It does not discover the switch in VLAN 50. It also does not discover the switch in VLAN 16 in the first column because the cluster command switch has no VLAN connectivity to it.

Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster command switch through their management VLAN. For information about discovery through management VLANs, see the “Discovery Through Different Management VLANs” section on page 6-8. For more information about VLANs, see Chapter 14, “Configuring VLANs.”



Note

For additional considerations about VLANs in switch stacks, see the “Switch Clusters and Switch Stacks” section on page 6-15.

Figure 6-3 Discovery Through Different VLANs

Discovery Through Different Management VLANs

Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches can discover and manage cluster member switches in different VLANs and different management VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. They do not need to be connected to the cluster command switch through their management VLAN. The default management VLAN is VLAN 1.




Note

If the switch cluster has a Catalyst 3750 or 2975 switch or has a switch stack, that switch or switch stack must be the cluster command switch.

The cluster command switch and standby command switch in [Figure 6-4](#) (assuming they are Catalyst 2960, Catalyst 2970, Catalyst 2975, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches) have ports assigned to VLANs 9, 16, and 62. The management VLAN on the cluster command switch is VLAN 9. Each cluster command switch discovers the switches in the different management VLANs except these:

- Switches 7 and 10 (switches in management VLAN 4) because they are not connected through a common VLAN (meaning VLANs 62 and 9) with the cluster command switch
- Switch 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Figure 6-4 Discovery Through Different Management VLANs with a Layer 3 Cluster Command Switch



Discovery of Newly Installed Switches


To join a cluster, the new, out-of-the-box switch must be connected to the cluster through one of its access ports. An access port (AP) carries the traffic of and belongs to only one VLAN. By default, the new switch and its access ports are assigned to VLAN 1.

When the new switch joins a cluster, its default VLAN changes to the VLAN of the immediately upstream neighbor. The new switch also configures its access port to belong to the VLAN of the immediately upstream neighbor.

The cluster command switch in [Figure 6-5](#) belongs to VLANs 9 and 16. When new cluster-capable switches join the cluster:

- One cluster-capable switch and its access port are assigned to VLAN 9.
- The other cluster-capable switch and its access port are assigned to management VLAN 16.

Figure 6-5 Discovery of Newly Installed Switches



HSRP and Standby Cluster Command Switches

The switch uses Hot Standby Router Protocol (HSRP) so that you can configure a group of standby cluster command switches. Because a cluster command switch manages the forwarding of all communication and configuration information to all the cluster member switches, we strongly recommend the following:

- For a cluster command switch stack, a standby cluster command switch is necessary if the entire switch stack fails. However, if only the stack master in the command switch stack fails, the switch stack elects a new stack master and resumes its role as the cluster command switch stack.
- For a cluster command switch that is a standalone switch, configure a standby cluster command switch to take over if the primary cluster command switch fails.

A *cluster standby group* is a group of command-capable switches that meet the requirements described in the “[Standby Cluster Command Switch Characteristics](#)” section on page 6-3. Only one cluster standby group can be assigned per cluster.

The switches in the cluster standby group are ranked according to HSRP priorities. The switch with the highest priority in the group is the *active cluster command switch* (AC). The switch with the next highest priority is the *standby cluster command switch* (SC). The other switches in the cluster standby group are the *passive cluster command switches* (PC). If the active cluster command switch and the standby cluster command switch become disabled *at the same time*, the passive cluster command switch with the highest priority becomes the active cluster command switch. For the limitations to automatic discovery, see the “[Automatic Recovery of Cluster Configuration](#)” section on page 6-13.



Note

The HSRP standby hold time interval should be greater than or equal to three times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds.

These connectivity guidelines ensure automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices. These topics also provide more detail about standby cluster command switches:

- [Virtual IP Addresses, page 6-11](#)
- [Other Considerations for Cluster Standby Groups, page 6-11](#)
- [Automatic Recovery of Cluster Configuration, page 6-13](#)

Virtual IP Addresses

You need to assign a unique virtual IP address and group number and name to the cluster standby group. This information must be configured on a specific VLAN or routed port on the active cluster command switch. The active cluster command switch receives traffic destined for the virtual IP address. To manage the cluster, you must access the active cluster command switch through the virtual IP address, not through the command-switch IP address. This is in case the IP address of the active cluster command switch is different from the virtual IP address of the cluster standby group.

If the active cluster command switch fails, the standby cluster command switch assumes ownership of the virtual IP address and becomes the active cluster command switch. The passive switches in the cluster standby group compare their assigned priorities to decide the new standby cluster command switch. The passive standby switch with the highest priority then becomes the standby cluster command switch. When the previously active cluster command switch becomes active again, it resumes its role as the active cluster command switch, and the current active cluster command switch becomes the standby cluster command switch again. For more information about IP address in switch clusters, see the [“IP Addresses” section on page 6-13](#).

Other Considerations for Cluster Standby Groups



Note

For additional considerations about cluster standby groups in switch stacks, see the [“Switch Clusters and Switch Stacks” section on page 6-15](#).

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

These requirements also apply:

- Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2960 switch, the standby cluster command switches must also be Catalyst 2960 switches. If the cluster command switch is a Catalyst 2960-S switch, the standby cluster command switches must also be Catalyst 2960-S switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

If your switch cluster has a Catalyst 2960 switch or a Cisco FlexStack (a stack that contains only 2960-S switches), it should be the cluster command switch.

- Only one cluster standby group can be assigned to a cluster. You can have more than one router-redundancy standby group.

- All standby-group members must be members of the cluster.




Note There is no limit to the number of switches that you can assign as standby cluster command switches. However, the total number of switches in the cluster—which would include the active cluster command switch, standby-group members, and cluster member switches—cannot be more than 16.

- Each standby-group member (Figure 6-6) must be connected to the cluster command switch through the same VLAN. In this example, the cluster command switch and standby cluster command switches are Catalyst 2970, Catalyst 3550, Catalyst 3560, or Catalyst 3750 cluster command switches. Each standby-group member must also be redundantly connected to each other through at least one VLAN in common with the switch cluster.

Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster standby group through their management VLANs. For more information about VLANs in switch clusters, see these sections:

- “Discovery Through Different VLANs” section on page 6-7
- “Discovery Through Different Management VLANs” section on page 6-8

Figure 6-6 VLAN Connectivity between Standby-Group Members and Cluster Members



Automatic Recovery of Cluster Configuration

The active cluster command switch continually forwards cluster-configuration information (but not device-configuration information) to the standby cluster command switch. This ensures that the standby cluster command switch can take over the cluster immediately after the active cluster command switch fails.

Automatic discovery has these limitations:

- This limitation applies only to clusters that have Catalyst 2950, Catalyst 3550, Catalyst 3560, and Catalyst 3750 command and standby cluster command switches: If the active cluster command switch and standby cluster command switch become disabled *at the same time*, the passive cluster command switch with the highest priority becomes the active cluster command switch. However, because it was a passive standby cluster command switch, the previous cluster command switch *did not* forward cluster-configuration information to it. The active cluster command switch only forwards cluster-configuration information to the standby cluster command switch. You must therefore rebuild the cluster.
- This limitation applies to all clusters: If the active cluster command switch fails and there are more than two switches in the cluster standby group, the new cluster command switch does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL cluster member switches. You must re-add these cluster member switches to the cluster.
- This limitation applies to all clusters: If the active cluster command switch fails and becomes active again, it does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL cluster member switches. You must again add these cluster member switches to the cluster.

When the previously active cluster command switch resumes its active role, it receives a copy of the latest cluster configuration from the active cluster command switch, including members that were added while it was down. The active cluster command switch sends a copy of the cluster configuration to the cluster standby group.

IP Addresses

You must assign IP information to a cluster command switch. You can assign more than one IP address to the cluster command switch, and you can access the cluster through any of the command-switch IP addresses. If you configure a cluster standby group, you must use the standby-group virtual IP address to manage the cluster from the active cluster command switch. Using the virtual IP address ensures that you retain connectivity to the cluster if the active cluster command switch fails and that a standby cluster command switch becomes the active cluster command switch.

If the active cluster command switch fails and the standby cluster command switch takes over, you must either use the standby-group virtual IP address or any of the IP addresses available on the new active cluster command switch to access the cluster.

You can assign an IP address to a cluster-capable switch, but it is not necessary. A cluster member switch is managed and communicates with other cluster member switches through the command-switch IP address. If the cluster member switch leaves the cluster and it does not have its own IP address, you must assign an IP address to manage it as a standalone switch.

For more information about IP addresses, see [Chapter 3, “Assigning the Switch IP Address and Default Gateway.”](#)

Hostnames

You do not need to assign a host name to either a cluster command switch or an eligible cluster member. However, a hostname assigned to the cluster command switch can help to identify the switch cluster. The default hostname for the switch is *Switch*.

If a switch joins a cluster and it does not have a hostname, the cluster command switch appends a unique member number to its own hostname and assigns it sequentially as each switch joins the cluster. The number means the order in which the switch was added to the cluster. For example, a cluster command switch named *eng-cluster* could name the fifth cluster member *eng-cluster-5*.

If a switch has a hostname, it retains that name when it joins a cluster and when it leaves the cluster.

If a switch received its hostname from the cluster command switch, was removed from a cluster, was then added to a new cluster, and kept the same member number (such as 5), the switch overwrites the old hostname (such as *eng-cluster-5*) with the hostname of the cluster command switch in the new cluster (such as *mkg-cluster-5*). If the switch member number changes in the new cluster (such as 3), the switch retains the previous name (*eng-cluster-5*).

Passwords

You do not need to assign passwords to an individual switch if it will be a cluster member. When a switch joins a cluster, it inherits the command-switch password and retains it when it leaves the cluster. If no command-switch password is configured, the cluster member switch inherits a null password. Cluster member switches only inherit the command-switch password.

If you change the member-switch password to be different from the command-switch password and save the change, the switch is not manageable by the cluster command switch until you change the member-switch password to match the command-switch password. Rebooting the member switch does not revert the password back to the command-switch password. We recommend that you do not change the member-switch password after it joins a cluster.

For more information about passwords, see the [“Preventing Unauthorized Access to Your Switch” section on page 9-1](#).

For password considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

SNMP Community Strings

A cluster member switch inherits the command-switch first read-only (RO) and read-write (RW) community strings with *@esN* appended to the community strings:

- *command-switch-readonly-community-string@esN*, where *N* is the member-switch number.
- *command-switch-readwrite-community-string@esN*, where *N* is the member-switch number.

If the cluster command switch has multiple read-only or read-write community strings, only the first read-only and read-write strings are propagated to the cluster member switch.

The switches support an unlimited number of community strings and string lengths. For more information about SNMP and community strings, see [Chapter 31, “Configuring SNMP.”](#)

For SNMP considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides specific to those switches.

Switch Clusters and Switch Stacks

A *switch cluster* can have one or more Catalyst 2960-S switch stacks. Each switch stack can act as the cluster command switch or as a single cluster member. [Table 6-2](#) describes the basic differences between switch stacks and switch clusters. For more information about switch stacks, see [Chapter 7, “Managing Switch Stacks.”](#)



Note

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Table 6-2 Basic Comparison of Switch Stacks and Switch Clusters

Switch Stack	Switch Cluster
Made up of Catalyst 2960-S switches only	Made up of cluster-capable switches, such as Catalyst 3750, Catalyst 3550, and Catalyst 2960-S switches
Stack members are connected through StackWise ports	Cluster members are connected through LAN ports
Requires one <i>stack master</i> and supports up to four other <i>stack members</i>	Requires 1 <i>cluster command switch</i> and supports up to 15 other <i>cluster member switches</i>
Can be a cluster command switch or a cluster member switch	Cannot be a stack master or stack member
Stack master is the single point of <i>complete</i> management for all stack members in a particular switch stack	Cluster command switch is the single point of <i>some</i> management for all cluster members in a particular switch cluster
Back-up stack master is automatically determined in case the stack master fails	Standby cluster command switch must be pre-assigned in case the cluster command switch fails
Switch stack supports up to four simultaneous stack master failures	Switch cluster supports only one cluster command switch failure at a time
Stack members (as a switch stack) behave and is presented as a single, unified system in the network	Cluster members are various, independent switches that are not managed as and do not behave as a unified system
Integrated management of stack members through a single configuration file	Cluster members have separate, individual configuration files
Stack- and interface-level configurations are stored on each stack member	Cluster configuration are stored on the cluster command switch and the standby cluster command switch
New stack members are automatically added to the switch stack	New cluster members must be manually added to the switch cluster

Recall that stack members work together to behave as a unified system (as a single switch stack) in the network and are presented to the network as such by Layer 2 and Layer 3 protocols. Therefore, the switch cluster recognizes switch stacks, not individual stack members, as eligible cluster members. Individual stack members cannot join a switch cluster or participate as separate cluster members. Because a switch cluster must have 1 cluster command switch and can have up to 15 cluster members, a cluster can potentially have up to 16 switch stacks, totalling 144 devices.

Cluster configuration of switch stacks is through the stack master.

These are considerations to keep in mind when you have switch stacks in switch clusters:

- If the cluster command switch is not a Catalyst 2960-S switch or switch stack and a new stack master is elected in a cluster member switch stack, the switch stack loses its connectivity to the switch cluster if there are no redundant connections between the switch stack and the cluster command switch. You must add the switch stack to the switch cluster.
- If the cluster command switch is a switch stack and new stack masters are simultaneously elected in the cluster command switch stack and in cluster member switch stacks, connectivity between the switch stacks is lost if there are no redundant connections between the switch stack and the cluster command switch. You must add the switch stacks to the cluster, including the cluster command switch stack.
- All stack members should have redundant connectivity to all VLANs in the switch cluster. Otherwise, if a new stack master is elected, stack members connected to any VLANs not configured on the new stack master lose their connectivity to the switch cluster. You must change the VLAN configuration of the stack master or the stack members and add the stack members back to the switch cluster.
- If a cluster member switch stack reloads and a new stack master is elected, the switch stack loses connectivity with the cluster command switch. You must add the switch stack back to the switch cluster.
- If a cluster command switch stack reloads, and the original stack master is not re-elected, you must rebuild the entire switch cluster.

For more information about switch stacks, see [Chapter 7, “Managing Switch Stacks,”](#)

TACACS+ and RADIUS

If Terminal Access Controller Access Control System Plus (TACACS+) is configured on a cluster member, it must be configured on all cluster members. Similarly, if RADIUS is configured on a cluster member, it must be configured on all cluster members. Further, the same switch cluster cannot have some members configured with TACACS+ and other members configured with RADIUS.

For more information about TACACS+, see the [“Controlling Switch Access with TACACS+”](#) section on page 9-10. For more information about RADIUS, see the [“Controlling Switch Access with RADIUS”](#) section on page 9-18.

LRE Profiles

A configuration conflict occurs if a switch cluster has Long-Reach Ethernet (LRE) switches that use both private and public profiles. If one LRE switch in a cluster is assigned a public profile, all LRE switches in that cluster must have that same public profile. Before you add an LRE switch to a cluster, make sure that you assign it the same public profile used by other LRE switches in the cluster.

A cluster can have a mix of LRE switches that use different private profiles.

Using the CLI to Manage Switch Clusters

You can configure cluster member switches from the CLI by first logging into the cluster command switch. Enter the **rcommand** user EXEC command and the cluster member switch number to start a Telnet session (through a console or Telnet connection) and to access the cluster member switch CLI. The command mode changes, and the Cisco IOS commands operate as usual. Enter the **exit** privileged EXEC command on the cluster member switch to return to the command-switch CLI.

This example shows how to log into member-switch 3 from the command-switch CLI:

```
switch# rcommand 3
```

If you do not know the member-switch number, enter the **show cluster members** privileged EXEC command on the cluster command switch. For more information about the **rcommand** command and all other cluster commands, refer to the switch command reference.

The Telnet session accesses the member-switch CLI at the same privilege level as on the cluster command switch. The Cisco IOS commands then operate as usual. For instructions on configuring the switch for a Telnet session, see the [“Disabling Password Recovery” section on page 9-5](#).

**Note**

The CLI supports creating and maintaining switch clusters with up to 16 switch stacks. For more information about switch stack and switch cluster, see the [“Switch Clusters and Switch Stacks” section on page 6-15](#).

Stacking is supported only on Catalyst 2960-S switches running the LAN base image.

Catalyst 1900 and Catalyst 2820 CLI Considerations

If your switch cluster has Catalyst 1900 and Catalyst 2820 switches running standard edition software, the Telnet session accesses the management console (a menu-driven interface) if the cluster command switch is at privilege level 15. If the cluster command switch is at privilege level 1 to 14, you are prompted for the password to access the menu console.

Command-switch privilege levels map to the Catalyst 1900 and Catalyst 2820 cluster member switches running standard and Enterprise Edition Software as follows:

- If the command-switch privilege level is 1 to 14, the cluster member switch is accessed at privilege level 1.
- If the command-switch privilege level is 15, the cluster member switch is accessed at privilege level 15.

**Note**

The Catalyst 1900 and Catalyst 2820 CLI is available only on switches running Enterprise Edition Software.

For more information about the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

Using SNMP to Manage Switch Clusters

When you first power on the switch, SNMP is enabled if you enter the IP information by using the setup program and accept its proposed configuration. If you did not use the setup program to enter the IP information and SNMP was not enabled, you can enable it as described in the “[Configuring SNMP](#)” section on page 31-6. On Catalyst 1900 and Catalyst 2820 switches, SNMP is enabled by default.

When you create a cluster, the cluster command switch manages the exchange of messages between cluster member switches and an SNMP application. The cluster software on the cluster command switch appends the cluster member switch number (*@esN*, where *N* is the switch number) to the first configured read-write and read-only community strings on the cluster command switch and propagates them to the cluster member switch. The cluster command switch uses this community string to control the forwarding of gets, sets, and get-next messages between the SNMP management station and the cluster member switches.



Note

When a cluster standby group is configured, the cluster command switch can change without your knowledge. Use the first read-write and read-only community strings to communicate with the cluster command switch if there is a cluster standby group configured for the cluster.

If the cluster member switch does not have an IP address, the cluster command switch redirects traps from the cluster member switch to the management station, as shown in [Figure 6-7](#). If a cluster member switch has its own IP address and community strings, the cluster member switch can send traps directly to the management station, without going through the cluster command switch.

If a cluster member switch has its own IP address and community strings, they can be used in addition to the access provided by the cluster command switch. For more information about SNMP and community strings, see [Chapter 31, “Configuring SNMP.”](#)

Figure 6-7 SNMP Management for a Cluster

