



CHAPTER 1

Overview

This chapter provides these topics about the Catalyst 3550 multilayer switch software:

- [Features, page 1-1](#)
- [Management Options, page 1-9](#)
- [Network Configuration Examples, page 1-10](#)
- [Where to Go Next, page 1-19](#)

In this document, IP refers to IP version 4 (IPv4). Layer 3 IP version 6 (IPv6) packets are treated as non-IP packets.

Features

The software supports the hardware listed in the release notes. This section describes the features supported in this release:



Note

All Catalyst 3550 Gigabit Ethernet switches ship with the IP services image, formerly known as the enhanced multilayer image (EMI), which provides Layer 2+ features, full Layer 3 routing, and advanced services. Catalyst 3550 Fast Ethernet switches can be shipped with either the IP base image, formerly known as the standard multilayer software image (SMI), or the IP services image installed. The IP base image provides Layer 2+ features and basic Layer 3 routing. You can order the IP services Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the IP base image to the IP services image.

Ease of Deployment and Ease of Use

The switch ships with these features to make the deployment and the use easier:

- Express Setup for quickly configuring a switch for the first time with basic IP information, contact information, switch and Telnet passwords, and Simple Network Management Protocol (SNMP) information through a browser-based program
- User-defined Smartports macros for creating custom switch configurations for simplified deployment across the network
- An embedded device manager for configuring and monitoring a single switch through a web browser. For information about launching the device manager, see the getting started guide. For more information about the device manager, see the switch online help.

- Network Assistant GUI for
 - Unified configuration, monitoring, authentication, and software upgrade of multiple switches (see the release notes for a list of eligible cluster members).
 - Automatic discovery of candidate switches and creation of clusters of up to 16 switches that can be managed through a single IP address.
 - Extended discovery of cluster candidates that are not directly connected to the command switch.
 - Downloading an image to a switch by using HTTP or TFTP.

Performance

- Autosensing of port speed and autonegotiation of duplex mode on all switch ports for optimizing bandwidth
- IEEE 802.3x flow control on all Ethernet ports
- EtherChannel for enhanced fault tolerance and for providing up to 8 Gbps (Gigabit EtherChannel) or 800 Mbps (Fast EtherChannel) full duplex of bandwidth between switches, routers, and servers
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links
- Per-port storm control for preventing broadcast, multicast, and unicast storms
- Port blocking on forwarding unknown unicast and multicast traffic
- Cisco Group Management Protocol (CGMP) server support and Internet Group Management Protocol (IGMP) snooping for IGMP versions 1, 2, and 3:
 - (For CGMP devices) CGMP for limiting multicast traffic to specified end stations and reducing overall network traffic
 - (For IGMP devices) IGMP snooping for limiting flooding of multicast traffic
- IGMP report suppression for sending only one IGMP report per multicast router query to the multicast devices (supported only for IGMPv1 or IGMPv2 queries)
- Multicast VLAN registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security reasons
- IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong
- IGMP throttling for configuring the action when the maximum number of entries is in the IGMP forwarding table
- IGMP leave timer for configuring the leave latency for the network.
- System Database Management (SDM) templates for allocating system resources to maximize support for user-selected features
- Web Cache Communication Protocol (WCCP) for redirecting traffic to local cache engines, for enabling content requests to be fulfilled locally, and for localizing web-traffic patterns in the network (requires the enhanced multilayer software image)

Manageability

- Cisco Intelligence Engine 2100 (IE2100) Series Cisco Networking Services (CNS) embedded agents for automating switch management, configuration storage and delivery.

- DHCP for automating configuration of switch information (such as IP address, default gateway, host name, and Domain Name System [DNS] and TFTP server names)
- DHCP server for automatic assignment of IP addresses and other DHCP options to IP hosts
- DHCP relay agent information (option 82) for subscriber identification and IP address management
- Directed unicast requests to a DNS server for identifying a switch through its IP address and its corresponding host name and to a TFTP server for administering software upgrades from a TFTP server
- Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding Media Access Control (MAC) address
- Unicast MAC address filtering to drop packets with specific source or destination MAC addresses
- Cisco Discovery Protocol (CDP) versions 1 and 2 for network topology discovery and mapping between the switch and other Cisco devices on the network
- Network Time Protocol (NTP) for providing a consistent timestamp to all switches from an external source
- Cisco IOS File System (IFS) for providing a single interface to all file systems that the switch uses
- Configuration logging to log and to view changes to the switch configuration
- Unique device identifier to provide product identification information through a **show inventory** user EXEC command display
- In-band management access through the embedded device manager over a Netscape Navigator or Internet Explorer session or through the Network Assistant application
- In-band management access through up to 16 simultaneous Telnet connections for multiple command-line interface (CLI)-based sessions over the network
- In-band management access for up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network
- In-band management access through SNMP versions 1, 2c, and 3 get and set requests
- Out-of-band management access through the switch console port to a directly attached terminal or to a remote terminal through a serial connection or a modem
- Secure Copy Protocol (SCP) feature to provide a secure and authenticated method for copying switch configuration or switch image files (requires the cryptographic version of the software)

**Note**

For additional descriptions of the management interfaces, see the [“Management Options” section on page 1-9](#).

Redundancy

- Hot Standby Router Protocol (HSRP) for command switch and Layer 3 router redundancy
- UniDirectional Link Detection (UDLD) and aggressive UDLD on all Ethernet ports for detecting and disabling unidirectional links on fiber-optic interfaces caused by incorrect fiber-optic wiring or port faults
- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks. STP has these features:
 - Per-VLAN spanning-tree plus (PVST+) for load balancing across VLANs

- Rapid PVST+ for load balancing across VLANs
- UplinkFast, cross-stack UplinkFast, and BackboneFast for fast convergence after a spanning-tree topology change and for achieving load balancing between redundant uplinks, including Gigabit uplinks and cross-stack Gigabit uplinks
- IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) for grouping VLANs into a spanning-tree instance, and providing for multiple forwarding paths for data traffic and load balancing
- IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the forwarding state
- Optional spanning-tree features available in PVST+, rapid-PVST+, and MSTP mode:
 - Port Fast for eliminating the forwarding delay by enabling a port to immediately transition from the blocking state to the forwarding state
 - BPDU guard for shutting down Port Fast-enabled ports that receive BPDUs
 - BPDU filtering for preventing a Port Fast-enabled port from sending or receiving BPDUs
 - Root guard for preventing switches outside the network core from becoming the spanning-tree root
 - Loop guard for preventing alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link

**Note**

The switch supports up to 128 spanning-tree instances.

- Flex Link Layer 2 interfaces to back up one another as an alternative to STP for basic link redundancy

VLAN Support

- Support for up to 1005 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth
- Support for VLAN IDs in the full 1 to 4094 range as allowed by the IEEE 802.1Q standard
- VLAN Query Protocol (VQP) for dynamic VLAN membership
- Inter-Switch Link (ISL) and IEEE 802.1Q trunking encapsulation on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources
- Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (IEEE 802.1Q or ISL) to be used
- VLAN Trunking Protocol (VTP) and VTP pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic
- Voice VLAN for creating subnets for voice traffic from Cisco IP Phones
- VLAN 1 minimization to reduce the risk of spanning-tree loops or storms by allowing VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received. The switch CPU continues to send and receive control protocol frames.

Security

- Web authentication to allow a supplicant (client) that does not support IEEE 802.1x functionality to be authenticated using a web browser.
- Password-protected access (read-only and read-write access) to management interfaces (device manager, Network Assistant, and CLI) for protection against unauthorized configuration changes
- Multilevel security for a choice of security level, notification, and resulting actions
- Static MAC addressing for ensuring security
- Protected port option for restricting the forwarding of traffic to designated ports on the same switch
- Port security option for limiting and identifying MAC addresses of the stations allowed to access the port
- Port security on trunk ports for limiting and identifying MAC addresses of the stations allowed to access the VLAN
- Port security aging to set the aging time for secure addresses on a port
- IP source guard to restrict traffic on nonrouted interfaces by filtering traffic based on the DHCP snooping database and IP source bindings.
- DHCP snooping to filter untrusted DHCP messages between untrusted hosts and DHCP servers
- Dynamic ARP inspection to prevent malicious attacks on the switch by not relaying invalid ARP requests and responses to other ports in the same VLAN
- Bridge protocol data unit (BPDU) guard for shutting down a Port Fast-configured port when an invalid configuration occurs
- Standard and extended IP access control lists (ACLs) for defining security policies in both directions on routed interfaces (router ACLs) and inbound on Layer 2 interfaces (port ACLs)
- Extended MAC access control lists for defining security policies in the inbound direction on Layer 2 interfaces
- VLAN ACLs (VLAN maps) for providing intra-VLAN security by filtering traffic based on information in the MAC, IP, and TCP/User Datagram Protocol (UDP) headers
- Source and destination MAC-based ACLs for filtering non-IP traffic
- IEEE 802.1x port-based authentication to prevent unauthorized devices (clients) from gaining access to the network
- IEEE 802.1x with per-user access control lists for providing different levels of network access and service to an IEEE 802.1x-authenticated user
- IEEE 802.1x with VLAN assignment for restricting IEEE 802.1x-authenticated users to a specified VLAN
- IEEE 802.1x with port security for controlling access to IEEE 802.1x multiple-host ports
- IEEE 802.1x with voice VLAN to permit an IP phone access to the voice VLAN irrespective of the authorized or unauthorized state of the port
- IEEE 802.1x with guest VLAN to provide limited services to non-IEEE 802.1x compliant users
- IEEE 802.1x with restricted VLAN to provide limited services to users who are IEEE 802.1x compliant, but do not have the credentials to authenticate via the standard IEEE 802.1x processes.
- IEEE 802.1x accounting to track network usage
- IEEE 802.1x with wake-on-LAN to allow dormant PCs to be powered on based on the receipt of a specific Ethernet frame

- Network Admission Control (NAC) features:
 - NAC Layer 2 IEEE 802.1x validation of the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access.
For information about configuring NAC Layer 2 IEEE 802.1x validation, see the [“Configuring NAC Layer 2 IEEE 802.1x Validation”](#) section on page 8-38.
 - NAC Layer 2 IP validation to validate the posture of endpoint systems or clients before granting the devices network access.
For information about configuring NAC Layer 2 IP validation, see the *Network Admission Control Software Configuration Guide*.
 - IEEE 802.1x inaccessible authentication bypass.
For information about configuring this feature, see the [“Configuring the Inaccessible Authentication Bypass Feature”](#) section on page 8-34.
 - Authentication, authorization, and accounting (AAA) down policy for a NAC Layer 2 IP validation of a host if the AAA server is not available when the posture validation occurs.
For information about this feature, see the *Configuring Network Admission Control Software Configuration Guide*.
- Network Admission Control (NAC) Layer 2 IEEE 802.1x validation to validate the antivirus condition or *posture* of endpoint systems or clients before granting the devices network access by using IEEE 802.1x port-based authentication on the network edge
- NAC Layer 2 IP validation to validate the posture of endpoint systems or clients before granting the devices network access by using UDP on the network edge
- TACACS +, a proprietary feature for managing network security through a TACACS server
- Kerberos security system to authenticate requests for network resources by using a trusted third party
- RADIUS, which provides detailed accounting information and flexible administrative control over authentication and authorization processes
- Secure Socket Layer (SSL) Version 3.0 support for the HTTP 1.1 server authentication, encryption, and message integrity, and HTTP client authentication to allow secure HTTP communications
- IEEE 802.1Q tunneling to allow customers with users at remote sites across a service provider network to keep VLANs segregated from other customers and Layer 2 protocol tunneling to ensure that the customer’s network has complete STP, CDP, and VTP information about all users
- Layer 2 point-to-point tunneling to facilitate the automatic creation of EtherChannels

Quality of Service (QoS) and Class of Service (CoS)

- Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues
- Classification
 - Classification on a physical interface or on a per-port per-VLAN basis
 - IP type-of-service/Differentiated Services Code Point (IP TOS/DSCP) and IEEE 802.1P CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications

- IP TOS/DSCP and IEEE 802.1P CoS marking based on flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance quality of service at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network
- Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain
- Trusted boundary for detecting the presence of a Cisco IP phone, trusting the CoS value received, and ensuring port security
- Policing
 - Policing on a physical interface or on a per-port per-VLAN basis
 - Traffic-policing policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow
 - Aggregate policing for policing traffic flows in aggregate to restrict specific applications or traffic flows to metered, predefined rates
 - Up to 128 policers on ingress Gigabit-capable Ethernet ports
Up to eight policers on ingress 10/100 ports
Up to eight policers per egress port (aggregate policers only)
- Out-of-Profile
 - Out-of-profile markdown for packets that exceed bandwidth utilization limits
- Egress Policing and Scheduling of Egress Queues
 - Four egress queues on all switch ports. These queues can either be configured with the Weighted Round Robin (WRR) scheduling algorithm or configured with one queue as a strict priority queue and the other three queues for WRR. The strict priority queue must be empty before the other three queues are serviced. You can use the strict priority queue for mission-critical and time-sensitive traffic.
 - Tail drop and Weight Random Early Detection (WRED) techniques for avoiding congestion on Gigabit Ethernet ports; tail drop for congestion avoidance on Fast Ethernet ports

Layer 3 Support

Some features and protocols require the enhanced multilayer software image.

- Hot Standby Router Protocol (HSRP) for Layer 3 router redundancy
- IP routing protocols for load balancing and for constructing scalable, routed backbones:
 - Routing Information Protocol (RIP) versions 1 and 2
 - Open Shortest Path First (OSPF)
 - Enhanced IGRP (EIGRP)
 - Border Gateway Protocol (BGP) Version 4
- IP routing between VLANs (inter-VLAN routing) for full Layer 3 routing between two or more VLANs, allowing each VLAN to maintain its own autonomous data-link domain
- Multiple VPN routing/forwarding (multi-VRF) instances in customer edge (CE) devices to allow service providers to support multiple virtual private networks (VPNs) and overlap IP addresses between VPNs.

- Policy-based routing (PBR) for configuring defined policies for traffic flows
- Fallback bridging for forwarding non-IP traffic between two or more VLANs
- Static IP routing for manually building a routing table of network path information
- Equal-cost routing for load balancing and redundancy
- Internet Control Message Protocol (ICMP) and ICMP Router Discovery Protocol (IRDP) for using router advertisement and router solicitation messages to discover the addresses of routers on directly attached subnets
- Protocol-Independent Multicast (PIM) for multicast routing within the network, allowing for devices in the network to receive the multicast feed requested and for switches not participating in the multicast to be pruned. Includes support for PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode.
- Distance Vector Multicast Routing Protocol (DVMRP) tunneling for interconnecting two multicast-enabled networks across non-multicast networks
- DHCP relay for forwarding UDP broadcasts, including IP address requests, from DHCP clients
- Nonstop forwarding (NSF) awareness to enable the Layer 3 switch to continue forwarding packets from an NSF-capable neighboring router during the interval while the primary route processor (RP) is crashing and the backup RP is taking over, or while the primary RP is manually reloaded for a nondisruptive software upgrade (requires the IP services image)

Monitoring

- Switch LEDs that provide port- and switch-level status
- Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) for traffic monitoring on any port or VLAN
- SPAN and RSPAN support of Intrusion Detection Systems (IDSs) to monitor, repel, and report network security violations
- Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis
- Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events
- MAC address notification for tracking users on a network by storing the MAC addresses that the switch has learned or removed
- Layer 2 traceroute to identify the physical path that a packet takes from a source device to a destination device

Power over Ethernet Support for the Catalyst 3550-24PWR Switch

- Ability to provide power to connected Cisco pre-standard and IEEE 802.3af-compliant powered devices from Power over Ethernet (PoE)-capable ports if the switch detects that there is no power on the circuit.
- Support for CDP with power consumption. The powered device notifies the switch of the amount of power it is consuming.

- Support for Cisco intelligent power management. The powered device and the switch negotiate through power-negotiation CDP messages for an agreed power-consumption level. The negotiation allows a high-power Cisco powered device to operate at its highest power mode.
- Automatic detection and power budgeting; the switch maintains a power budget, monitors and tracks requests for power, and grants power only when it is available.
- Fan-fault and over-temperature detection through the device manager and Network Assistant

Management Options

The switch is designed for plug-and-play operation: you need to configure only basic IP information for the switch and connect it to the other devices in your network. If you have specific network needs, you can configure and monitor the switch—on an individual basis or as part of a switch cluster—through its various management interfaces.

Management Interface Options

You can configure and monitor individual switches and switch clusters by using these interfaces:

- An embedded device manager—The device manager is a GUI that is integrated in the software image. You use it to configure and to monitor a single switch. For more information about the device manager, see the switch online help.
- Network Assistant—Network Assistant is a GUI that can be downloaded from Cisco.com. You use it to manage a single switch or a cluster of switches. For more information about Network Assistant, see the *Getting Started with Cisco Network Assistant*, available on Cisco.com.
- CLI—The switch Cisco IOS software supports desktop- and multilayer-switching features. You can access the CLI either by connecting your management station directly to the switch console port or by using Telnet from a remote management station.

For more information about the CLI, see [Chapter 2, “Using the Command-Line Interface.”](#)

- IE2100—Cisco Intelligence Engine 2100 Series Configuration Registrar is a network management device that works with embedded CNS Agents in the switch software. You can automate initial configurations and configuration updates by generating switch-specific configuration changes, sending them to the switch, executing the configuration change, and logging the results.

For more information about IE2100, see [Chapter 4, “Configuring Cisco IOS CNS Agents.”](#)

- SNMP—SNMP provides a means to monitor and control the switch and switch cluster members. You can manage switch configuration settings, performance, security, and collect statistics by using SNMP management applications such as CiscoWorks2000 LAN Management Suite (LMS) and HP OpenView.

You can manage the switch from an SNMP-compatible management station that is running platforms such as HP OpenView or SunNet Manager. The switch supports a comprehensive set of MIB extensions and four RMON groups.

For more information about using SNMP, see [Chapter 27, “Configuring SNMP.”](#)

Advantages of Using Network Assistant and Clustering Switches

Using Network Assistant and switch clusters can simplify and minimize your configuration and monitoring tasks. You can use Cisco switch clustering technology to manage up to 16 interconnected, supported Catalyst switches through one IP address. This can conserve IP addresses if you have a limited number of them. Network Assistant is the easiest interface to use and makes switch and switch cluster management accessible to authorized users from any PC on your network.

By using switch clusters and Network Assistant, you can

- Manage and monitor interconnected Catalyst switches (see the release notes for a list of supported switches), regardless of their geographic proximity and interconnection media, including Ethernet, Fast Ethernet, Fast EtherChannel, Cisco GigaStack Gigabit Interface Converter (GBIC), Gigabit Ethernet, and Gigabit EtherChannel connections.
- Accomplish multiple configuration tasks from a single Network Assistant window without needing to remember CLI commands to accomplish specific tasks.
- Apply actions from Network Assistant to multiple ports and multiple switches at the same time. Here are some examples of configuring and managing multiple ports and switches:
 - Port configuration such as speed and duplex settings
 - Port and console port security settings
 - NTP, STP, VLAN, and QoS configurations
 - Inventory and statistic reporting and link- and switch-level monitoring and troubleshooting
 - Group software upgrades
- View a topology of interconnected devices to identify existing switch clusters and eligible switches that can join a cluster. You can also use the topology to quickly identify link information between switches.
- Monitor real-time status of a switch or multiple switches from the LEDs on the front-panel images. The system, redundant power system (RPS), and port LED colors on the images are similar to those used on the physical LEDs.
- Use an interactive mode that takes you step-by-step through configuring complex features such as VLANs, ACLs, and QoS.
- Use a wizard that prompts you to provide only the minimum required information to configure complex features such as QoS priorities for video traffic, priority levels for data applications, and security.

For the Network Assistant software and browser requirements, and for more information about clustering, see *Getting Started with Cisco Network Assistant*, available on Cisco.com. For clustering requirements, including supported Cisco IOS releases, see the release notes for this release.

Network Configuration Examples

This section provides network configuration concepts and includes examples of using the switch to create dedicated network segments and interconnecting the segments through Fast Ethernet and Gigabit Ethernet connections.

- [“Design Concepts for Using the Switch” section on page 1-11](#)
- [“Small to Medium-Sized Network Using Mixed Switches” section on page 1-14](#)
- [“Large Network Using Only Catalyst 3550 Switches” section on page 1-15](#)

- [“Multidwelling Network Using Catalyst 3550 Switches” section on page 1-17](#)
- [“Long-Distance, High-Bandwidth Transport Configuration” section on page 1-18](#)

Design Concepts for Using the Switch

As your network users compete for network bandwidth, it takes longer to send and receive data. When you configure your network, consider the bandwidth required by your network users and the relative priority of the network applications they use.

[Table 1-1](#) describes what can cause network performance to degrade and how you can configure your network to increase the bandwidth available to your network users.

Table 1-1 *Increasing Network Performance*

Network Demands	Suggested Design Methods
Too many users on a single network segment and a growing number of users accessing the Internet	<ul style="list-style-type: none"> • Create smaller network segments so that fewer users share the bandwidth, and use VLANs and IP subnets to place the network resources in the same logical network as the users who access those resources most. • Use full-duplex operation between the switch and its connected workstations.
<ul style="list-style-type: none"> • Increased power of new PCs, workstations, and servers • High bandwidth demand from networked applications (such as e-mail with large attached files) and from bandwidth-intensive applications (such as multimedia) 	<ul style="list-style-type: none"> • Connect global resources—such as servers and routers to which the network users require equal access—directly to the high-speed switch ports so that they have their own high-speed segment. • Use the EtherChannel feature between the switch and its connected servers and routers.

Bandwidth alone is not the only consideration when designing your network. As your network traffic profiles evolve, consider providing network services that can support applications for voice and data integration, multimedia integration, application prioritization, and security. [Table 1-2](#) describes some network demands and how you can meet those demands.

Table 1-2 Providing Network Services

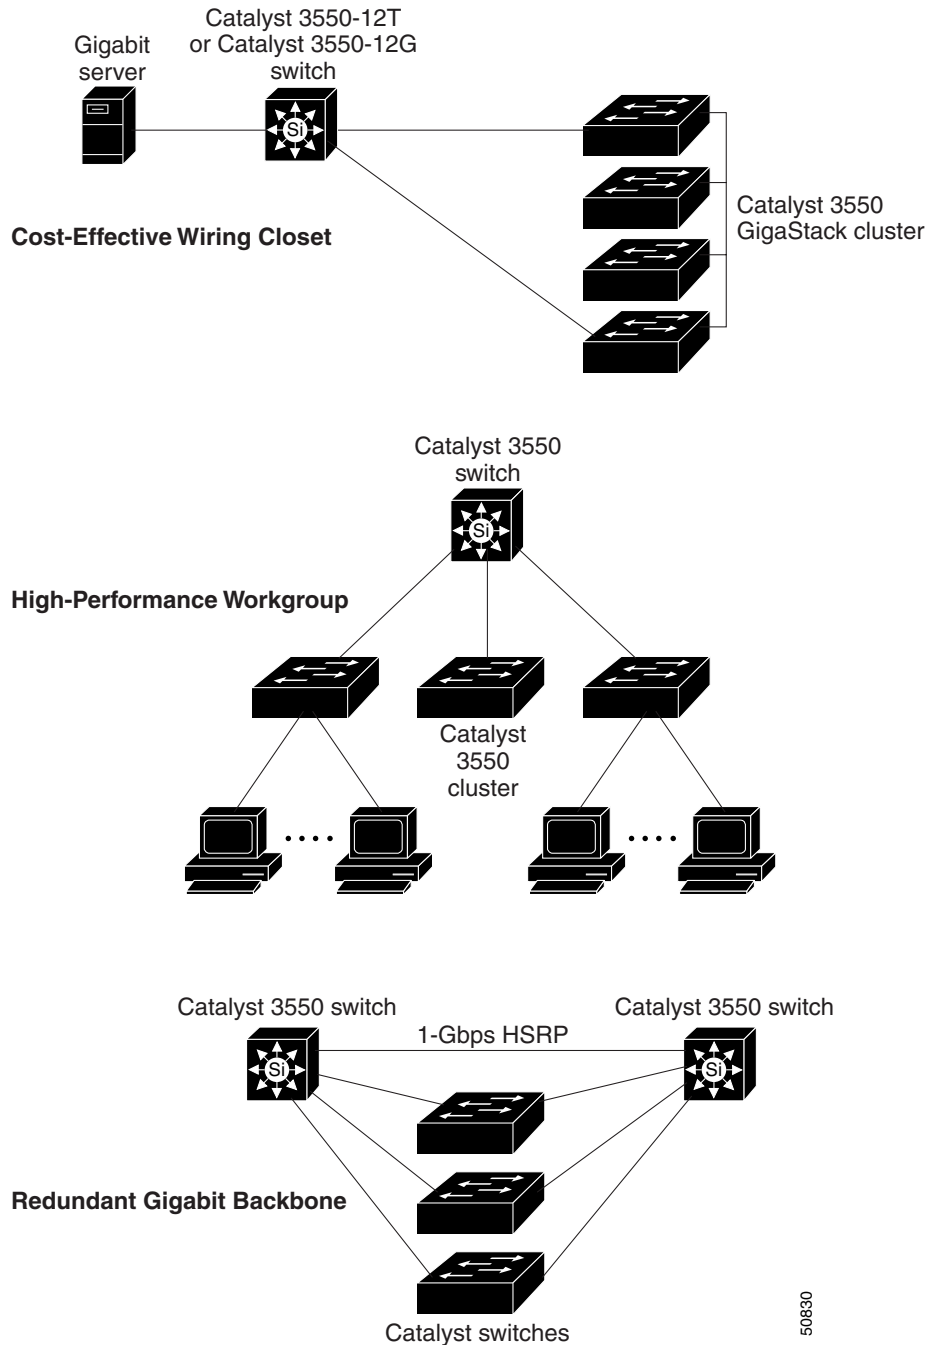
Network Demands	Suggested Design Methods
Efficient bandwidth usage for multimedia applications and guaranteed bandwidth for critical applications	<ul style="list-style-type: none"> • Use IGMP snooping to efficiently forward multimedia and multicast traffic. • Use other QoS mechanisms such as packet classification, marking, scheduling, and congestion avoidance to classify traffic with the appropriate priority level, thereby providing maximum flexibility and support for mission-critical, unicast, and multicast and multimedia applications. • Use optional IP multicast routing to design networks better suited for multicast traffic. • Use MVR to continuously send multicast streams in a multicast VLAN, but to isolate the streams from subscriber VLANs for bandwidth and security reasons.
High demand on network redundancy to provide <i>always on</i> mission-critical applications	<ul style="list-style-type: none"> • Use HSRP for router redundancy. • Use VLAN trunks, cross-stack UplinkFast, and BackboneFast for traffic-load balancing on the uplink ports so that the uplink port with a lower relative port cost is selected to carry the VLAN traffic.
An evolving demand for IP telephony	<ul style="list-style-type: none"> • Use QoS to prioritize applications such as IP telephony during congestion and to help control both delay and jitter within the network. • Use switches that support at least two queues per port to prioritize voice and data traffic as either high- or low-priority, based on IEEE 802.1P/Q. • Use voice VLAN IDs (VVIDs) on the Catalyst 2900 XL and 3500 XL switches to provide separate VLANs for voice traffic.
A growing demand for using existing infrastructure to transport data and voice from a home or office to the Internet or an intranet at higher speeds	<p>Use the Catalyst 2900 LRE XL and Catalyst 2950 LRE switches to provide up to 15 Mb of IP connectivity over existing infrastructure, such as existing telephone lines.</p> <p>Note Long-Reach Ethernet (LRE) is the technology used in the Catalyst 2900 LRE XL and Catalyst 2950 LRE switches. See the switch documentation sets about these switches and the LRE technology.</p>

Figure 1-1 shows three configuration examples of using Catalyst switches to create the following:

- Cost-effective wiring closet—A cost-effective way to connect many users to the wiring closet is to connect a Catalyst switch cluster of up to nine Catalyst 3550 XL switches (or with a mix of Catalyst 3550, Catalyst 2950, Catalyst 3500 XL, and Catalyst 2900 XL switches) through GigaStack GBIC connections. To preserve switch connectivity if one switch in the stack fails, connect the bottom switch to the top switch to create a GigaStack loopback, and enable cross-stack UplinkFast on the cross-stack Gigabit uplinks.

You can have redundant uplink connections, using Gigabit GBIC modules, from the GigaStack cluster to a Gigabit backbone switch such as the Catalyst 3550-12T or Catalyst 3550-12G switch. You can also create backup paths by using Fast Ethernet, Gigabit, or EtherChannel links. If one of the redundant connections fails, the other can serve as a backup path. You can configure the Catalyst 3550-12T or Catalyst 3550-12G switch as a switch cluster manager to manage stack members through a single IP address. The Catalyst 3550-12T or Catalyst 3550-12G switch can be connected to a Gigabit server through a 1000BASE-T connection.

Figure 1-1 Example Configurations



- **High-performance workgroup**—For high-speed access to network resources, you can use Catalyst 3550 switches in the access layer to provide Gigabit Ethernet to the desktop. To prevent congestion, use QoS DSCP marking priorities on these switches. For high-speed IP forwarding at the distribution layer, connect the Catalyst 3550 switches in the access layer to a Gigabit multilayer switch (such as the Catalyst 3550 multilayer switch) in the backbone.

Each switch in this configuration provides users with a dedicated 1-Gbps connection to network resources in the backbone. Compare this with the switches in a GigaStack configuration, where the 1-Gbps connection is shared among the switches in the stack. Using these Gigabit GBIC modules also provides flexibility in media and distance options:

- 1000BASE-T GBIC: copper connections of up to 328 feet (100 m)
- 1000BASE-SX GBIC: fiber-optic connections of up to 1804 feet (550 m)
- 1000BASE-LX/LH GBIC: fiber-optic connections of up to 32,808 feet (6 miles or 10 km)
- 1000BASE-ZX GBIC: fiber-optic connections of up to 328,084 feet (62 miles or 100 km)
- Redundant Gigabit backbone—Using HSRP, you can create backup paths between two Catalyst 3550 multilayer switches to enhance network reliability and load balancing for different VLANs and subnets. Using HSRP also provides faster network convergence if any network failure occurs. You can connect the Catalyst switches, again in a star configuration, to two Catalyst 3550 multilayer backbone switches. If one of the backbone switches fails, the second backbone switch preserves connectivity between the switches and network resources.

Small to Medium-Sized Network Using Mixed Switches

Figure 1-2 shows a configuration for a network of up to 500 employees. This network uses Catalyst 3550 multilayer switches to aggregate up to ten wiring closets through high-speed uplinks. For network reliability and load balancing, this network includes two routers and two Catalyst 3550 multilayer switches, all with HSRP enabled. This ensures connectivity to the Internet, WAN, and mission-critical network resources if one of the routers or Catalyst 3550 multilayer switches fails.

The wiring closets have a mix of switches such as the Catalyst 3550, Catalyst 3500 XL, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switches. These switches are connected to workstations, Cisco IP Phones, and local servers. You can cluster these switches into multiple clusters, as shown, or into a single cluster. You can manage a cluster through the IP address of its primary and secondary command switches, regardless of the geographic location of the cluster members.

This network uses VLANs to segment the network logically into well-defined broadcast groups and for security management. Data and multimedia traffic are configured on the same VLAN. Voice traffic from the Cisco IP Phones are configured on separate VVIDs. You can have up to four VVIDs per wiring closet. If data, multimedia, and voice traffic are assigned to the same VLAN, only one VLAN can be configured per wiring closet. For any switch port connected to Cisco IP Phones, IEEE 802.1P/Q QoS gives voice traffic forwarding-priority over data traffic.

Cisco IP Phones are connected—using standard straight-through, twisted-pair cable with RJ-45 connectors—to the 10/100 PoE ports on the Catalyst 3550-24PWR switches and to the 10/100 ports on the Catalyst 3550 switches. These multiservice switch ports automatically detect any IP phones that are connected. Cisco CallManager controls call processing, routing, and IP phone features and configuration. Users with workstations running Cisco SoftPhone software can place, receive, and control calls from their PCs. Using Cisco IP Phones, Cisco CallManager software, and Cisco SoftPhone software integrates telephony and IP networks, and the IP network supports both voice and data.

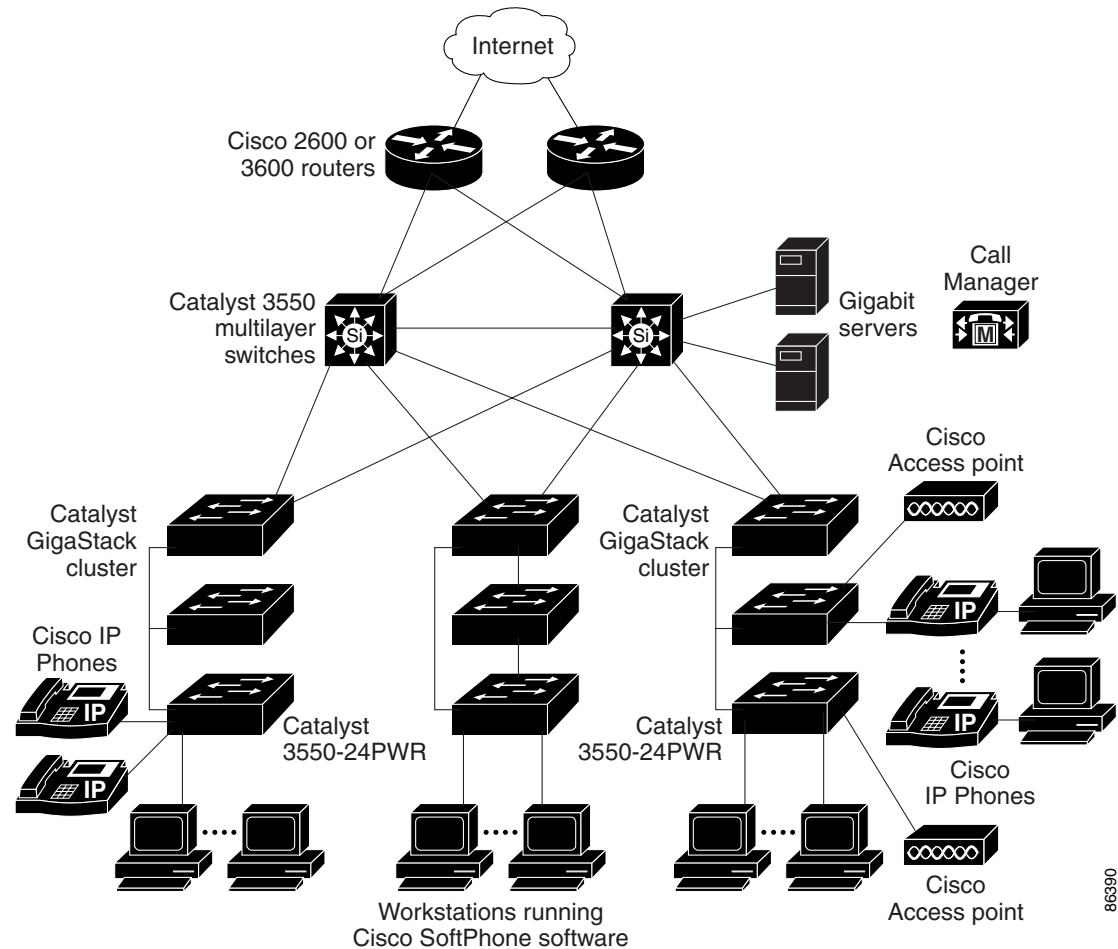
Each 10/100 PoE port on the Catalyst 3550-24PWR switches provides 15.4 W per port. The IP phone can receive redundant power when it is also connected to an AC power source. IP phones not connected to the Catalyst 3550-24PWR switches receive power from an AC power source.

When an end station in one VLAN needs to communicate with an end station in another VLAN, a router or multilayer switch routes the traffic to the appropriate destination VLAN. In this network, the Catalyst 3550 multilayer switches provide inter-VLAN routing. VLAN access control lists (VLAN maps) on the Catalyst 3550 switches provide intra-VLAN security and prevent unauthorized users from accessing critical pieces of the network.

In addition to inter-VLAN routing, the Catalyst 3550 multilayer switches provide QoS mechanisms such as DSCP priorities to prioritize the different types of network traffic and to deliver high-priority traffic in a predictable manner. If congestion occurs, QoS drops low-priority traffic to allow delivery of high-priority traffic.

With the Catalyst 3550 multilayer switches providing inter-VLAN routing and other network services, the routers focus on firewall services, Network Address Translation (NAT) services, voice-over-IP (VoIP) gateway services, and WAN and Internet access.

Figure 1-2 Catalyst 3550 Switches in a Collapsed Backbone Configuration



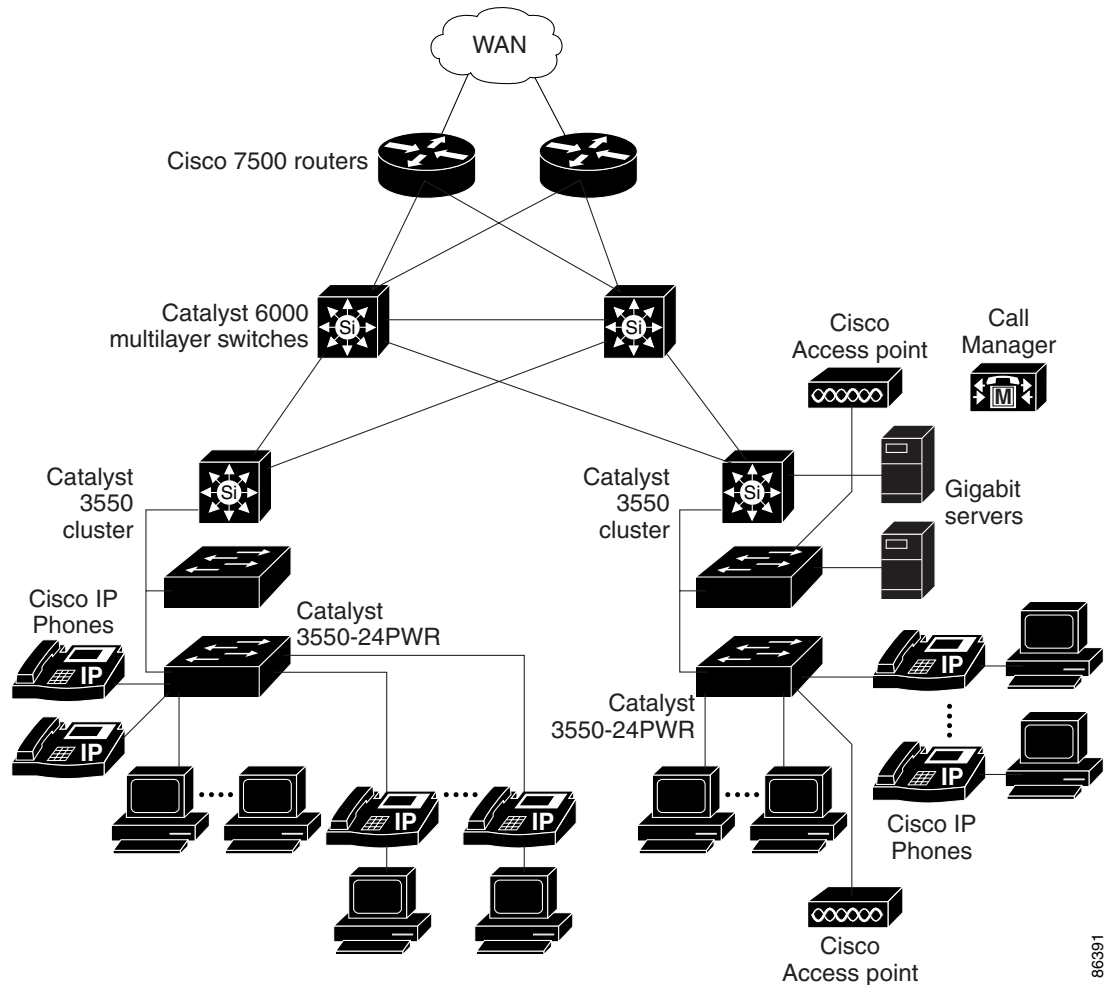
86390

Large Network Using Only Catalyst 3550 Switches

Switches in the wiring closet have traditionally been Layer 2-only devices, but as network traffic profiles evolve, switches in the wiring closet are increasingly employing multilayer services such as multicast management and traffic classification. Figure 1-3 shows a configuration for a network exclusively using Catalyst 3550 multilayer switches in the wiring closets and a Catalyst 6000 switch in the backbone to aggregate up to ten wiring closets.

In the wiring closet, each Catalyst 3550 switch has IGMP snooping enabled to efficiently forward multimedia and multicast traffic. QoS ACLs that either drop or mark nonconforming traffic based on bandwidth limits are also configured on each switch. VLAN maps provide intra-VLAN security and prevent unauthorized users from accessing critical pieces of the network. QoS features can limit bandwidth on a per-port or per-user basis. The switch ports are configured as either trusted or untrusted. You can configure a trusted port to trust the CoS value, the DSCP value, or the IP precedence. If you configure the port as untrusted, you can use an ACL to mark the frame in accordance with the network policy.

Figure 1-3 Catalyst 3550 Switches in Wiring Closets in a Backbone Configuration



Within each wiring closet is a Catalyst 3550 multilayer switch for inter-VLAN routing. These switches provide proxy ARP services to determine IP and MAC address mapping, thereby removing this task from the routers and lessening this type of traffic on the WAN links. These switches also have redundant uplink connections to the backbone switches, with each uplink port configured as a trusted routed uplink to provide faster convergence in case of an uplink failure.

The routers and Catalyst 6000 multilayer backbone switches have HSRP enabled for load balancing and redundant connectivity to guarantee mission-critical traffic.

The Catalyst 6000 switch provides the workgroups with Gigabit access to core resources. The server farm includes a call-processing server running Cisco CallManager software. Cisco CallManager controls call processing, routing, and IP phone features and configuration.

Multidwelling Network Using Catalyst 3550 Switches

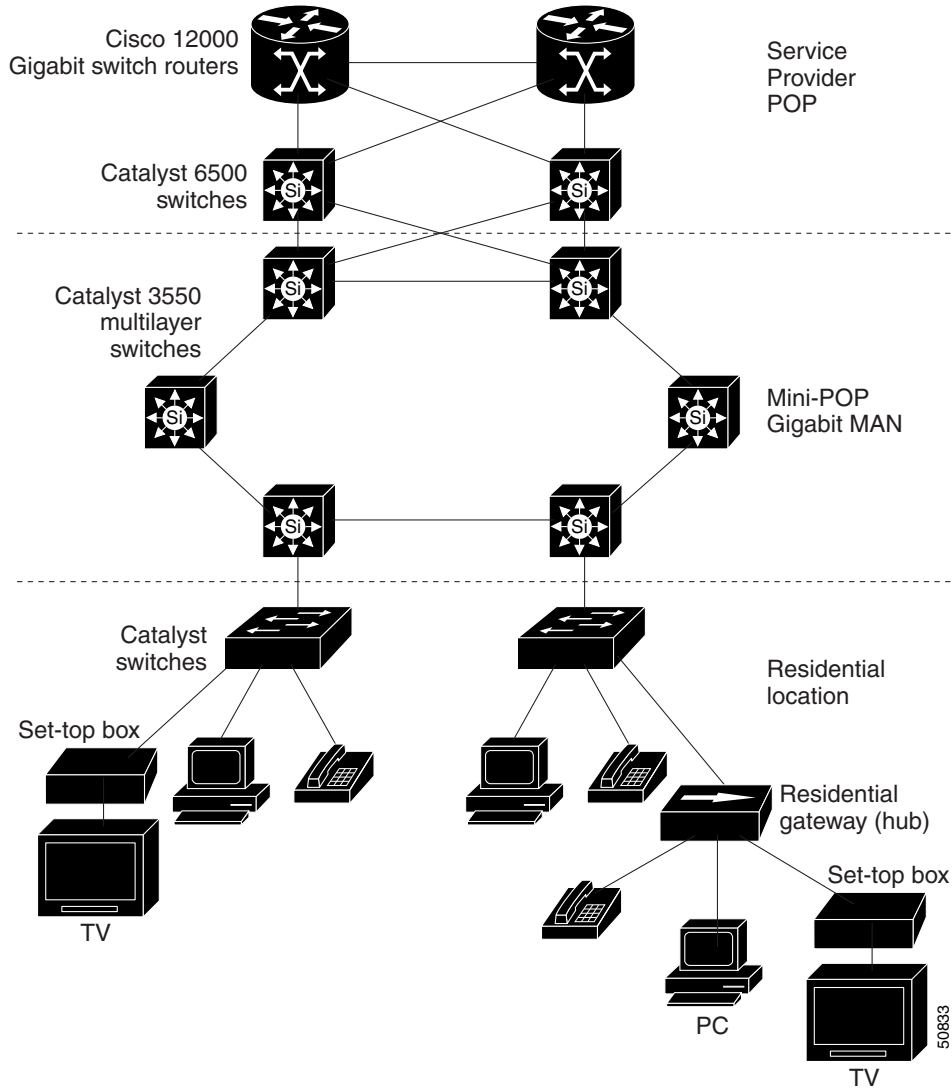
A growing segment of residential and commercial customers are requiring high-speed access to Ethernet metropolitan-area networks (MANs). [Figure 1-4](#) shows a configuration for a Gigabit Ethernet MAN ring using Catalyst 3550 multilayer switches as aggregation switches in the mini-point-of-presence (POP) location. These switches are connected through 1000BASE-X GBIC ports.

The resident switches can be Catalyst 3550 switches, providing customers with high-speed connections to the MAN. Catalyst 2900 LRE XL or 2950 LRE Layer 2-only switches also can be used as residential switches for customers requiring connectivity through existing phone lines. The Catalyst LRE switches can then connect to another residential switch or to an aggregation switch.

All ports on the residential Catalyst 3550 switches (and Catalyst LRE switches if they are included) are configured as IEEE 802.1Q trunks with protected port and STP root guard features enabled. The protected port feature provides security and isolation between ports on the switch, ensuring that subscribers cannot view packets destined for other subscribers. STP root guard prevents unauthorized devices from becoming the STP root switch. All ports have IGMP snooping or CGMP enabled for multicast traffic management. ACLs on the uplink ports to the aggregating Catalyst 3550 multilayer switches provide security and bandwidth management.

The aggregating switches and routers provide services such as those described in the previous examples, [“Small to Medium-Sized Network Using Mixed Switches”](#) section on page 1-14 and [“Large Network Using Only Catalyst 3550 Switches”](#) section on page 1-15.

Figure 1-4 Catalyst 3550 Switches in a MAN Configuration



Long-Distance, High-Bandwidth Transport Configuration

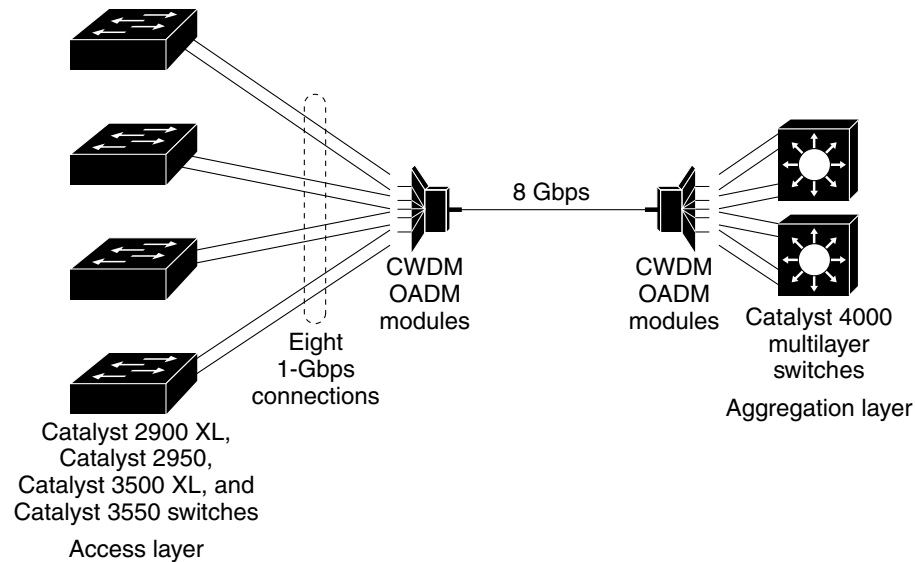
Figure 1-5 shows a configuration for transporting 8 Gigabits of data over a single fiber-optic cable. The Catalyst switches have Coarse Wave Division Multiplexer (CWDM) fiber-optic GBIC modules installed. Depending on the CWDM GBIC module, data is sent at wavelengths from 1470 nm to 1610 nm. The higher the wavelength, the farther the transmission can travel. A common wavelength used for long-distance transmissions is 1550 nm.

The CWDM GBIC modules connect to CWDM optical add/drop multiplexer (OADM) modules over distances of up to 393,701 feet (74.5 miles or 120 km). The CWDM OADM modules combine (or *multiplex*) the different CWDM wavelengths, allowing them to travel simultaneously on the same fiber-optic cable. The CWDM OADM modules on the receiving end separate (or *demultiplex*) the different wavelengths.

Using CWDM technology with the switches translates to farther data transmission and an increased bandwidth capacity (up to 8 Gbps) on a single fiber-optic cable.

For more information about the CWDM GBIC modules and CWDM OADM modules, see the *Installation Note for the CWDM Passive Optical System*.

Figure 1-5 Long-Distance, High-Bandwidth Transport Configuration



74089

Where to Go Next

Before configuring the switch, review these sections for start up information:

- [Chapter 2, “Using the Command-Line Interface”](#)
- [Chapter 3, “Assigning the Switch IP Address and Default Gateway”](#)
- [Chapter 4, “Configuring Cisco IOS CNS Agents”](#)

