# LAN Design Guide for the Midmarket

## Introduction

With the rapid adoption of Internet technologies and increasing consolidation of data, voice, and video on a common network infrastructure, networks have become a critical component of a company's business infrastructure. Scalable switching solutions based on Ethernet, Fast Ethernet, and Gigabit Ethernet have been introduced to handle the rising demand for more network bandwidth and higher performance. Proper network designs are required to achieve the business objectives and to make the network a competitive advantage.

In this paper, we present a structured approach to designing standards-based local area networks (LANs) for midmarket companies using the Catalyst® family of stackable switches. Midmarket companies are defined as having 100 to 1000 users. While price/performance, ease of management, and scalability continue to remain key requirements, midmarket networks have evolved to supporting mission-critical applications such as e-commerce, virtual private networks (VPNs), enterprise resource planning (ERP), and IP Telephony. These applications are driving the deployment of intelligent network services for high availability and high performance—features that in the past were typically only seen in larger campus network designs. These new-generation applications have positioned Internet Protocol (IP) as the dominant protocol in midmarket networks. In addition, due to their attractive price-performance and scalability, Ethernet LAN solutions have become the solutions of choice in these midmarket networks.

This document is targeted at systems engineers, network design engineers, and network managers catering to midmarket customers. We assume that the reader has a basic knowledge of an Ethernet LAN. Design concepts and examples outlined here are the result of extensive Cisco experience in building switched networks that provide reliability, performance, security, and manageability, and have been deployed by numerous customers worldwide.

In this paper, we discuss the following major topics.

- *Applications Driving Differentiated Services Networks*—Summarizes key end-to-end requirements for mission-critical applications and introduce differentiated services networks.
- *Structured Network Design*—Presents structured design techniques and define hierarchical requirements at the wiring closet and core (backbone) layer.
- *Intelligent Network Services*—Outlines intelligent network services such as quality of service (QoS), high availability, security, and management services required to deliver differentiated services.
- *Building Blocks for Structured Design and Design Examples*—Presents several design examples to cover an array of customer requirements in terms of price, performance, network size, network services, and scalability. We develop these sample designs based on building blocks with well-defined goals and characteristics.
- *Summary and Quick Reference*—Provides a quick reference guide to select the most appropriate design that meets your business requirements.

### Applications Driving Differentiated Services Networks

Different business applications require different network characteristics in terms of throughput, delay, jitter, and packet loss in order to make the user experience acceptable. Unacceptable network delays for business applications such as SAP or Oracle transactions can hamper the productivity of a company and decrease its agility. Hence, a thorough understanding of application requirements and projected traffic growth are crucial to designing scalable networks with built-in future proofing.

Embracing new technologies such as IP Telephony and business-to-business integration (for example, supply chain management), requires the network to be able to discriminate between these different types of traffic and dynamically configure its behavior to meet the business goals. This has led to the deployment of "differentiated services" in IP networks. For example, network throughput, packet loss and latency are important considerations for ERP applications. For multimedia and real-time voice applications, minimizing network latency and jitter become critical considerations.

Differentiated services must be instrumented end-to-end. Layer 2 switches support up to eight levels of differentiation using class of service (CoS). Layer 3 switches support 8 to 64 levels of differentiated service using IP Precedence or Differentiated Services Code Point (DSCP). Figure 1 shows packet classification available in Layers 2 and 3 to provide differentiated services.

Figure 1 Packet Classification Methods Available in LAN Switches.

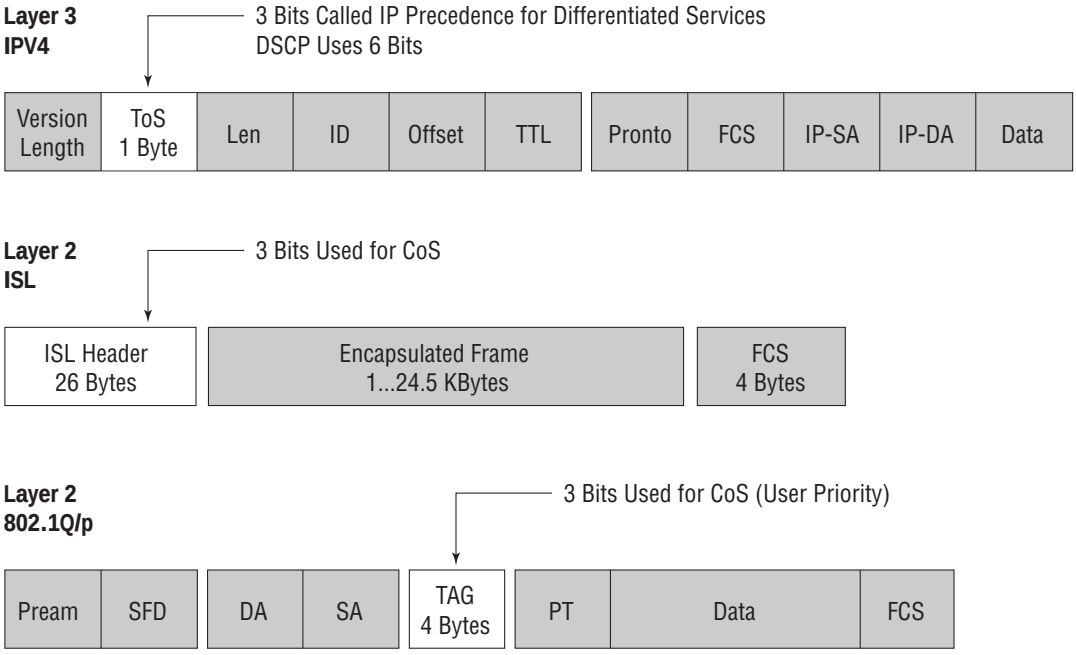


### Structured Network Design

The development of Layer 2 switching in hardware has provided rapid migration from shared to switched networks. Technologies such as virtual LANs (VLANs) were developed to increase scalability and to control broadcast domains. Now, with Layer 3 switching hardware, deployment of cost-effective, high-performance LAN backbones with rich QoS to support critical IP applications has become a reality.

Working with customers all over the world, Cisco has developed structured campus LAN design methodology that will enable customers to build scalable, high-performance, reliable networks. This methodology also offers the following additional benefits:

- It simplifies new network designs and makes it easy to replicate, debug and troubleshoot networks
- It provides for rapid migration from legacy networks and technologies as it preserves existing addressing schemes
- It provides support for critical applications such as IP Telephony, multimedia, and e-commerce
- It supports most common network protocols
- It provides room for network growth without fundamentally changing the architecture

At a high level, as shown in Figure 2, midmarket LANs can be partitioned into two regions—the access or wiring closet layer and the LAN backbone layer. Larger campus networks may also have an intermediate "distribution layer" to provide additional scalability. We will only consider the two-layer design methodology in this document.

The requirements and price/performance at these layers vary depending on the applications supported and the breadth of network services deployed.

Figure 2 A Hierarchical, Two-Layer Architecture for Midmarket Networks

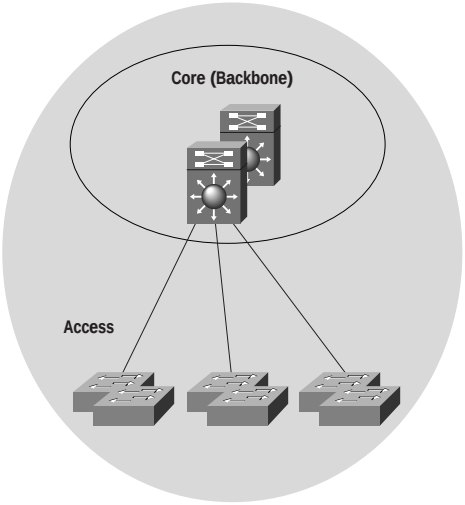


**Access layer**

The access layer is where user computers are attached and workgroups are created. Major considerations for access layer switches include:

- Scalable port density to cover from small workgroups to 1000-user midsized network or enterprise branch. 10/100 auto sensing Layer 2 switches suffice for most applications. To support very high-bandwidth applications on desktop PCs, there may be a need for Gigabit Ethernet speeds at the access layer.
- Superior price performance is essential as hundreds of users access the network through this layer.
- Multiple, high-performance Gigabit uplinks for redundant connections to the LAN backbone layer. Because the access layer supports hundreds of users, redundant uplink configuration is important to the overall network availability.
- Intelligent Layer 2 functions such as VLANs, fast convergent Spanning Tree, and other high-availability features. In the event of topology changes, the uplinks should also converge fast enough to avoid disrupting user access to servers in the network core.
- LAN edge QoS to prioritize applications such as IP Telephony during congestion and help control both delay and jitter within the access layer. At minimum, the switches should support at least two queues per port to be able to prioritize traffic as either high or low priority based on 802.1p/Q class of service (CoS).
- LAN-based telephone power to provide power supply to IP phones attached to workgroups. The switch port incorporates auto-sensing features to detect if telephones are attached to the ports before enabling line power.
- Traffic storm control and pruning to prevent flooding does transmit forwarding data only where required.
- The ability to create stacks with a high-performance stacking bus.
- Scalable embedded, Web-based multidevice, management architecture for simplified management. The embedded software offers LAN management without requiring any client-side application installation.
- Plug-and-play switching, which is critical to minimizing configuration and management efforts. Hence, most access layer switches are intelligent Layer 2 switches. In special cases where elaborate application layer filtering and policies need to be applied, Layer 3 switches can be used at the access layer. However, this may entail additional network configuration.

**Core (Backbone) Layer**

The LAN core layer forms the backbone for the network. All major network services and application servers are attached to the backbone. Virtually all traffic between application servers and clients traverses the backbone. With the advent of wire speed Layer 3 Gigabit switching devices, LAN backbones have migrated to switched Gigabit architectures which provide all the benefits of routing with wire speed packet forwarding.

- A Layer 3 switching core in a structured design reduces (or terminates) the scope of Spanning-Tree domains. It also provides inter-VLAN and inter-subnet routing, thus offloading the router to focus on critical WAN routing functions.
- A Layer 3 core is not susceptible to loops (due to the time-to-live parameter in IP) unlike the Layer 2 network that is susceptible to loops due to the slow spanning tree convergence.
- Use of Layer 3 routing protocols, such as Enhanced IGRP (EIGRP) and open shortest path first (OSPF) help handle difficult tasks such as load balancing, redundancy, and recovery in the backbone.
- Layer 3 designs are highly scalable because they exchange only summary reachability information rather than maintaining path information for every host.
- A Layer 3 switching core also provides rich QoS using IP Precedence or DSCP.

General requirements for the core layer include the following:

- Wire speed Layer 3 Gigabit switching performance as provided by Layer 3 switches such as Catalyst 4908G-L3. In cases where Gigabit performance is not required, Layer 3 Fast Ethernet switches such as Catalyst 2948G-L3 can provide the LAN Core services.
- Redundancy support with rapid failover to maintain high core availability.
- Support for large number of queues per port (at least four) with intelligent scheduling algorithms to manage congestion and to prioritize traffic.
- Intelligent network services such as QoS, security, and load balancing.
- Rich Layer 2 to Layer 4 QoS to support a large number of applications.
- Support for important routing protocols such as routed information protocol (RIP), OSPF, EIGRP, and multicast protocols such as protocol independent multicast (PIM).

## Intelligent Network Services

To help design highly available networks, Cisco has incorporated several intelligent network services into its switches, routers, and other network products. These services deliver high network uptime with high performance for mission-critical applications. Major Cisco IOS® network services available in the Catalyst family of stackable switches are listed below and are discussed in the following sections.

- High-Availability Services
- QoS Services
- Scalable, Multidevice Management Services
- Multicast Services
- Traffic Storm Control
- Security Services

### High-Availability Services

In this section we will highlight the following features:

- Spanning Tree enhancements for rapid Layer 2 topology convergence
  - PortFast, UplinkFast, and BackboneFast features
- Per VLAN Spanning Tree to control the scope of Spanning Tree and increase scalability
- Redundancy and Load Balancing
  - Redundant Gigabit uplinks from Catalyst switches
  - Port aggregation and load balancing using Cisco EtherChannel® technology
  - Switch redundancy using Hot Standby Redundancy Protocol (HSRP)

#### Spanning Tree Enhancements

Layer 2 switches run Spanning Tree Protocol (STP) to discover the Layer 2 topology and to prevent network loops. Spanning Tree enables LAN switches to be connected in redundant manner without creating loops. One drawback of Spanning Tree is that it is recomputed every time there is a change in topology or there is a link or switch failure resulting in loss of connectivity. In typical networks the minimum convergence time for the Spanning Tree is about 50 seconds based on the default values of Spanning Tree timers. Cisco has incorporated several enhancements to the STP to achieve network convergence in the order of 2-5 seconds thus increasing the network availability. Examples of such enhancements include UplinkFast, PortFast and BackboneFast.

UplinkFast helps achieve convergence in the order of 2 seconds upon failure of an active redundant uplink to the LAN backbone. With UplinkFast, the switch instantaneously cuts over to the new Spanning Tree root because the standby redundant link goes directly to the forwarding state from the blocked state. For each downstream MAC address the switch knows how to reach, it sends a proxy multicast frame on the

new root port at a configured update rate to rapidly restore reachability information. The upstream switches update their tables to use the new path.

PortFast minimizes the initialization time for access ports with respect to the STP states. With PortFast enabled, access ports bypass the listening and learning states and directly proceed to forwarding state from blocking state.

BackboneFast enables a switch attached to a backbone to react quickly to indirect faults without really lowering the max_age timer of STP. Connectivity to the backbone is rapidly restored by detecting a link or switch failure using a special Protocol Data Unit (PDU) and then unblocking a blocked uplink as if the max_age timer had expired.

### Per VLAN Spanning Tree

Cisco switches offer Per VLAN Spanning Tree (PVST) while many switches in the industry offer only a single Spanning Tree per switch. PVST increases network scalability by balancing load across different VLANs for connections from the access layer to the LAN backbone. In addition, PVST convergence time may be smaller than that for the larger STP topology for the entire switch. Note that PVST with Cisco enhancements (PVST+) is interoperable with switches that use only a single Spanning Tree per switch.

A detailed example of the benefits of PVST is provided in Section 5.

### Redundancy and Load Balancing

Catalyst switches offer redundancy at link as well as at switch level. The Catalyst 3500 XL family of switches offers two gigabit interface converter (GBIC) uplink ports per switch. These can be connected to two backbone layer switches to provide redundant paths to the services in the core. Customers can also create a Gigabit EtherChannel connection from the two uplinks to aggregate uplink performance as well as provide load balancing.

EtherChannel technology provides aggregation of multiple ports with similar characteristics into a single logical port to provide both higher bandwidth and link redundancy. EtherChannel technology is supported for both Fast Ethernet and Gigabit Ethernet. Typically, it is used in creating higher-bandwidth uplinks between access and backbone layers. EtherChannel technology can also be used to achieve high-speed connections to the servers. When all the ports in the port group are operational, load balancing is achieved. In the event of a port/link failure, the other available ports carry the load.

Using HSRP, switch level redundancy can be obtained at Layer 3 for the LAN Backbone. Both Catalyst 4908G-L3 and Catalyst 2948G-L3 switches support HSRP. A virtual IP address is created for the two switches connected in a redundant fashion. When one of the switches fails, the other seamlessly takes over with no change to the assigned virtual IP address. HSRP is used to determine the status of the switches and detect any failures.

Figure 3 Some Connectivity Examples Using Gigabit EtherChannel and Fast EtherChannel Technologies

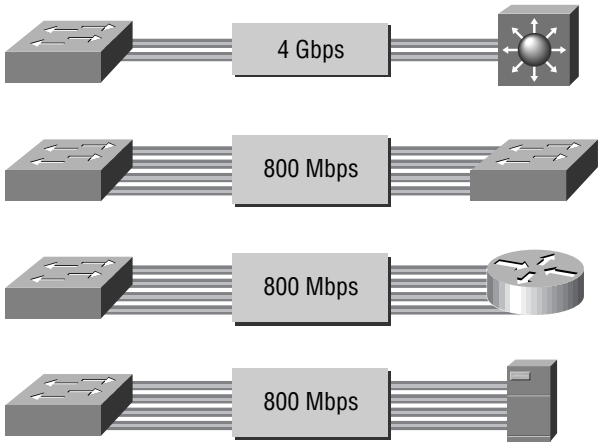

**Quality of Service (QoS) Services**

Intelligent QoS is essential to classify traffic based on business importance and application priority. QoS is an end-to-end network architecture, and this is especially true for applications such as IP Telephony. The Catalyst 2900/3500 XL family of switches provides the LAN edge QoS that can integrate with the QoS features of campus core products such as Catalyst 4908G-L3 or Catalyst 6500 to achieve the end-to-end QoS.

LAN edge QoS functions typically involve inspecting the incoming packet, classifying the packet, queuing it based on the number of queues per port supported by the edge switch, and scheduling the packet for forwarding. In a QoS-aware network, preferential forwarding is provided for mission-critical traffic such as IP Telephony. Campus core switches, such as the Catalyst 4908G-L3 and Catalyst 6500, provide a rich set of QoS functions to prioritize mission-critical applications. Features include IP Precedence or DSCP, multiple queues with sophisticated scheduling techniques such as Weighted Round Robin (WRR), and congestion avoidance techniques such as Weighted Random Early Detection (WRED).

**Scalable, Multidevice Management Services**

Innovative Switch Clustering technology from Cisco provides unparalleled single IP address management of groups of switches (up to 16 switches in each cluster) while eliminating all the deficiencies of traditional stacking. Cisco Switch Clustering, available in Catalyst 2900 and 3500 XL series, eliminates the geographical boundaries of wiring closets, adds embedded, Web-based management with no complex commands to learn, and spans different types of equipment and networks. Clustering dramatically minimizes the administrative efforts by providing cluster-wide abstraction for configuration, monitoring, and troubleshooting of the switches.

Figure 4 Cisco Switch Clustering Technology Provides Single IP Address Management for Catalyst 2900/3500 Workgroups

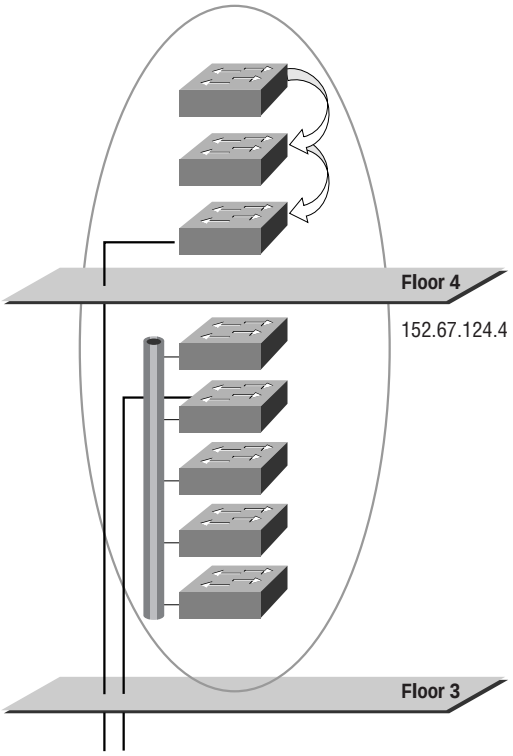


A master switch (called the command switch) provides the proxy and redirection services for single IP address management per cluster. All cluster management commands are targeted to the command switch IP address. For redundancy, a second switch can be assigned an IP address and the overall cluster can then be managed using a single virtual IP address. In the event the primary command switch fails, the back up or secondary command switch seamlessly takes over the management of the cluster while the user still accesses the cluster via the virtual IP address. Figure 4 illustrates the concept of cluster management supported by Catalyst 2900 and 3500 XL Layer 2 switches.

The Cisco Cluster Management Suite, Web-based management software embedded in Catalyst 2900/3500 XL switches, also integrates with campus management applications such as CiscoWorks2000 to provide end-to-end management.

### Multicast Services

IP Multicast enabled applications such as training, distance learning, and conferencing are becoming increasingly vital to business operations. A well-designed network does not flood these types of traffic throughout the network; instead, it dynamically sends the traffic only to those clients that have registered. Protocols such as Cisco group management protocol/Internet group management protocol (CGMP/IGMP) with features such as Fast Leave and Join help deliver efficient multicast through dynamic client registration and intelligent pruning (Figure 5). This saves network bandwidth and client CPU cycles, thus increasing multicast scalability.

Cisco Layer 3 switches such as Catalyst 4908G-L3 and Catalyst 2948G-L3, support other multicast protocols such as IGMP and PIM to provide end-to-end network interoperability.

Figure 5 Example of Dynamic Registration and Intelligent Pruning for Multicast

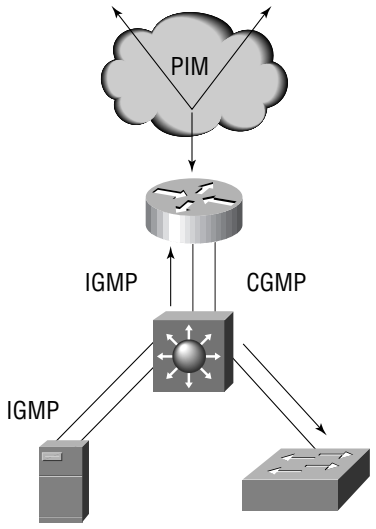


### Traffic Storm Control

Catalyst switches support broadcast, multicast, and traffic storm on a per port basis. In the event of traffic storms the erring ports are disabled based on user defined thresholds. This will prevent network downtime and increase bandwidth available for mission- critical applications.

### Port Security and TACACS+

These services, offered by Catalyst switches, help control access to network. Using port security, administrators can lock down a switch port to a secure MAC address. This will prevent unauthorized systems from getting on to the network.

With TACACS+, access to device command line interface (CLI) can be controlled. A log is also maintained of such access.

## Building Blocks for Structured Design

In this section we present three key building blocks with well-defined performance characteristics and configuration guidelines. Using these building blocks, we illustrate several midmarket designs that meet a broad range of performance, cost, and availability criteria. Overall architecture can be preserved while providing network growth. Note that several of the intelligent network services and redundancy features discussed earlier come standard in the Catalyst switches at no additional cost.

### Building Block #1: Cost-Effective Wiring Closet Using GigaStack Technology

Figure 6 Cost-Effective Wiring Closet Using Catalyst 3500 XL Switches

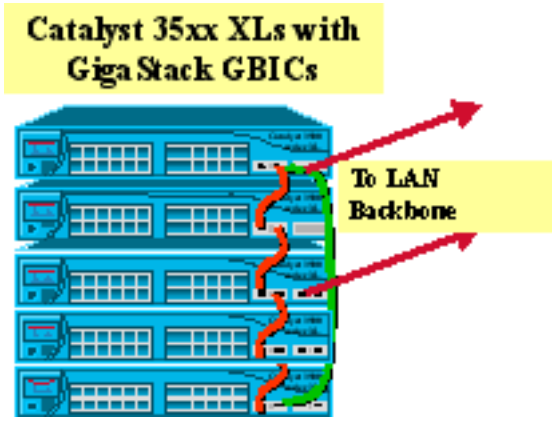


Key Benefits

- Cost-effective; can stack up to nine 3500 XL switches per stack; with Catalyst 3548XL, a stack can support up to 432 users
- 8 Mpps maximum switch performance
- One Gbps stack performance
- Two GBIC ports per switch for cost-effective redundancy; one is used for stacking and the other can be used for uplink to the LAN backbone
- Redundant stack loop back for high availability
- QoS based on 802.1p/Q and user-defined per port priority
- Building block can be clustered for single IP management

Configuration Guidelines

- PortFast = ON
- UplinkFast = ON for the workgroup switches
- STP Max_Age = 2 Seconds, FWD_DELAY = 5 seconds for the stack
- In general, limit the number of VLANs per switch and per stack
- For edge QoS, select CoS values 4-5 for high priority traffic; use values 6 and 7 for network control traffic—this will provide maximum benefits when an upstream device maps CoS to L3 DSCP

**Building Block #2: High-Performance Workgroup in Star Configuration**

Figure 7 A High-Performance Workgroup with Star Configuration to Layer 3 Backbone

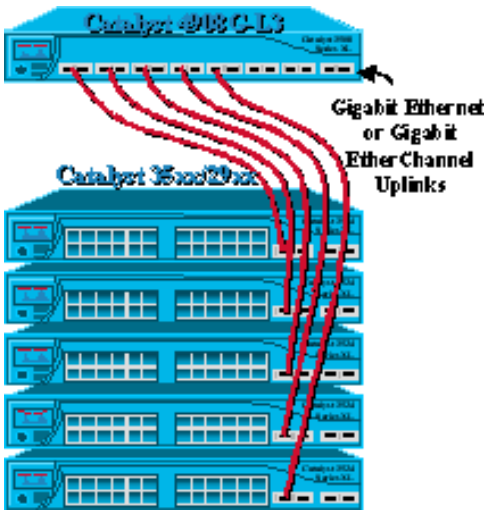


**Key Benefits**

- High-performance workgroup; 2 Gbps full-duplex uplink performance; the Layer 2 switches provide 10 Gbps switch performance and the Catalyst 4908G-L3 provides 22 Gbps performance; for networks with only a few subnets, this design becomes the lowest-cost Layer 3 configuration
- Two GBIC uplink ports per switch for cost-effective redundancy
- Gigabit and Gigabit EtherChannel performance and load balancing from aggregation layer; supports up to 288 users with 2 Gbps EtherChannel uplink to the core or 192 users with 4 Gbps EtherChannel uplink
- Wire-speed Layer 3 Gigabit backbone performance and inter-VLAN routing provided by the Catalyst 4908G-L3; this offloads the router to focus on WAN functionality
- QoS to support delay-sensitive applications such as IP Telephony
- Workgroup 2900/3500 switches can be clustered for single IP management
- GBIC-based gigabit ports also provide flexibility in media and distance options; Catalyst 2900/3500 XL provide support for GigaStack copper (1m), SX fiber (550 m), LX fiber (10 km), and ZX fiber (100 Km) GBICs

Configuration Guidelines

- All items under 5.1.2
- Create Gigabit EtherChannel connections to the core layer for higher performance, redundancy, and load balancing

**Building Block #3: LAN Backbone with High Availability**

Figure 8 Gigabit Backbone Using Catalyst 4908G-L3 Switches in Redundant Configuration

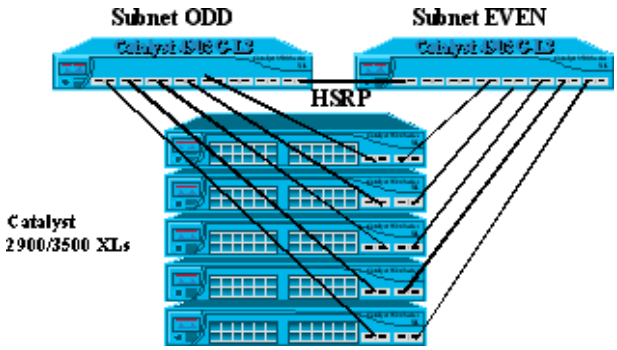


Key Benefits

- All the benefits of 5.2.1
- Tolerant to single link and single switch failure at the aggregation layer
- Layer 2 redundancy provided by redundant links from the workgroup 3500 XL switches; Layer 3 redundancy between the two 4908G-L3 provided by HSRP
- Gigabit and Gigabit EtherChannel performance, redundancy, and load balancing from aggregation layer; supports up to 288 users with 2 Gbps uplink capacity to the core or up to 240 users with two 2 Gbps EtherChannel uplink for a total capacity of 4 Gbps

Configuration Guidelines

- For higher availability configure HSRP for Layer 3 redundancy between the two Catalyst 4908G-L3 switches; in the event of a failure of one of the two switches, the other switch takes over
- Distribute inter-VLAN routing between the two Catalyst 4908G-L3 backbone switches, based on even and odd subnets; this provides scalability and load balancing
- For end-to-end QoS, select CoS values 4–5 for high priority traffic at the edge; use values 6 and 7 for network control traffic

## Design Examples

### Cost-Effective, High-Performance, 100-User Business Network

Figure 9 A Low-Cost, 100-User Network Using Catalyst 3500 XL Switches

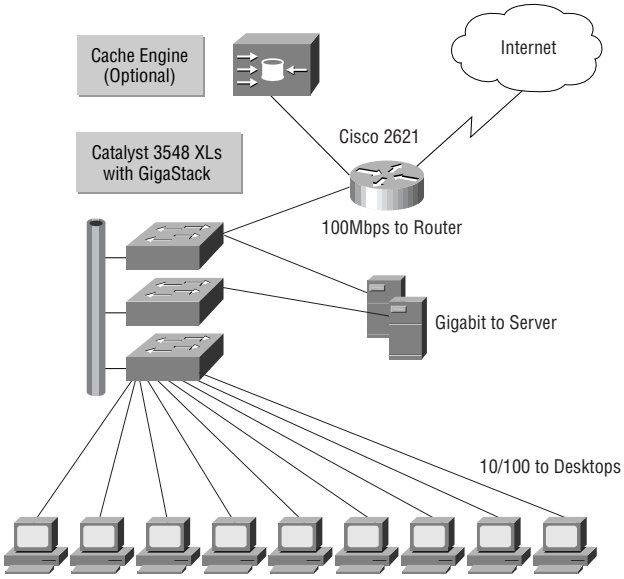


#### Business Goals

Company ABC has 100 employees. The network is mainly used for e-mail, file sharing, and customer database access. Internet access is used for market and competitive research and browsing. Low initial network cost and investment protection are important design requirements.

#### Example Design

This design is optimized for cost using the GigaStack™ building block. However, the features of Catalyst 3500 XL also provide high performance. One of the two Gigabit GBIC ports on each of the 3548 XL is used to create a GigaStack configuration with a performance of one Gbps. GigaStack loop back is enabled to create a high availability stack. The other Gigabit GBIC port can be used to attach to the servers. The cost of the Gigabit NIC card on the server is distributed over the total number of workgroup users.

A Fast Ethernet port is connected to the router to provide Internet access. Firewall, VPN, and Network Address Translation (NAT) features are enabled at the router. If there is a lot of Web (HTTP) activity, we recommend that a Cisco CacheEngine be added to the network to save on WAN costs.

If Gigabit performance is not required from the server, Catalyst 2900 XL switches with 10/100 ports can be used. Note that Catalyst 3500 XL offers two gigabit uplink ports for increased performance and scalability.

### A Low-Cost, Layer 3, Collapsed Back Bone Design

#### Business Goals

Company XYZ has predominantly Fast Ethernet network and has suddenly grown to 500 employees due to a recent acquisition. There is significant overhead on the router in the head quarters due to the increase in number of VLANs. LAN performance has also deteriorated due to the increase in the number of users.

In addition, the company would like to save on telephone costs for communication with the smaller branches.

#### Example Design

Using a Catalyst 4908G-L3, the company creates a high-performance Gigabit LAN backbone. Workgroups are created using Catalyst 3500 XL switches. GigaStack is used for workgroups with Gigabit uplink to the Backbone. Common application servers and the CallManager server are upgraded with Gigabit network interface cards and directly connected to the backbone.

The Catalyst 4908G-L3 provides inter-VLAN routing and offloads the router to focus on WAN access. The optional Cache Engine helps save WAN bandwidth for Web (HTTP) traffic.

The Cisco CallManager provides the call-processing software component of the Cisco AVVID (Architecture for Voice, Video and Integrated Data) network. Cisco CallManager runs on a dedicated server and extends telephony features and functions to packet telephony network devices such as IP phones, software phones, and IP Telephony gateways.

Figure 10 Collapsed Backbone Design Using Catalyst 4908G-L3 Gigabit Switch

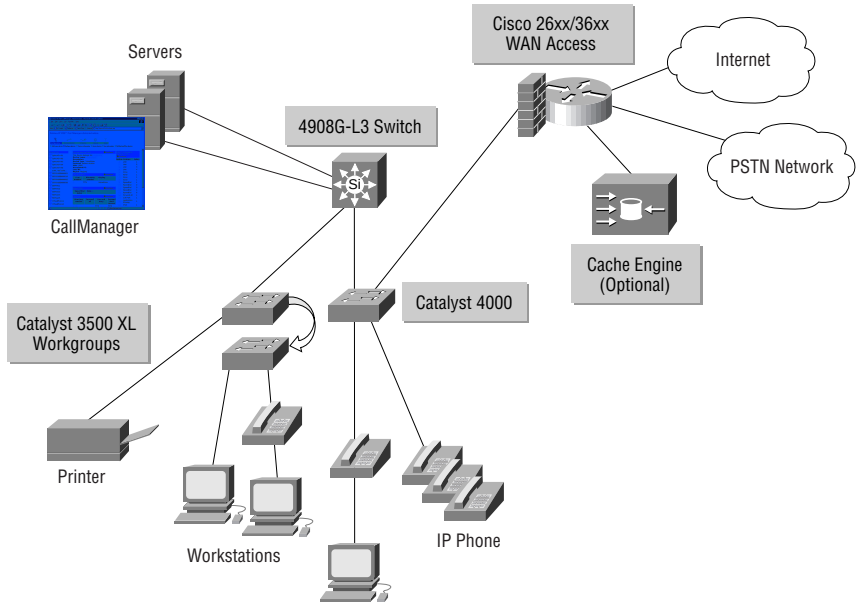


In midmarket designs where gigabit performance is not required, the Catalyst 2948G-L3 can be used to create a Layer 3 backbone with Fast Ethernet performance, as shown in Figure 11. Servers are connected using Fast Ethernet or Fast EtherChannel technology. The Catalyst 2948G-L3 performs inter-VLAN routing, thus, offloading the WAN router. The workgroups are created using Catalyst 2900 XL 10/100 switches.

Figure 11 A Midmarket Network with Fast Ethernet Collapsed Backbone Design

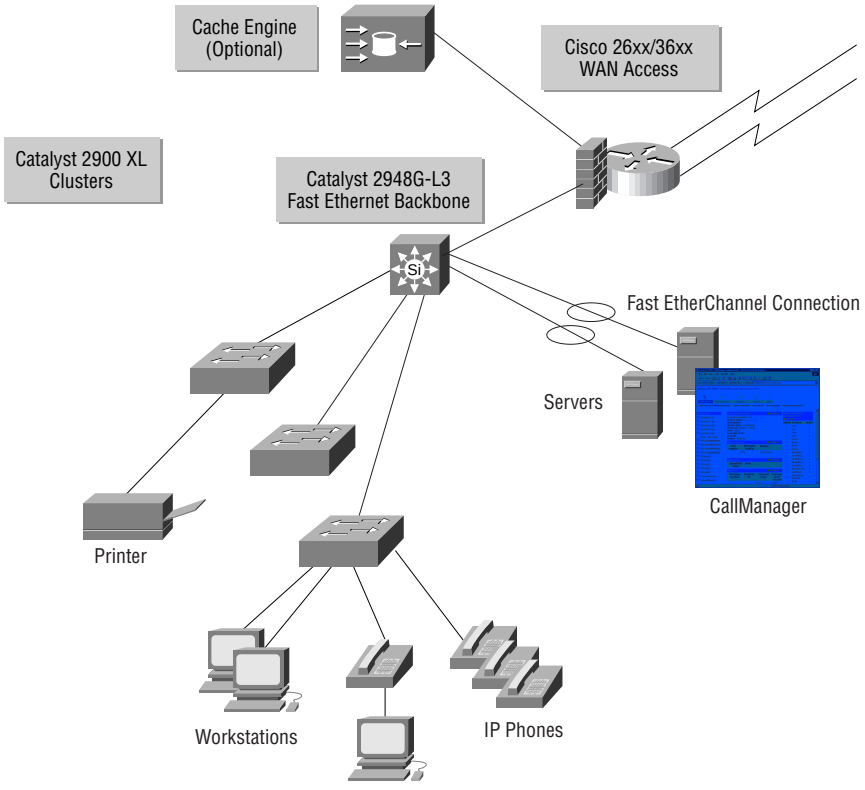

### A Highly Available Midmarket Company Network

#### Business Goals

MyCompany has more than 800 employees at the head quarters and more than 100 sales people in branches in major cities. Many supply chain partners and sales personnel regularly access applications and databases at the head quarters; hence, it is critical that the network be available at all times. Any network outage directly translates into lost revenue. MyCompany is actively deploying Internet and IP Telephony technologies.

#### Example Design

In this design, the main focus is high availability from the WAN edge all the way to the servers. The routers and firewall are configured in redundant manner using HSRP. Firewall functionality can be obtained by either using dedicated devices such as a Cisco PIX™ Firewall, or by enabling IOS® firewall features in the routers. The LAN backbone switches Catalyst 4908G-L3s are connected in a redundant fashion using HSRP. To achieve greater scalability and load balancing, inter-VLAN routing is distributed between the two core switches, for example, based on odd and even numbered subnets. One-to-one mapping between Layer 2 VLANs and Layer 3 subnets is established to integrate the workgroups and to minimize the number of hops.

GigaStack with stack redundancy is used where one-Gigabit uplink performance is sufficient for the workgroups. For high performance workgroups, star configuration of Catalyst 3500 XL switches using gigabit or Gigabit EtherChannel technology is created. IP Telephony is used for telephony between the head quarters and the major branches.

A similar design can also be used for an e-commerce company or enterprise branch network where network robustness and redundancy are critical to the business.

Figure 12 A Fully Redundant Design for Very High Availability Branch Network

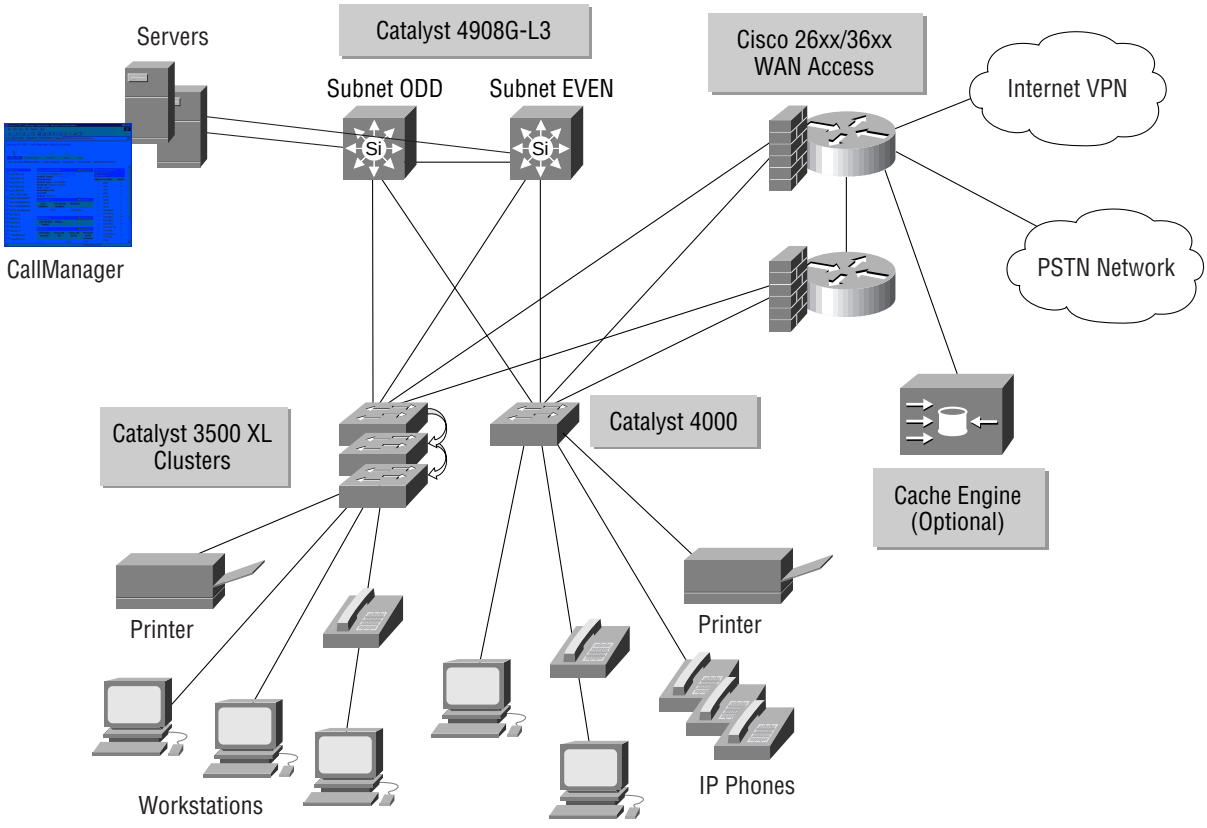

### A Campus LAN Supporting More than 1000 Users

Large campus network design is not the focus of this paper. However, for midmarket LAN backbones that aggregate a large number of gigabit ports we recommend the Catalyst 6500 multilayer switch, which can aggregate up to 130 gigabit ports. The design is similar to that in Section 6.2 except that the collapsed backbone is created using the Catalyst 6500. The WAN router can also be a 7200 or 7500 series depending on the number of interfaces required.

Workgroups created using the Catalyst 4000 or Catalyst 3500 are connected to the core using gigabit or Gigabit EtherChannel uplinks. Multiple gigabit uplinks can be used for redundancy. The core Catalyst 6500 switch can also be connected in redundant fashion for high availability.

Please refer to documentation on Catalyst 6500 family of switches for additional information on designing large campus networks.

Figure 13 A Large Campus Network with End-to-End QoS

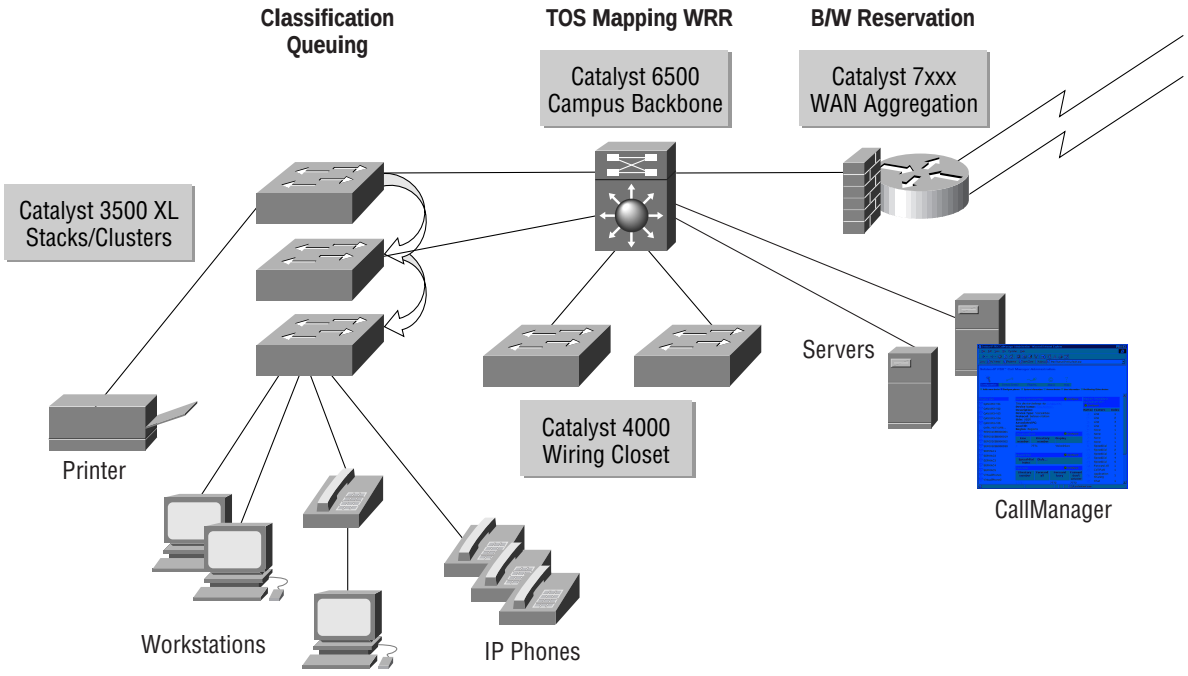


### Summary and Quick Reference

The following table summarizes the product features from the Catalyst family of switches used in the design examples presented in this document. The design examples outlined here provide scalable performance from switched Ethernet to switched Fast Ethernet to switched Gigabit Ethernet, common IOS software architecture, intelligent network services such as QoS and Access Control Lists and industry-leading management with CiscoWorks2000. These features provide a solid network infrastructure for midmarket customers to build simple, hierarchical, easily managed, highly available networks that will not only meet current business needs but also support emerging applications such as multimedia, multicasting, and IP Telephony.

Detailed product information can be obtained at www.cisco.com or from your Cisco Reseller.

Table 1 Product Features Summary

| Feature | Catalyst 2900 XL Series | Catalyst 3500 XL Series | Catalyst 2948G-L3 | Catalyst 4000 | Catalyst 4908G-L3 |
|---|---|---|---|---|---|
| **Product Focus in Midmarket Designs** | Access Layer Workgroups | Access Layer Workgroups | Fast Ethernet Layer 3 Backbone | Access Layer Workgroups | Gigabit Ethernet Layer 3 Backbone |
| **Number of switched 10/100 Ethernet Ports** | 12, 24 (Layer 2) | 12, 24, 48 (Layer 2) | 48 (Layer 3) | 240 | None |
| **Number of switched Gigabit Ethernet Ports** | Up to 2 in 2924M | 2 (Layer 2) | 2 (Layer 3) | 32 | 8 (Layer 3) |
| **GBIC Options** | GigaStack connection, SX, LX, LH | GigaStack connection, SX, LX, LH, ZX | SX, LX, LH, ZX | SX, LX, LH, ZX | SX, LX, LH, ZX |
| **Switching Performance** | 3.2-Gbps Backplane<br>3 Mpps | 5-Gbps Backplane<br>8 Mpps | 22-Gbps Backplane<br>10 Mpps | 24-Gbps Backplane<br>18 Mpps | 22-Gbps Backplane<br>10 Mpps |
| **IOS Software** | Yes | Yes | Yes | Yes (Optional) | Yes |
| **QoS** | 2 queues/port with priority scheduling L2 CoS | 2 queues/port with priority scheduling L2 CoS | 4 queues/port with WRR scheduling | 2 queues/port with priority scheduling | 4 queues/port with WRR scheduling |
| **Routing Protocols** | NA | NA | RIP, OSPF, IGRP, EIGRP | RIP, OSPF, IGRP, EIGRP (Optional) | RIP, OSPF, IGRP, EIGRP |
| **CGMP** | Yes | Yes | Yes | Yes | Yes |
| **IGMP** | No | No | Yes | Yes | Yes |
| **Access Control Lists** | No | No | Yes | Yes | Yes |
| **TACACS+ Support** | Yes | Yes | Yes | Yes | Yes |
| **CiscoView Device Management** | Yes | Yes | Yes | Yes | Yes |
| **CiscoWorks2000 Support** | Yes | Yes | Yes | Yes | Yes |


**Corporate Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

**European Headquarters**
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
http://www-europe.cisco.com
Tel: 33 1 69 18 61 00
Fax: 33 1 69 28 83 26

**Americas Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

**Asia Headquarters**
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001


**Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at http://www.cisco.com/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela



01/00 BW5835