



LAN Baseline Architecture Overview—Branch Office Network

This document provides guidance on how to design a local area network (LAN) for a Business Ready Branch or autonomous Business Ready Office where corporate services such as voice, video, and data are converged onto a single office network.

This document provides an overview of LAN architecture. Because of the numerous combinations of features, platforms, and customer requirements that make up an office design, this version of the design guide focuses on various LAN design discussions for voice and data services without making specific design recommendations.

LAN Services Overview

LAN services provide connectivity to end devices into the corporate network within the office. With the convergence of services onto a single network infrastructure, devices such as computers, telephones, surveillance cameras, cash registers, kiosks, and inventory scanners all require connection to the corporate network via the LAN. This assortment of devices requires simplified connectivity tailored to the demands of each device. For example, devices such as IP telephones or cameras may be powered via the LAN switch, automatically assigned an IP address, and be placed in a virtual LAN (VLAN) to securely segment them from the other devices. Wireless access points may be used to provide secure mobile access for laptop computers, scanning devices, wireless IP phones, or kiosks. These are just a few examples of the LAN services that are used in the Business Ready Branch or Office solution.

In addition to providing the integrated voice, video and data services for the employees, branch offices also require guest network access, and in some cases should support demilitarized zones (DMZs). The guest access can be for partners or customers, and guest access includes both wired and wireless access.

Regardless of the presence of DMZ, security in branch offices is a key element of branch LAN services. The LAN must be protected against malicious attacks, and the users accessing the corporate network must be authorized/authenticated.



Americas Headquarters:

Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2007 Cisco Systems, Inc. All rights reserved.

Branch LAN Design Considerations

Branch LAN infrastructure provides connectivity to the end devices to access the corporate network. In a small office and even a medium-sized branch office, the resources are typically located at the corporate headquarters and accessed through a wide area network (WAN) of varying bandwidth. For certain branch offices, a limited amount of end user connectivity is desired, and these end users access the computational resources at the corporate headquarters. However, it is also desired that the computational resources be deployed in certain branch offices. In such a case, in addition to providing connectivity to the corporate headquarters, the branch LAN must meet additional requirements. Based on these computational and connectivity requirements, branch offices can be categorized into the following categories:

- Small branch (up to 50 users)
- Medium branch (up to 100 users)
- Large branch (up to 200 users)

The small branch office is typically characterized by small number of users, usually less than 50 users. The medium branch office is up to 100 users. The large branch office should accommodate up to 200 users. Typically, secure connectivity to the corporate headquarters is the main focus for small- and medium-sized branch offices. In a small- and medium-sized office, the following issues must be considered when deploying the LAN:

- Coverage considerations for wireless LAN (WLAN) users in a branch office
- Distance considerations from the closet to the desk for wired clients
- Inline power requirements for all IP phone users in the branch office
- Security, and manageability considerations

For the large branch office, several services and computational resources must be provided as well as end user connectivity to the corporate office. These services are typically handled by well-defined entities in campus environments. These entities have their own LAN design and tie into the campus core. The following services are expected to be provided in large branch office designs in addition to the services mentioned above for small and medium sized branch offices.:

- DMZ and small server farm
- Wide area file services
- Local authentication (survivability) for users
- Security services such as intrusion detection/prevention
- High availability and scalability

Deployment of the above features/services means increased switching capabilities for the LAN. The network must not only be designed to meet current requirements, but should scale and be able to accommodate value-added services without having to redesign the entire network.

These additional requirements for a large branch office LAN are met by a multilayer LAN architecture. The following section provides more details about the considerations and capabilities of a multilayer branch LAN architecture.

Multilayered Branch Architecture

Typically, the branch LAN infrastructure is logically similar to the campus LAN infrastructure. However, because of the differences in scalability, high availability, manageability, and cost considerations, the network devices deployed can be different in branch and campus environments. Even when some of the low-end devices that are used in both branch and campus LAN environments are the same, the devices upstream that aggregate the traffic are different, and the ways in which the network is designed to accommodate the branch requirements are significantly different from the campus LAN environment.

The following are the main design criteria for designing a branch office LAN:

- **High availability**—A redundant path should be provided for the traffic in case of device or link failure.
- **Scalability**—The architecture should accommodate the addition of more users and services without major changes to the infrastructure.
- **Security**—The network should be secure to exclude unauthorized users and prevent malicious attacks.
- **Manageability**—The network should be simple to deploy, troubleshoot, and manage without compromising high availability, security, and scalability.

Multilayered architecture provides several strengths. The layers are clearly defined, providing modularity; each device in a layer performs the same function, thereby making the configuration simpler in a modular design. The multilayered design also makes it easier to troubleshoot network problems, and provides scalability and high availability. Specifically, with a limited number of Layer 2 versus Layer 3 ports available on the router, the multilayered architecture provides support for more users, and also helps in providing a good integration point with the edge router. The multilayered architecture also provides traffic separation between layers and reduces CPU utilization on the router; for example, by transferring some of the functions from the edge to the distribution, the CPU on the router is freed from performing those functions. If required, this architecture also provides an integration point for various technologies without the need to redesign.

The benefits of multilayered architecture can be summarized as follows:

- Simplifies configuration
- Provides modularity
- Facilitates troubleshooting
- Scales well
- Provides traffic separation
- Provides CPU load sharing
- Provides a hook to add additional services without having to redesign the network

A multilayered branch LAN architecture can be divided into the following layers:

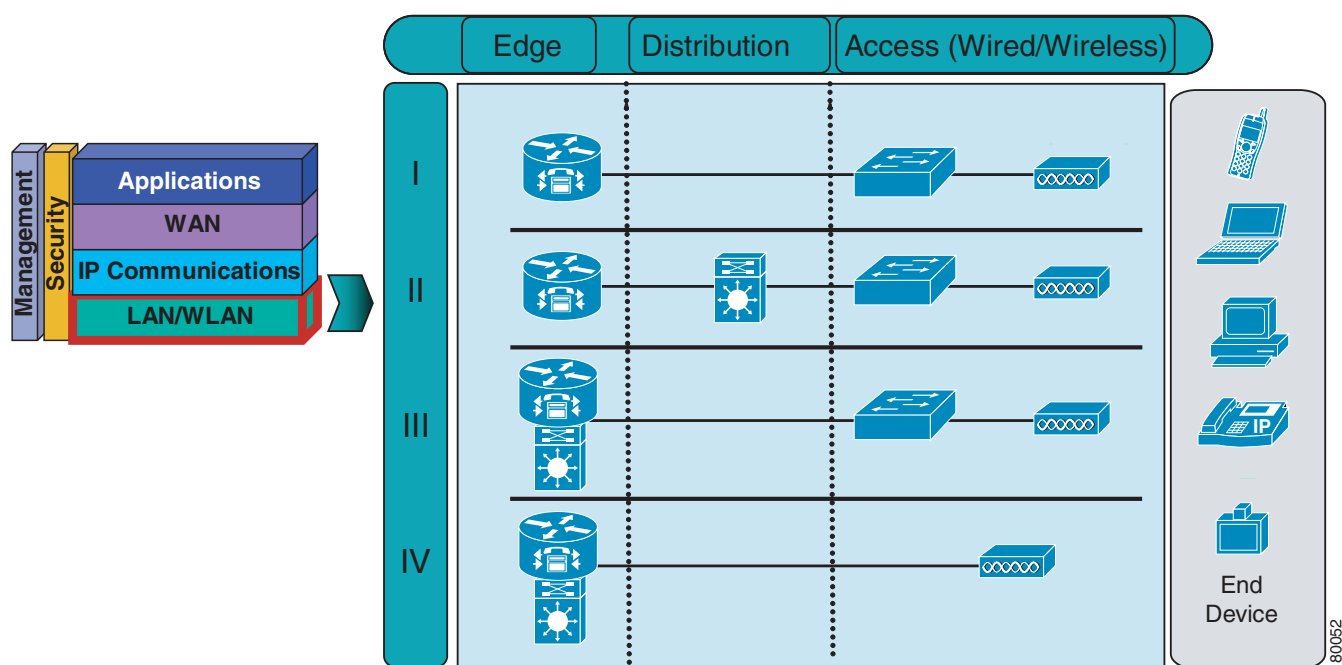
- **Access layer**—Provides connectivity to end users, either via wireless or wired network. L2 security, authentication, and wireless services are also addressed at this layer.
- **Distribution layer**—Provides DHCP, routing, and policy-based routing (PBR) while migrating to advanced services such as segmentation or guest access.
- **Edge layer**—Provides WAN, firewall, intrusion protection system (IPS), voice services, L3-type traffic and an exit point to the rest of the network. Only integration to the edge layer is discussed in this design guide.

Figure 1 shows the various layers of a branch multilayered architecture, and also shows various ways in which a branch office network can be designed.

The architecture should be highly available as well as scalable. Based on the products available, and the scalability and high availability requirements, the architecture can be modified without losing the distinct services offered by each layer. The various possibilities are shown in Figure 1. The most flexible option is the second option (II) in Figure 1, which provides high availability as well as scalability. The number of access switches supported can be scaled easily, thereby increasing the number of users.

The distribution layer can be collapsed into the edge, or the distribution and access layers into the edge, based on high availability or scalability requirements.

Figure 1 Layers of a Multilayered Branch Architecture




Note

Small branch LAN offices can use integrated switching at the edge, and might not have to resort to a multilayer architecture, depending on the number of users and the size of the office. Also, some of the integrated switches for ISR, do not provide the advanced spanning tree and security features that are important for quick convergence in case of switch or link failure in a highly available branch office architecture. High availability and scalability requirements are met by adopting a multilayered architecture. Medium and large branch offices must adopt some variety of multilayer architecture.

Services

Figure 2 shows the services at various layers of the branch architecture.

Figure 2 Services at Various Layers of a Branch Architecture

Edge layer services include WAN, firewall, intrusion detection and prevention, and voice. Edge layer services and details about the edge design are not covered in this document. Only the integration of the edge with the LAN is covered in this document.

Distribution layer services include DHCP, routing, and if required, PBR, while migrating to advanced services such as segmentation or guest access. The distribution layer can be used to add additional services if required. Examples of these services include LAN Controller and wireless domain services (WDS) for WLANs, and appliance-based firewalls or IDS/IPS.

The access layer provides wired and wireless connectivity to end users. The access layer mainly provides Layer 2 security, authentication, and wireless services. Details of the access and distribution services are provided in the following sections. The design options are described in the *LAN Baseline Architecture Branch Office Network Reference Design*

Guide <http://www.cisco.com/en/US/docs/solutions/Enterprise/Branch/Design.html>.

Access Layer

The user connects to the network via the access layer using either a wired or wireless connection. The access layer can also provide the following value-added services:

- Voice and data VLANs to segregate voice and data traffic
- Layer 2 security to protect against malicious attacks
- Quality of service (QoS) to prioritize traffic and also to protect against denial of service attacks and worm mitigation
- Authentication services such as dot1X and IBNS


- Guest services or guest VLANs at the access layer
- Network Admission Control (NAC) to protect against viruses

With many of these services provided at the access layer, the best design practice should integrate all these services seamlessly either at Layer 2 or Layer 3 access. The following sections provide more details of the considerations that go into the design of an access layer and the various elements of the access layer.

Layer 2 versus Layer 3 at Access Layer

There are two options for the switches in the access layer. The first option is to use Layer 2 at the access layer, and the second option is to enable routing and to use VLANs to place users in different groups at the access layer. These two options are shown in [Figure 3](#).

Figure 3 *Layer 2 versus Layer 3 at the Access Layer*




Layer 2 Access

Traditionally, the switches deployed at the access layer operate at Layer 2, which can result in the following two spanning tree issues for some customers:

- Troubleshooting is more difficult
- Convergence in high availability designs can take longer in case of switch or link failure

These problems arise in a traditional, highly-available architecture. In a traditional design, two distribution switches and an access switch are involved with a Layer 2 loop, as shown in [Figure 4](#).


Figure 4 *Traditional Highly-Available LAN Design*



The Layer 2 access switch is connected to both the distribution switches, and the distribution switches are connected together by a trunked EtherChannel. Typically, the Layer 2 topology is designed in such a way that the spanning tree blocks predetermined links so that the traffic takes a deterministic path under both normal and failure circumstances. The convergence problem is addressed by Rapid Spanning Tree, which converges in the sub-second range under failure conditions. Misconfiguration always causes problems when troubleshooting, but by following the appropriate design guide, this should not be a problem for a trained engineer.

Figure 5 provides a different approach for using Layer 2 at the access layer.

Figure 5 *Highly-Available LAN Design with No Layer 2 Loops*




This topology uses stackable switches at the distribution layer instead of two distribution switches running Hot Standby Routing Protocol (HSRP). This topology is highly available and scalable. In this topology, the Layer 3 redundancy is built into the stack. High availability between access and distribution is provided by using EtherChannels. This topology has no Layer 2 loops. However, spanning tree should be enabled and configured to mitigate any accidental Layer 2 loops. Layer 2 at the access layer also makes the integration of various technologies easier, and also provides more flexibility. If appliance devices have to be used rather than service capabilities on the ISR for higher throughput reasons, the hierarchical design with Layer 2 at the access layer provides more flexibility to integrate the appliances.

Layer 3 at the Access

Layer 3 in the access brings a different perspective to the solution. In this solution, routing is enabled on the access switch but still provides the capability to put end users in different VLANs. Routing on the access switch implies using a platform that supports routing and switching. [Figure 6](#) provides some details of a Layer 3 access solution. The access switch provides equal cost multiple paths to the core/edge device. Under failure circumstances, the convergence can very well be in sub-seconds with EIGRP.

Figure 6 **Layer 3 at the Access Layer**



Note

It is quite possible that sub-second convergence under failover scenarios is achievable with both EIGRP and OSPF routing protocols. The testing has not been done.

Layer 3 at the access is not recommended in the branch office designs because of the following reasons:

- Higher costs involved with deploying such a solution.
- The Layer 2 access solution provides a platform to seamlessly integrate all the various services discussed in the previous sections.
- The distribution layer ties all the services together and is analogous to the campus core.

Adopting routing at the access layer creates a very thin Layer 2 domain on the switch. This Layer 2 domain provides the necessary VLANs for the end devices. With this solution, there are no Layer 2 loops, and spanning tree influence is diminished such that it can be disabled to make troubleshooting easier. From a Layer 2 perspective, this is a low maintenance and quick convergence solution. Although it might increase the cost in some situations because the low-end Layer 2 switches cannot be used in this solution, it is a viable solution if cost is not a factor.

VLANs and Spanning Tree Protocol

VLANs are Layer 2 broadcast domains. The traffic in a VLAN is confined to the VLAN until it is routed either into a different VLAN or into a traditional Layer 3 network. A VLAN consists of several end systems, either hosts or network equipment (such as switches and routers), all of which are members of a single logical broadcast domain. The network devices can be members of different VLANs as well. Traffic between two switches that are members of different VLANs is carried on a common link between the two switches while maintaining the broadcast domain. These links are called trunks. Several trunking protocols form a trunk between two switches.

VLANs help to segregate the traffic from different endpoints. For example, voice, video, and data can be segregated by putting the devices into different VLANs. VLANs are also widely used to segregate different users. The proliferation of VLANs results in the various types of spanning tree protocols; Spanning Tree Protocol (STP) prevents loops from being formed when switches or bridges are interconnected by multiple links. Spanning tree protocols such as IEEE 802.1D spanning tree, which was initially designed to protect from Layer 2 loops, has evolved. Multiple types of spanning tree protocols can currently be deployed when using VLANs, including the following:

- Common Spanning Tree
- Per VLAN Spanning Tree (PVST)
- Per VLAN Spanning Tree Plus (PVST+)
- Multiple Instance Spanning Tree (MISTP/802.1S)
- Rapid Spanning Tree (RSTP/802.1W)

**Note**

More information on STP can be found at the following URL:

http://www.cisco.com/en/US/partner/tech/tk389/tk621/tsd_technology_support_protocol_home.html

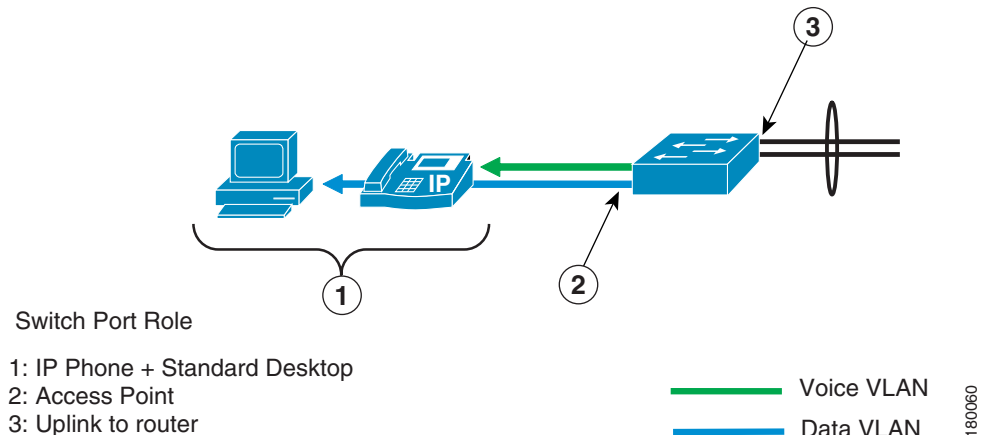
The limitations of 802.1D standard spanning tree protocol, such as slow convergence after a topology change, are eliminated by RSTP. RSTP also supports Cisco specific features such as PortFast, UplinkFast, and BackboneFast for faster network convergence. With RSTP, the convergence time is reduced to a few hundred milliseconds as opposed to the standard 30 to 40 seconds with the 802.1D standard. Cisco switches that support 802.1s/w spanning tree protocols should be deployed to achieve quick convergence.

Voice and Data VLANs

Convergence of voice and data into a single infrastructure lowers the overall cost of ownership of a network, and simplifies administration and maintenance through the elimination of separate voice and data infrastructures.

Convergence also implies that to provide the reliability and quality for the voice and data applications, the traffic type has to be identified at the edge of the network so that appropriate QoS parameters can be applied to the traffic.

Cisco switches allows both the voice and data devices to be connected to a single physical port. On Cisco switches, the concept of access port has been extended, and it is possible to configure a voice and data VLAN. The switch can now receive traffic on two VLANs, as shown in [Figure 7](#).

Figure 7 Data and Voice VLAN on a Switch Port

The first VLAN, called the data VLAN, is sent and received untagged. The second VLAN, called the voice VLAN, is sent tagged with a dot1q header and a voice VLAN to which it belongs. However, the switch port is not a trunk port. The tagged packet comes from the IP phone. The data device that is connected to the IP phone receives and transmits only untagged packets and belongs to the native VLAN.

Security

Security is one of the most important considerations while designing the network. Malicious users can use tools available freely on the Internet to launch an attack if the access switches and ports are not secured; the attacker must simply gain physical access to these unsecured ports, and the entire network is wide open for an attack. This emphasizes the need to protect the internal ports against possible attacks.

In addition to protection from attacks, additional layers of security can be added to authenticate and authorize users trying to get access to the port, as well as enforcing policies on the edge of the network to ensure that users meet the policy requirements before accessing the network.

The following layers of security services can be deployed at the access layer:

- Layer 2 security
- IBNS and 802.1x
- Network Admission Control

Layer 2 security and user authentication are described in the following sections. Network Admission Control will be added in the future.

Layer 2 Security

Protecting against snooping and denial of service (DoS) attacks can be achieved simply by turning on the security features embedded in the Cisco switches. Layer 2 security plays an important role in the branch office to mitigate internal threats. The possibility of lack of tight physical security and monitoring in a branch office is a compelling reason to incorporate some of these security features into the design.

These features have become standard features in almost all Cisco switches. Catalyst Integrated Security Feature is a security feature that can be enabled on the access switches. CISF can be turned on to protect against various internal threats. The various kinds of attacks and which features can be used to mitigate the attacks have been documented extensively.

IBNS and 802.1x

Security can be further enhanced by authenticating and authorizing users before letting them on the network. Such a mechanism is inherent in wireless technologies. Authentication and authorization can also be enforced on the wired LAN ports by using Cisco Identity-Based Networking Services (IBNS) at the following URL:

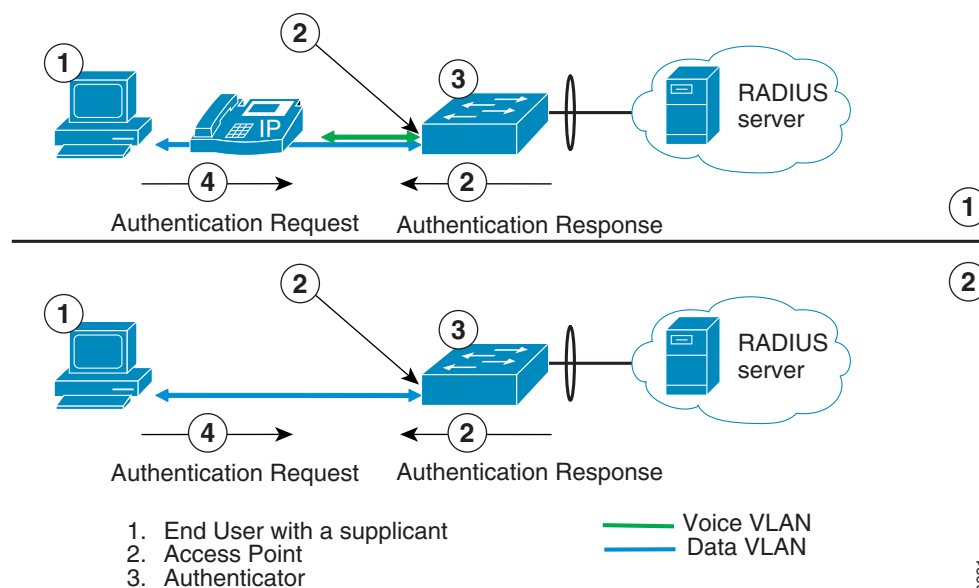
http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns75/networking_solutions_sub_sub_solution_home.html.

The Cisco IBNS solution is based on standard RADIUS and 802.1x implementations.

Cisco IBNS interoperates with all IETF authentication servers that comply with the RADIUS (RFC 2865, 2866, and 2868) and Extensible Authentication Protocol (RFC 2284) standards. Cisco has enhanced its Cisco Secure ACS to provide a tight integration across all Cisco switches.

802.1x is a standardized framework defined by the IEEE, designed to provide port-based network access. Using 802.1x, users are authenticated using information unique to the client and with credentials known only to the client. Figure 8 provides the basic framework used to authenticate the end users.

Figure 8 **Authenticating the End User using 802.1x**



The authentication process consists of exchanges of Extensible Authentication Protocol (EAP) messages between the supplicant and the authentication server. The authenticator (Cisco switch) relays the exchange between the server and the client transparently.

The switch can also enforce a policy dynamically, provided by the RADIUS server based on client credentials during the authentication phase. This policy dictates how the user accesses the network. Policies that can be enforced include putting the client into a specific VLAN and applying ACLs on the specific port.

By providing flexible port-based access control and policy enforcement capabilities at the network edge, this solution provides an important addition to the tools available for securing your network.

To deploy these solutions, it is important that end user machines have supplicants. Without the supplicants on the user machine, the user can be placed in a VLAN with very limited access. Microsoft provides these supplicants on some operating systems by default; on others, it must be downloaded and configured. Other third-party vendors also provide supplicants to various operating systems.

Network Admission Control

NAC preserves enterprise resilience by auditing and enforcing adherence to corporate endpoint security policies when accessing the network. While most *users* are authenticated, their endpoint devices (laptops, PCs, PDAs, etc.) are *not* checked for security policy compliance. NAC helps ensure the health of endpoints before they are granted network access. NAC works with software installed on workstations wishing to access the network to assess their condition (including operating system version, security patches, anti-virus, CSA, and other installed software), called the posture, of a client prior before allowing it to access the network. NAC also ensures that a network client has an up-to-date virus signature set prior to gaining access to the network. If the client requires a signature update, NAC directs it to the appropriate resources to complete the update. One example in which NAC accomplishes this is through placing the client into a quarantined network segment until disinfection is completed. More details will be documented in a future document.

QoS

During congestion in the network, traffic is delivered on a best effort basis. The switches and routers in the network do not differentiate between packets. With the converged network, it is important that traffic be prioritized so that packets that belong to certain applications get preferential treatment. A lot has been discussed and written about QoS. This document takes the QoS recommendations and applies them to the Branch infrastructure.



Note

See [References, page 18](#) for QoS reference pointers.

As per the QoS design principles provided in *End-to-End QoS Network Design*, following are some of the design considerations:

- Voice, video, and data applications should be classified and marked as close to their sources as possible.
- Unwanted traffic should be policed as close to its source as possible and dropped.
- QoS should be done in hardware; the complexity of the QoS policies to be deployed close to the source dictates the hardware requirements.

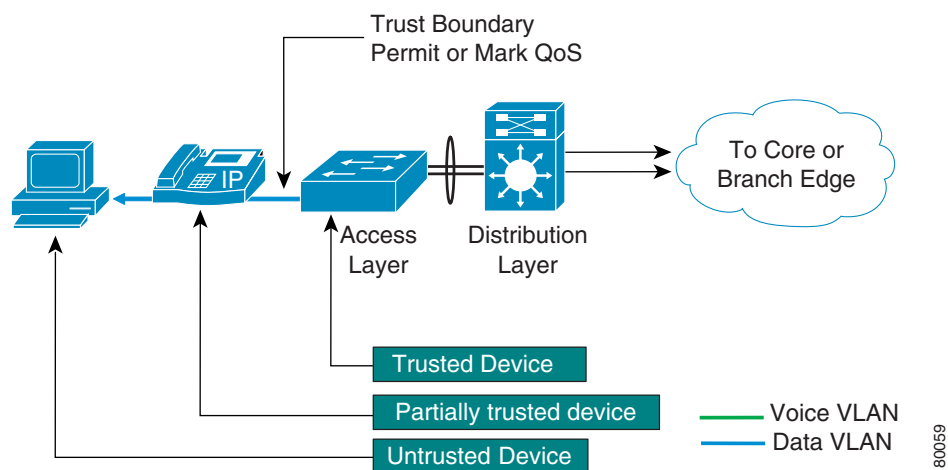
End points are capable of marking class of service (CoS) and Differentiated Services Code Point (DSCP) values. However, it is a matter of policy whether these end points can be trusted. Trusting the device means accepting the markings by these end devices and prioritizing traffic based on those values. If the end devices cannot be trusted, the device closest to the end point can be used to mark the CoS and DSCP values, and also police and rate limit traffic. This closest trusted device that marks the CoS and DSCP values creates a trusted boundary. All these functions require significant CPU time if done in software. Performing these tasks in hardware by ASICs relieves the CPU to do other tasks. As such, the granularity of policing and rate limiting might dictate the use of specific hardware.

By defining a trust boundary in the network, the device at the boundary can permit or remark the QoS values. In addition to trusted devices, there are devices that are partially trusted or conditionally trusted. Devices such as Cisco IP phones provide Ethernet ports to connect additional devices. The Cisco IP phone in this scenario is a partially trusted device because it provides connectivity to other devices as well. In such a case, the traffic originating from the Cisco IP phone can be permitted, and the rest of the

traffic can be marked at the trust boundary. The access layer is the closest layer to the end points, and the QoS policies can be defined at the access layer. The access switches then forms the trust boundary. At this trust boundary, the traffic is marked or remarked depending on the trust worthiness of the device. It is good practice to let traffic on voice VLANs through without remarking if it is being originated from a Cisco IP phone (Cisco Discovery Protocol running on the access switches determines whether the device is a Cisco IP phone). All other traffic has to be marked or remarked at the access switch or the trusted boundary.

The Cisco press book discusses the various models in depth. The trust boundary is shown in [Figure 9](#) for convenience. For more information about the trust models, trusted/untrusted/conditionally-trusted endpoints, see the Cisco press book.

Figure 9 Trust Boundary



Distribution Layer

The distribution layer provides the following services:

- High availability
- Scalability
- An aggregation point to deploy additional services if required

High Availability

Typically, high availability designs at the Layer 3 level or distribution level involve two switches running HSRP. The topology for a traditional design is shown in [Figure 4](#). Although this has some advantages, such as providing active-active distribution switches that share the load between two distribution switches, the complexity of deployment and troubleshooting in this design is worth considering.

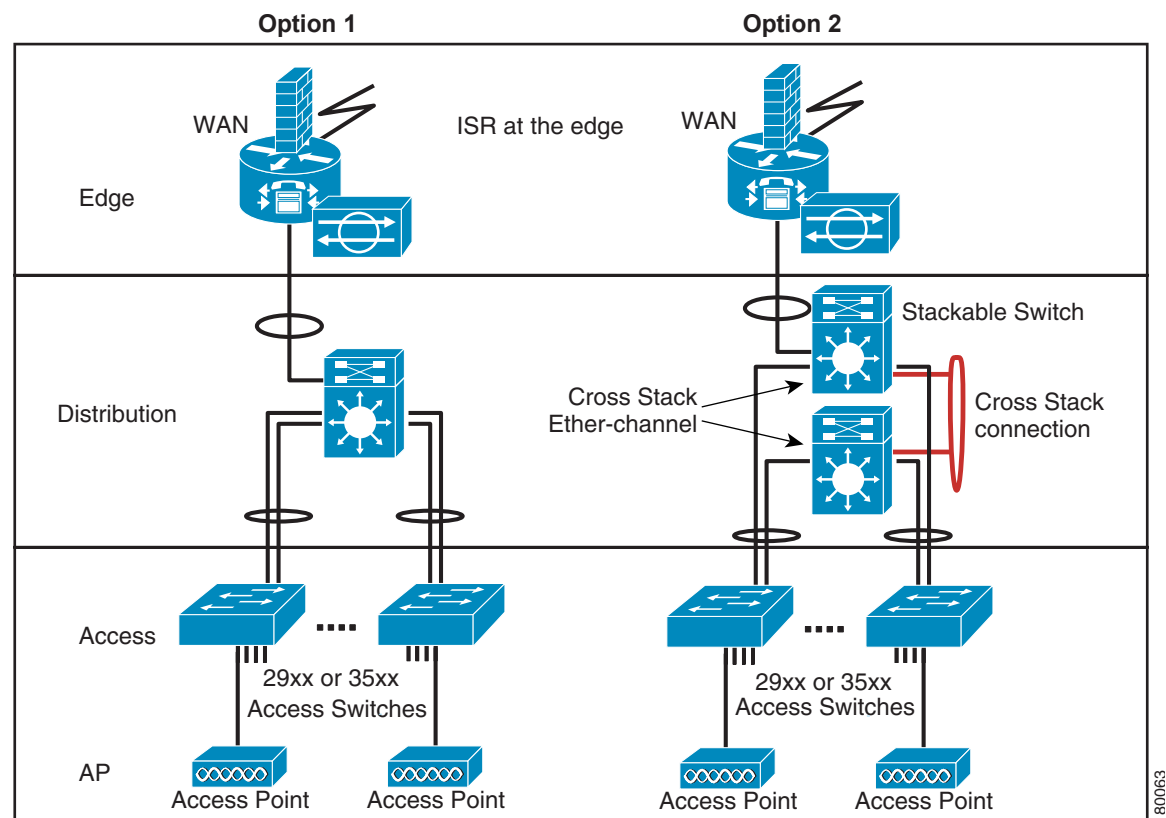
High availability design ensures network availability under failure conditions. High availability is achieved by providing redundant links and backup devices. High availability design also means additional complexity in the design to address failover and convergence. In a branch network, more often the cost is also a contributing factor for design considerations. High availability in the branch networks

can be achieved by avoiding dual homing of access layer switches to the distribution active and backup switches, and thereby avoiding spanning tree and the problems associated with troubleshooting and convergence. The loop-free topology can be achieved in three ways:

- Using a chassis-based switch, such as a Catalyst 4500 with redundant supervisor and redundant power supplies, to protect against device and power failure; and using EtherChannels/Link Aggregation Control Protocol (LACP)/Port Aggregation Protocol (PAgP) between the distribution and the access layers.
- Using an external stackable switch at the distribution to protect against device failure, and using EtherChannels to protect against link failure between the distribution and the access layers.
- Using an EtherSwitch Services Module for the ISR.

Figure 10 shows the options using external distribution switches.

Figure 10 *Multilayered Branch Architecture using External Distribution Switches*



As shown in Figure 10, using external distribution switches, Option 1 is a single chassis solution at the distribution layer that can also be used with redundant supervisor cards and redundant power supplies. This option has the following two distinct advantages:

- It can provide sub-second convergence for both Layer 2 and Layer 3 failures.
- It can provide LACP/PAgP/EtherChannel support.

The disadvantage is the lack of flexibility.

As shown in Figure 10, Option 2 is more cost effective and flexible for branch offices. If IP Base image is used, static routing can be used at the distribution layer without sacrificing the port scalability and high availability.

**Note**

Two kinds of binary files are available for the switches: IP base image is a standard multilayer image, and EMI is enhanced multilayer image. Enhanced image supports a host of Layer 3 features in addition to Layer 2 features of the standard image.

If dynamic routing is desired at the distribution layer, the stackable switch can be reconfigured with an EMI image. The stackable switch behaves like a single switch with line cards. More switches can be added to the stack to meet the growth requirements without sacrificing the high availability. The scalability is limited to the number of switches that can be stacked. There are the following two disadvantages of this option:

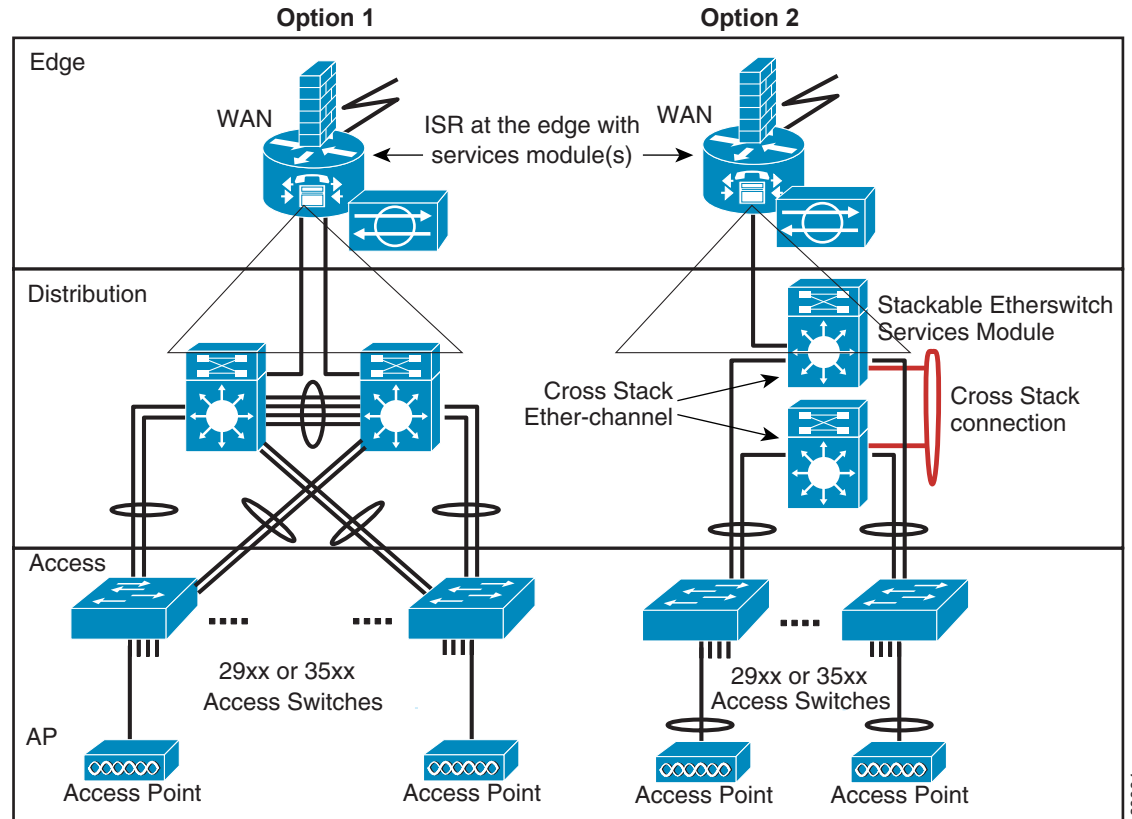
- The stackable switches do not support cross stack LACP/PagP at the time of testing. EtherChannel was the only option to achieve high availability.
- Layer 3 failure can take up to three seconds to converge under failure conditions. If the applications used need significantly less convergence times, then other options (option 1) have to be considered.

In addition to high availability, it is possible to configure the switches to load balance the traffic on the EtherChannels. Contrary to an L3 design where per packet load balancing can be achieved on equal cost multiple paths, load sharing is done based on source and destination MAC addresses. This implies that for a specific host, the traffic always uses the same link of the EtherChannel.

When using the Cisco EtherSwitch Services Module that fits into a slot of the Cisco ISR, the topology looks similar to Option 1 in [Figure 10](#) if a single EtherSwitch Services Module is used. The main reason that the topology looks similar is because the ISR interfaces with the EtherSwitch Services Module through a gigabit Ethernet interface. If high availability is desired for the distribution layer, the following two options can be used:

- Stack the services module with an external switch (topology similar to Option 2 in [Figure 10](#)).
- Use a dual services module with a trunk in between for high availability.

These two topologies are shown in [Figure 11](#).

Figure 11 Using EtherSwitch Services Module(s) at the Distribution Layer

180064

Scalability

In terms of the number of LAN/switch ports supported, scalability is achieved by putting the distribution switch. The distribution switch has the required port density to fan out more access switches without compromising high availability. Without the distribution switch, scalability is limited to the number of LAN ports available on the edge router. The Ethernet interfaces embedded in the ISR do not support switched virtual interfaces (SVIs). In addition, EtherChannels, LACP, and PAGP are not supported on the embedded Ethernet interfaces on the ISR. EtherChanneling and SVIs are supported only on network module-based Ethernet switches, which plug into ISRs to provide Catalyst switch features. Incorporating high availability in the design means sacrificing scalability and sacrificing a network module slot. In the event that a network module-based Ethernet switch is used with the ISR, scalability is limited to the number of LAN ports available on the network module Ethernet switch.

Additional Services

Because of the inherent Layer 2 and Layer 3 services within the switch, the distribution layer can be used to deploy additional services. The distribution layer can also help customers to migrate to use advanced services without having to redesign the entire branch office network. It is also possible to deploy appliance devices at the distribution layer if required, and to migrate towards an architecture where these advanced services are integrated into the distribution switch software images. The distribution layer provides great flexibility without compromising high availability and scalability.

Some of the areas that can benefit are the security and WLAN. Specifically, the following are some of the services that can be deployed on the distribution switches:

- VRF on the distribution layer switches
- Policy-based routing
- DHCP for IP address management
- Firewall services for the DMZ servers
- DMZ services, including Wide Area File System
- Intrusion detection/prevention
- Wireless LAN management using mini WLSE
- Cisco WLAN Controller

Conclusion

The next generation branch office should be able to add services as the branch office grows. Providing advanced services requires a baseline architecture onto which these advanced services can be added without having to re-architect the network. Keeping this in mind, the various architectures discussed in this document take into consideration the growth, high availability, security, and deployment of advanced services without having to redesign the network.

In this document, Layer 2 at the access layer is recommended. It is difficult to meet all the requirements with one single box. The layered architecture provides the required flexibility to meet all the requirements of the next generation branch office. In some cases, this layered architecture might be housed in a single box but still provides the required high availability and scalability to meet the branch office requirements. At the same time, the layered architecture must be easily deployable. With that in mind, either a topology with no Layer 2 loops can be deployed, or if more control over traffic paths and failover times is desirable, other architectures can be deployed. With either a loop-free or looped topology, a layered architecture with Layer 2 at the access provides more flexibility for adding services.

From a security perspective, providing layered security in the branch office is desirable. For example, Layer 2 security is supported in all Cisco access layer switches, and provides a strong obstacle against some denial of service attacks. Also, users are authenticated and authorized before logging on to the network, when they are connected either directly or via the Cisco IP phone. Additional services can be deployed or enabled as they become available without having to redesign the network for the foreseeable future.

References

- Cisco Campus Network Design Guide—
http://www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html
- LAN Baseline Architecture Branch Office Network Reference Design Guide <http://www.cisco.com/en/US/docs/solutions/Enterprise/Branch/Design.html>.
- Configuration guides
 - <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3560/>
 - <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750/>
 - <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2970/>

- <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2950/>
- Additional useful links
 - http://www-tac.cisco.com/Training/bootcamps/advanced_lanswitching_bootcamp/lectures/module5/common-issues-day5.pdf

