CHAPTER 8

Enterprise Branch

The enterprise branch, along with the enterprise WAN edge, provides users at geographically disperse remote sites access to the same rich network services as users in the main site. The availability and overall security of the branch, the WAN edge, and the WAN transit, is thus critical to global business operations.

The challenge, from a security perspective, is enabling the enterprise to confidently embrace and extend these rich global services and remote collaboration capabilities to all locations. This is achieved through a defense-in-depth approach to security that extends and integrates consistent end-to-end security policy enforcement and system-wide intelligence and collaboration across the entire enterprise network.

The aim of this chapter is to illustrate the role of the enterprise branch in this end-to-end security policy enforcement, including how to apply, integrate, and implement the SAFE guidelines. See Figure 8-1.

Figure 8-1 Enterprise Branch

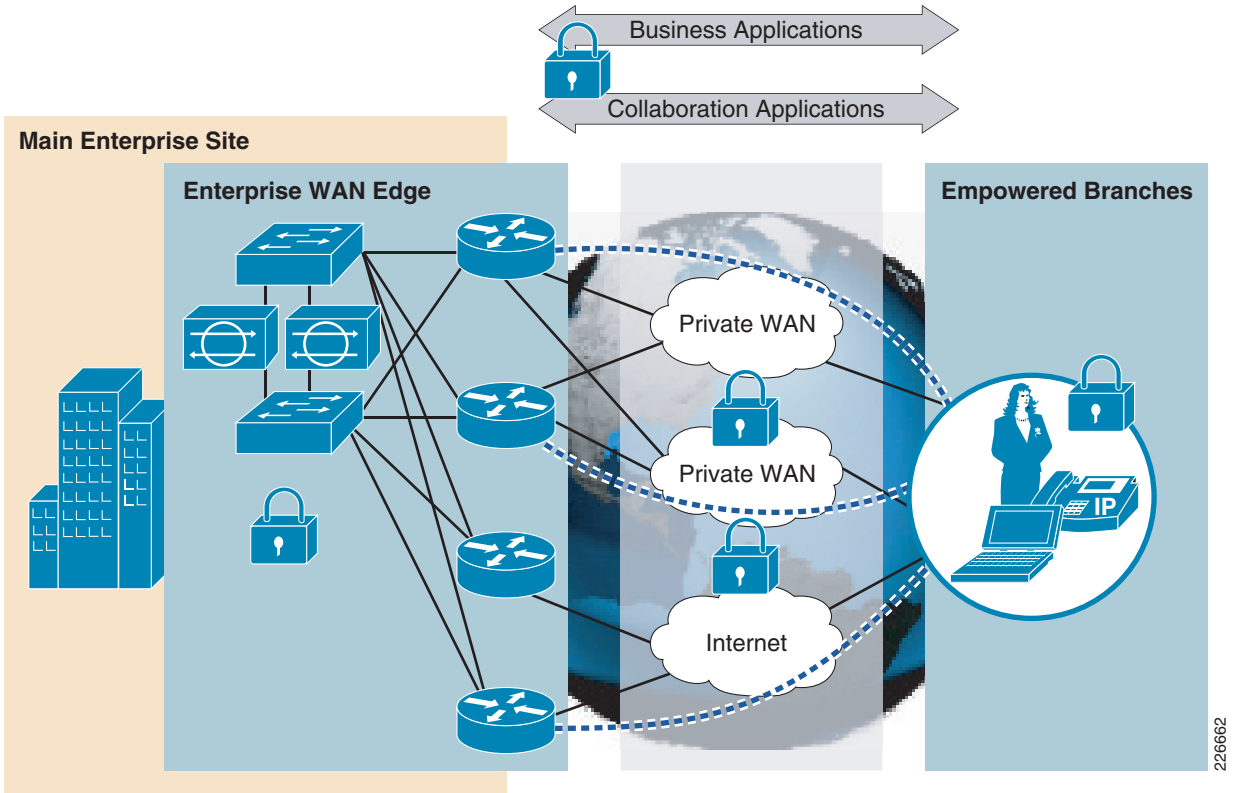

From a functional perspective, an enterprise branch typically includes the following:

- WAN edge device

  Terminates the WAN link and may offer additional services, such as site-to-site VPN, local Internet access, security, application optimization, and voice services. The device may be owned by the enterprise or it may be owned and managed by a Service Provider (SP). The WAN link may be a private network, the Internet, wireless, or a combination thereof.

- Switching infrastructure

  Provides wired LAN access to clients and LAN distribution for local services.

- Local services infrastructure

  Based on particular branch business needs, additional infrastructure may be deployed to support local services such as voice, video, application servers, and wireless LAN, as well as security services such as VPN, encryption, IPS, and firewall.

Remote access, teleworker, partner, customer, and Internet access are addressed in Chapter 6, "Enterprise Internet Edge," along with their related security guidelines.

Two typical enterprise branch architectures are illustrated in Figure 8-2.

Figure 8-2 Enterprise Branch Architecture

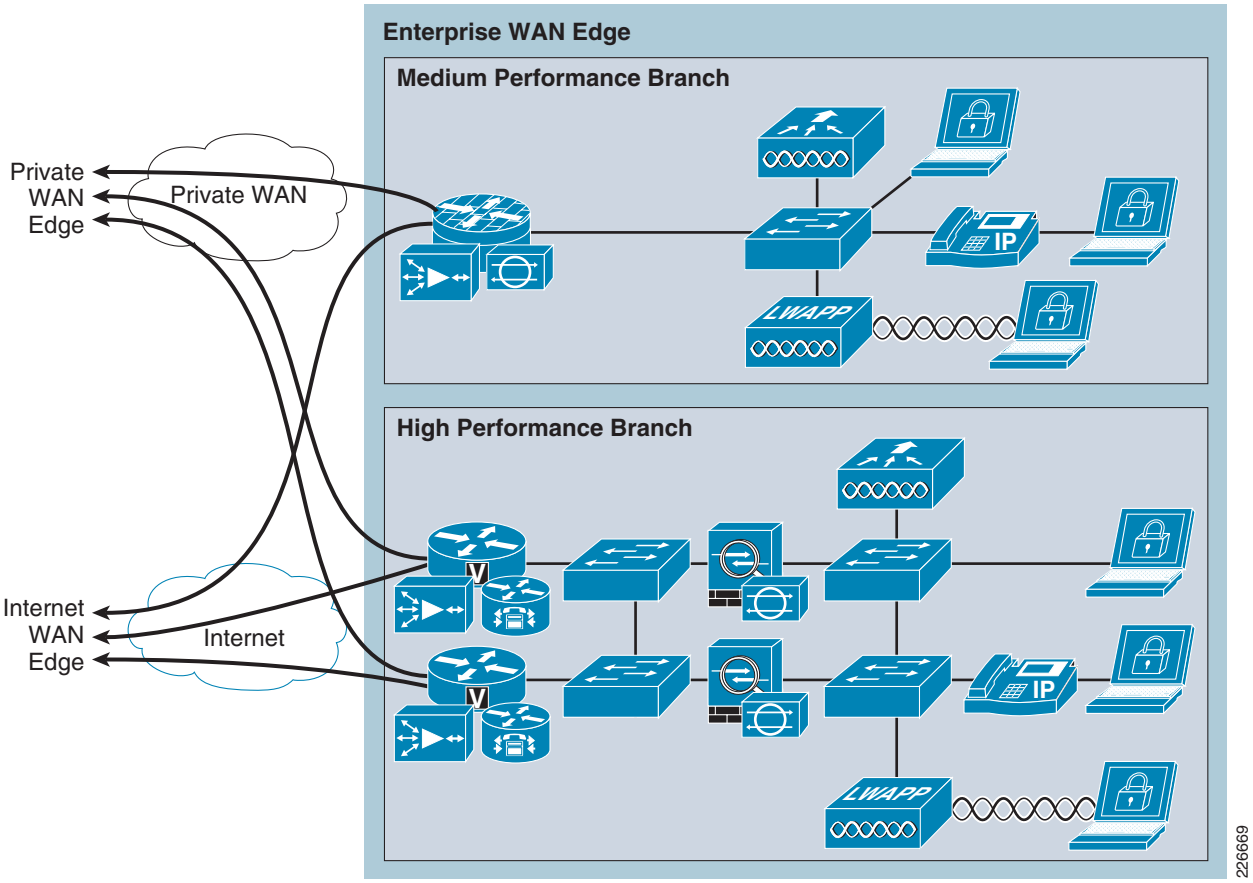

For more information about SAFE for the enterprise WAN edge, see Chapter 7, "Enterprise WAN Edge."

The following section describes the threats that the branch is exposed to and the security technologies integrated to address them.

Key Threats in the Enterprise Branch

The threats addressed in the branch of an end-to-end enterprise architecture are focused on the following key areas:

- Malicious activity by branch clients, including malware proliferation, botnet detection, network and application abuse, and other malicious or non-compliant activity.
- WAN transit vulnerabilities such as sniffing and man-in-the-middle (MITM) attacks.
- Attacks against the infrastructure itself, such as unauthorized access, privilege escalation, and denial-of-service (DoS) attacks.

Web and E-mail threats posed to branch clients, such as malicious web sites, compromised legitimate web sites, spam, and phishing, are addressed as part of a centralized deployment in Chapter 6, "Enterprise Internet Edge." This assumes that the branch is not using split-tunneling. If a branch does use split-tunneling, whereby there is local Internet access directly from the branch, web security must be implemented locally.

The particular threat focus of an enterprise branch, and the specific security objectives and integration elements to mitigate these threats, are shown in Table 8-1.

Table 8-1 Key Threats in the Enterprise Branch

Threat Focus	Threats Mitigated	Security Objectives	Security Integration
Malicious branch client activity	Malware proliferation, botnets, worms, viruses, Trojans Application and network abuse	Detect and mitigate threats	• IPS Integration • Endpoint Security • Web Security[1] • E-mail Security[1]
WAN transit threats	Unauthorized access to network and data such as through sniffing and man-in-the-middle (MITM) attacks	Isolate and secure WAN data and access	• Secure WAN Connectivity
Attacks against the infrastructure	Unauthorized access to devices, network and data Reconnaissance DoS	Deliver resilient and highly available services	• Routing Security • Service Resiliency • Network Policy Enforcement • Switching Security • Secure Device Access • Telemetry

1. Addressed as part of a centralized deployment in Chapter 6, "Enterprise Internet Edge," assuming a non-split-tunneling policy at remote sites.

The design and integration of each of these security elements into the branch is addressed in the following section.

Design Guidelines for the Branch

Security integration in the branch addresses the key threat areas locally through the following:

As mentioned earlier, web and E-mail security are addressed as part of a centralized deployment in the Chapter 6, "Enterprise Internet Edge." If localized web security is required, due to the use of split-tunneling, Cisco offers a number of web and content security options, including the Cisco IronPort S-Series, Cisco ASA.5500 Series, and Cisco IOS Content Filtering. For more information, see the Web Security section of Appendix A, "Reference Documents."

Secure WAN Connectivity in the Branch

The branch is reliant upon its WAN connectivity for access to centralized corporate services and business applications, as well, in many cases, Internet services. As such, the WAN is critical to service availability and business operations. Consequently, the WAN must be properly secured to protect it against compromise, including unauthorized access, and data loss and manipulation from sniffing or man-in-the-middle (MITM) attacks.

The security objective being to provide confidentiality, integrity and availability of data as it transits the WAN.

The design and implementation of secure WAN connectivity is addressed as an end-to-end system, incorporating both the branch and the WAN edge. The key design recommendations and considerations are presented in "Secure WAN Connectivity in the WAN Edge" section on page 7-5, but must be developed in conjunction with the branch WAN design and tie together to provide an end-to-end, secure WAN.

The recommendation for secure WAN connectivity includes the following:

- VPN for traffic isolation over the WAN

  There are a number of VPN options and the choice will vary based on specific customer requirements. DMVPN, for example, offers support for VPN over both a private WAN and the Internet, as well as multicast and dynamic routing. Consequently, DMVPN can be integrated to enable a common VPN implementation if both of these WAN types are deployed at remote sites.

- Public Key Infrastructure (PKI) for strong tunnel authentication

  PKI provides secure, scalable, and manageable authentication that is critical to large-scale VPN deployments. PKI also features the dynamic renewal and revocation of certificates that enables the dynamic commissioning and decommissioning of branches with ease.

- Advanced Encryption Standard (AES) for strong encryption

  Data over the Internet is vulnerable to sniffing; therefore, encryption is critical to data confidentially and integrity. Data over a private WAN can also be encrypted for maximum security or for compliance reasons.

For more information on VPN technologies and PKI, refer to the WAN Design section of Appendix A, "Reference Documents."

Routing Security in the Branch

Routing in the branch is critical to service availability, and as such, it must be properly secured to protect it against compromise, including unauthorized peering sessions and DoS attacks that may attempt to inject false routes, and remove or modify routes.

The security of the routing is particularly important in the branch as it features a key network border, supporting both an external and an internal routing domain. Consequently, it is critical, not only that the external peering interface is properly secured, but that the routing information is properly filtered to ensure that only necessary routes are advertised out and that only valid routes are propagated into the internal routing table.

There are two routing domains to consider:

- External routing domain

  Maximum routing security, including strict routing protocol membership, routing domain termination and route redistribution filtering to ensure only the necessary routes are advertised.

- Internal routing domain

  Routing security for internal interfaces is typically less stringent though should, at a minimum, include neighbor authentication. In addition, if dynamic routing is not being used within the branch itself, the routing domain should be terminated on the edge device.

The areas of focus, objectives and implementation options for routing security in the branch are outlined in Table 8-2.

Table 8-2 Routing Security in the Branch

Routing Security Focus	Routing Security Objectives	Implementation
Restrict Routing Protocol Membership	Restrict routing sessions to trusted peers and validate the origin and integrity of routing updates	• Routing peer definition • Neighbor authentication • BGP TTL Security Hack (BTSH) • Default passive interface
Control Route Propagation	Ensure only legitimate networks are advertised and propagated	• Terminate the external routing domain on the WAN edge (e.g., using EIGRP stub routing)[1] • Only advertise required routes to the external routing domain • Terminate the internal routing domain if dynamic routing not required in the branch (e.g., using EIGRP stub routing)[1] • Advertise branch routes over the VPN to the internal routing domain
Log Neighbor Changes	Detect neighbor status changes that may indicate network connectivity and stability issues, due to an attack or general operations problems	• Neighbor logging on all routing domains

1. Stub routing is not supported in OSPF, thus outbound filters should be enforced to restrict route propagation.

A sample implementation of secure routing in the branch module is shown below and it integrates the SAFE guidelines to:

- Authenticate all routing peers.
- Disable routing on all interfaces by default.
- Explicitly enable the internal routing domain on the VPN tunnels.
- Explicitly enable the external routing domain on interfaces to the private WAN.
- Only permit distribution of the directly-connected and summary branch subnets over the internal routing domain.
- Limit router participation in the external router domain to learning only, preventing the distribution of any routes from the branch into that routing domain.
- Enable neighbor logging on all routing domains.

```
! Internal Routing Domain
router eigrp 1
! By default disables routing on all interfaces
passive-interface default
! Enables internal routing on the VPN tunnels
 no passive-interface Tunnel0
 no passive-interface Tunnel1
 network 10.0.0.0
 no auto-summary
! EIGRP stub routing to only advertise directly connected networks and summarized routes
```

```
 eigrp stub connected summary
! Enables neighbor logging
 eigrp log-neighbor-changes
!
! External Routing Domain
router eigrp 100
! By default disables routing on all interfaces
 passive-interface default
! EGP is permitted on interfaces to the Private WAN
 no passive-interface GigabitEthernet0/1
 network 192.168.0.0 0.0.255.255
 no auto-summary
! Router only accepts, but does not explicitly advertise, any routes.
 eigrp stub receive-only
! Enables neighbor logging
 eigrp log-neighbor-changes
!
! Authenticate internal routing peers
key chain eigrp-auth
 key 10
   key-string <strong-key>
!
! Authenticate external routing peers
key chain eigrp-auth-egp
 key 11
   key-string <strong-key>
!
interface Tunnel0
 description Private WAN Tunnel
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 eigrp-auth
!
interface Tunnel1
 description Internet Tunnel
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 eigrp-auth
!
interface GigabitEthernet0/1
 description Private WAN
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 eigrp-auth-egp
!
```

Note that neighbor logging is enabled by default in EIGRP; therefore, the **eigrp log-neighbor-changes** command does not appear explicitly in the configuration.

For more information on routing security, including design and configuration guidelines , see Chapter 2, "Network Foundation Protection."

Design Considerations

- Neighbor authentication and BTSH require identical configuration on routing peers in order to accept routing updates. Consequently, the enterprise should work with their service provider to enable these security features.
- If neighbor logging is enabled by default for a particular routing protocol, the logging commands will not appear in the configuration. For example, this is currently the case for EIGRP.

Service Resiliency in the Branch

The resiliency of services provided by the branch infrastructure itself is critical to the operation of the remote site. The branch edge is also a key network border. Consequently, all infrastructure devices and links must be resilient to targeted, indirect, malicious and unintentional attacks, as well as general failure scenarios. This is particularly important since the edge devices have external interfaces.

Possible attacks include DoS attacks based on unauthorized and authorized protocols, distributed DoS (DDoS) attacks, flood attacks, reconnaissance and unauthorized access. General failure scenarios include power outages, physical link failures, and device failures.

Service resiliency in the branch involves the following three key design areas:

- Device resiliency
- Central site service availability
- High availability

The areas of focus, objectives , and implementation options for service resiliency in the branch are outlined in Table 8-3.

Table 8-3 Service Resiliency in the Branch

Service Resiliency Focus	Service Resiliency Objectives	Implementation
Restrict attack surface	Disable unnecessary services Address known vulnerabilities	• Disable unnecessary services on all infrastructure devices • Patch infrastructure devices with updated software
Harden the device	Protect device resources from exhaustion attacks by limiting, filtering and rate-limiting traffic destined to the control plane.	• Memory protection • Port security on access ports • Limit and rate-limit control plane traffic, including service-specific considerations (e.g., DHCP snooping and DAI on access switches[1)] • Implement CoPP/CoPPr, if available
Preserve and optimize remote site services	Ensure any limited resources at a remote site, such as a low bandwidth WAN link or a low performance platform, are not overwhelmed, and optimize their utilization.	• QoS: Ingress and egress QoS on the access switch. Egress QoS on the WAN link. • Application optimization
Implement redundancy[1]	Deploy device, link and geographical diversity to eliminate single points of failure	• Redundant devices • Redundant links • Redundant WAN providers • Geographically diverse locations

1. The level of redundancy implemented in a branch is typically dependent on the size, business need, and budget associated with a particular site. In some cases, a decision may be taken to accept the risk associated with a particular failure scenario, in lieu, for example, of cost-saving.

The particular considerations for QoS in the branch are covered below. For more information on the other service resiliency techniques, including design and configuration guidelines, see Chapter 2, "Network Foundation Protection."

Service resiliency should be complemented by network policy enforcement techniques that filter traffic at the network edges, permitting only authorized services and originators to directly address the infrastructure. These techniques restrict accessibility to the infrastructure in order to reduce exposure to unauthorized access, DoS, and other network attacks. For more information, refer to the "Network Policy Enforcement in the Branch" section on page 8-11.

QoS in the Branch

QoS is critical to the optimal performance and availability of business-critical services in a branch, even under adverse network conditions, such as high data rates and worm outbreaks. In addition, since some service control and all remote management are in-band, it is critical that QoS is employed to accurately classify, prioritize, control, and management traffic.

The fundamental principles are to accurately classify and mark traffic at the access edge, then police and schedule traffic at key network borders, particularly on links with limited resources that are subject to congestion.

In a branch QoS implementation, the primary elements are as follows:

- Ingress QoS on the access ports

  Accurately identify, classify, mark and police ingress traffic.

- Egress QoS on the access ports

  Queuing according to the defined service classes of the enterprise QoS policy.

- Egress QoS on the access switch uplink

  Queuing according to the defined service classes of the enterprise QoS policy. The DCSP markings can be trusted as the ingress QoS policy has been enforced to mark or remark QoS settings.

- Egress QoS on the WAN link

  Queuing according to the defined service classes of the enterprise QoS policy to optimize usage of the, typically limited, WAN resources. The DCSP markings can be trusted as the ingress QoS policy has been enforced to mark or remark QoS settings. See Figure 8-3.

Figure 8-3 QoS in the Branch

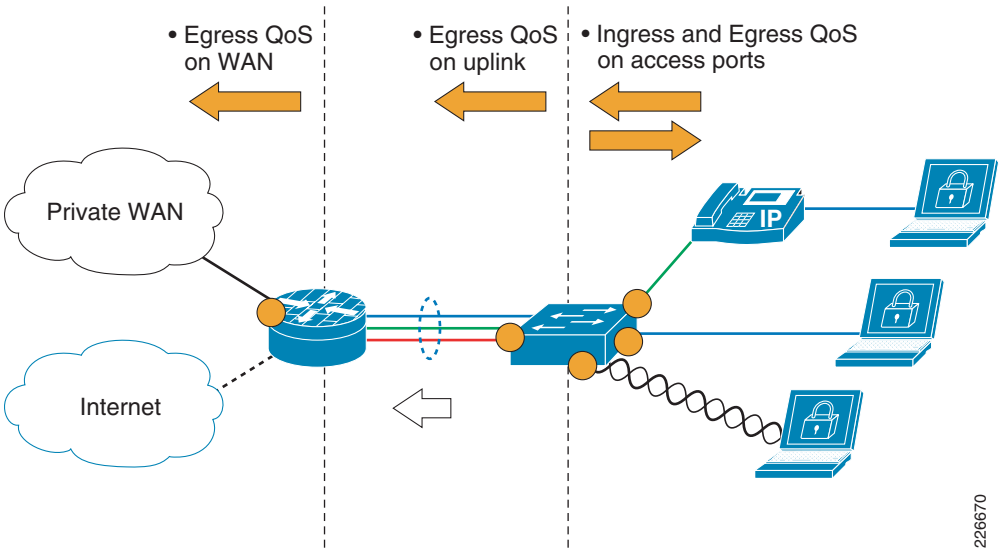

QoS in the branch is just one element of an end-to-end QoS implementation, including per-branch egress QoS on the WAN edge and consistent ingress classification and marking across all access edges of the enterprise network.

The following is a sample branch QoS configuration:

```
! Define the Egress QoS policy
! Prioritize voice, interactive video, call signaling and control traffic
policy-map WAN-Edge-QoS
 class Voice
    priority percent 18
 class Interactive-Video
    priority percent 15
 class Call-Signaling
    bandwidth percent 5
 class Network-Control
    bandwidth percent 5
 class Critical-Data
    bandwidth percent 27
     random-detect dscp-based
 class Bulk-Data
    bandwidth percent 4
     random-detect dscp-based
 class Scavenger
    bandwidth percent 1
 class class-default
    bandwidth percent 25
     random-detect
!
! Enforce the QoS policy on this traffic types
class-map match-all Bulk-Data
 match ip dscp af11  af12
class-map match-all Interactive-Video
 match ip dscp af41  af42
class-map match-any Network-Control
 match ip dscp cs6
 match ip dscp cs2
class-map match-all Critical-Data
 match ip dscp af21  af22
class-map match-any Call-Signaling
 match ip dscp cs3
 match ip dscp af31
class-map match-all Voice
 match ip dscp ef
class-map match-all Scavenger
 match ip dscp cs1
!
! Enforce the policy on traffic over the VPN tunnel interfaces
interface Tunnel0
 qos pre-classify
!
interface GigabitEthernet0/1
 description Private WAN
 service-policy output WAN-Edge-QoS
```

Design Considerations

- Do not trust traffic on access ports.
- Perform ingress QoS as close to the source as possible.
- Perform QoS in hardware.
- Implement per-branch egress QoS on the WAN edge.
- Enforce a consistent ingress QoS policy across all access edges of the enterprise network.
- Complement QoS on the WAN with application optimization to maximize application performance and WAN utilization.

For more information on QoS, see the QoS Design section of Appendix A, "Reference Documents."

Network Policy Enforcement in the Branch

The branch features two key network borders: a WAN edge and an access edge. Thus, it is critical to enforce a strong network policy on both these network borders. This includes restricting the incoming traffic that is permitted, blocking unauthorized access and validating the source IP address of traffic on both the WAN interfaces , and the access edge. Anomalous traffic is discarded as close to the edge of the network as possible, thereby minimizing the risk of exposure.

Possible threats include unauthorized access and IP spoofing that can be used to anonymously launch an attack, bypass network access and policy enforcement controls, and snoop data through MITM attacks that combine IP and ARP spoofing.

The areas of focus, objectives, and implementation options for network policy enforcement in the enterprise branch are listed in Table 8-4.

Table 8-4 Network Policy Enforcement in the Branch

Network Policy Enforcement Focus	Network Policy Enforcement Objectives	Implementation
Filter Incoming Traffic	Restrict incoming traffic to authorized sources and for authorized services only	• WAN edge ACLs applied inbound on WAN interfaces • Access edge iACLs applied inbound on the access edge • Firewall integration
IP Spoofing Protection	Ensure traffic is topologically valid, (i.e., sourced from a valid address that is consistent with the interface it is received on)	• uRPF loose mode on WAN interfaces • IP Source Guard on access ports or uRPF on Layer 3 access edge

The particular considerations for WAN edge ACLs, access edge iACLs and firewall integration in a branch are covered in detail below. For more information on the other network policy enforcement techniques, including design and configuration guidelines, see Chapter 2, "Network Foundation Protection."

Additional Security Technologies

If the enterprise security posture assessment determines that additional access and policy control technologies are a required element of the corporate security policy, then this must be extended and integrated to the branch.

Additional technologies may include 802.1X, Identity-Based Network Networking Services (IBNS) and Network Admission Control (NAC). For more information, see the Identity-Based Network Services section of Appendix A, "Reference Documents."

Design Considerations

- Consistent network policy enforcement on all key network borders

  A consistent network policy must be enforced on all key network borders, including the WAN edge, the Internet edge and the access edge. For more information on edge policy enforcement for WAN Edge, refer to Chapter 7, "Enterprise WAN Edge." For more information on edge policy enforcement for the Internet edge, refer to Chapter 6, "Enterprise Internet Edge."

- Address space planning

  Careful planning of the corporate address space facilitates the definition and maintenance of traffic filtering that is used in many areas of security policy enforcement, including ACLs, firewalls, route filtering, and uRPF. It is recommended that a rational, summarized, or compartmentalized IP address scheme be used across the enterprise network, enabling a manageable and enforceable security policy, offering a significant benefit to overall network security

  For more information on address space planning, see Chapter 2, "Network Foundation Protection."

- IP Spoofing Protection

  Enforce IP spoofing protection on the access edge to enable generic and consistent access edge iACLs to be enforced across the enterprise, thereby minimizing operational overhead. IP spoofing protection removes the need to specify the particular access subnet as the source address in the ACL entries, since the source IP address has already been validated.

  For more information on IP spoofing protection, see Chapter 2, "Network Foundation Protection."

WAN Edge ACLs

The primary objective of WAN edge ACLs is to restrict incoming traffic on the WAN links only to the minimum required traffic and services, and *only* from authorized originators. This typically involves permitting only the necessary routing updates from defined external routing peers, along with VPN access to the WAN edge.

In addition, standard ingress edge filtering is enforced, per Bridge Control Protocol (BCP) 38 and RFC2827, denying traffic with illegitimate, invalid, or reserved source addresses.

A WAN edge ACL for a branch with site-to-site VPN only, thus typically features the following elements:

- Deny fragments
- Deny the corporate address space originating from external sources
- Deny RFC1918 private address space (10/8, 172.16/12, 192.168/16)
- Deny RFC3330 special use IPv4 addressing (0.0.0.0, 127/8, 192.0.2/24, 224/4)
- Permit routing updates from authorized, external peers

- Permit VPN to WAN edge
- Permit ping and traceroute for troubleshooting

For more information on traffic filtering, see the Edge Filtering section of Appendix A, "Reference Documents."

Access Edge iACLs

The primary objective of iACLs on the access edge is to restrict client access to the network infrastructure, thereby reducing the risk exposure of these devices. Consequently, direct traffic to the infrastructure address space is blocked across all access edges of the enterprise.

In a branch, access edge iACLs are typically enforced on the branch edge router or, if a dedicated firewall is deployed on the access edge, they are integrated into the firewall policy. This ensures ease of operational management.

The following is a sample configuration (guidelines) for creating an iACL:

```
access-list 125 remark Client Access Edge iACL
! Permit Clients to perform ping and traceroute
access-list 125 remark Permit Client ping and traceroute
access-list 125 permit icmp any any ttl-exceeded
access-list 125 permit icmp any any port-unreachable
access-list 125 permit icmp any any echo-reply
access-list 125 permit icmp any any echo
! Permit VPN to Mgmt FW for local operational staff
access-list 125 remark Permit VPN to Mgmt FW
access-list 125 permit udp any host <mgmt-fw> eq isakmp
access-list 125 permit esp any host <mgmt-fw>
! Deny Client Access to Network Infrastructure Address Space
access-list 125 remark Deny Client to OOB Mgmt Network
access-list 125 deny ip any <subnet> <inverse-mask>
access-list 125 remark Deny Client to NoC & Core
access-list 125 deny ip any <subnet> <inverse-mask>
access-list 125 remark Deny Client to Internet Edge
access-list 125 deny ip any <subnet> <inverse-mask>
access-list 125 remark Deny Client to WAN Edge
access-list 125 deny ip any <subnet> <inverse-mask>
access-list 125 remark Deny Client to VPN Tunnels
access-list 125 deny ip any <subnet> <inverse-mask>
access-list 125 remark Deny Client to Branch Infra
access-list 125 deny ip any <subnet> <inverse-mask>
! Permit All Other Client Traffic
access-list 125 remark Permit Client Non-Infra traffic
access-list 125 permit ip any any
!
! Enforces ACL on Branch access port
interface GigabitEthernet0/0.10
 description Wired Clients
 ip access-group 125 in
```

Note that it is not necessary to specify the particular access subnet as the source address in the ACL entries if IP source address validation is already being enforced; for example, through IP Source Guard on the access ports. This enables generic and consistent iACLs to be deployed across the enterprise access edge, thereby minimizing the operational overhead.

For more information on iACLs, see Chapter 2, "Network Foundation Protection."

Design Considerations

- IP Spoofing Protection

  Enforce IP spoofing protection on the access edge to enable generic and consistent access edge iACLs to be enforced across the Enterprise, thereby minimizing operational overhead. IP spoofing protection removes the need to specify the particular access subnet as the source address in the ACL entries, since the source IP address has already been validated.

Firewall Integration in the Branch

Firewall integration in the branch enables the segmentation and enforcement of different security policy domains. This provides enhanced protection from unauthorized access that may be required if, for example, the branch features a point-of-sale (PoS) segment for credit card processing that requires PCI compliance, if split-tunneling is being employed, if there is an unsecured wireless network or if there are multiple user groups with different security policies to be enforced.

In addition, firewall integration offers more advanced, granular services , such as stateful inspection and application inspection and control on Layer 2 through Layer 7. These advanced firewall services are highly effective of detecting and mitigating TCP attacks and application abuse in HTTP, SMTP, IM/P2P, voice, and other protocols.

The two common design criteria for firewall integration in a branch are cost and administrative domains (i.e., who manages the infrastructure devices (NetOps, SecOps, SP)). A combination of these two factors typically dictates the platform selection.

To meet the deployment criteria of each customer, Cisco offers two key firewall integration options:

- IOS Firewall

  Cost-effective, integrated firewall that is typically implemented in the branch edge router. Cisco IOS Firewall is offered as a classic, interface-based firewall or as a zone-based firewall (ZBFW) that enables the application of policies to defined security zones.

- Adaptive Security Appliance (ASA) Series

  Dedicated firewall enabling a highly scalable, high performance, high availability and fully featured deployment on a range of platforms. It also supports distinct administrative domains, including a separate NetOps and SecOps model, as well as deployments where the edge router is SP-owned and managed.

  For more information on firewall integration using either a Cisco IOS firewall or a Cisco ASA, see the Firewall section of Appendix A, “Reference Documents.”

IOS Zone-based Firewall (ZBFW) Integration in a Branch

IOS ZBFW enables the creation of different security zones and the application of particular network policies to each of these defined zones. The first design step is thus to determine the zones required, based on the different network policies to be enforced. IOS ZBFW features an implicit deny for traffic between zones and so zones need only be created for zones with traffic flows that will be permitted between zones.

Typical security zones for a branch are as follows:

- VPN

  Tunnel interfaces to WAN edge hubs

- Clients

Client VLAN interfaces

- Infrastructure

  Management VLAN and integrated modules, such as switch, WLAN infrastructure, and IPS

Since there is an implicit deny, unless a branch is hosting externally-accessible services, the definition of a WAN zone is not required, because traffic from the WAN is not, by default, permitted to pass through the ISR to internal interfaces.

A sample baseline IOS ZBFW design for a branch, illustrating some sample zones and the associated permit policies to be enforced, is shown in Figure 8-4.

Figure 8-4 Sample IOS ZBFW Integration in a Branch

	VPN	Infrastructure	Clients
VPN		Permit	Permit
Infrastructure	Permit		Permit
Clients	Permit	Deny	

• Implicit deny for non-permitted flows

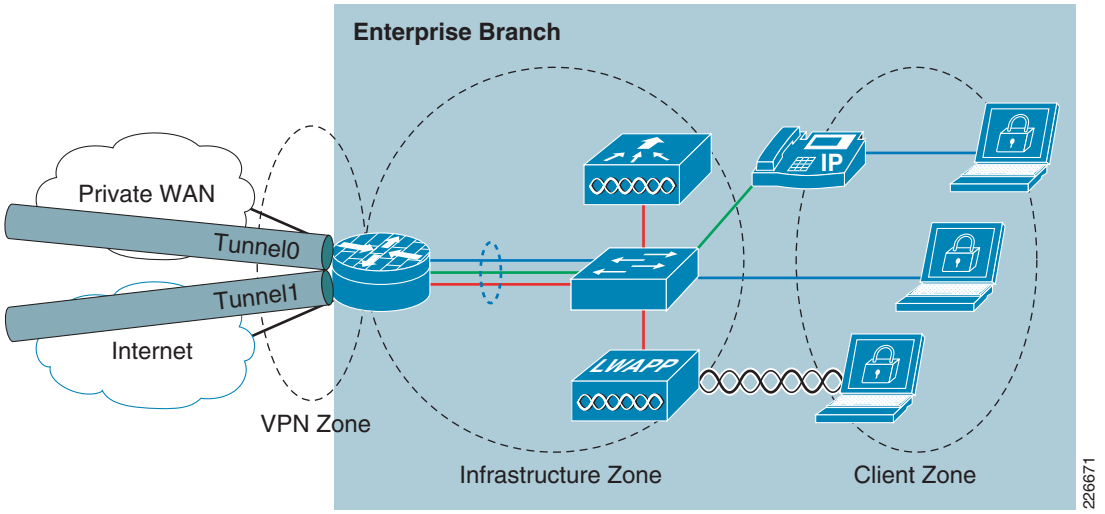

The following sample configuration illustrates the use of ZBFW on the branch:

```
zone security vpn
 description VPN to Corporate
zone security infra
 description Infrastructure Devices
zone security clients
 description Clients
zone-pair security clients-hq source clients destination vpn
 service-policy type inspect clients-hq-policy
zone-pair security infra-hq source infra destination vpn
 service-policy type inspect infra-hq-policy
zone-pair security hq-clients source vpn destination clients
 service-policy type inspect hq-clients-policy
zone-pair security hq-infra source vpn destination infra
 service-policy type inspect hq-infra-policy
zone-pair security infra-clients source infra destination clients
 service-policy type inspect infra-clients-policy
policy-map type inspect infra-clients-policy
 class type inspect frm-infra-class
```

```
  inspect
 class class-default
  drop
policy-map type inspect infra-hq-policy
 class type inspect frm-infra-class
  pass
 class class-default
  drop
policy-map type inspect hq-infra-policy
 class type inspect to-infra-class
  pass
 class class-default
  drop
policy-map type inspect clients-hq-policy
 class type inspect frm-clients-class
  pass
 class class-default
  drop
policy-map type inspect hq-clients-policy
 class type inspect to-clients-class
  pass
 class class-default
  drop
 class-map type inspect match-any to-infra-class
  match access-group 104
 class-map type inspect match-any to-clients-class
  match access-group 103
 class-map type inspect match-any frm-infra-class
  match access-group 102
 class-map type inspect match-any frm-clients-class
  match access-group 101
!
access-list 101 remark Client Source
access-list 101 permit ip 10.200.1.0 0.0.0.255 any
access-list 102 remark Infra Source
access-list 102 permit ip 10.201.1.0 0.0.0.255 any
access-list 103 remark Clients Dest
access-list 103 permit ip any 10.200.1.0 0.0.0.255
access-list 104 remark Infra Dest
access-list 104 permit ip any 10.201.1.0 0.0.0.255
```

Design Considerations

- An implicit deny applies as soon as a single zone is created on the device. Consequently, even if an interface is not placed in a zone, traffic will, by default, be denied.
- Policies are, by default, only applied to traffic flowing through the device, not to traffic directed to the device itself. This behavior can be modified by defining policies for what is referred to as the *self* zone.
- Once a baseline IOS zone-based firewall (ZBFW) design has been developed, advanced firewall inspection can easily be integrated by simply modifying the policies being enforced.

For more information on IOS ZBFW, see the Firewall section of Appendix A, “Reference Documents.”

ASA Integration in a Branch

The Adaptive Security Appliance (ASA) is a dedicated, fully featured firewall device enabling a scalable, high performance and high availability design, depending on a particular branch needs. It also provides support for separate management domains. See Figure 8-5.

The ASA is placed logically inline, between the branch access edge and the branch edge router, as well as between any additional security domains, to enforce network policy at the branch. The ASA access policy also typically includes access edge iACLs and, if the branch is hosting externally-accessible services or the branch edge router is not owned and managed by the enterprise, WAN edge ACLs.

Figure 8-5 Sample ASA Integration in a Branch

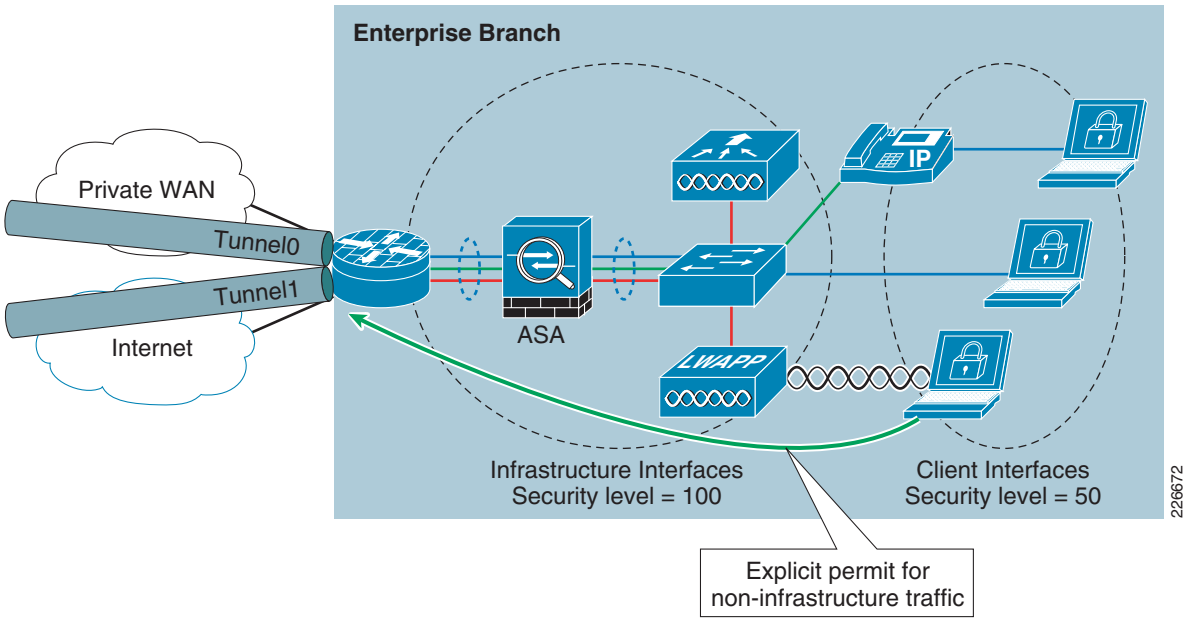

The Cisco ASA enforces network access policies based on the security level of an interface, with a default network access policy of an implicit permit for interfaces of the same security level, and an implicit permit from a higher to a lower security level interface. It is recommended, however, to enforce explicit policies as this provides maximum visibility into traffic flows.

For more information on the Cisco ASA, see the Firewall section of Appendix A, "Reference Documents."

A typical ASA deployment in a branch, as shown in Figure 8-5, illustrates the following:

- Access interfaces with a lower security level than the infrastructure interfaces. This ensures that, by default, clients are not permitted access to other interfaces.
- Access interface network access rules can also integrate the access edge iACLs.
- An explicit permit must be defined for client access non-infrastructure addresses.
- Infrastructure interfaces with a high security level permits traffic flows, by default, to other interfaces.
- There are no external interfaces on this ASA but, if they exist, they should be assigned the lowest security level and a strong network policy enforced.
- As with standard ACLs, a final, explicit deny should be enforced to provide maximum visibility into traffic being denied.

A sample ASA branch configuration is shown in Figure 8-6, illustrating the network access policy being enforced on a client access interface, including access edge iACLs.

Figure 8-6 Sample ASA Network Access Policy for a Client Access Interface

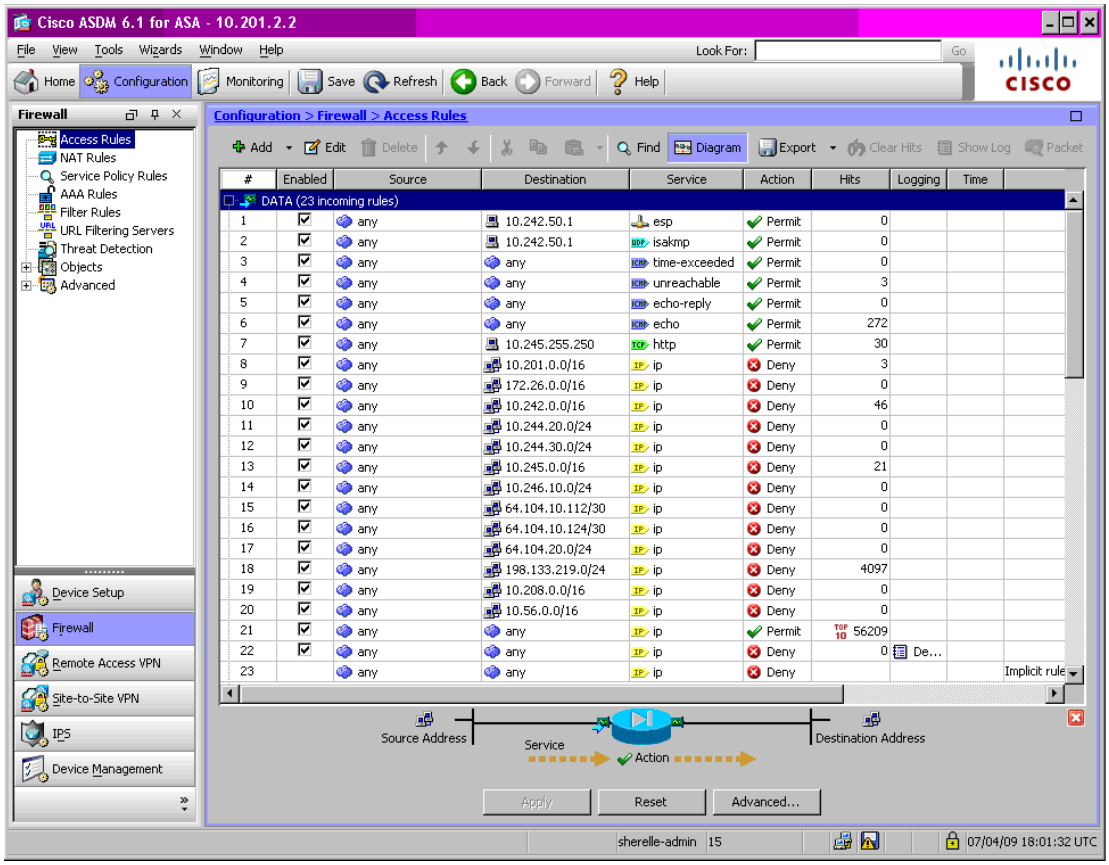

Additional policies can be enforced on the ASA, including application layer protocol inspection, IPS, and QoS. For more information on the ASA, see the Firewall section of Appendix A, "Reference Documents."

IPS Integration in the Branch

Cisco IPS provides signature and reputation-based threat detection and mitigation for threats such as worms, spyware, adware, network viruses, and application abuse. Its integration in a branch enables the localized detection and mitigation of malicious and anomalous activity. This is highly effective at enabling threats to be detected and mitigated in a timely manner and as close to the source as possible, thereby reducing the possible impact on the rest of the corporate network.

In addition, Cisco IPS collaboration with other Cisco devices provides enhanced visibility and control through system-wide intelligence. This includes host-based IPS collaboration with CSA, reputation-based filtering and global correlation using SensorBase, automated threat mitigation with the WLAN controller (WLC), multi-vendor event correlation and attack path identification using Cisco Security Monitoring, Analysis, and Response System (CS-MARS), and common policy management using Cisco Security Manager (CSM). For more information on Cisco security collaboration, see Chapter 10, "Monitoring, Analysis, and Correlation," and Chapter 11, "Threat Control and Containment."

The general IPS design considerations of deployment mode, scalability, availability and maximum threat coverage apply to a branch but the branch also, typically, introduces cost and management considerations that must be taken into account.

The Cisco IPS includes a wide range of platform options that enable customers to select a platform that is appropriate to their deployment criteria, as shown in Figure 8-7.

Figure 8-7 Cisco IPS Deployment Options

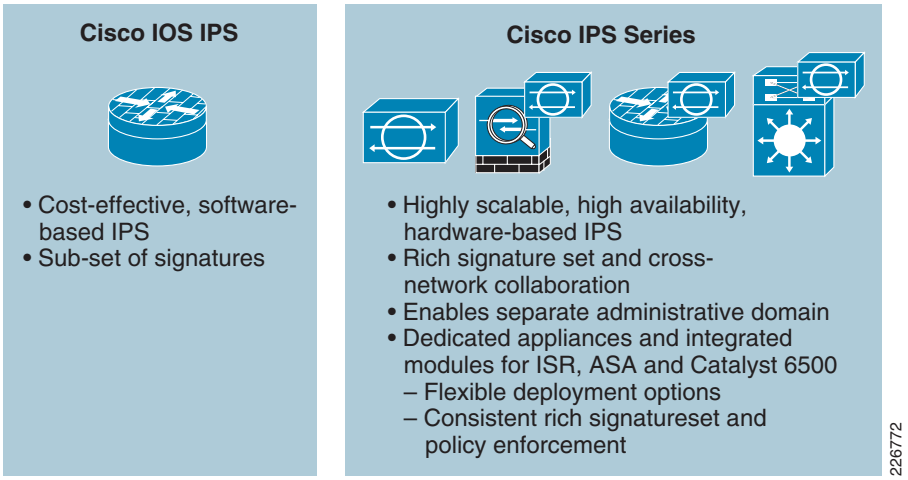

This wide range of IPS platforms shares a common set of signatures and can all be remotely managed by a central management platform, such as Cisco Security Manager (CSM). This enables consistent, rich signature and policy enforcement across the entire enterprise network, facilitating IPS tuning and operations, while at the same time accommodating the particular design criteria of diverse locations.

In large branch locations that require high scalability and availability, multiple IPS can be deployed.

Design Considerations

- IOS IPS currently only supports a sub-set of the signatures supported by the IPS devices and modules. In addition, IOS IPS does not currently support collaboration with other Cisco devices.
- The IPS modules provide a cost-effective IPS solution that maintains a consistent, rich signature set across all enterprise network IPS.
- IPS modules enable operational staff to easily migrate from promiscuous to inline mode, through a simple configuration change on the host platform.
- IPS modules offer limited scalability and availability that must be taken into account in a design.
- Symmetrical traffic flows offer a number of important benefits, including enhanced threat detection, reduced vulnerability to IPS evasion techniques and improved operations through reduced false positives and false negatives. Consequently, leveraging the Cisco IPS Normalizer engine is a key design element. If multiple IPS exist in a single flow, for instance, in multiple edge routers, maintaining symmetric flows requires careful consideration of the IPS integration design.
- It is recommended that IPS monitoring is performed on internal branch interfaces only in order to focus threat detection and mitigation on internal threats. This avoids the local IPS and the centralized monitoring station from being inundated with alerts that do not necessarily indicate a risk.

- IPS can, alternately, be integrated in the enterprise WAN edge as a centralized IPS deployment. This enables a scalable, highly available and cost-effective design, that also offers ease of management advantages, since it typically features a smaller number of devices. The threat coverage offered by this type of deployment must, however, be considered, since only traffic passing through the WAN edge distribution block will be monitored. For more information on a centralized IPS deployment, see Chapter 7, "Enterprise WAN Edge."
- A combination of centralized and distributed IPS enables the appropriate deployment model to be chosen according to the needs of a particular branch, while maintaining consistent policy enforcement.
- IPS signature tuning enables the automated response actions taken by Cisco IPS to be tuned and customized according to the customer environment and policy.

For more information on IPS design considerations as well as high scalability and availability IPS designs, see "IPS Integration in the WAN Edge Distribution" section on page 7-19.

Implementation Option

- IPS Promiscuous Mode

  Cisco IPS can also be deployed in promiscuous mode. In promiscuous mode, the IPS performs passive monitoring, with traffic being passed to it through a monitoring port. Upon detection of anomalous behavior, management systems are informed of an event and operational staff can subsequently decide what action, if any, to take in response to an incident. The time between threat detection and mitigation may thus be extended.

IPS Module Integration in a Cisco ISR

IPS integration in a small, cost-sensitive branch can leverage an IPS module integrated in the branch edge ISR. Integration of an IPS module enables a consistent, rich signature set across all enterprise network IPS.

IPS module integration is very simple to implement, with the IPS receiving traffic over the backplane of the ISR. Once the module is installed, it is simply a case of enforcing IPS monitoring on the desired interfaces. See Figure 8-8.

Figure 8-8 IPS Module Integration in a Cisco ISR

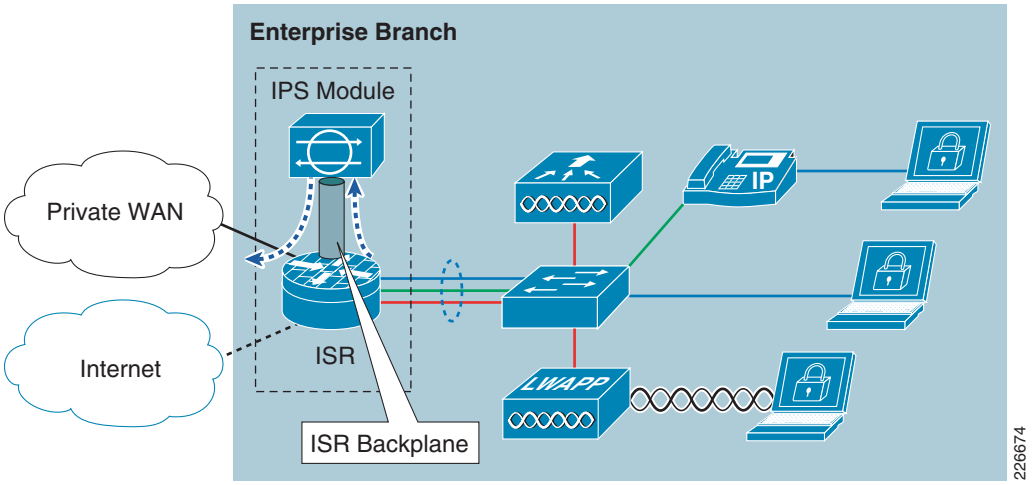

IPS monitoring is enforced on the ISR on a per-interface basis, as shown in the following example:

```
!
interface GigabitEthernet0/0.10
 description Wired Clients
 encapsulation dot1Q 10
 ip address 10.200.1.1 255.255.255.128
ids-service-module monitoring inline
!
```

The IPS configuration is a standard, consistent IPS policy that is enforced across the enterprise, as shown in Figure 8-9.

Figure 8-9 IPS Module in ISR Configuration

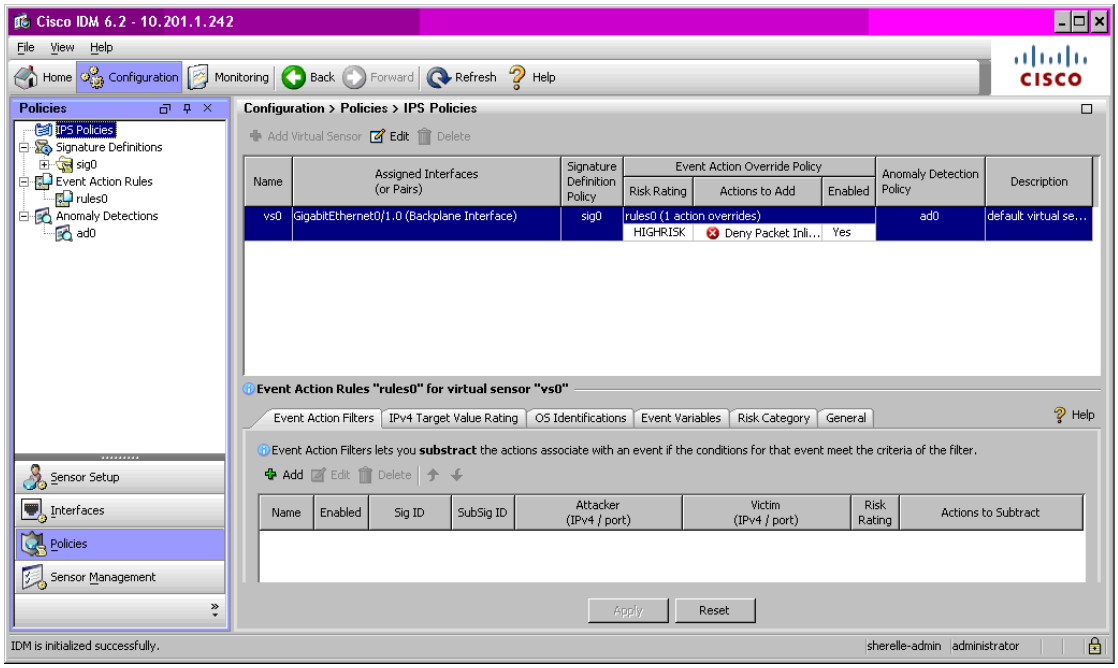

IPS Module Integration in a Cisco ASA

A branch with a Cisco ASA can integrate an IPS module in this ASA to provide a cost-effective, integrated solution. Integration of an IPS module enables a consistent, rich signature set across all enterprise network IPS.

IPS module integration is very simple to implement, with the IPS receiving traffic over the backplane of the ASA. Once the module is installed, it is simply a case of enforcing IPS monitoring on the desired interfaces. See Figure 8-10.

Figure 8-10 IPS Module Integration in a Cisco ASA

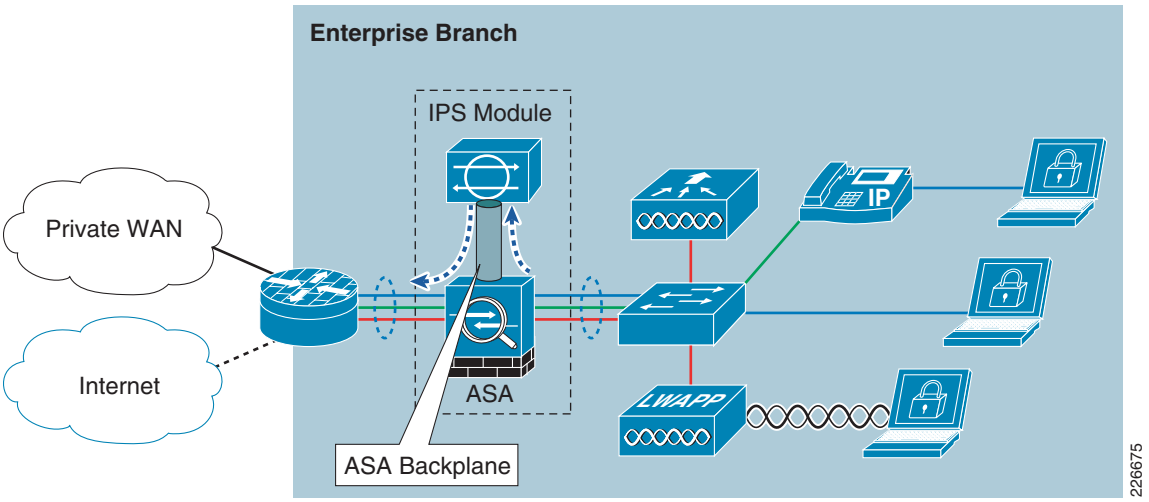

IPS monitoring is enforced on the ASA by applying a service policy on a per-interface basis, as illustrated in Figure 8-11.

Figure 8-11 IPS Policy Enforcement on the ASA

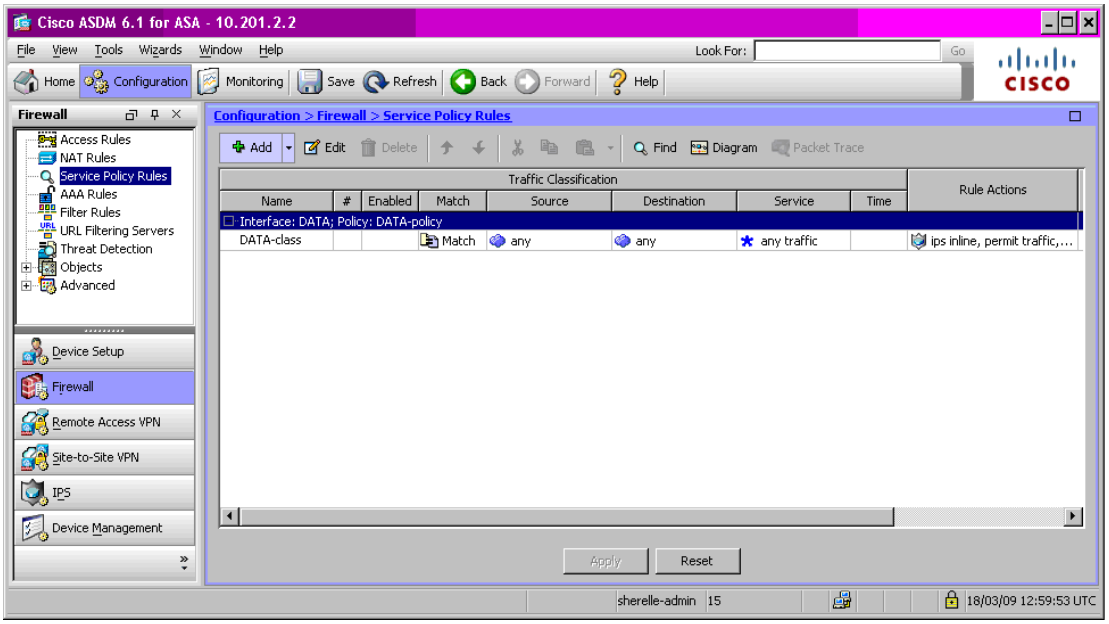

The IPS configuration is a standard, consistent IPS policy that is enforced across the enterprise, as shown in Figure 8-12.

Figure 8-12 IPS Module in ASA Configuration

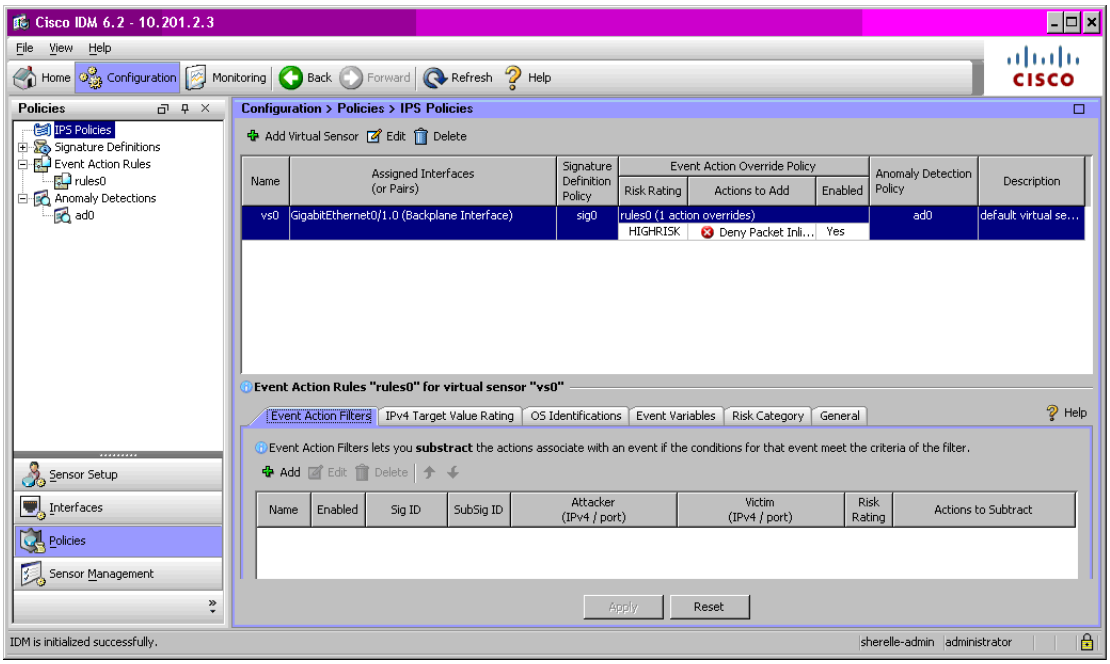

For more information on ASA integration in a branch, see the "ASA Integration in a Branch" section on page 8-17.

Switching Security in the Branch

Switching in the branch is critical to network services; therefore, its security and resiliency is vital to business operations. Consequently, the switching infrastructure and services themselves must be resilient to attacks against them, and they must offer protection to users and devices against attacks within their Layer 2 domain.

Possible attacks include DoS attacks through Spanning Tree Protocol (STP) manipulation, MAC flooding, DHCP starvation and unknown, multicast and broadcast frame flooding, reconnaissance, unauthorized access and MITM attacks through DHCP server spoofing, ARP spoofing, VLAN hopping, and STP manipulation. These attacks may be targeted or indirect, malicious or unintentional, conducted by authorized or unauthorized users, or performed by malware.

Consistent policies should be enforced across all enterprise switches including those in the campus, branch, data center, Internet edge and WAN edge, though taking into consideration the role of each switch, for example, if it is an access edge switch, a distribution switch or a core switch. Consequently, switching security in the branch involves extending and applying these same switching security policies to the branch switches.

The areas of focus, objectives, and implementation options for switching security in a branch access edge switch are listed in Table 8-5.

Table 8-5 Switching Security in the Branch Access Edge Switch

Switching Security Focus	Switching Security Objectives	Implementation
Restrict Broadcast Domains	Limit the Layer 2 domain in order to minimize the reach and possible extent of an incident	• Restrict each VLAN to an access switch.[1]
Spanning Tree Protocol (STP) Security	Restrict STP participation to authorized ports only	• Rapid Per-VLAN Spanning Tree (PVST) • BPDU Guard • STP Root Guard
DHCP Protection	Prevent rogue DHCP server and DHCP starvation attacks	• DHCP Snooping on access VLANs
ARP Spoofing Protection	Prevent ARP spoofing-based MITM attacks	• Dynamic ARP Inspection (DAI) on access VLANs[2]
IP Spoofing Protection	Ensure traffic is topologically valid, (i.e., sourced from a valid address that is consistent with the interface it is received on)	• IP Source Guard on access ports or uRPF on Layer 3 access edge
MAC Flooding Protection	Prevent switch resource exhaustion attacks that can cause flooding of a Layer 2 domain	• Port security on access ports
VLAN Best Common Practices	Apply VLAN security guidelines across the infrastructure	• Define a port as a trunk, access or voice port rather than enabling negotiation • VTP transparent mode • Disable unused ports and place in an unused VLAN • Use all tagged mode for the native VLAN on trunks • Traffic storm control

1. Keeping VLANs unique to an access switch is an enterprise campus BCP. For more information on enterprise campus design see the Campus Design section of Appendix A, "Reference Documents."
2. ARP spoofing attacks are often conducted in combination with IP spoofing, in order to avoid traceability. Consequently, IP spoofing protection should also be enforced on a switch. For more information on IP spoofing protection, see the Chapter 2, "Network Foundation Protection."

The following is a sample secure switching configuration:

```
Global Configuration
! Spanning Tree Security
spanning-tree mode pvst
spanning-tree portfast bpduguard default
! Enable DHCP Snooping on Access VLANs
ip dhcp snooping vlan 10,20
no ip dhcp snooping information option
ip dhcp snooping
! Enable DAI on Access VLANs with ARP ACLs for Default Gateway
ip arp inspection vlan 10,20
ip arp inspection filter staticIP vlan 10,20
 arp access-list staticIP
  permit ip host 10.200.1.1 mac host 0015.622e.8c88
  permit ip host 10.200.1.129 mac host 0015.622e.8c88
! VTP Transparent Mode
```

```
vtp mode transparent
vlan 10
 name DATA
!
vlan 20
 name VOICE
!
vlan 201
 name Mgmt
!
vlan 2000
 name Unused
!
! Native VLAN Tagging
vlan dot1q tag native
!
! Access Port
interface FastEthernet1/0/2
 switchport access vlan 10
 switchport mode access
 switchport voice vlan 20
 switchport port-security maximum 3
 switchport port-security maximum 2 vlan access
 switchport port-security maximum 1 vlan voice
 switchport port-security
 switchport port-security aging time 1
 switchport port-security violation restrict
 switchport port-security aging type inactivity
 ip arp inspection limit rate 100
 load-interval 60
 priority-queue out
 no snmp trap link-status
 storm-control broadcast level pps 1k
 storm-control multicast level pps 2k
 storm-control action trap
 no cdp enable
 spanning-tree portfast
 spanning-tree bpduguard enable
 spanning-tree guard root
 ip verify source
 ip dhcp snooping limit rate 15
!
! Trunk Port to ISR
interface GigabitEthernet1/0/1
 description Trunk to ISR
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,20,201
 switchport mode trunk
 ip arp inspection trust
 load-interval 60
 ip dhcp snooping limit rate 15
 ip dhcp snooping trust
!
! Unused Port
interface GigabitEthernet1/0/4
 switchport access vlan 2000
 shutdown
 no cdp enable
```

The particular considerations for DHCP protection and ARP spoofing protection in a branch are covered in detail below. For more information on the other switching security techniques, including design and configuration guidelines, see Chapter 2, “Network Foundation Protection.”

Design Considerations

- Rate-limiting must be enabled for both DHCP snooping and DAI in order to ensure that these features do not create a DoS vector.
- If automatic recovery of an interface after a security violation is not enabled, the interface will remain in an error-disabled state until it is manually recovered with a shutdown and no shutdown.
- DAI is highly effective for ARP spoofing protection in DHCP environments but must be bypassed for any device that does not use DHCP. This is achieved either by explicitly defining the interface it is connected to as trusted, or by creating an ARP inspection ACL to permit the source MAC and IP address of that device.
- The operational management of the switching infrastructure can be greatly enhanced my leveraging Smartports macros on a Cisco switch. Smartports macros enable customized port templates to be defined according to corporate policy and applied to ports on an as-needed basis. This ensures consistent policy enforcement, eases operations and avoids misconfiguration. For more information on Smartports macros, see the Switching Security section of Appendix A, "Reference Documents."

DHCP Protection

DHCP protection is critical to ensure that a client on an access edge port is not able to spoof or accidentally bring up a DHCP server, nor exhaust the entire DHCP address space by using a sophisticated DHCP starvation attack.

Both these attacks are addressed with the Cisco IOS DHCP snooping feature that performs two key functions to address these attacks:

- Rogue DHCP Server Protection

  If reserved DHCP server responses (DHCPOFFER, DHCPACK, and DHCPNAK) are received on an untrusted port, the interface is shut down.

- DHCP Starvation Protection

  Validates that the source MAC address in the DHCP payload on an untrusted interface matches the source MAC address registered on that interface.

In addition, DHCP snooping rate-limiting must be enabled to harden the switch against a resource exhaustion based DoS attack.

For more information on the DHCP Snooping feature, see the Switching Security section of Appendix A, "Reference Documents."

ARP Spoofing Protection

ARP spoofing protection ensures that a client on an access edge port is not able to perform a MITM attack by sending a gratuitous ARP that presents its MAC address as that associated with a different IP address, such as that of the default gateway.

This attack is addressed with the Cisco IOS Dynamic ARP Inspection (DAI) feature that validates that the source MAC and IP address in an ARP packet received on an untrusted interface matches the source MAC and IP address registered on that interface.

DAI is enabled on a per-VLAN basis and all interfaces in that VLAN are untrusted by default. Consequently, for a device that does not use DHCP, such as the default gateway, ARP inspection must be bypassed by either explicitly defining the interface it is connected to as trusted, or creating an ARP inspection ACL to permit the source MAC and IP address of that device.

In addition, DAI rate-limiting must be enabled to harden the switch against a resource exhaustion-based DoS attack.

For more information on the DAI feature, see the Switching Security section of Appendix A, "Reference Documents."

Endpoint Security in the Branch

Enterprise clients and servers are exposed to a range of threats, including malware, botnets, worms, viruses, trojans, spyware, theft of information, and unauthorized access. Consequently, the vulnerability of any particular endpoint can impact the security and availability of an entire enterprise network.

Endpoint security is thus a critical element of an integrated, defense-in-depth approach to security, protecting both the endpoint itself and the corporate network to which it connects. Consistent policies should be enforced across all enterprise clients and servers, and so endpoint security in the branch involves extending and applying these same security policies to branch endpoints.

Endpoint security must not only harden the endpoint against the initial attack, compromise and subsequent activity, but also general malicious and non-compliant client activity. This is generally referred to as host-based IPS and the key elements are as follows:

- Protection against known attacks

  Signature-based threat detection and mitigation, such as known worms and viruses.

- Protection against zero-day or unpatched attacks

  Behavioral-based threat detection and mitigation, such as attempts to load an unauthorized kernel driver, buffer overflow, capture keystrokes, create rogue services, modify system configuration settings and inset code into other processes.

- Policy enforcement

Host-based IPS functionality to support all these key elements is provided by the Cisco Security Agent (CSA). CSA is deployed on all endpoints and centralized policy management and enforcement is performed by the CSA Management Center (CSA-MC), enabling a common and consistent policy enforcement across all enterprise endpoints.

For more information on CSA and endpoint security in the enterprise, see Chapter 5, "Enterprise Campus."

Design Considerations

- Endpoint security protects the client even when they are not connected to the corporate network, such as at a hotel, a hotspot or home.
- CSA on enterprise endpoints is typically managed by a centralized CSA MC. CSA continues to protect the endpoint, even when the CSA MC is not accessible and so a branch WAN outage is not an issue.
- CSA policies that are enforced based on location-aware policies must take into consideration the fact that the CSA MC may not always be reachable by branch endpoints. For instance, a policy that blocks local network access if the CSA MC is not reachable may not be viable in a branch.
- User behavior is a key factor in endpoint and overall network security. This is becoming even more critical as attacks evolve to focus on social engineering and targeted attacks. CSA can be leveraged to reinforce user education and training by, for instance, advising users when they perform an anomalous action, outlining the risks it presents and the associated corporate policy, before allowing them to permit the action, along with, perhaps, a justification.

Complementary Technology

Additional endpoint security includes the following:

- Cisco Security Services Client (CSSC)

  The CSSC is a software supplicant that enables identity-based access and policy enforcement on a client, across both wired and wireless networks. This includes the ability to enforce secure network access controls, such as requiring the use of WPA2 for wireless access and automatically starting a VPN connection when connected to a non-corporate network.

  For more information on CSSC, see the Endpoint Security section of Appendix A, "Reference Documents."

- System and application hardening

  It is critical that the operating system and applications running on an endpoint are hardened and secured in order to reduce the attack surface and render the endpoint as resilient to attacks as possible. This involves implementing a secure initial configuration, the regular review of vulnerabilities and the timely application of any necessary updates and security patches.

- User education and training

  End-users should receive ongoing education and training to make them aware of the role they play in mitigating existing and emerging threats, including how to be secure online, protecting corporate data, and minimizing their risk. This should be presented in a simple, collaborative way to reinforce the defined corporate policies.

Secure Device Access in the Branch

Access to all infrastructure devices in the branch must be secured. If infrastructure device access is compromised, the security and management of the entire network can be compromised. Consequently, it is critical to establish the appropriate controls in order to prevent unauthorized access to infrastructure devices.

There will be some variations in the actual implementation of secure device access, based on the particular device and software release, but all the fundamental objectives must be applied:

- Restrict device accessibility

  Limit the accessible ports and access services, restrict access to authorized services from authorized originators only, enforce session management and restrict login vulnerability to dictionary and DoS attacks

- Present legal notification

  Display legal notice, developed in conjunction with company legal counsel, for interactive sessions.

- Authenticate access

  Ensure access is only granted to authenticated users, groups, and services.

- Authorize actions

  Restrict the actions and views permitted by any particular user, group, or service.

- Ensure the confidentiality of data

  Protect locally stored sensitive data from viewing and copying. Consider the vulnerability of data in transit over a communication channel to sniffing, session hijacking, and MITM attacks.

- Log and account for all access

  Record who accessed the device, what occurred, and when for auditing purposes.

For more information on secure device access, including design and configuration guidelines for the areas outlined above, see Chapter 2, "Network Foundation Protection."

Design Considerations

- Remote Management

  Remote management typically occurs in-band for remote sites; therefore, it is critical that management access and management traffic be properly isolated, secured, and resilient to challenging network events, such as high data rates and worm outbreaks. This is achieved through service resiliency measures designed to ensure remote manageability even under adverse circumstances, including device hardening and QoS. For more information on service resiliency, see the "Service Resiliency in the Branch" section on page 8-8.

- Integrated IPS Module Access

  It is generally recommended that outgoing access is not permitted on a device, unless explicitly required. One example of such a requirement is access to an IPS module integrated in an ISR. Console access to the integrated IPS module is only available through a reverse telnet from the host ISR. Consequently, outgoing telnet must be permitted on the host ISR console and VTY lines.

Telemetry in the Branch

Telemetry must be extended to the branch in order to provide visibility into its status, as well as activity on both the local network and its external interfaces. This enables the timely and accurate detection and mitigation of anomalies. Consequently, telemetry is enabled across all infrastructure devices in the branch and integrated with a centralized management system for event monitoring, analysis, and correlation. The key elements include:

- Synchronize time

  Synchronize all network devices to the same network clock by using Network Time Protocol (NTP) to enable accurate and effective event correlation.

- Monitor system status information

  Maintain visibility into overall device health by monitoring CPU, memory, and processes.

- Implement CDP best common practices

  Enable CDP on all infrastructure interfaces for operational purposes but disable CDP on any interfaces where CDP may pose a risk, such as external-facing interfaces.

- Enable remote monitoring

  Use syslog and SNMP to a centralized server, such as CS-MARS, for cross-network data aggregation. This enables detailed and behavioral analysis of the data which is key to traffic profiling, anomaly-detection and attack forensics, as well as general network visibility and routine troubleshooting.

For more information on telemetry, including design and configuration guidelines for the areas outlined above, see Chapter 2, "Network Foundation Protection."

For more information on remote monitoring, analysis and correlation, including syslog, SNMP and NetFlow, see Chapter 10, "Monitoring, Analysis, and Correlation."

Design Considerations

- CDP is enabled by default in Cisco IOS and should be disabled on all external-facing interfaces. This can be verified on a per interface basis using the command show cdp interface.
- CDP may pose a risk on access ports but it should be noted that some services, such as Cisco UC phones, require CDP. Thus, care must be taken when disabling CDP.
- As with secure device access, the isolation of management access and management traffic is recommended using an out-of-band (OOB) network in order to provide an extra degree of security. This is typically employed using an OOB network that is physically independent of the data network and features limited and strictly controlled access. For more information on the implementation of a management network, refer to Chapter 9, "Management."
- One key element to consider is that, since these are remote sites, telemetry and remote management are typically in-band. It is thus critical that QoS is employed to accurately classify and prioritize control and management traffic. This will ensure continuing service availability and remote management even under adverse network conditions, such as high data rates and worm outbreaks.

Threats Mitigated in the Enterprise Branch

Table 8-6 Enterprise Branch Threat Mitigation Features

		Botnets	DoS	Unauthorized Access	Phishing, Spam	Malware, Spyware	Application, Network Abuse	Data Leakage	Visibility	Control
Secure WAN Connectivity				Yes				Yes		Yes
Routing Security			Yes	Yes					Yes	Yes
Service Resiliency	**Device Hardening** **QoS** **Redundancy**		Yes	Yes						Yes
Network Policy Enforcement	**WAN Edge ACLs** **Access Edge iACLs** **Cisco Firewall** **IP Source Guard or uRPF**	Yes		Yes			Yes	Yes		Yes
Cisco IPS Integration		Yes				Yes	Yes			Yes
Switching Security			Yes	Yes			Yes	Yes		
Endpoint Security		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Secure Device Access				Yes				Yes	Yes	Yes
Telemetry		Yes	Yes	Yes			Yes		Yes	

Web and E-mail threats, such as malicious web sites, compromised legitimate web sites, spam, and phishing, are addressed as part of a centralized deployment in the enterprise Internet edge (see Chapter 6, "Enterprise Internet Edge.") If a branch employs split-tunneling, whereby there is local Internet access direct from the branch, web security must be implemented locally. This can be achieved using the Cisco IronPort S-Series, Cisco ASA 5500 Series, or Cisco IOS Content Filtering. For more information, see the Web Security section of Appendix A, "Reference Documents."