Internet Edge Solution Overview

December 2, 2010

Contents

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

System Overview

This document describes a set of best practices to successfully designing, implementing, and maintaining an Internet edge. While this document is written with medium and large size enterprises in mind, most of the principles presented here are universal and therefore applicable to other environments as well.

The Internet edge is the network infrastructure that provides connectivity to the Internet, and that acts as the gateway for the enterprise to the rest of the cyber space. The Internet edge serves the other many building blocks (referred by Cisco as places in the network (PINs)) present in a typical enterprise network. This modular building block approach enables flexibility and customization in network design to meet the needs of different size customers and their business models.

As Figure 1 shows, users at the campus access the Internet through the Internet edge; the enterprise website and other public resources are accessible to clients and partners through the Internet edge, mobile and home-based employees may access corporate resources and applications through the Internet edge; and the Internet edge can also provide backup access to remote and branch offices in case the primary WAN links fail.

Figure 1 Internet Edge Topology

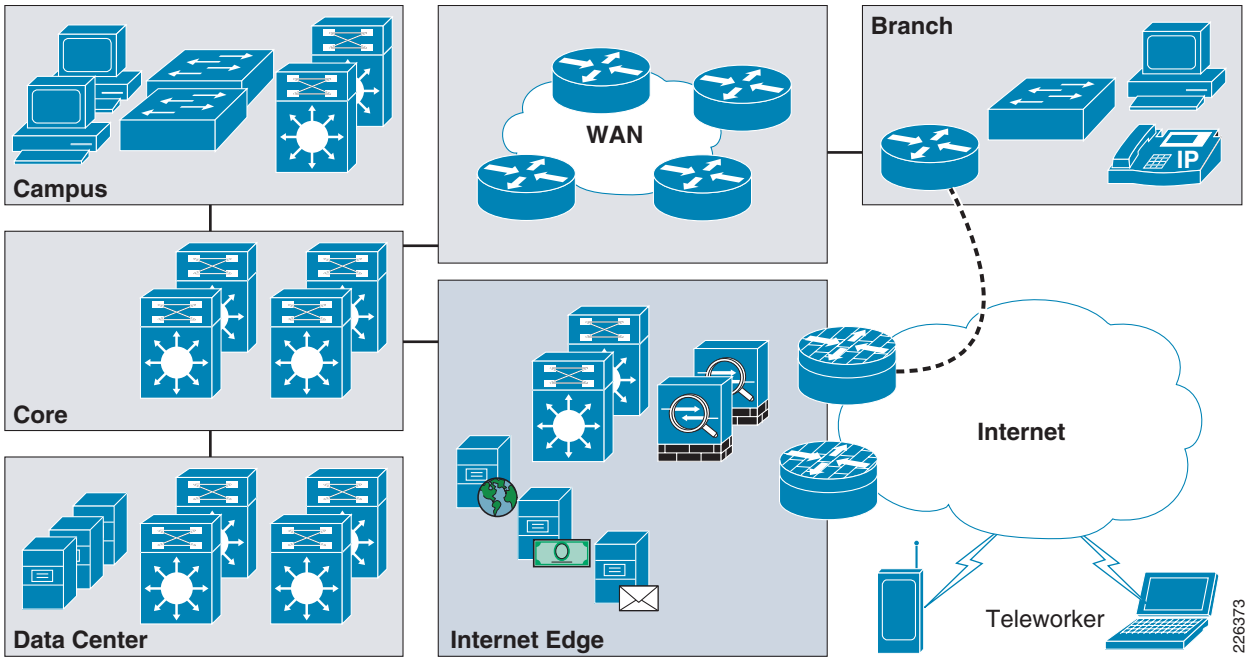

As the gateway to the Internet, the Internet edge infrastructure plays a critical role in supporting the services and activities that are fundamental to the operation of the modern enterprise. For this reason, the Internet edge has to be designed to provide service availability and resiliency, to be compliant with regulations and standards, to provide flexibility in accommodating new services and adapt with the time, to be secure, and to facilitate administration (reducing OPEX).

Service Availability and Resiliency

The disruption of e-commerce portals, corporate websites, and communication channels with partners, are all examples of events that could severely inhibit the productivity and even halt the business operation of a corporation. The Internet edge design here proposed incorporates several layers of redundancy to eliminate single points of failure and to maximize the availability of the network infrastructure. The design also uses a wide set of features destined to make the network more resilient to attacks and network failures.

Regulatory Compliance

Standards such as the Payment Card Industry Data Security Standard (PCI DSS) for the payment card industry and regulations like Health Insurance Portability and Accountability Act (HIPAA) for the health industry impose a series of requirements to be followed by organizations, and for which noncompliance may lead to the revocation of licenses, stiff penalties, and even legal actions. The Internet edge design includes a security baseline built-in as intrinsic part of the network infrastructure. The security baseline incorporates a rich set of security practices and functions commonly required by regulations and standards, and provides a solid platform to achieve regulatory compliance.

Modularity and Flexibility

The Internet edge follows a modular design where all components are described by functional roles rather than point platforms. This results in added flexibility when it comes to selecting the best platform for a given functional role, enabling the network to fit your business model and grow with your business. At the same time, this modular design facilitates the implementation of future services and roles, extending the useful life of existing equipment and protecting previous capital investment (CAPEX).

Security

The rise in organized crime use of the Internet, cyber espionage, growing data theft, the increasing sophistication of network attacks, are all examples of the real threats faced by organizations these days. As a key enabler of the business activity, networks need to be designed with security in mind, and to ensure the confidentiality, integrity and availability of applications, endpoints and the network itself. The Internet edge design incorporates security as an intrinsic component of the network architecture, where a rich set of security technologies and capabilities are deployed in a layered approach, but under a common strategy. The selection of technologies and capabilities is driven by the application of the Cisco Security Framework (CSF), a methodology that aims at achieving complete visibility and total control.

Operational Expenditures

As operational expenditures continue to rise, and as the cost of hiring and training personnel increases, designing networks that facilitate operations becomes a fundamental requirement for cost reduction. The Internet edge is designed to accommodate operations, right from deployment and throughout the operational life cycle. In addition to guiding the design and initial deployment, this guide presents an implementation roadmap, allowing users to start with a subset of the design, and systematically implement the remaining technologies and capabilities as they see fit. With a focus on operations, tools and procedures are provided to verify the effectiveness and the proper operation of each network element in the design.

Customer Use Cases

Medium and large size enterprises with more than 500 users onsite typically require Internet access to serve externally-facing data centers, campus users, mobile users, and to provide backup for remote offices.

Public Services DMZ

Traditionally, public-facing services were typically placed on a demilitarized zone (DMZ) for security and control purposes. The DMZ acts as a middle stage between the Internet and organization's private resources, preventing external users from direct access to internal servers and data. In today's network, most public services such as email and web serverfarms are located inside in the data center. DMZs in today's network normally provide network services such as DNS, FTP, NTP, etc. Other services implemented at a DMZ often include the email and web security appliances. See Figure 2.

Figure 2 DMZ Topology

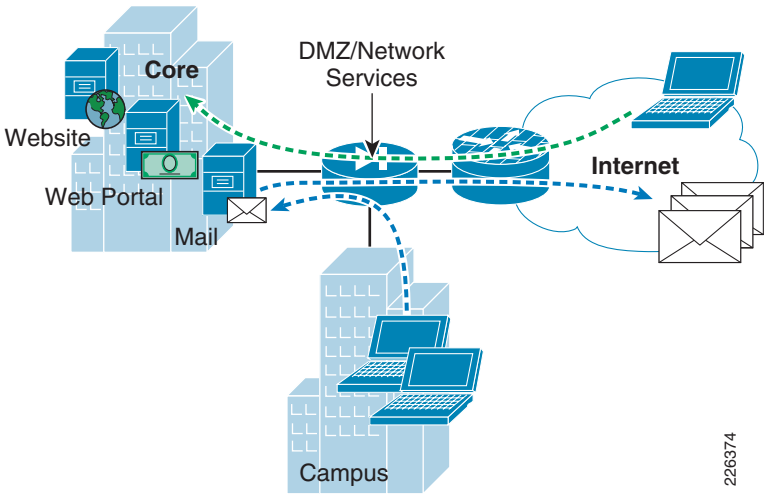

The following are some of the key attributes to be expected in the DMZ design:

- Service availability and resiliency
- Regulatory compliance
- Security: prevent intrusions, data leakage and fraud, and ensure user confidentiality and data integrity

Corporate Internet Access

Users at the campuses access email, instant messaging, web browsing, and other common services through the existing Internet edge infrastructure at the headquarters or regional offices. Depending on the organization's policies, users at the branches may also be forced to access the Internet over a centralized Internet connection, typically at the headquarters. These cases are represented in Figure 3.

Figure 3 Internet Access

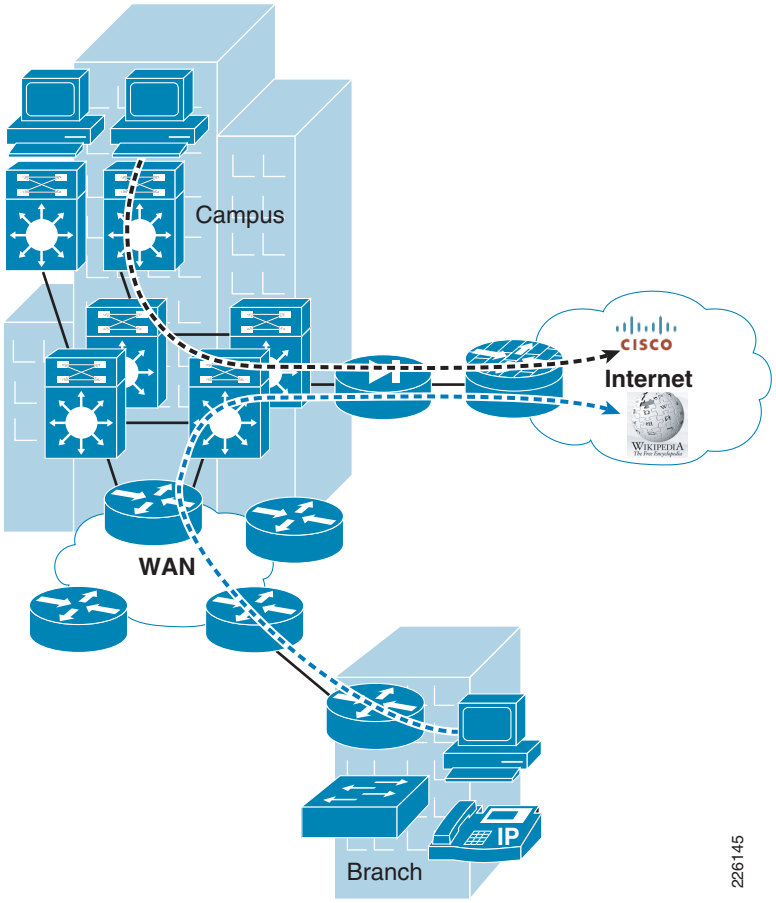

The following are some of the key attributes to be expected in the Internet access design:

- Service availability and resiliency
- Regulatory compliance
- Security: prevent network abuse, intrusions, data leakage and fraud, and ensure user confidentiality and data integrity

Teleworker

The Internet edge infrastructure may also provide mobile users and teleworkers with access to private applications and data residing in the organization's network.

This sort of remote access is authenticated and secured with SSL or IPSec VPNs. Access control policies may also be enforced to limit access to only the necessary resources and according to the user's role. Typical services provided to mobile users and teleworkers include email, access to intranet websites, business applications, video on demand, IP telephony, instant messaging, etc. See Figure 4.

Figure 4 Teleworker Topology

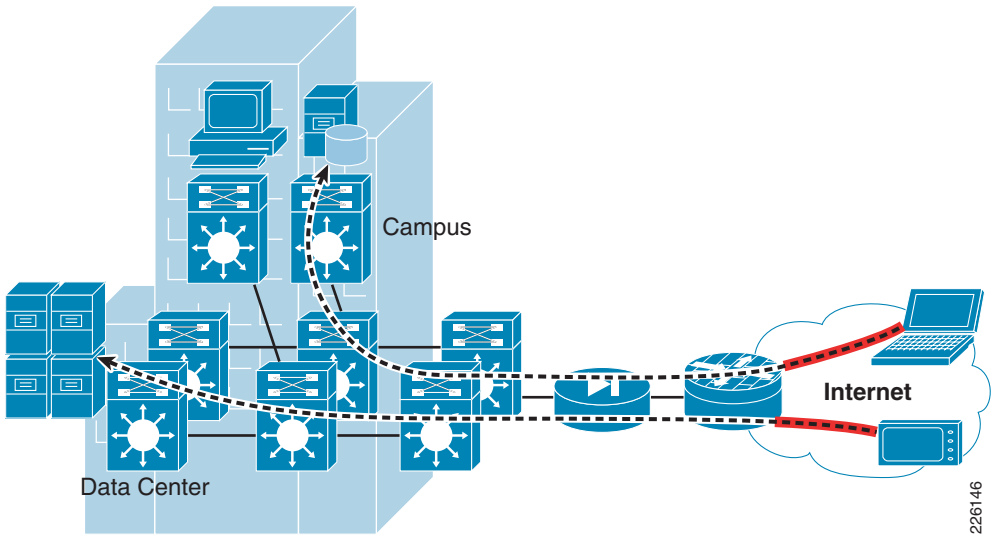

The following are some of the key attributes to be expected in the teleworker design:

- Service availability and resiliency
- Regulatory compliance
- Security: prevent network abuse, intrusions, data leakage and fraud, and ensure user confidentiality, data integrity, user segmentation.

The design proposed in this document presents a firewall based teleworker solution. Separate firewalls will be used to segment teleworker traffic from other traffic flows.

Branch Internet Connectivity

Under normal conditions, if the branch has not implemented split-tunneling, all Internet-bound traffic from the branch will have to go through the headend. The Internet-bound traffic will pass through the WAN-edge and out through the Internet edge. It is therefore imperative that all Internet traffic from the branches is treated in a similar fashion to Internet traffic from corporate users. This implies that all monitoring, threat mitigation tools and enforcement policies has to apply to branch-originated Internet-bound traffic.

WAN Backup

To ensure business continuity and service availability, remote offices may implement an Internet connection to be used as a backup of the primary WAN links (see Figure 5). Since the Internet is a public medium, communications to headquarters or regional offices are to be secured with a Virtual Private Network (VPN) technology like IPSec. In this scenario, VPN backup connections are terminated and aggregated at the headquarters or at regional offices. For the same reason, branch routers and other Internet-facing equipment need to be properly hardened.

In case centralized controls and security measures are favored, the organization may enforce a policy to prevent branch users from accessing the Internet directly. In this case, Internet access may be provided centrally at the headquarters or regional offices.

Depending on the bandwidth available, branch users may be limited to a subset of applications during the failover of the primary WAN links.

Figure 5 Internet WAN Backup Topology

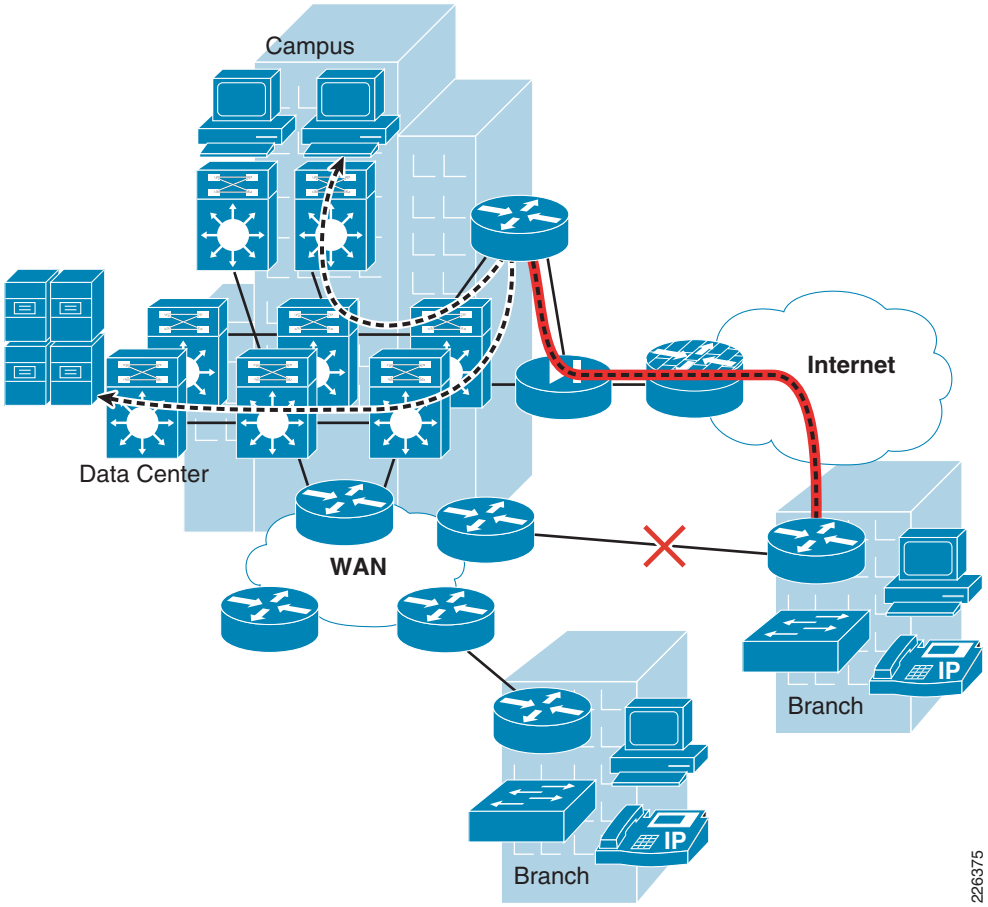

The following are some of the key attributes to be expected in the Internet WAN backup design:

- Service availability and resiliency
- Regulatory compliance
- Security: prevent network abuse, intrusions, data leakage and fraud, and ensure user confidentiality and data integrity.

Systems Architecture

Integrated Services Model and Appliance Model

The Internet edge architecture can be implemented using either the integrated model or the appliance model. In the integrated model, various capabilities and functional blocks are integrated within the same appliance.

There are several benefits of integrating security functions on a router or a switch.

1. First, everyone uses routers for routing purposes. Adding security to them reduces the number of boxes one must support and maintain in the network. This significantly reduces the real estate required for equipment and often reduces CAPEX.

2. Second, because the router already takes an active role in the overall network (routing and switching traffic), adding security functions can usually be done without impacting the network design.
3. Third, smaller sites which may not have any administrators exclusively to manage the security can now use network operators to maintain security as well. An integrated router platform series can provide that viable alternative where a router can be used for core IP routing services such as IGP/BGP, QoS, v4/v6 Multicast, NAT, NetFlow, GRE, RTP compression ISSU, and many others. At the same time, it can perform advanced technology function such as firewall, encryption, participating in WAN optimization functions by way of PfR or WCCPv2 along with Cisco WAAS solution. All this can be done at multiple speeds with 5, 10, or 20 Gbps within a single platform.

The benefits of an integrated model are as follows:

- Simplified operations and reduced cost
- Consolidation and service aggregation
- Faster and reliable service deployments
- Reduced carbon footprint-efficient power consumption

On the other hand, many large enterprises prefer to keep routing and advanced services in separate platforms for multiple reasons including but not limited to high availability, use of best-of-breed products, enforcing the boundary of responsibilities between Net-operations (NetOP) and Sec-operations (SecOP), and perhaps sometimes just the preference of the user interface that the existing staff is comfortable with.

The benefits of an appliance model are as follows:

- Organizational NetOP and SecOP boundaries
- Separate feature domains that isolates failures
- Separate feature domains easing management and troubleshooting
- Increased availability
- Better scaling for high-end large implementations

Note The first phase of the design of this Internet solution architecture emphasizes on the appliance model. The second phase of this project will use the integrated model as the main approach in designing the solution.

Common Infrastructure

The use cases described earlier in this document may be implemented under a common shared infrastructure or on independent Internet edges. The later implies building duplicated infrastructures, which results in higher capital expenditures and operational expenses. For this reason, most organizations implement a common Internet edge infrastructure to satisfy the different use cases.

Figure 6 depicts the common elements present in a shared infrastructure.

Figure 6 Common Infrastructure

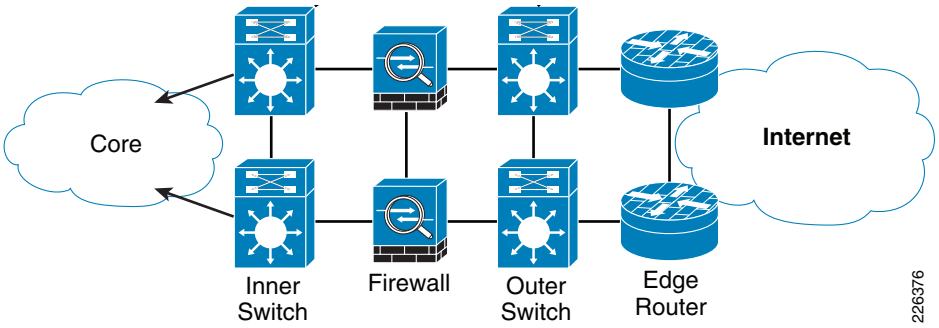

Common elements in a shared infrastructure include the following:

- Edge routers

  The primary function of the edge routers is to route traffic between the organization's network and the Internet. They provide connectivity to the Internet through one or more Internet service providers (ISPs). Edge routers may also provide QoS and rate-limiting. In terms of security, the edge routers act as the first line of defense against external attacks. Access control lists (ACLs), uRPF, and other filtering mechanisms are implemented for antispoofing and to block invalid packets. NetFlow, syslog, SNMP are used to gain visibility on traffic flows, network activity and system status. In addition, edge routers are secured following the practices discussed in the "Baseline Security" section on page 15. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the dynamic exchange of routing information.

- Outer switches

The outer switches provide data link layer (Layer 2) connectivity between the edge routers and the firewalls. The outer switches are secured following the principles explained in the "Baseline Security" section on page 15. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching infrastructure.

- Firewall

  The Internet edge design implements firewalls and uses their stateful access control and deep packet inspection to accomplish the following:

  - Protect the organization's internal resources and data from external threats by preventing incoming access from the Internet
  - Protect public resources served by the DMZ by restricting incoming access to the public services, and by limiting outbound access from DMZ resources out to the Internet
  - Control user's Internet-bound traffic

  In addition, firewalls may provide address translation (NAT/PAT) and may be used to terminate VPN tunnels. This design will use separate firewalls for teleworker access and corporate connectivity.

  Administrative access to the firewalls is also secured following the same principles described in the "Baseline Security" section on page 15.

- Inner switches

The inner switches provide network layer (Layer 3) and data link layer (Layer 2) connectivity between the Internet edge and the rest of the enterprise network, typically through the core. The inner switches are configured with a routing process that routes information between the VLANs that connect to the core switches and the firewall inside VLAN.

The function of the inner and outer switches can be collapsed in a single pair of switches. In that case, the inside and outside segments of the firewall need to be properly segmented with VLANs.

The inner switches are secured following the principles explained in the "Baseline Security" section on page 15. This includes restricting and controlling administrative access, protecting the management and control planes, and securing the switching infrastructure.

Routing and Switching

As illustrated in Figure 7, the outer switches implement a single Layer 2 segment or VLAN that provides connectivity between the firewalls and the edge routers. This VLAN is the firewall's outside segment. As firewall failover requires all firewall interfaces to be Layer 2 adjacent, the outside VLAN needs to be carried by both outer switches.

Each interface of the edge routers is to be configured as Layer 3 segments. In case a redundant ISP connection is implemented, each Layer 3 interface of the router is likely to be configured with different IP subnets.

The firewalls are configured in hot-standby mode, therefore logically the active and standby units share the same IP addresses. Nevertheless, all firewall interfaces need to be Layer 2 adjacent, which means the VLANs and Layer 2 segments on the firewall need to be trunked across the primary and secondary switches.

Figure 7 Layer 3 and Layer 2 Design

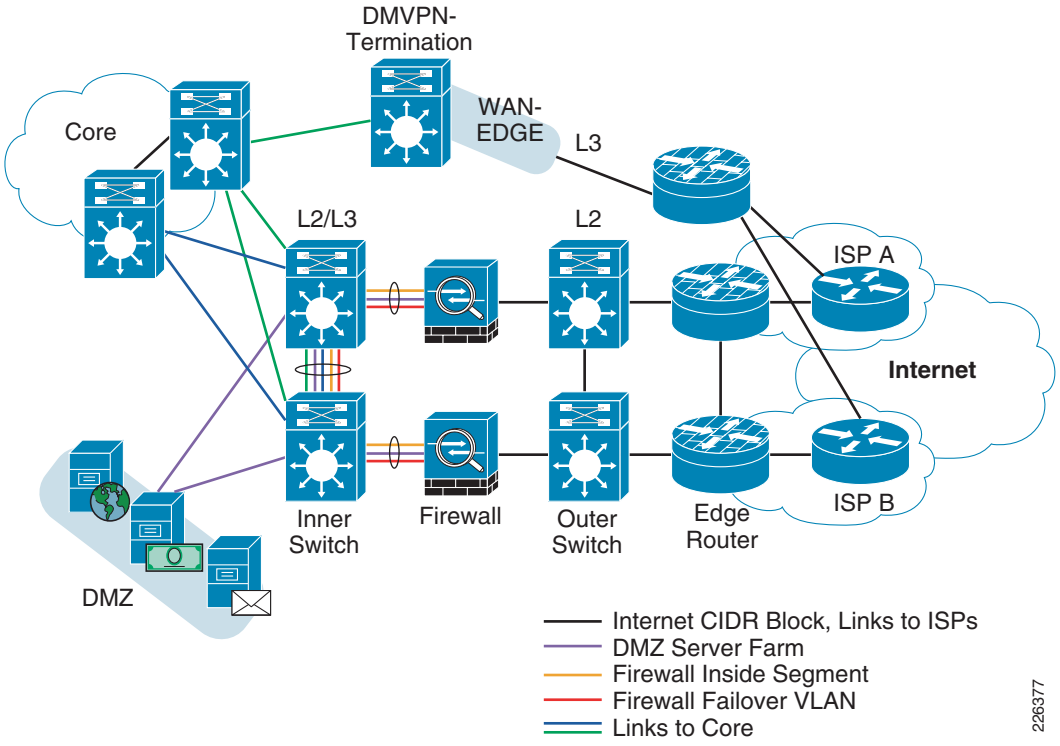

The inner switches are configured with multiple VLANs supporting the DMZ services. All DMZ VLANs converged on the firewalls and do not participate on the switch's routing process. In addition, these switches provide Layer 3 connectivity to the core of the enterprise network through a set of VLANs connecting to the core switches. All VLANs are trunked between the primary and secondary switches.

Note that the functions provided by the inner and outer switches can be collapsed on a single pair of redundant switches. In this case, proper segregation between the inside and outside firewall segments needs to be ensured.

Dynamic routing is implemented by using a combination of BGP and an IGP protocol such as OSPF or EIGRP. The edge routers run external BGP sessions to the ISP routers. These are shown in Figure 8. Assuming the organization owns a CIDR block, the same address space may be announced by the ISP (or ISPs in case of having two providers).

Figure 8 Routing Processes

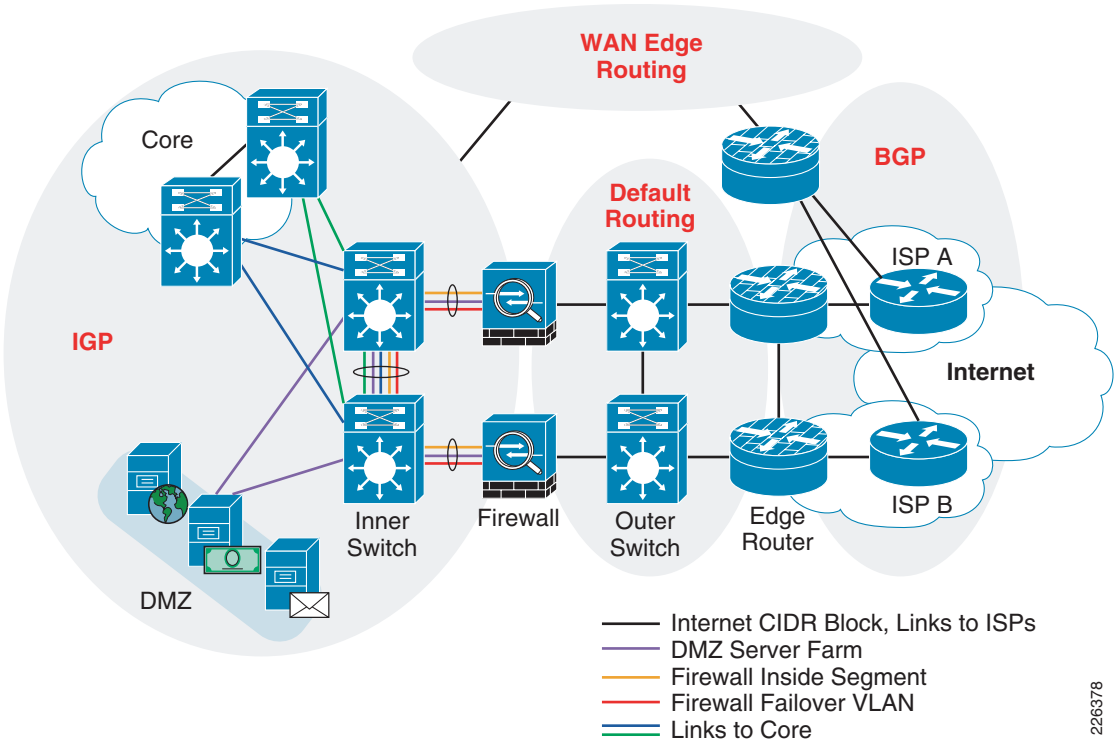

The firewalls may be configured with static routing or an IGP to inject a default route to the interior of the network. In case of an IGP is used, it is important not to use the same routing process as inside the firewall. As a best practice, different routing processes should be implemented for subnets inside and outside the firewall to allow for better control and reduce the possible effects of an attack on the routing infrastructure.

Internally, an IGP may be used for dynamic routing. This routing process can be configured on the inside interfaces and DMZ interfaces of the firewall. This allows for the dynamic propagation of route information. As described in the "Baseline Security" section on page 15, dynamic routing is secured by the enforcement of neighbor authentication, route filtering and device hardening.

High Availability

The Internet edge is built out of many platforms and components that may fail or that may be subject to attack. Designing a redundant architecture helps eliminate single points of failure, therefore improving the availability and resiliency of the network.

The Internet edge is designed with several layers of redundancy including redundant interfaces, standby devices, and topological redundancy.

Redundant Interfaces

The design includes the adoption of redundant interfaces at various points of the architecture to provide alternative paths. Dynamic routing protocols are used for path selection. The design allows for the use of multiple ISP connections, each served by different interfaces.

Standby Devices

The Internet edge design implement redundant firewalls and routers by using the existing firewall failover mechanisms and router redundancy protocols.

The design implements firewall redundancy with a pair of units deployed in stateful active/standby failover mode. In this configuration, at any given time one of the firewalls is active while the other one remains idle in standby mode. Under normal operation, only the active firewall processes network traffic. Firewall configuration is maintained synchronized between the active and standby units. In addition, and thanks to the stateful nature, the active unit shares the connection state and flow information with the stand by unit. When a failure occurs, the standby firewall becomes active and starts processing network traffic.

Firewall redundancy may also be implemented in active/active failover mode, in which case all units process network traffic. This failover mode requires the separation of traffic flows across the active firewalls. This is achieved by defining multiple firewall contexts, which is out of the scope of this document.

The Internet edge design also makes use of a First Hop Redundancy Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP) or Gateway Load Balancing Protocol (GLBP) to allow for redundant routers. These protocols provide transparent failover of the first-hop IP router in segments not configured with a dynamic routing protocol (i.e., firewall outside segment). In this configuration, two routers are set up together in a group, sharing a single IP address that other systems in the segment use as the next hop. One of the two routers is elected as the active router, and it is responsible for handling all traffic sent to the IP address. In the event the active router fails, the standby router takes over.

Topological Redundancy

The Internet edge implements redundant links and devices to extend availability and to make the network more resilient to attacks and errors. Topological redundancy is implemented at both the network as well as the date link level. At the network level, it is implemented by using redundant routers and network links and by using a dynamic routing protocol. At the data link level, redundant switches multiple paths are implemented in conjunction with a spanning tree protocol.

Figure 9 illustrates the high availability design.

Figure 9 Redundancy

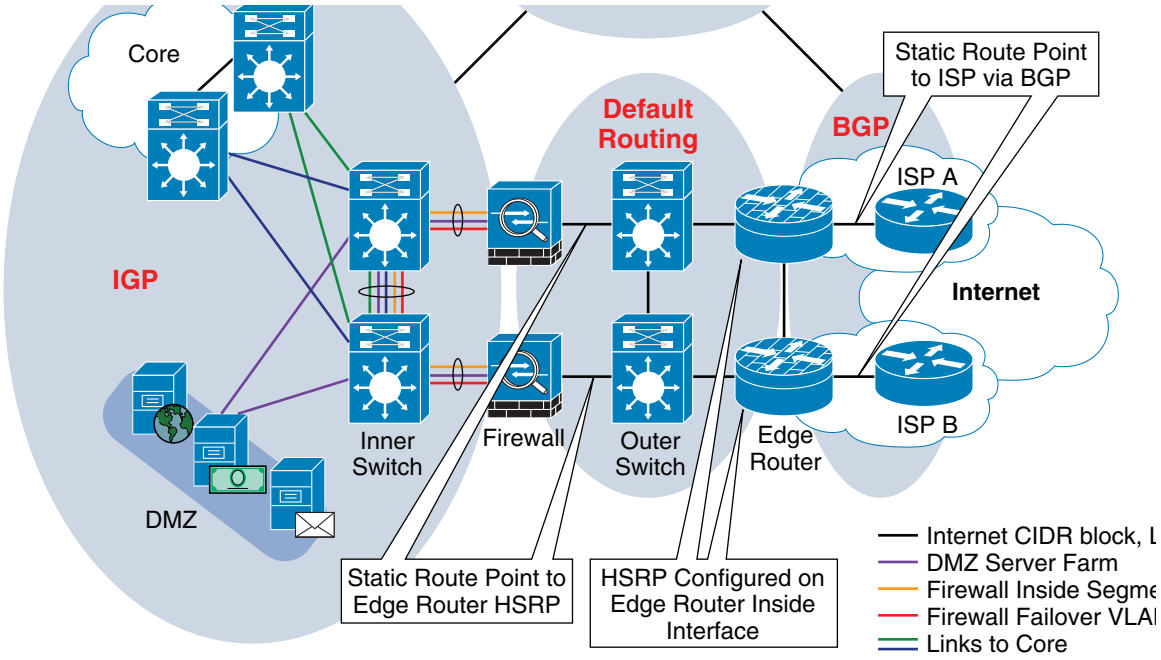

Management Network

The Internet edge design includes a management network dedicated to carrying control and management plane traffic such as NTP, SSH, SNMP, syslog, etc. The management network combines out-of-band management and in-band management as illustrated in Figure 10.

Figure 10 Management Network

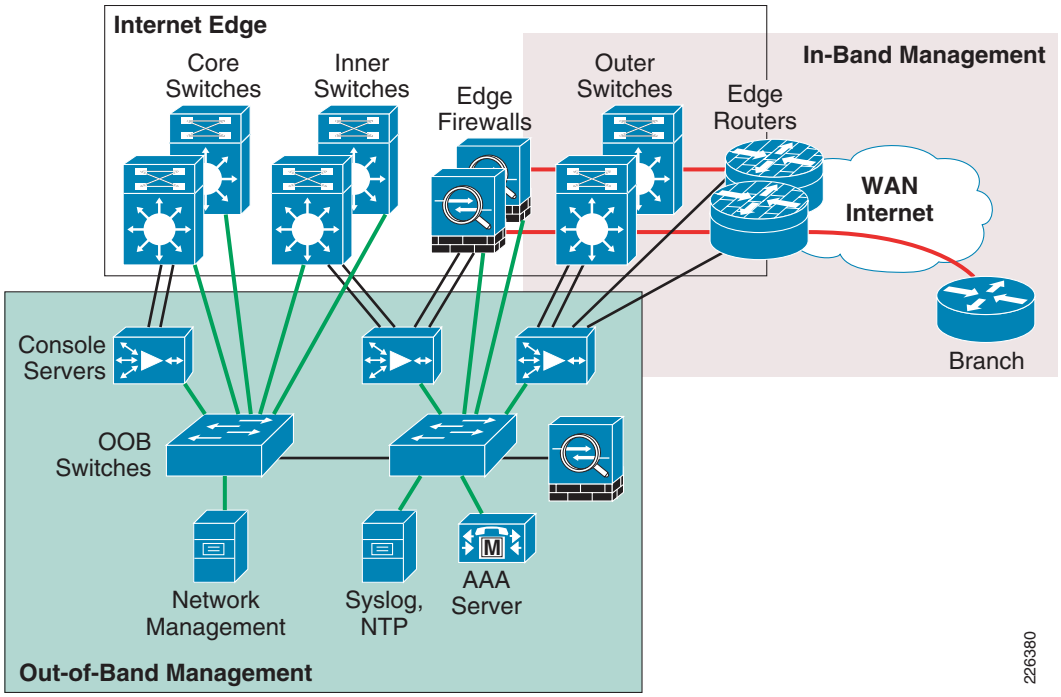

At the headquarters, where the Internet edge resides, an out-of-band (OOB) management network is implemented by using dedicated switches that are independent and physically disparate from the data network. Routers, switches, and other network devices connect to the OOB network through dedicated management interfaces. The OOB network hosts console servers, network management stations, AAA servers, analysis and correlation tools, NTP, FTP, syslog servers, and any other management and control services. This OOB management network may serve the other places in the network at the headquarters.

Any device outside the edge firewalls are managed in-band, using the same physical and logical infrastructure as the data traffic. Despite being deployed at the headquarters, the outer switches and edge routers are located outside the edge firewall, therefore they are managed in-band. The edge firewalls are responsible of securing the OOB network by permitting control and management connections only from the expected devices. Connecting the outer switches or the edge routers directly to the OOB network is highly discouraged, as it would facilitate the bypass of the firewall protection. Devices residing at the branches are also to be managed in-band, but over a secure VPN connection.

Baseline Security

Effective network security demands the implementation of various security measures in a layered approach and guided under a common strategy. The Internet edge is designed with security in mind, where multiple security technologies and capabilities are strategically deployed throughout the network to complement each other and to collaborate. Under a common strategy, security measures are positioned to provide maximum visibility and control.

This section describes the best practices for securing the infrastructure itself, the control and management planes, setting a strong foundation on which more advanced methods, and techniques can subsequently be built on. Later in this document, each use case will be presented with the additional security design elements required to enhance visibility and control and to secure the data plane.

The following are the key areas of baseline security:

- Infrastructure device access
- Routing infrastructure
- Device resiliency and survivability
- Network telemetry
- Network policy enforcement
- Switching infrastructure

In order to ensure a comprehensive solution, the selection of technologies and capabilities follows the Cisco Security Framework (CSF). CSF provides a method of assessing and validating the security requirements of a system, guiding the selection of security measures to be considered for each particular contextual area.

Infrastructure Device Access

Securing the network infrastructure requires securing the management access to these infrastructure devices. 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.

Network infrastructure devices often provide a range of different access mechanisms, including console and asynchronous connections, as well as remote access based on protocols such as Telnet, rlogin, HTTP, and SSH. Some mechanisms are typically enabled by default with minimal security associated with them; for example, Cisco IOS software-based platforms are shipped with console and modem access enabled by default. For this reason, each infrastructure device should be carefully reviewed and configured to ensure only supported access mechanisms are enabled and that they are properly secured.

The key steps to securing both interactive and management access to an infrastructure device are as follows:

- Restrict device accessibility—Limit the accessible ports, restrict the permitted communicators, and the permitted methods of access.
- 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 man-in-the-middle (MITM) attacks.
- Log and account for all access—Record who accessed the device, what occurred, and when for auditing purposes.

Table 1 CSF Assessment —Infrastructure Device Access

Total Visibility		
Identify	**Monitor**	**Correlate**
• AAA enforcement – Centralized AAA and local fallback – Administrator access – Privileged level access • SNMP accounts (community strings or auth/privacy policy) • Device management best common practices – Strong password policy – Per-user accounts • Remove default accounts and passwords	• Logging – Syslog – SNMP – AAA server-based accounting • Configuration change notification and logging	
Total Control		
Harden	**Isolate**	**Enforce**
• SNMP • SSH/Telnet • HTTP/HTTPS • Restrict device accessibility – Transport types – VTY ACLs – SNMP ACLs – IOS login enhancements	• Console • Dedicated management interface • Management network • SSH/Telnet • HTTP/HTTPS • ACLs • Out-of-band (OOB) management	• Banners (MOTD, EULA) • Local password protection – Password encryption – Secrets • File transfer and verification – FTP – TFTP – SCP – IOS image verification • Session management • Device management best common practices • Minimum access privileges

Routing Infrastructure

Routing is one of the most important parts of the infrastructure that keeps a network running, and as such, it is absolutely critical to take the necessary measures to secure it. There are different ways routing can be compromised, from the injection of illegitimate updates to DoS specially designed to disrupt routing. Attacks may target the router devices, the peering sessions, and/or the routing information.

The Internet edge design uses the following measures to effectively secure the routing plane:

- Restrict routing protocol membership—Limit routing sessions to trusted peers, validate origin, and integrity of routing updates.
- Control route propagation—Enforce route filters to ensure only valid routing information is propagated. Control routing information exchange between routing peers and between redistributing processes.
- Log status changes—Log the status changes of adjacency or neighbor sessions.

Table 2 CSF Assessment —Routing Infrastructure

Total Visibility		
Identify	**Monitor**	**Correlate**
• Neighbor authentication • Routing peer definition • Route redistribution filtering	• Neighbor adjacency logging • Logging (syslog, SNMP)	
Total Control		
Harden	**Isolate**	**Enforce**
• BGP TTL security check • Standby interfaces • Standby devices • Element redundancy • Topological redundancy		• Prefix filtering • Maximum prefix filtering • Route redistribution filtering • Stub routing • iACL • Control plane policing

Device Resiliency and Survivability

Routers and switches may be subject to attacks designed to or that indirectly affect the network availability. Possible attacks include DoS based on unauthorized and authorized protocols, DDoS, flood attacks, reconnaissance, unauthorized access, and more.

The Internet edge design uses the following best practices to ensure the resiliency and survivability of routers and switches:

- Disable unnecessary services—Disable default enabled services that are not required.
- Restrict access to the Infrastructure address space—Deploy ACLs at the network edges to shield the infrastructure from unauthorized access, DoS, and other network attacks.
- Protect control plane—Filter and rate limit traffic destined to the control plane of routers and switches.
- Control switch Content Addressable Memory (CAM) usage—Restrict the MAC addresses that are allowed to send traffic on a particular port.
- Implement redundancy—Eliminate single points of failure using redundant interfaces, standby devices and topological redundancy.

Table 3 CSF Assessment—Device Resiliency and Survivability

Total Visibility		
Identify	**Monitor**	**Correlate**
• Port Security	• Logging (Syslog, SNMP)	
Total Control		
Harden	**Isolate**	**Enforce**
• Control plane policing • Standby interfaces • Standby devices (stateful failover, HSRP/GLBP) • Topological redundancy		• Disable unnecessary services • iACL

Network Telemetry

In order to operate and ensure availability of a network, it is critical to have visibility and awareness into what is occurring on the network at any one time. Network telemetry offers extensive and useful detection capabilities that can be coupled with dedicated analysis systems to collect, trend, and correlate observed activity.

This section highlights the baseline forms of telemetry recommended for network infrastructure devices:

- Time synchronization—Implement Network Time Protocol (NTP) to ensure dates and times in logs and alarms are synchronized.
- Maintain local device traffic statistics—Leverage device global and interface traffic statistics.
- Maintain system status information—Leverage memory, CPU and process status information.
- System logging—Log and collect system status, traffic statistics and device access information.
- Log and account for all access—Record who accessed the device, what occurred, and when for auditing purposes.
- Packet capture—Establish the mechanisms to allow the capture of packets in transit for analysis and correlation purposes.
- URL monitoring—Establish mechanisms to allow for statistical analysis of web usage and web policy enforcement.

Table 4 CSF Assessment—Network Telemetry

Total Visibility		
Identify	**Monitor**	**Correlate**
• CDP • SNMP • Syslog	• NTP • Local device statistics • System status information – Memory/CPU/processes – CPU and memory threshold notification • CDP best common practices logging (syslog, SNMP, accounting, configuration change notification and logging) • Packet capture (SPAN/RSPAN, copy/capture VACLs)	

Network Policy Enforcement

Baseline network policy enforcement is primarily concerned with ensuring that traffic entering a network conforms to the network policy, including the IP address range and traffic types. Anomalous packets should be discarded as close to the edge of the network as possible, thereby minimizing the risk of exposure.

The Internet edge design implements the following measures:

- Access edge filtering—Control traffic destined to the infrastructure space.
- IP anti-spoofing—Implement packet filters and other dynamic mechanisms to block packets with spoofed IP addresses.
- Web URL monitoring and filtering—Web monitoring and filtering can be achieved by using web security appliances. One can monitor, authenticate and block users web access to previously defined websites

Table 5 CSF Assessment – Network Policy Enforcement

Total Visibility		
Identify	**Monitor**	**Correlate**
	• Logging (syslog, SNMP)	
Total Control		
Harden	**Isolate**	**Enforce**
	Access control lists (ACLs)	• Access edge filtering – iACLs • IP spoofing protection – uRPF – IP source guard – DHCP/ARP enforcement • URL filtering – Reputation analysis – Web URL monitoring – Authentication

Switching Infrastructure

Baseline switching security is concerned with ensuring the availability of the Layer 2 switching network. To that end the Internet edge design implements:

- Broadcast domain restriction—Design the Layer 2 infrastructure limiting the size of the broadcast domains.
- Spanning Tree Protocol (STP) Security—Leverage existing features to secure STP.
- VLAN best common practices

Table 6 CSF Assessment—Switching Infrastructure

Total Visibility		
Identify	**Monitor**	**Correlate**
	• Logging (syslog, SNMP)	
Total Control		

Table 6 CSF Assessment—Switching Infrastructure (continued)

Harden	Isolate	Enforce
	• Restrict broadcast domains • VLAN • Layer 3 hierarchical design	• STP security – Disable dynamic trunking – PVST – BPDU guard – Root guard • VLAN best common practices – Disable unused ports and put them in an unused VLAN – Use a dedicated VLAN ID for all trunk ports – Avoid using default VLAN 1 – Disable auto-trunking on user facing ports – Explicitly configure trunking on infrastructure ports – Use all tagged mode for the native VLAN on trunks

Architecture

Having explained the common elements building the Internet edge infrastructure, next we describe the functions and detailed designs to satisfy each one of the use cases.

Public Services (DMZ)

The DMZ is the piece of the Internet edge infrastructure that hosts public services. As illustrated in Figure 11, the DMZ design consists of several network devices and functions, which are explained next.

Figure 11 DMZ Design

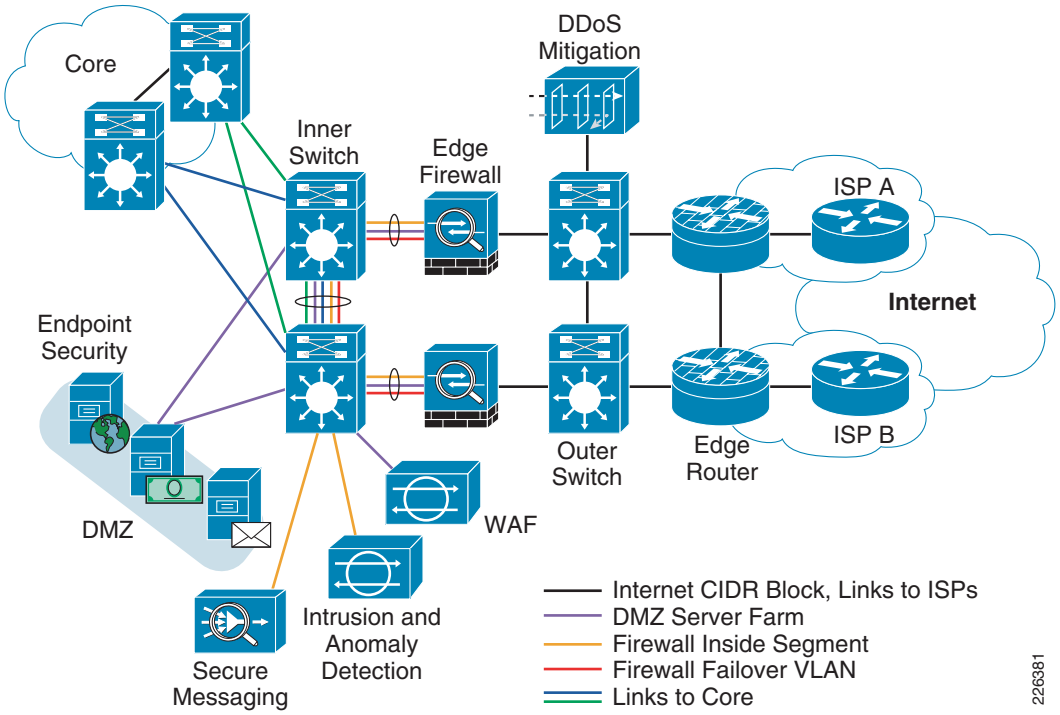

In this design, the edge routers are responsible for providing a reliable communication path to and from the Internet. The edge routers also serve as the first line of protection against spoofing and other network attacks. The edge routers also act as a collection point for network flow and event information useful for analysis and correlation purposes. Redundancy is achieved by deploying two routers and by configuring a First Hop Redundancy Protocol (FHRP) is configured in their inner interfaces.

The outer switches provide Layer 2 connectivity between the Layer 2 firewalls and the edge routers. Two switches are deployed for redundancy. Since there are no Layer 2 loops in the design, STP is not required. The outer switches provide connectivity to the DDoS mitigation component.

Public services are hosted on a DMZ, and depending on the services offered the DMZ may be implemented as several isolated Layer 2 segments (i.e., one hosting the website and another one for the mail services). The DMZ Layer 2 segments are implemented at the inner switches, which are configured to ensure the isolation of the various Layer 2 segments. Two switches are deployed for redundancy. Servers at the DMZ are protected with endpoint security software that works in conjunction with the intrusion protection system (IDS) and the monitoring and analysis system. Services and applications hosted at the DMZ are protected by a line of stateful firewalls, an intrusion and anomaly detection system, and a DDoS mitigation system.

The edge firewalls secure the DMZ by controlling and inspecting all traffic entering and leaving the DMZ segments. This includes traffic between DMZ, as well as traffic between the DMZ and the Internet and the internal network. For redundancy, firewalls are deployed in stateful active/standby failover mode. Complementary, a secure messaging system is deployed at the DMZ to inspect incoming and outgoing emails and eliminate threats such as e-mail spam, viruses and worms.

The DMZ implements an IDS with the intention to identify and alert on well-known attacks and suspicious activity. Alert and alarm generated by the IDS and the endpoint security software is processed by a monitoring and analysis system for analysis and correlation purposes. Complementary to the IDS, an anomaly detection system is also deployed at the DMZ. This system is responsible of identifying DDoS and other network-based attacks, and it works in conjunction of the DDoS mitigation system

deployed at the outer switches. Once an ongoing attack is confirmed on one of the public services (i.e., HTTP), the DDoS mitigation system triggers a traffic redirection to put itself on the traffic path for the affected service, and starts inspecting traffic ensuring only good traffic reaches the server.

Table 7 illustrates how all these components play together as part of the same security strategy.

Table 7 CSF Assessment—Public Services DMZ

Total Visibility		
Identify	**Monitor**	**Correlate**
• Firewall deep packet inspection • Secure messaging	• Intrusion detection system (IDS) • Anomaly detection system (ADS) • Network management • Network flow data collection • Packet capture • Endpoint monitoring • Event monitoring	Event analysis and correlation
Total Control		
Harden	**Isolate**	**Enforce**
• Baseline security • Stateful firewall • Link and system redundancy	• VLANs • Firewall access control policies	• Firewall access control • Endpoint security • Secure messaging

Corporate Internet Access

The Internet edge infrastructure provides users at the headquarter or regional offices access to the Internet. Optionally, the same infrastructure may serve users at the branches that are mandated to access the Internet over a centralized connection. The components and functions present in the architecture are depicted in Figure 12.

Figure 12 Internet Access Design

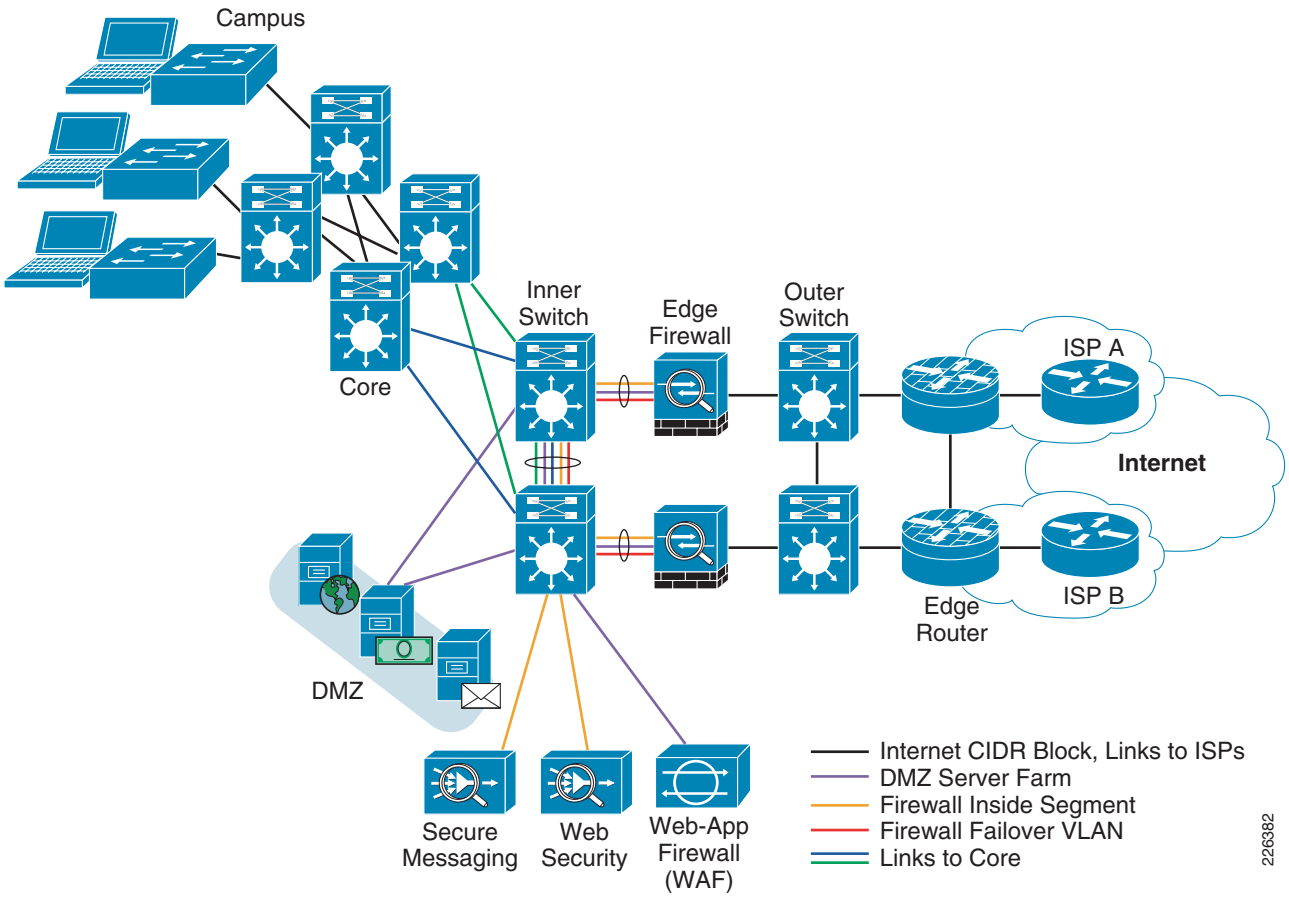

The edge routers are responsible for providing a reliable communication path to and from the Internet. To that end, redundancy is achieved by deploying two routers and by configuring an FHRP in their inner interfaces. The outer switches provide Layer 2 connectivity and are deployed in pair for redundancy.

The edge firewalls protect the interior of the network from the Internet, and are also responsible of enforcing the Internet access policy for internal users. To that end, firewalls enforce access policies, keep track of connection status, and inspect packet payloads. As part of this, the firewalls may be configured to enforce policies destined to limit or block instant messaging and peer-to-peer applications to mitigate network abuse. The edge firewalls also control traffic destined to the DMZ. In addition, the firewalls are in charge of performing port address translation (PAT) for the traffic bound to the Internet.

An intrusion detection and prevention system (IDS/IPS) is deployed at the level of the inner switches to inspect inside traffic and to alert on known attacks or malicious activity. The IDS/IPS may be configured to alert on activities considered as network abuse or to block traffic that is deemed to be malicious. Alarm information generated by the IDS/IPS is forwarded to a monitoring and analysis system for analysis and correlation purposes.

A web security system is deployed at the level of the inner switches to inspect web traffic bound to the Internet. This web security appliance is in charge of blocking spyware, malware, and other known threats, to provide content filtering and optionally to authenticate user requests.

E-mail communications are inspected by the secure messaging system deployed at the DMZ hosting the mail server. This email security appliance is responsible for analyzing email payloads and eliminating threats such as e-mail spam, viruses, and worms.

The web security firewall (WAF) is responsible for blocking Internet-based web application attacks such as XML attacks at Layer 7. These type of attacks take advantage of Layer-7 vulnerability in the web servers that may go undetected using traditional firewalls. WAFs are especially recommended for corporations implementing e-commerce services.

Endpoints residing at the campuses and other places in the network may be protected with endpoint security software that works in conjunction with the IDS and with the monitoring and analysis system. This collaboration allows for a better calculation of the risk level associated with an event, and the dynamic enforcement of watch lists for systems believed to have been compromised.

Table 8 illustrates how all these components play together as part of the same security strategy.

Table 8 CSF assessment – Corporate Internet Access

Total Visibility		
Identify	**Monitor**	**Correlate**
• Firewall deep packet inspection • Web security • Content filtering • Secure messaging	• Intrusion Detection System • Network management • Network flow data collection • Packet capture • Endpoint monitoring • Event monitoring	Event analysis and correlation
Total Control		
Harden	**Isolate**	**Enforce**
• Baseline security • Stateful firewall • Link and system redundancy	• VLANs • Firewall access control policies	• Firewall access control • Intrusion prevention • Endpoint security • Content filtering • Secure messaging

Teleworker

One of the objectives of the Internet edge architecture is to provide mobile users and teleworkers with secure access to applications and data residing at the corporate network. Other services may include IP telephony. Figure 13 shows the infrastructure supporting this function.

Figure 13 Teleworker Access Design

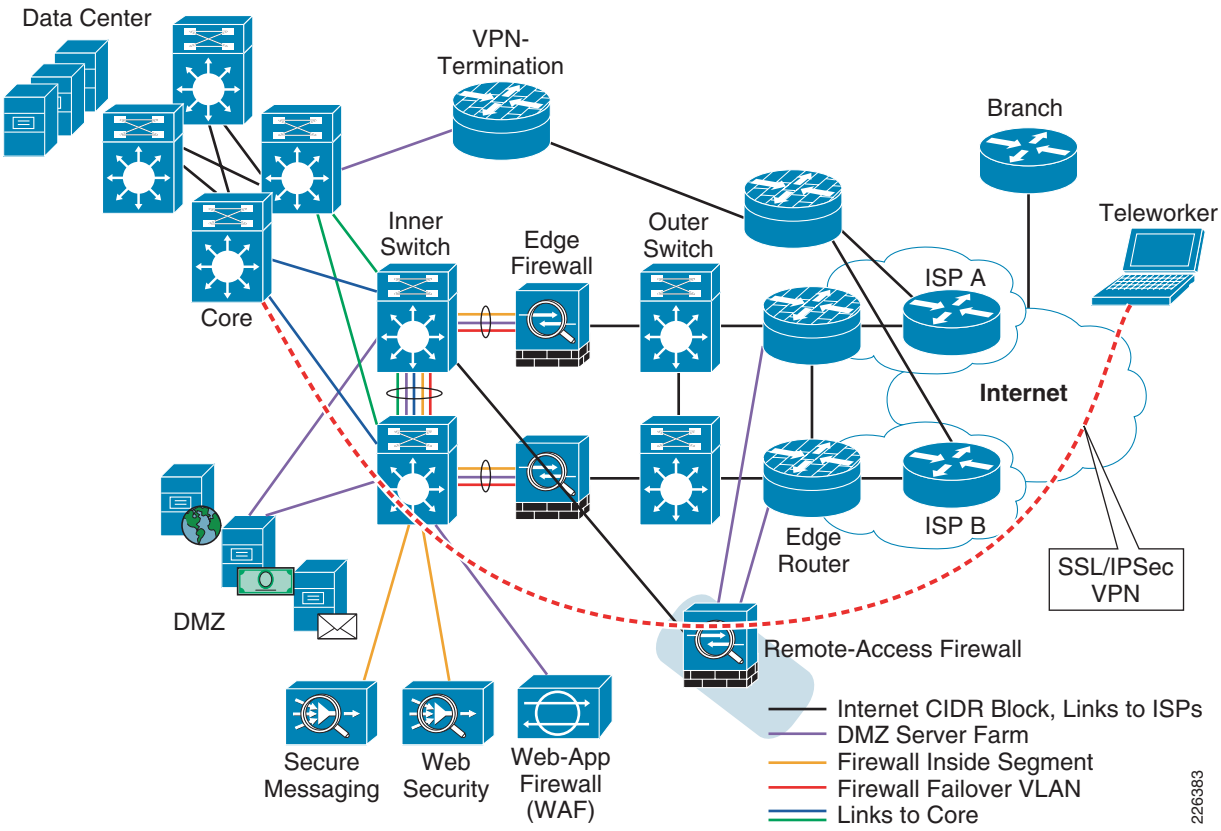

Teleworker and mobile user access is authenticated and encrypted with either SSL or IPSec. These VPN tunnels are terminated on the edge firewalls. The firewalls are not only responsible of authenticating users and terminating the VPN sessions, but also of enforcing per-user or per-group access policies that restrict access to only the necessary resources.

Complementary, an IDS deployed at the firewall inside segment inspects traffic coming and going to the remote users. All alarm information generated by the IDS is forwarded to a monitoring and analysis system for analysis and correlation purposes.

In case remote users use the central e-mail service, the secure messaging system deployed close to the mail server inspects all e-mail communications, analyzing email payloads and eliminating threats such as e-mail spam, viruses, and worms.

In case the organization's policy forces all Internet access throughout a central location, web communications from remote users can also be secured by the web security system is deployed at inner switches level. In addition, remote users may be protected with endpoint security software that works in conjunction with the IDS and with the monitoring and analysis system. This collaboration allows for a better calculation of the risk level associated with an event, and the dynamic enforcement of watch lists for systems believed to have been compromised.

Table 9 illustrates how all these components play together as part of the same security strategy.

Table 9 CSF assessment – Teleworker Access

Total Visibility		
Identify	**Monitor**	**Correlate**
• Firewall deep packet inspection • VPN authentication • Web security* • Content filtering* • Secure messaging*	• Intrusion detection system (IDS) • Network management • Endpoint security posture • Event monitoring	Event analysis and correlation
Total Control		
Harden	**Isolate**	**Enforce**
• Baseline security • Stateful firewall • VPN redundancy • Link and system redundancy	• Per user/group firewall policies • VPN	• Firewall access control • Intrusion prevention • Endpoint security • Content filtering* • Secure messaging*

* In case access to the Internet is only allowed throughout a central location.

WAN Backup

The Internet edge design includes the option of using the Internet as a backup for the WAN links connecting the remote offices (see Figure 14). This requires an Internet connection at the branch offices.

Under this configuration, when the WAN links to a branch fail, traffic is automatically redirected over the Internet and over an authenticated and encrypted VPN connection. The VPN connection can either be permanently established or triggered on demand after a failure is detected. In this design, a dynamic routing protocol is used to identify failures and redirect traffic over the VPN tunnels. The dynamic routing protocol is transported over both the primary WAN links and the VPN tunnels.

The VPN tunnels are terminated at a pair of dedicated routers residing at the Internet edge of the head quarter or regional office. For configuration simplicity, Dynamic Multicast VPN (DMVPN) is used in this design.

For enhanced security and control, the VPN routers are placed behind the edge firewalls and distinct interfaces for encrypted traffic and traffic in clear. The interfaces dedicated to handling the encrypted traffic connect to the firewalls, which shield the VPN routers from a variety of threats by only accepting VPN packets initiated from the known IP addresses assigned to the branches.

The interfaces dedicated to handling the traffic in clear (after decryption) connect directly to the inside segment of the firewall. This assumes that traffic coming inside the VPN tunnels is to be trusted, and therefore no additional controls are required.

In addition, traffic coming from the backup VPN connections can be inspected by the IDS deployed at the firewall inside. All alarm information generated by the IDS is forwarded to a monitoring and analysis system for analysis and correlation purposes.

In case the organization mandates all Internet access to be centralized, the secure messaging system and web security system deployed at the Internet edge may also be used to secure e-mail and web communications.

Finally, remote users may be protected with endpoint security software that works in conjunction with the IDS and with the monitoring and analysis system. This collaboration allows for a better calculation of the risk level associated with an event, and the dynamic enforcement of watch lists for systems believed to have been compromised.

Figure 14 Internet WAN Backup Design

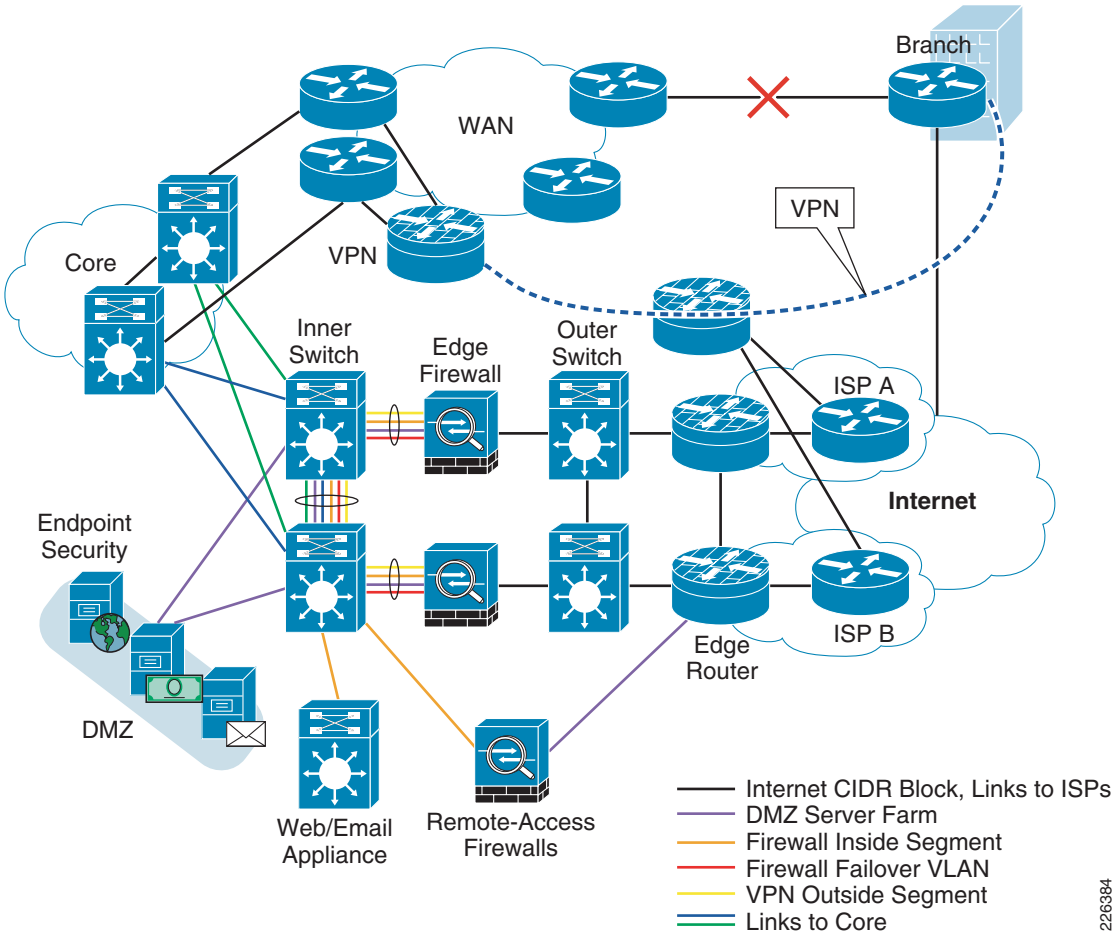

Table 10 illustrates how all these components play together as part of the same security strategy.

Table 10 CSF Assessment—Internet WAN Backup

Total Visibility		
Identify	**Monitor**	**Correlate**
• VPN authentication • Firewall deep packet inspection* • Web security* • Content filtering* • Secure messaging*	• Intrusion detection system (IDS) • Network management • Endpoint security posture • Event monitoring	Event analysis and correlation
Total Control		
Harden	**Isolate**	**Enforce**
• Baseline security • Stateful firewall • VPN redundancy • Link and system redundancy	• Per user/group firewall policies • VPN • VLANs	• Firewall access control* • Intrusion prevention • Endpoint security • Content filtering* • Secure messaging*

* In case access to the Internet is only allowed throughout a central location.

References

- Network Security Baseline

  http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/Baseline_Security/securebasebook.html
- Data Center Design Guides

  http://www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html
- WAN/Branch Design Guides

  http://www.cisco.com/en/US/netsol/ns816/networking_solutions_program_home.html