Enterprise Data Center Topology

This chapter provides a detailed description on how to harden and modify enterprise data center topologies for data center security. It includes the following sections:

Enterprise Data Center Topology Overview

A typical large enterprise network often consists of multiple data centers, each with a responsibility for supporting key functions. In general, these data centers can be classified into three types:

- Internet
- Extranet
- Intranet

The Internet data center, which is used by external clients connecting from the Internet, supports the servers and devices required for business-to-consumer transaction-based web applications (e-commerce).

The extranet data center provides support and services for external, business-to-business (B2B) partner transactions. These services are often accessed over secure VPN connections or private WAN links between the partner network and the enterprise extranet.

The intranet data center houses applications and services accessed by clients with connectivity to the internal enterprise network. The applications and services housed in the intranet data center often support functions for manufacturing, marketing, HR, research and development, payroll, and other core business services.

Figure 2-1 shows a common design for enterprise data centers. As illustrated, business transactions from the service providers (ISP1 and ISP2) enter the intranet server farm through a set of firewalls. These transactions might require load balancing to the DMZ servers to the presentation tier of the business-to-consumer (B2C) applications. The DMZ servers also include DNS servers and SMTP servers and they can equally benefit from the network load balancing.

Figure 2-1 Enterprise Data Center Network with Internet and Intranet Server Farms

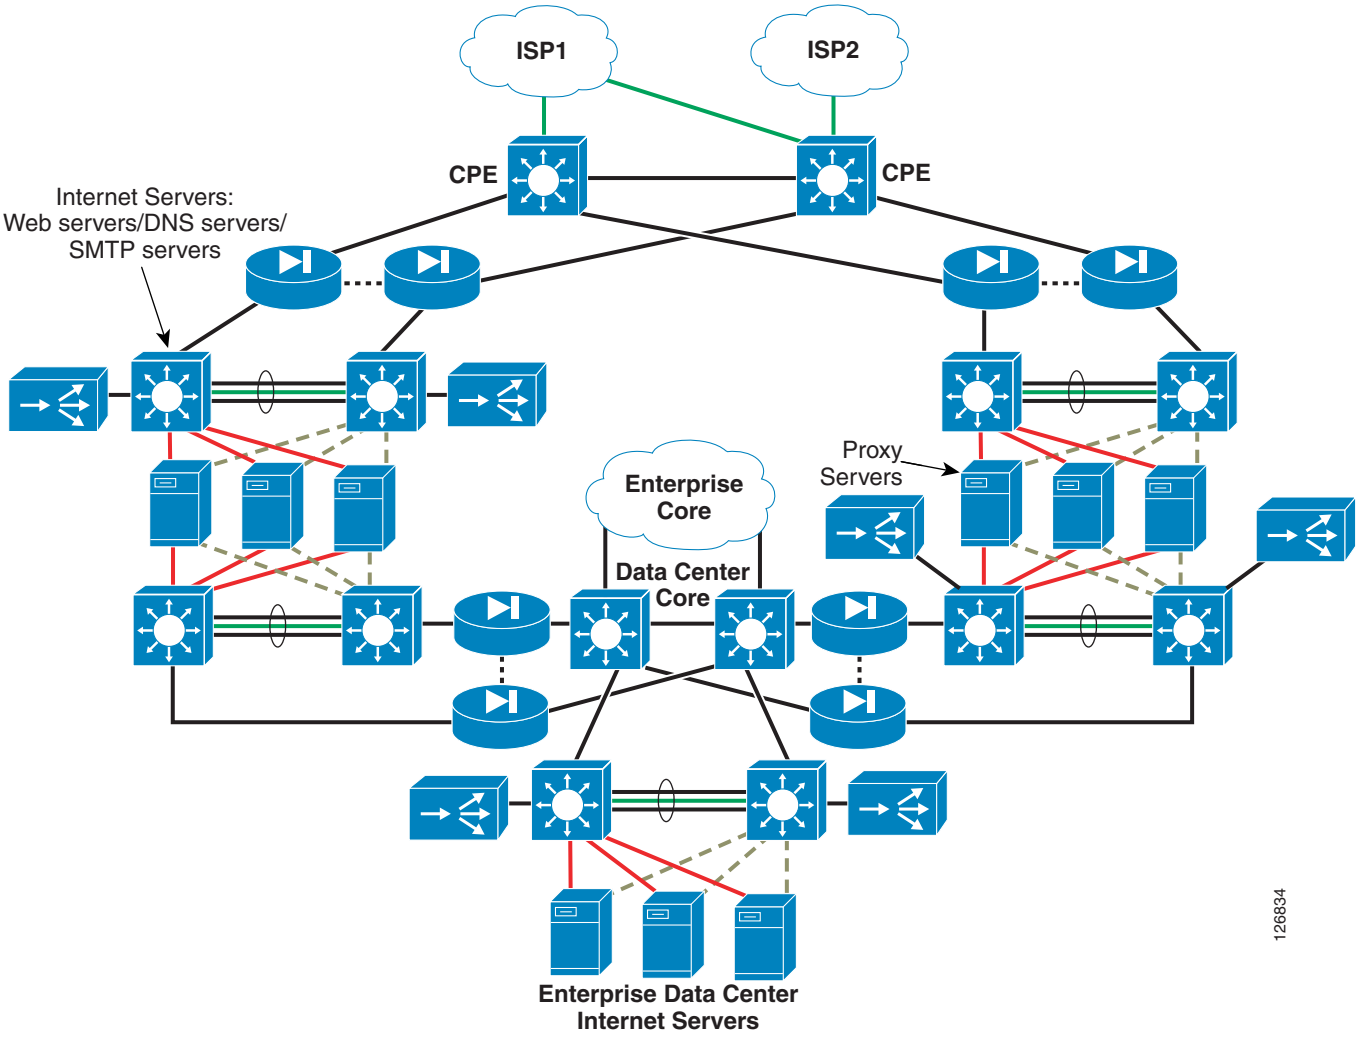

The B2C servers can be dual-homed using two NICs, with the public NIC used for transaction exchange and the private NIC used to communicate with the application and/or the database servers. Figure 2-1 does not illustrate the application and database servers. The figure shows only that the back-end NIC gives the intranet servers connectivity to the data center core through a pair of firewalls.

Figure 2-1 shows the proxy servers, which provide campus network users with connectivity to the Internet. In the illustration, the intranet data center connects to the data center core through redundant Layer 3 links. The data center core simplifies connectivity among the various data center environments such as B2C, business-to-business (B2B), intranet server farms, and so on.

Some data center implementations completely isolate the Internet servers from the rest of the network at the physical level. This means that a separate set of non-routable links connect these servers directly to the intranet data center with no physical path available to any other part of the network.

Network Design for Multi-tier Applications

This section analyzes the network design of the Internet and/or intranet server farm and provides some additional details. The same model can be used for the B2B server farm. This section includes the following topics:

- Network Design for B2B and B2X Server Farms
- Using Firewalls, Cisco IOS ACLs, and VACLs
- Virtual Firewalls
- Preventing VLAN Hopping

Network Design for B2B and B2X Server Farms

Server farms are often built with physical separation between application tiers, as shown in Figure 2-2.

Figure 2-2 Typical B2C Architecture with Physical Separation Between Application Tiers

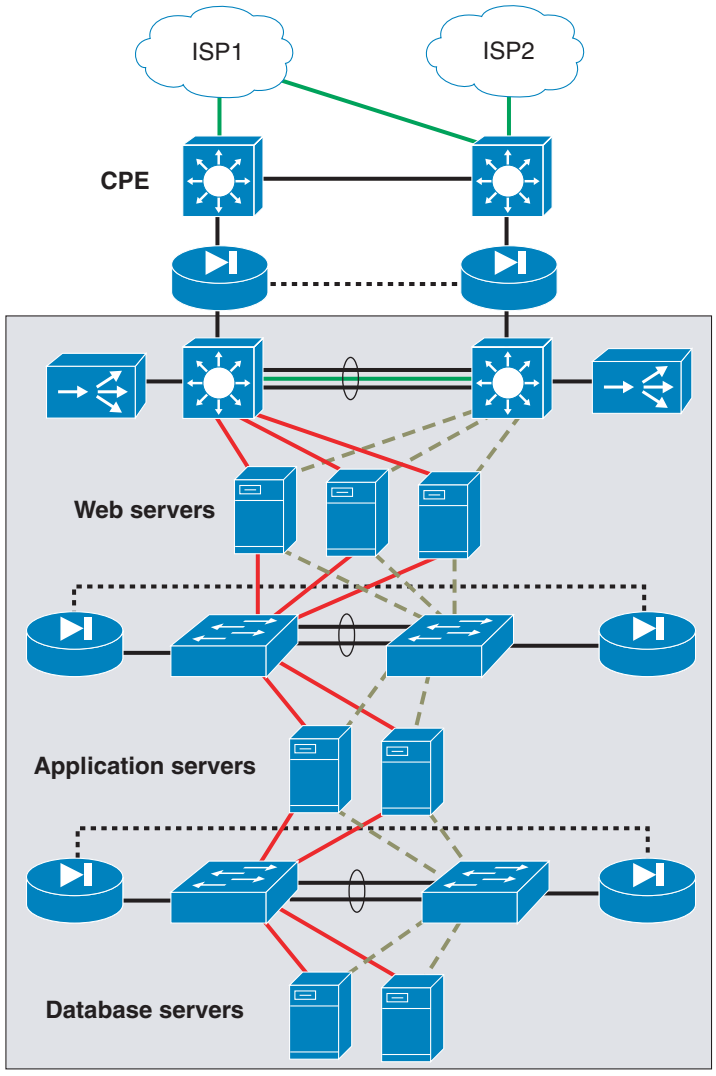

In this example, the B2C environment consists of a first tier of web servers, each of which has at least two NIC cards: a public interface and a private interface. The public interface may use either a public IP address or a private address with a firewall or load balancer providing Network Address Translation (NAT). The private interface uses a private address and gives access to the application servers through a pair of firewalls. The application servers, in turn, have at least two NICs: one for communication with the web servers and one for communication with the database servers.

Note B2X generically refers to the e-commerce, business-to-business, and intranet server farms.

The current trend for consolidated data centers is to simplify the network infrastructure by reducing the number of network devices (see Figure 2-3).

Figure 2-3 Consolidated B2C Architecture Topologies

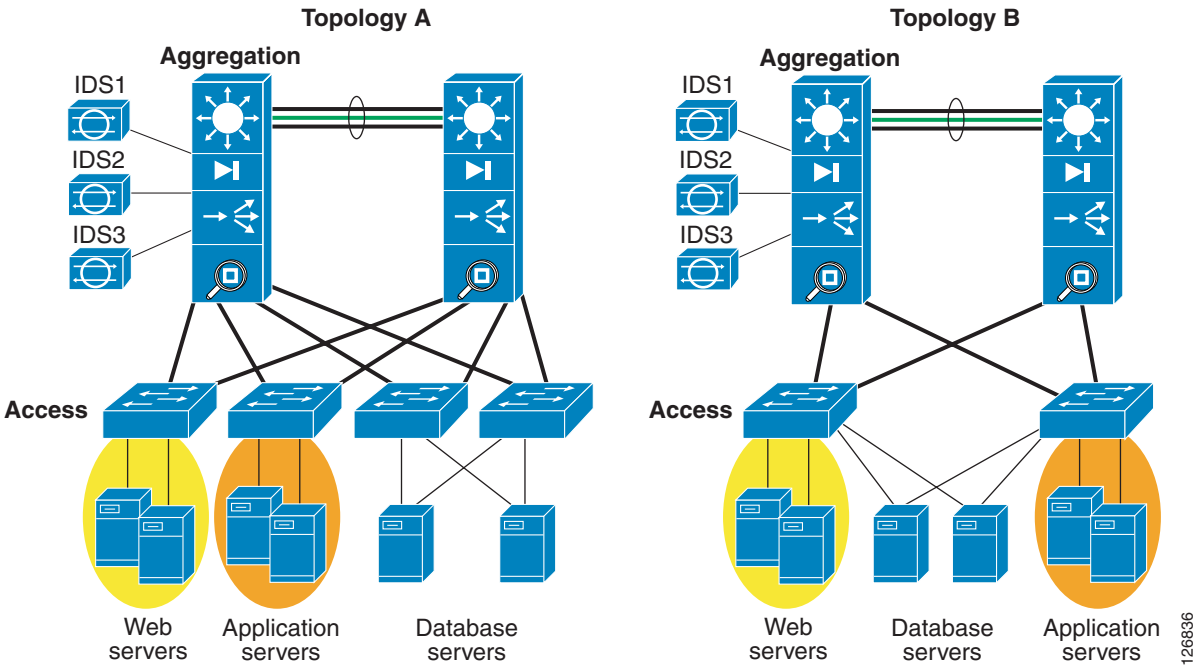

In Topology A, each server of a different type is connected to a physically separate access switch. Web servers are connected to one switch, application servers are connected to a different switch, and database servers are connected to a pair of access switches for increased availability. The traffic from these access switches is aggregated by a pair of Cisco Catalyst 6500 switches with service modules. Segmentation between these servers is ensured by the use of VLANs and/or virtual firewall contexts.

Topology B shows a more consolidated infrastructure in which web, database, and application servers connect to a single pair of access switches. At the access layer, VLANs provide segmentation between these servers. At the aggregation layer, segmentation is provided by VLANs and virtual firewall contexts.

The aggregation layer in both Topology A and Topology B provides the following functions:

- Firewalling
- Load balancing
- Network analysis

- SSL offloading services
- Intrusion detection

These services can be either integrated into a single aggregation chassis or some services can be offloaded to a separate layer of switches, referred to as service switches (see Figure 2-4). Each hardware accelerated service takes one slot in the chassis. Service switches are useful for consolidating multiple security and load balancing services to provide high density of ports for the servers.

To optimize security, only offload load balancing and SSL offloading to the service switches. Keep the firewall and network analysis modules in the aggregation layer switch.

Note This design guide does not include the data center design that uses service switches.

Figure 2-4 B2C Topology with Service Switches

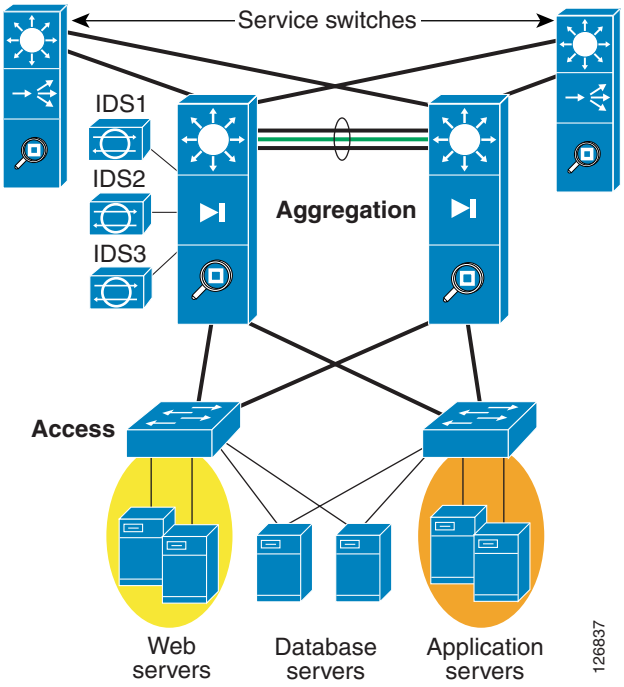

Segmentation by means of VLANs combined with access control lists (ACLs) and/or firewalls within a consolidated infrastructure allows servers belonging to different organizations to be kept logically separate for security reasons, while remaining physically connected to the same device.

Using Firewalls, Cisco IOS ACLs, and VACLs

Deploying ACLs in the data center is most beneficial for limiting access to and from devices (for example, subnet segregation) through basic Layer 3 and Layer 4 packet filtering. ACLs can be set to filter by the Layer 4 port, but they are not capable of providing upper-layer application protection. ACLs do not support stateful packet inspection, which is a key benefit of using a firewall. Stateful packet inspection allows a firewall to perform packet-by-packet inspection of connection-oriented requests, and to deny incomplete or malformed requests.

The Cisco Firewall Services Module (FWSM) is an integrated firewall for the Cisco Catalyst 6000 Series switches. The FWSM is configured like a Cisco PIX Firewall and therefore can be deployed to perform stateful packet inspection for both inbound and outbound traffic, as well as server-to-server communications. The FWSM module provides packet inspection throughput at 5 Gbps.

In a multi-tier architecture, filtering is recommended and should be performed in front of the presentation tier, between the presentation and application tiers, and between the application and database tiers. Packet filtering may also be performed between servers residing in the same tier. The packet filtering recommendations are dependent on the type of architecture deployed. For the physical multi-tier server farm, Cisco recommends that you filter at each layer because this provides optimum security.

With a traditional appliance-based firewall, filtering at each layer requires a minimum of two firewalls at each tier. This in turn adds to the complexity of physical connections, management, and high availability. Figure 2-2 shows a typical multi-tier server farm architecture with appliance-based firewalls deployed at each tier.

You can configure separate VLANs on the FWSM for each layer with routing and packet filtering performed between each tier. This allows all traffic between VLANs to pass through the FWSM, therefore centralizing packet filtering services on a single physical pair of firewalls. A single FWSM pair can be “virtualized” into multiple logical firewalls. This virtualization allows you to create separate logical firewalls per tier, and, if desirable, per customer.

Whether it is better to use ACL packet filtering or firewalling between server farm tiers depends on the application. Firewalls have been optimized for transactional protocols such as HTTP, DNS, SMTP, and SQL as well as typical DMZ applications. As a result, the presentation tier for typical web-based transactional applications benefits most from a firewall.

Server-to-server protocols that negotiate dynamic ports and keep connection idle for a long time typically are handled better with pure packet filtering. These protocols include IIOP, RMI, and DCOM.

Virtual Firewalls

Security between server farms of different types and residing in separate VLANs can be provided by partitioning a single FWSM into multiple virtual firewalls, known as security contexts. Each context is an independent system with its own security policy, interfaces, and administrators. Multiple contexts are equivalent to having multiple standalone firewalls. Each context has its own configuration that identifies the security policy, interfaces, and almost all the options that are configurable on a standalone firewall. If desired, individual context administrators can implement the security policy for each context. To prevent one context from inadvertently affecting other contexts, some resources, such as VLANs and system resources, are controlled by the overall system administrator.

Figure 2-5 shows the resulting topology in a consolidated server farm where each firewall context protects the application tiers. VLAN segmentation enforces traffic from the web to the application tier through the firewall context protecting the application tier.

Figure 2-5 Data Center Topology with Virtual Firewalls

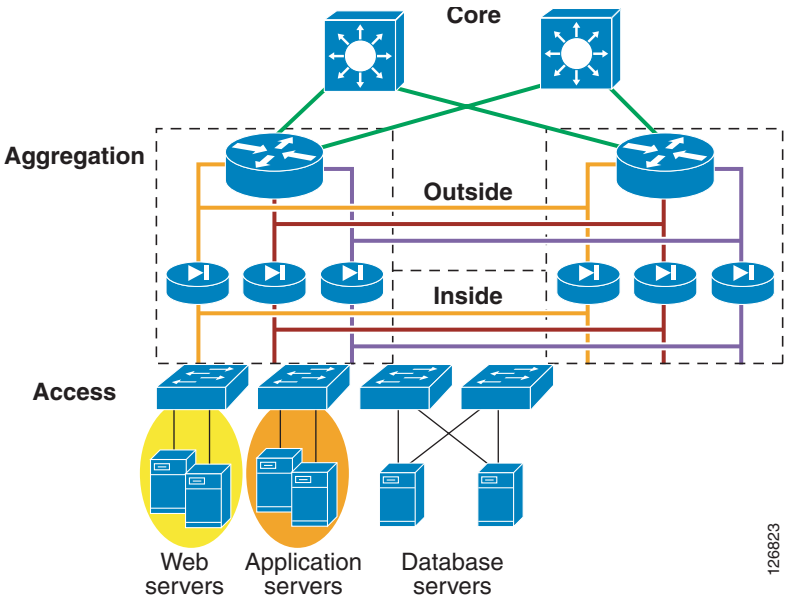

Several variations in this design are possible. For example, servers often have two NIC cards: one for the public-facing network and one for the web-to-application communication. In this case, the NIC might be placed on the same subnet on the outside VLAN of the next-tier firewall. Better yet, it can be placed in its own subnet, routed only to the application tier subnet, and without any public access.

The same concepts can be used to provide security for applications that belong to different departments within the same organization.

Preventing VLAN Hopping

The data center access switches are typically connected to the aggregation switches through 802.1q trunk ports. By default, a trunk carries all VLANs when it is first configured. When using a Layer 2 access, each access switch that supports more than one server VLAN must have a trunk connection to the data center aggregation switch, as shown in Figure 2-6.

Figure 2-6 802.1q Trunking in the Data Center

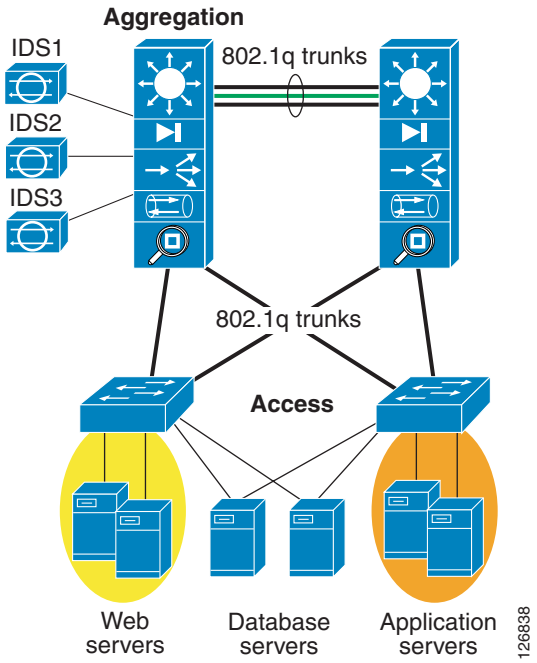

By default, all trunks carry VLAN 1 and all ports reside in VLAN 1. Cisco Discovery Protocol (CDP) and VLAN Trunking Protocol (VTP) control messages are also carried on VLAN 1, by default. Even if VLAN 1 is cleared from a trunk interface, the control messages are still sent over VLAN 1 even though no data traffic is forwarded on VLAN 1.

It is theoretically possible for an attacker to craft frames with two VLAN 802.1q tags. In Figure 2-6, this would make it possible for a compromised web server to gain access to the application servers, bypassing the firewall protection, because the native VLAN on the access switch is set to VLAN 1. The attacker simply double encapsulates the packet with two VLAN tags. The first tag is VLAN 1, and the second tag is the target VLAN where the application servers reside (for example, VLAN 10 in Figure 2-6).

When the switch receives the double-encapsulated packet from the attacker, it strips off the first VLAN tag (native VLAN) and forwards the packet to VLAN 10. In this case, the port to which the attacker connected does not have to carry VLAN 10 for the attacker to reach VLAN 10. It is only necessary for the attacker to install software on the server for crafting a packet with two 802.1q tags.

However unlikely this attack is, it is important to make sure that the design of the data center eliminates any possibility. VLANs are used as the segmentation mechanism in a consolidated environment and firewalls operate on VLANs to apply security policies to traffic that goes from one server farm to the other.

Several steps can be taken to prevent VLAN hopping:

- First, clear the native VLAN from all trunk ports. The control protocols may still be carried over the native VLAN, but no data traffic will be transmitted over it.
- If the native VLAN cannot be cleared from the trunk port, pick an unused VLAN to use as the native VLAN and use it for nothing else.
- Tag all VLANs on the trunks including the native VLAN.

To change the native VLAN in Catalyst IOS, enter the following command:

```
access> (enable) set vlan 2 1/1
access> (enable) sh trunk 1/1
```

```
* - indicates vtp domain mismatch
Port      Mode         Encapsulation  Status        Native vlan
--------  -----------  -------------  ------------  -----------
 1/1        auto         negotiate     trunking          2
```

To change the native VLAN in Cisco IOS software, enter the following command:

```
agg(config-if)#switchport trunk native vlan 2
agg#sh int gig 2/8 trunk
Port      Mode         Encapsulation  Status        Native vlan
Gi2/8     desirable    negotiate      not-trunking  2
```

Also, disable DTP on server ports. If the port is left with DTP auto-configured, which is the default on many switches, an attacker can connect and arbitrarily cause the port to start trunking and consequently pass all VLAN information.

To disable DTP in CatOS, enter the following command:

```
Access> (enable) set trunk 3/47 off
Port(s)  3/47 trunk mode set to off.
```

To disable DTP in Cisco IOS software, enter the following command:

```
agg(config-if)#switchport mode access
```

To avoid this problem, you should choose to not use any access VLAN as the native VLAN of a trunk, or you can make sure you tag all the VLANs carried on a trunk by using the **vlan dot1q tag native** command.

A recent series of tests performed on the Catalyst product line by @Stake were specifically directed at testing the vulnerability of VLANs in these switches. The tests found that when VLAN security configuration guidelines were properly followed, they were not able to hop or bypass VLANs on these switches using a variety of well-known attacks.

Note To see the @Stake security document, see the following URL: http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/stake_wp.pdf

Network Design for DoS Protection

The objective of a denial of service (DoS) attack is to deny access for legitimate users to enterprise resources. It is an attack that floods the network with useless traffic, and that usually exploits limitations in the TCP/IP stack. Synchronization (SYN) floods, Ping-Of-Death, and Teardrop attacks are common DoS methods used by attackers. The Catalyst 6500, load balancers, and firewalls provide protection against DoS attacks. The data center design can also be optimized to provide maximum protection against DoS attacks.

This section describes methods used to prevent DoS attacks, the impact on performance of these methods, and recommended designs. It includes the following topics:

- TCP Intercept
- SYN Cookies
- Performance Considerations
- Design Models

TCP Intercept

The TCP Intercept feature protects downstream servers using TCP (FTP, HTTP, SMTP, and so on) from SYN flood DoS attacks. In a SYN flood attack, one or more machines controlled by an attacker bombard a server with a barrage of requests for connection. Because these request messages have unreachable return addresses, the connections cannot be established. The resulting volume of unresolved open connections eventually overwhelms the server and can cause it to deny service to valid requests. This prevents legitimate users from connecting to websites, accessing e-mail, using FTP service, and so on.

TCP Intercept on the Catalyst 6500

The TCP Intercept feature mitigates SYN floods by intercepting and validating TCP connection requests, and operates in two modes:

- Intercept mode
- Watch mode

In intercept mode, TCP Intercept software matches and intercepts TCP SYN packets from clients to servers using an extended access list. The software establishes a connection with the client on behalf of the server and, if successful, establishes a connection with the server on behalf of the client. Finally, the software knits the two connections together. The entire process is transparent to the client and server. In this way, connection attempts from unreachable hosts never reach the server. The software continues to intercept and forward packets throughout the duration of the connection. In intercept mode, TCP options that are normally negotiated during the handshake between client and server, such as RFC 1323 (window scaling), are not negotiated.

In watch mode, the TCP Intercept software is not an active participant in the connection process. Rather, it passively observes connection requests. Those that fail to fully establish a connection in a configurable time interval are terminated with a TCP Reset (RST) issued by the software to the server in question. The aggressive timeouts on half-open connections and the thresholds on TCP connection requests protect servers from illegitimate traffic while still allowing valid requests.

When establishing a security policy using TCP Intercept, all client/server traffic may be selected for interception, or traffic can be selected based on the source or destination network. Connection rate and threshold values for half-open connections are also configurable.

When using Supervisor 2 or Supervisor 720, TCP Intercept can be hardware assisted. However, when using intercept mode with timeout, all traffic belonging to the given connection is handled in the software, which may overwhelm the CPU. For other modes of TCP Intercept, after the connection is successfully established, the software installs a hardware shortcut to switch the rest of the flow in hardware.

TCP Intercept on the FWSM

Before Release 2.3, the FWSM implemented DoS protection against SYN flooding using TCP Intercept. A SYN flood consists of a large number of independent connection requests. The FWSM employing the TCP Intercept feature validates incoming connection requests and replies to the client SYN with an acknowledgement (SYN-ACK) on behalf of the destination device, which is usually a server. If the client responds with the appropriate acknowledgement (ACK), the FWSM establishes a connection with the destination device on behalf of the client and then weaves the two connections together. This process prevents illegitimate connection requests from consuming the limited resources of enterprise endpoints, which thwarts the DoS attack.

The FWSM TCP Intercept feature employs an embryonic limit, which is a threshold that defines the number of "incomplete" connections the FWSM permits before intercepting further connection requests (SYN packets). The definition of an incomplete connection is a client that has not responded to the SYN-ACK sent by the destination device protected by the FWSM. When the embryonic limit is surpassed, the FWSM begins intercepting incoming connection requests. The embryonic limit may be set using either the **static** or **nat** commands. In the examples described in this design guide, the embryonic limit is set using the **static** command.

SYN Cookies

In Release 2.3, FWSM uses SYN cookies. SYN cookies are an effective mechanism to protect a server farm from DoS attacks. By using this technology, the Cisco CSM and FWSM can withstand attacks of hundreds of thousands of connections per second while preserving legitimate user connections.

The SYN cookie mechanism protects the SYN queue of the TCP/IP stack of a device (either a network device or a server) by selecting an ISN (the cookie value) based on a Message Digest 5 (MD5) hash of the source and destination IP addresses and port numbers with a rotating secret key. When a certain threshold in the queue is reached, a SYN/ACK is still sent but connection state information is no longer kept. The connection request information sent by the host can be reconstructed from the cookie. If the final ACK for the three-way handshake is received, the server recalculates the original information that had come with the initial SYN.

Note When using SYN cookies, TCP options such as large windows and selective acknowledgement cannot be supported.

SYN Cookies on the CSM

Starting from Release 3.2, the CSM implements the SYN cookie technology to mitigate the impact of a SYN flood attack on the enterprise and its endpoints. Using SYN cookies on the CSM requires setting a threshold of embryonic or half-open connections. If this threshold is exceeded, a 32-bit number (cookie) is created by the CSM from a cryptographic function performed on the received SYN information. The CSM creates the SYN cookie using the following SYN information:

- Maximum segment size (MSS)
- Source IP address and port
- Destination IP address and port
- Secret key local to the CSM

The cookie is sent back to the host in the SYN-ACK as the initial sequence number (ISN). CSM resources are not required to maintain the connection request information sent by the host because this information exists in the cookie. If the host responds with an ACK, the cookie is available to the CSM in the acknowledgement number field. The CSM reconstructs the original SYN information from the cookie (acknowledgement number field -1) by reversing the hash operation. The CSM initiates a backend connection to a server only when it receives a data packet from the host. The CSM then manages the connections between the host and server.

Note SYN cookies limit the use of some TCP options because of the 32-bit restriction of the ISN field. For example, the Window Scale Option is not contained in the cookie.

To use SYN cookies on the CSM, enable the termination service for each virtual server requiring DoS protection. The default embryonic threshold is 5000. To modify this threshold, set the SYN_COOKIE_THRESHOLD variable to any number between 0–1000000. For example, to use SYN cookies for all connections requests, set the threshold to 0 (zero).

SYN cookie technology requires a secret key, which the hashing algorithm uses to generate the cookie. The CSM generates a new key every three seconds by default. Use the SYN_COOKIE_INTERVAL variable to modify the key generation period from 1 to 60 seconds. The CSM commits to memory the current key and one previous secret key to allow it to reverse the SYN cookie hash. This implies that the host has twice the time set in the SYN_COOKIE_INTERVAL to send an ACK to the CSM for validation. The CSM drops invalid host acknowledgements (ACKs).

SYN Cookies on the FWSM

Starting with Release 2.3, SYN cookies are used to protect against SYN flood attacks. When the embryonic connection threshold of a connection is crossed, the FWSM acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the FWSM receives an ACK from the client, it can then authenticate the connection and allow the connection to the server. The configuration is the same as described in TCP Intercept on the FWSM, page 2-10.

Performance Considerations

The average size of each frame in a SYN flood can be calculated as follows:

20 bytes (IP header) + 20 bytes (TCP header) + 8 bytes (Options) = 48 bytes.

This means that the number of SYN packets that can be carried by a DS-3 channel is as follows:

DS-3 channel = 51.84 Mbps / (48 * 8) = ~135,000 SYNs/s

An OC-3 link (3 DS-3 channels) can carry approximately 403,000 SYNs/s and an OC-12 link can carry about 1.6 M SYNs/s. A single host connected to the Internet at 1 Mbps can generate approximately 2,600 SYNs/s.

The performance implications of using the SYN flood protection technologies described in the previous section are as follows:

- TCP Intercept on sup2 (MSFC2)—In watch mode, this can sustain the DoS attack of a single host connected to the Internet at 1 Mbps. It can sustain bursts of higher levels of SYN floods but this causes higher CPU utilization.
- TCP Intercept on sup720 (MSFC3)—This can sustain about three times the performance of a sup2 (MSFC2). The performance is unrelated to the ASIC performing the sequence number adjustment, but rather is limited by the router processor performance in setting up new connections.
- TCP Intercept on FWSM (releases before 2.3)—This can sustain an attack of approximately 45,000 SYNs/s (approximately 15 x 1 Mbps-connected hosts attacking in parallel) without a noticeable effect to the user in terms of transactions per second. The performance degradation of legitimate HTTP transactions is about 10 percent, which means that legitimate transactions still complete, but the connection setup rate goes down. The performance impact of the DoS attack becomes significant at DS-3 rates.
- TCP SYN cookies on CSM (starting from Release 3.2)—This can sustain a DS-3 level of DoS attack with no visible impact to a user on HTTP transactions. The performance degradation is about 10 percent, which means that legitimate transactions still complete, but the connection setup rate goes

down. The performance degradation becomes significant (30–40 percent) at about 300,000 SYNs/s of SYN flood. At that level, HTTP transactions still complete, but the setup rate for legitimate transactions is significantly reduced.

- TCP SYN cookies on FWSM (starting from Release 2.3)—The performance is superior to the performance of the CSM with SYN cookies.

Design Models

Knowing the technology that performs best against DoS attacks can help when choosing the most effective design. When configuring a data center with load balancers and firewalls, two main models of deployment are possible (see Figure 2-7).

Figure 2-7 In-line Designs with Firewall and Load Balancers

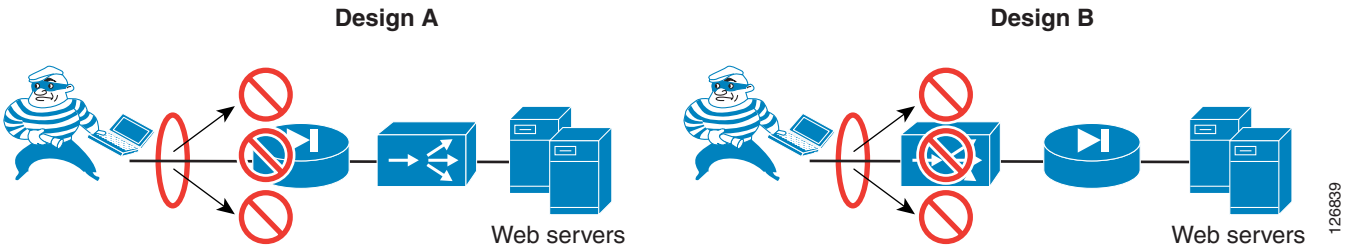

In Design A, the firewall stops the DoS attack, and the load balancer does not even see it. The level of DoS protection that this design provides equals the DoS performance provided by the firewall.

Design B, in which the load balancing function precedes the firewall function, may be preferable when the load balancer provides better DoS protection than the firewall. For example, Cisco CSM Release 3.2, which uses SYN cookies, provides much better DoS protection than Cisco FWSM Release 2.2, which uses TCP Intercept. The design that works better depends entirely on the release of software available for each product at the time of deployment.

A third design combines the DoS capabilities of both products (see right side of Figure 2-8). As illustrated, this design uses a CSM one-arm design combined with the FWSM placed between the MSFC and the servers.

Note There is no technical reason for using a firewall to protect a load balancer.

Figure 2-8 Cisco Data Center Solution—FWSM and CSM for DoS Protection

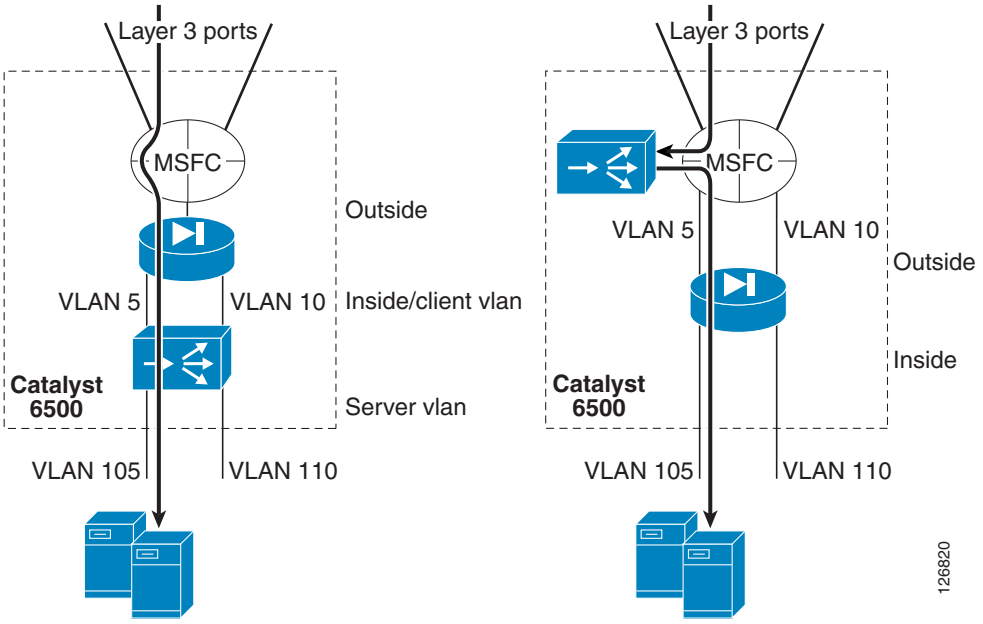

Figure 2-8 shows two designs:

- The design on the left represents one of the inline designs: MSFC–FWSM–CSM–servers (inline CSM)
- The design on the right represents the one-arm design: MSFC–FWSM+MSFC–CSM (one-arm)

With the CSM one-arm design, traffic that is load balanced to virtual IP addresses (VIPs) is directed to the CSM, while other traffic goes directly through the FWSM, bypassing the CSM. If an attacker launches a SYN flood against a VIP, the CSM is the first device that is hit. If the attacker launches a SYN flood against an IP address that does not require load balancing, the FWSM sees the traffic first.

The benefit of this design is that the DoS protection capabilities of the CSM and FWSM are combined:

- The CSM protects against DoS attacks directed at a VIP.
- The FWSM protects against DoS attacks directed at non-load balanced servers.

The FWSM can be configured for DoS protection in both routed mode and transparent mode. Currently, the FWSM does not perform NAT when operating in transparent mode. NAT can be applied to outside-facing servers that do not require any load balancing by using the FWSM in routed mode. If you use the FWSM in transparent mode to protect outside-facing servers that do not require load balancing, these servers should be assigned valid, public IP addresses. On outside-facing servers that are assigned private IP addresses and that require load balancing, NAT can be applied by the CSM.

Network Design for Intrusion Detection

This section describes the reference architecture for capturing traffic for network intrusion detection or other anomaly detection, as shown in Figure 2-9. This section includes the following topics:

- Topology
- VSPAN and PSPAN
- Locally Switched Traffic and Routed Traffic

Figure 2-9 Network IDS Capture Architecture

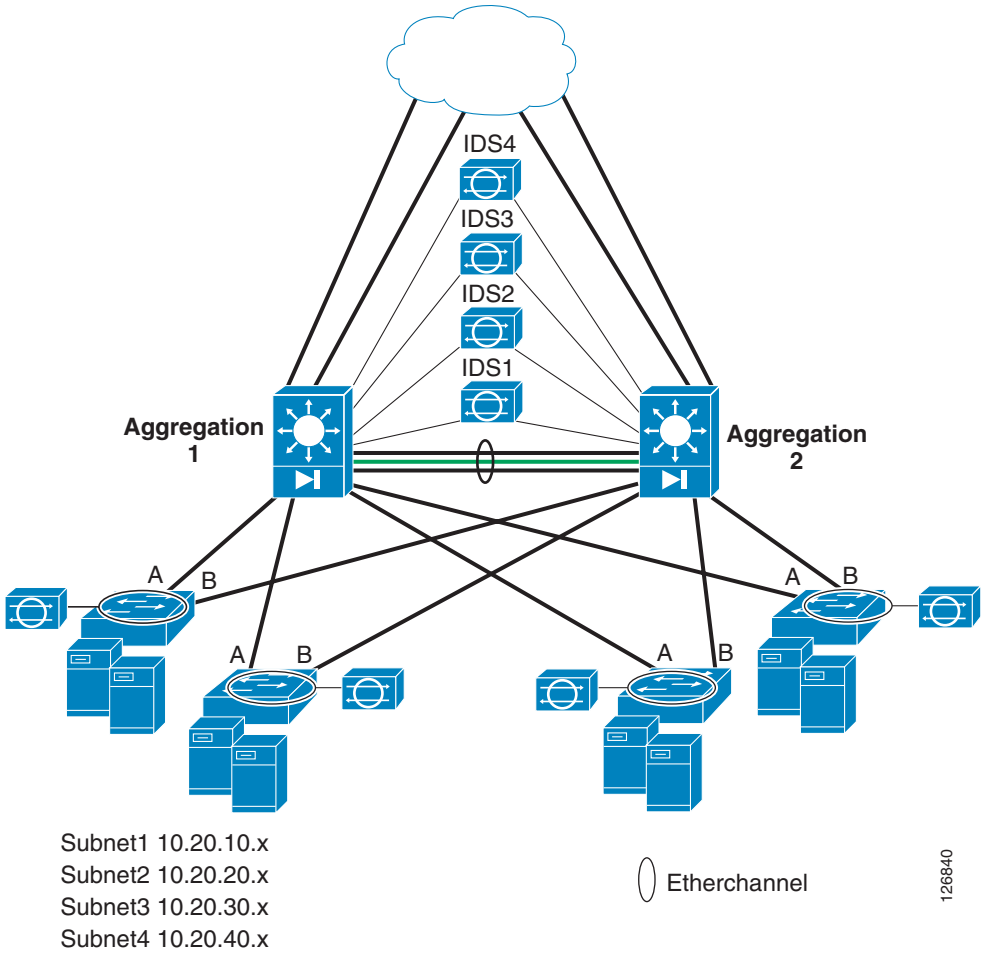

Topology

The design shown in Figure 2-9 is a fully redundant data center topology with access and aggregation layers. The aggregation layer is built with Catalyst 6500s with an IDS sensor attached to each aggregation switch to capture TCP streams that take asymmetric paths, and with an optional FWSM in each aggregation switch. The same configuration present on Aggregation1 is also present on Aggregation2 so that a given flow can take one aggregation switch in its inbound direction and Aggregation2 in the outbound direction. The IDS sensors are able to correlate the directions of the traffic as part of the same connection or flow.

Optionally, the IDS sensors can be attached to a single Catalyst 6500 because the mirrored traffic from Aggregation2 can be carried on the RSPAN VLAN to Aggregation1.

This topology has a number of subnets but no assumption is made on where these subnets reside in the access switches. For example, each data center subnet can be monitored respectively by IDS1, IDS2, IDS3, or IDS4, regardless of where (on which access switches) these subnets reside in the data center.

Several techniques can be used to differentiate the traffic on multiple sensors. The most recent and powerful techniques include the following:

- RSPAN combined with VACL redirects

- Virtual SPAN sessions

The user must establish policies on what traffic IDS1, IDS2, IDS3, and IDS4 need to monitor. For example, IDS1 could monitor HTTP traffic, IDS2 could monitor DNS traffic, and IDS3 could monitor SMTP traffic, and so on. The policy can be modified whenever the user desires without impacting the way traffic is forwarded on the network.

VSPAN and PSPAN

Whether to use VLAN SPAN (VSPAN) or Port SPAN (PSPAN) depends on the specific configuration of the system, and design guidance regarding this topic is available in the following chapters of this guide:

- Chapter 7, "Traffic Capturing for Granular Traffic Analysis"
- Chapter 9, "Deployment of Network-Based IDS Sensors and Integration with Service Modules"

If an FWSM module is present in the architecture, it is important to realize that TCP sequence numbers are normally randomized, so you need to be careful with where to place the SPAN. Two solutions are possible:

- Let the FWSM randomize the sequence numbers and configure the VSPAN outside or inside the FWSM.
- Use PSPAN on the ports surrounding the Catalyst 6500 switches and disable sequence number randomization on the FWSM.

Whether to use VSPAN outside or inside the FWSM depends on a number of factors that are not strictly related to security, including the following:

- Generation of duplicate frames
- Need to see both directions of the traffic (this is easy to do by using VSPAN outside)
- Requirement to protect the IDS sensors from seeing DoS traffic that is stopped by the FWSM or the CSM

The use of VSPAN inside or outside is described in detail in Chapter 9, "Deployment of Network-Based IDS Sensors and Integration with Service Modules."

Locally Switched Traffic and Routed Traffic

Traffic monitoring at the aggregation layer can use RSPAN with VACL redirect (indicated by the green line in Figure 2-9). This provides the maximum flexibility in monitoring all data center traffic and assigning IDS1, 2, 3, and 4 to different user-definable traffic categories. Also, RSPAN with VACL redirect is used at the aggregation layer because it poses no restrictions to monitor any-to-any routed or switched traffic.

The access layer is implemented at Layer 2; no traffic routing occurs on the access switches. For simplicity, traffic monitoring at the access layer uses VACL capture with IDS sensor directly connected to the access switches.

Optionally, you can perform RSPAN on the access layer switches and trunk the RSPAN VLAN to the aggregation layer. This allows the sensors at the aggregation layer to monitor locally switched traffic at the access layer. The topology for the RSPAN VLAN does not need redundancy, and it can be mapped to the links that normally do not forward any traffic, the Spanning-Tree blocking links in Figure 2-10.

Figure 2-10 Using RSPAN to Monitor Access Switches Traffic

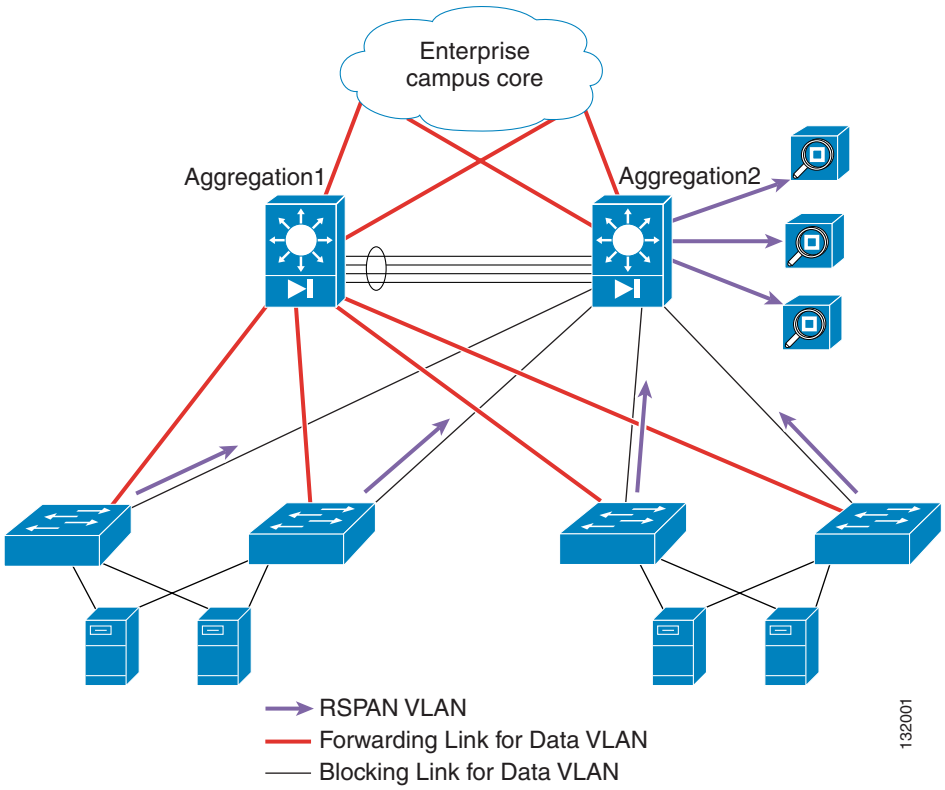

In Figure 2-10, the RSPAN VLAN is present only on uplinks to Aggregation2. These links are the backup links for the data traffic (Spanning-Tree blocks these links for the data VLAN).