Web-Site Security and Denial-of-Service Protection

Background

The challenge for successful public Web sites is to encourage access to the site, while eliminating undesirable or malicious traffic, and providing the necessary levels of sufficient security without creating constraining site limitations in performance or scalability.

Disruption of service caused by denial-of-service (DoS) attacks is the "kiss of death" for Web-driven enterprises such as portals and e-commerce sites. The 1999 Computer Crime and Security Survey found that system penetration by outsiders increased for the third year in a row, with 30 percent of respondents reporting intrusions. Those reporting their Internet connection as a frequent point of attack rose for the third straight year, from 37 percent of respondents in 1996 to 57 percent in 1999.

Preventing DoS attacks is critical for most Web sites. These attacks are specifically designed to bring down a Web site using methods that appear to be normal network traffic—until it is too late. Web-site administrators have used packet filtering in their IP routers to provide basic access control, but often this slows router performance to an unacceptable point and fails to eliminate many common types of DoS attacks.

Traditional firewalls can act as a boundary for IP addresses, using Network Address Translation (NAT), preventing DoS by limiting traffic using specific TCP ports, limiting traffic coming from specific network addresses, or even scanning traffic for viruses or undesirable applications. However, these solutions were designed to prevent access to systems, a concept that is incompatible with today's Web. Further, traditional firewalls do not scale well for today's extremely high-traffic Web environments.

Cisco CSS 11000 series content services switches provide comprehensive Web-site and back-end system security capabilities designed to provide the right level of security without sacrificing scalability or performance, including (See Figure 1):

Site-Level Security

- *DoS attack prevention*—Cisco CSS Web switches validate every session flow at initial flow setup time and eliminate all connection-based DoS attacks and other attempted malicious or abnormal connections, with no impact on the performance of the Cisco CSS Web switch.
- *FlowWall security*—Cisco CSS Web switches provide firewall services including high-speed access control lists (ACLs) that block particular content requests by IP address, TCP port, host tag, complete URL, or file type.
- *Network Address Translation*—Wire-speed NAT capabilities on Cisco CSS Web switches effectively hide the IP addresses for all devices located behind the Web switch, such as Web servers and caches, eliminating the ability for hackers to attack servers directly by using explicit IP addresses.

CISCO SYSTEMS

Security for Back-End Systems

- *Firewall load balancing*—When full firewall security is needed, either in the path from the Internet or to protect mission-critical back-end systems or networks, Cisco CSS 11000 series switches can prevent bottlenecks and eliminate single points of failure by distributing traffic among multiple load-balanced firewalls.

Figure 1 Web-Site Security DoS Attacks

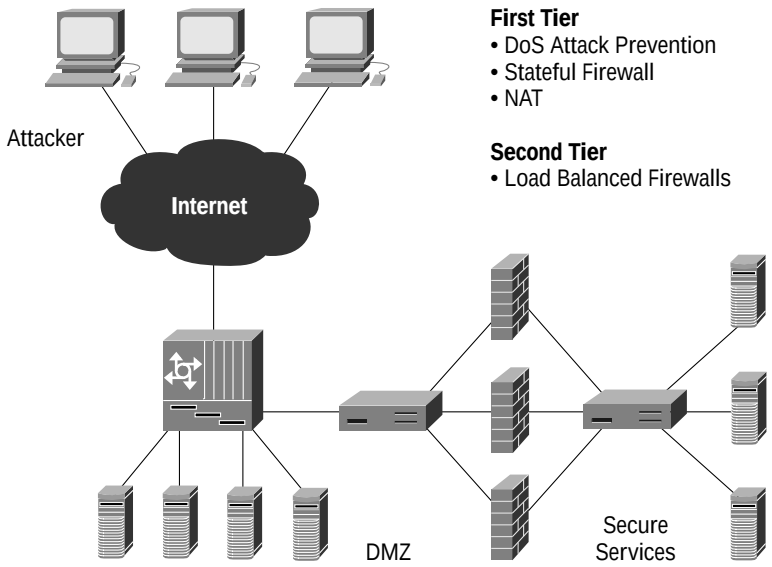

- *DoS attacks*—A DoS attack typically involves the misuse of standard protocols or connection processes with the intent to overload and disable the target Web servers. For example:
- *TCP SYN floods*—These attacks are created by sending repeated TCP connection requests with no subsequent completion, causing the target system to allocate TCP control blocks until it runs out of resources.
- *"Smurf" and "fraggle" attacks*—These attacks involve sending a large number of Internet Control Message Protocol (ICMP) echo (ping) messages to an IP broadcast address, with the forged source address of the intended victim. The routing device forwarding traffic to those broadcast addresses performs the IP broadcast to Layer 2 broadcast function, whereupon most network hosts will each take the ICMP echo request and issue an echo reply, multiplying the traffic by the number of hosts responding. A fraggle is similar to a smurf except that it uses User Datagram Protocol (UDP) echo messages. On a multiaccess broadcast network, potentially hundreds of machines could reply to each packet.
- *UNIX process table DoS attacks*—These attacks entail sending repeated open-connection requests to a UNIX server. Subprograms—such as Internet Daemon, Secure Shell Daemon, and Internet Message Access Protocol Daemon—are written to automatically answer every connection and carry out requests. But if the connection is initiated with no request, most daemons keep the line open, using resources from the server process table, which is equipped to handle between 600 and 1500 simultaneous tasks. Repeated connections can very quickly overload the process table and crash the server.
- *Finger of death*—These attacks involve sending a finger request to a specific computer every minute, but never disconnecting. Program failure to terminate the connection can quickly overload a UNIX server "process tables" and bring the Internet service provider's (ISP's) services to a standstill for hours.

Cisco CSS 11000 series switches eliminate all of the above DoS attacks, as well as any malicious connection requests attempting to go through the switch, with no impact on the Web switch itself.

General Protection for All Traffic

The Cisco CSS 11000 series switches will discard frames if:

- Length is too short
- Frame is fragmented
- Source IP address = IP destination (LAND attack)
- Source address = Cisco addresses, or the source is a subnet broadcast
- Source address is not a unicast address
- Source IP address is a loop-back address
- Destination IP address is a loop-back address
- Destination address is not a valid unicast or multicast address

For Layers 4 and 5

For Hypertext Transfer Protocol (HTTP) flows (directed to Virtual IP address [VIP] with Layer 5 rules in the switch), the Web switch must receive a valid content frame within 16 seconds of starting the flow or it will discard the frames and tear the flow down. Servers are never contacted until the Web switch receives a valid content frame for them. Thus, there is no danger of dangling TCP state blocks on a server front ended by an Cisco CSS 11000 series switch.

For TCP flows (directed to a VIP with Layer 4 rules in the switch), Web switches must receive a return ACK for the three-way TCP handshake within 16 seconds, or it will tear down that TCP flow. This eliminates any process-table DoS attacks that attempt to make repeated open request connections to a server.

For any flow that has tried an initial SYN more than eight times, the Web switch will kill the flow and cease processing any more SYNs from that source having the same initial sequence numbers and source and destination address and port pairs. This eliminates SYN flood DoS attacks.

Network Address Translation

NAT hides the IP addresses for all devices located behind the switch, allowing unlimited use of thousands of private IP addresses (10.xxx.xxx.xxx) to be mapped to one or more external VIPs that are based on assigned, globally unique, IP addresses. NAT is also important for managing existing allocated IP addresses, a scarce resource, and reduces the need to acquire additional addresses as the network expands.

NAT is based on an industry-standard implementation described in RFC 1631, but only Cisco CSS 11000 series switches allow full bidirectional NAT on any port, and at wire speed. The switch also supports source group NAT, which provides NAT for server-initiated flows going back to the client (port-based dynamic File Transfer Protocol [FTP]) or server-initiated flows going to locations other than the client. As a result, higher levels of security are achieved with no negative impact on site performance.

FlowWall Security

Upon detection of a new flow, the Web switch firewall rules, including ACL and flow admission control, are invoked. After these policies are verified as part of the flow setup, all traffic for the duration of that flow is switched at wire speed. Cisco provides comparable functionality to the Cisco PIX™ Firewall—only faster—but Cisco is not replacing conventional software-based firewalls, which inspect every packet. For instance, FlowWall does not scan for Java and Active-X traffic that is common in the Internet but may not be acceptable for secure enterprise networks.

In addition to several important traditional firewall rules, the Web switch can also be configured for policies for any particular domain name or Universal Resource Locator (URL). In fact, the Cisco CSS 11000 series switches allow ACLs to be configured that specify an action (include/bypass/block) for any request matching some or all of the following:

- Source IP address
- Destination IP address
- TCP port
- Host tag
- URL
- File extension

A “blocking” rule will prevent any requests for specific content from going to the origin server or cache. This capability could be used on a switch which is front ending cache servers located in a point of presence (PoP) or cable headend, to block access to specific content in the Internet.

ACLs can also be used to enable advanced transparent caching policies, for example:

- An “include” rule can be used to direct all requests for particular content to the cache and to bypass the cache for everything else, optimizing the delivery of specific customer content.
- A bypass rule can be used to bypass the cache for all requests for particular content, and to direct everything else to the cache. Some customers may have technical, legal, or philosophical reasons for having their traffic cached, and this capability provides granular control to exclude all or part of a particular customer’s content from being diverted to the cache.

Firewall Load Balancing

Cisco CSS 11000 series switches can load balance traffic among multiple firewalls, eliminating performance bottlenecks and single points of failure for securing Web sites, back-end databases, networks, or other resources.

Cisco CSS 11000 series switches are deployed in front and in back of the firewalls being load balanced. Physically unique switches are deployed on each side of the firewalls, although ports on each switch not being used for firewall load balancing can be configured for other uses, as shown in Figure 2.

Figure 2 Enterprise Firewall Load-Balancing Example

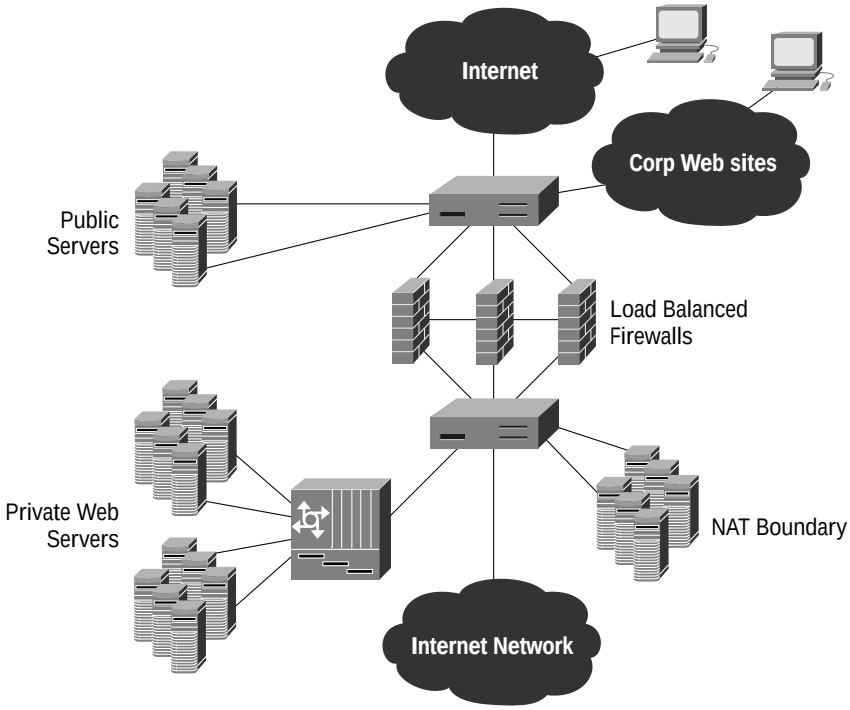

Cisco firewall load balancing is designed to distribute traffic among stateful firewalls that have their own IP addresses, including any Cisco PIX Firewall, CheckPoint FirewallOne, or Nokia firewall products.

Figure 3 Load-Balanced Firewall Configuration

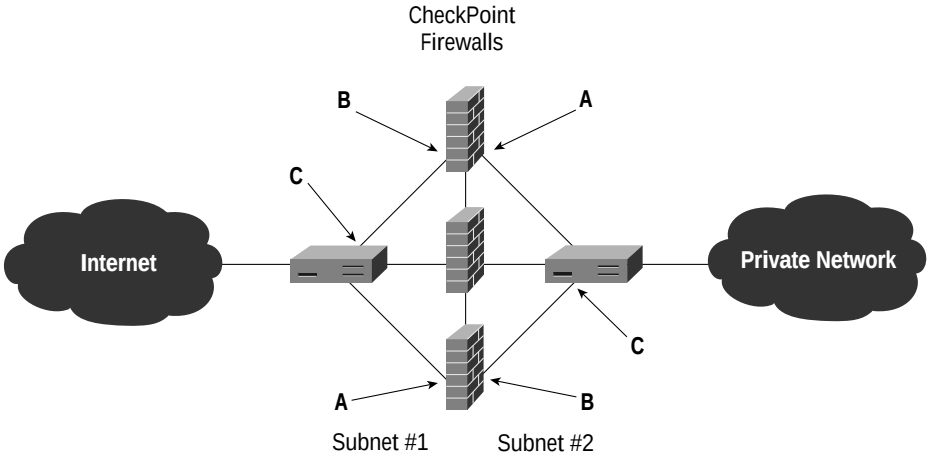

Cisco CSS 11000 series switches ensure that all traffic for a given Web flow between a pair of IP addresses, in either direction, will traverse the same firewall. As shown in Figure 3, this is accomplished by configuring static IP routes on each switch with the IP address of the adjacent firewall port (A), the IP address of the remote firewall port (B), as well as the IP address of the remote Web switch port (C). In addition, all ports on each side of the load-balanced firewall are utilizing a different IP subnet address. In this configuration, the firewalls cannot be configured to perform NAT—rather all NAT processing is done on the Web Switch on the private network side of the firewall.

Each of the Cisco CSS 11000 series switches secures the paths through the firewall by utilizing health checks to verify that each stage on the paths is responding, including its adjacent firewall port, its remote firewall port, and its remote Web switch port. If any part of the path becomes unavailable, the Web switches reroute all traffic through the surviving paths. Cisco CSS 11000 series switches can be configured to send an alarm or log an event to notify the system administrators of faults. If firewalls are configured to share session-state information and have physical connections to each other, then user sessions will not be interrupted in the case of the failure in any path or firewall.

For additional redundancy, multiple Cisco CSS 11000 series switches can be deployed on both sides of the firewalls being load balanced in an active/passive configuration. This eliminates single points of failure for both firewalls and Web switches, as shown in Figure 4. In this configuration, the passive Web switches monitor the health of the active Cisco CSS 11000 series switches, and upon detecting a switch or an uplink failure, the Virtual Router Redundancy Protocol (VRRP) is used to transfer control to the passive switch pair.

Figure 4 Fully Redundant Firewall Load-Balancing Example

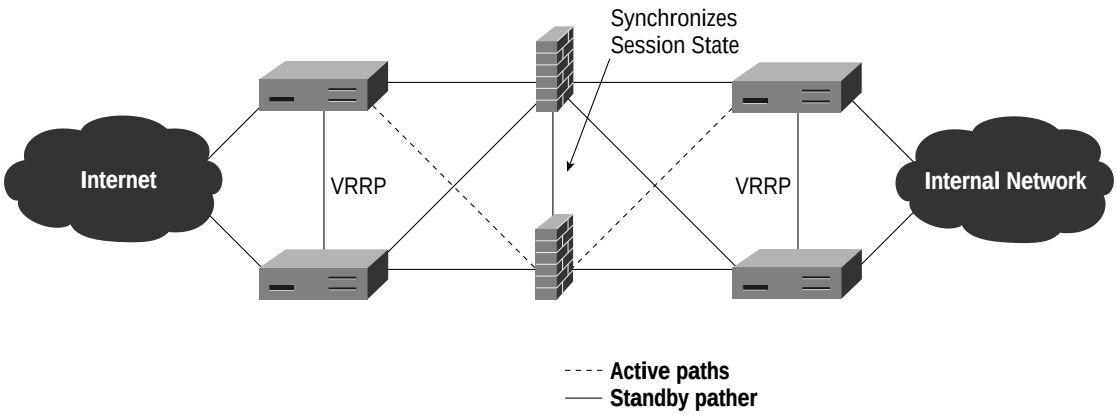

Conclusions

Cisco CSS 11000 series switches are designed from the ground up to provide comprehensive solutions for all aspects of Web-site security without compromising performance or scalability. They combine the inherent intelligence to eliminate sophisticated DoS attacks with the flexibility to configure custom security policies for each individual Web site, ensuring the constant availability of the site for real customers and legitimate users.

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

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

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

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco.com Web site at www.cisco.com/go/offices.

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

7/00 KW/SP