WHITE PAPER

PROTECTION FOR THE CISCO CATALYST 6500 SERIES SWITCHES AGAINST DENIAL-OF-SERVICE ATTACKS

Denial-of-service (DoS) attacks continue to be a serious threat to enterprise and service provider networks. They can disrupt mission-critical services, prevent data transfer between devices, and decrease overall productivity. The Cisco® Catalyst® 6500 Series Supervisor Engine 32 and Supervisor Engine 720 include hardware-based mechanisms that can effectively protect against DoS attacks on the Cisco Catalyst 6500 Series switches.

This paper assumes the reader is familiar with the basic forwarding operation of the Cisco Catalyst 6500 Series and does not review those details. This paper will leave a reader with a good understanding of all DoS mitigation mechanisms available on the Cisco Catalyst 6500 Series as of Cisco Catalyst OS Release 8.4 and Cisco IOS® Software Release 12.2(18)SXE1. Additional information can be found in the Cisco IOS Software 12.2SX DoS protection configuration guide at:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080435872.html

INTRODUCTION

Denial of service has become a growing concern, especially when considering the associated costs of such an attack. DoS attacks can decrease the performance of networked devices, disconnect the devices from the network, and cause system crashes. When network services are unavailable, enterprises and service providers suffer the loss of productivity and sales.

The objective of a DoS attack is to deprive a user or organization access to services or resources. If a Website is compromised by a DoS attack, millions of users could be denied access to the site. DoS attacks do not typically result in intrusion or the illegal theft of information. Instead of providing access to unauthorized users, DoS attacks can cause much aggravation and cost to the target customer by preventing authorized access. Distributed DoS (DDoS) attacks amplify DoS attacks in that a multitude of compromised systems coordinate to flood targets with attack packets, thereby causing denial of service for users of the targeted systems.

Common forms of DoS attacks include “SYN flood attacks,” “land attacks,” “smurf attacks,” viruses, and worms. Some of these attacks can be hard to defend against because DoS packets may look exactly like normal packets. However, while these DoS attacks all exploit various system and network vulnerabilities, they are similar in the way they spread and impact the network infrastructure. Most DoS attacks rely on spoofing and flooding techniques. Intrinsic network impacts include the resource exhaustion of the media’s bandwidth capacity, switch forwarding capacity, or switch CPU capacity. Understanding common attack vectors and network impact are key elements to deploying effective DoS mitigation techniques.

DOS PROTECTION FOR THE CISCO CATALYST 6500 SERIES

Because DoS attacks can be extremely costly and destructive to internetworking environments, networks must be proactively designed and configured with preventative means to counteract these attacks. This paper introduces the mechanisms available on the Cisco Catalyst 6500 Series switches to mitigate the threat of DoS attacks. Once enabled, these features minimize the potential for DoS attacks to impact the switch and network devices. The goal of these features is to maximize network availability in the context of a DoS attack.

All contents are Copyright © 1992–2005 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

The Cisco Catalyst 6500 Series offers several mechanisms that allow DoS protection. Crucial to this DoS protection and mitigation process is the inherent separation of the Cisco Catalyst 6500 Series data-plane and control- and management-plane traffic: data-plane traffic is typically hardware-switched and does not affect control-plane traffic; control-plane traffic needs to be processed by the CPU. This paper distinguishes between DoS attacks targeted at end systems, which impact the switch indirectly, and DoS attacks targeted at the switch itself. DoS attacks through the switch can be mitigated using data-plane DoS protection mechanisms. DoS attacks directed at the switch can be mitigated using control- and management-plane DoS protection mechanism.

Available mechanisms to counteract DoS attacks can be further categorized into the following generic mitigation mechanisms:

- Data-plane protection
- Control- and management-plane protection
- DoS detection
- Specialized DoS protection

It is important to note that the supervisor engine DoS mitigation mechanisms listed in this paper are hardware-based features unless noted otherwise: turning any of these features on will not impact the performance of the switch itself. So not only does the Cisco Catalyst 6500 Series mitigate DoS attacks, it also protects against line-rate DoS attacks with its superior hardware-based architecture. As an example of line-rate DoS attack mitigation, consider a DoS attack with Internet Control Message Protocol (ICMP) echo requests destined to the switch CPU, and a control-plane policy that limits ICMP traffic destined to the CPU in the order of 100 Kbps. No matter how much ICMP traffic is destined to the switch CPU (100 Kbps, 1 Mbps, 1 Gbps, 10 Gbps, etc.), the CPU utilization remains constant with no more than 100 Kbps of ICMP traffic reaching the CPU and the rest of the traffic being dropped in hardware.

DATA-PLANE PROTECTION

Because most traffic travels through the data plane, it is important to make sure that valid user traffic gets its fair share of the switch data-plane capacity. A DoS attack could directly or indirectly impact the switch's data-plane resources by flooding illegitimate or unnecessary traffic, hence depleting the finite hardware resources. Such attacks can make the network unavailable to legitimate users or critical traffic and may include the following symptoms:

- Abnormal bandwidth utilization and link saturation
- Forwarding capacity saturation
- Port buffer overflow
- Content Addressable Memory (CAM) table overflow
- Indiscriminate drops of incoming and outgoing packets

The Cisco Catalyst 6500 Series supports robust anti-spoofing and anti-flooding techniques to mitigate DoS attacks that directly or indirectly affect the data plane.

Anti-Spoofing Mechanisms

To protect against DoS attacks targeted at the switch itself or at hosts, it is important to restrict possible attack vectors. This can be done through source-address verification mechanisms validating the traffic originator such that attackers cannot use spoofed source addresses. Administrators should also restrict network access to allowed traffic and applications only. Anti-spoofing mechanisms play an important role in preventing DoS attacks. Dynamic Host Configuration Protocol (DHCP) Snooping, Dynamic Address Resolution Protocol (ARP) Inspection, and IP Source Guard provide anti-spoofing protection and man-in-the-middle attack prevention in the Layer 2 access layer. Unicast Reverse Path Forwarding check and access-lists mechanisms can be deployed with no performance cost throughout the network.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Anti-Spoofing and Man-in-the-Middle Attack Prevention in the Access Layer

Figure 1 shows how deploying DHCP Snooping, Dynamic ARP Inspection (DAI), and IP Source Guard in the access layer can prevent man-in-the-middle attacks. A malicious user (Host D) tries to act as a DHCP server to redirect all network traffic through itself by specifying a wrong default gateway. Another possibility is to release all IP addresses on the network by sending DHCP release messages to all clients on one or multiple subnets and therefore prevent legitimate users to access network. Enabling DHCP Snooping on the switch stops these types of attacks and allows the creation of a binding table listing IP address to MAC address bindings per source interface. The same malicious user (Host D) tries to poison the switch's ARP cache and intercept all the traffic from Host A by declaring itself the gateway. Enabling DAI stops this type of attack by comparing incoming traffic on untrusted ports with information stored in the DHCP Snooping binding table. The same malicious user (Host D) tries to carry out a DoS attack using Host A's source IP to carry the attack. Enabling IP Source Guard stops this type of attack by comparing incoming source IP address on untrusted ports with the source IP address stored in the DHCP Snooping binding table for a particular incoming interface.

Figure 1. Man-in-the-Middle Attack Prevention with DHCP Snooping, Dynamic ARP Inspection, and IP Source Guard

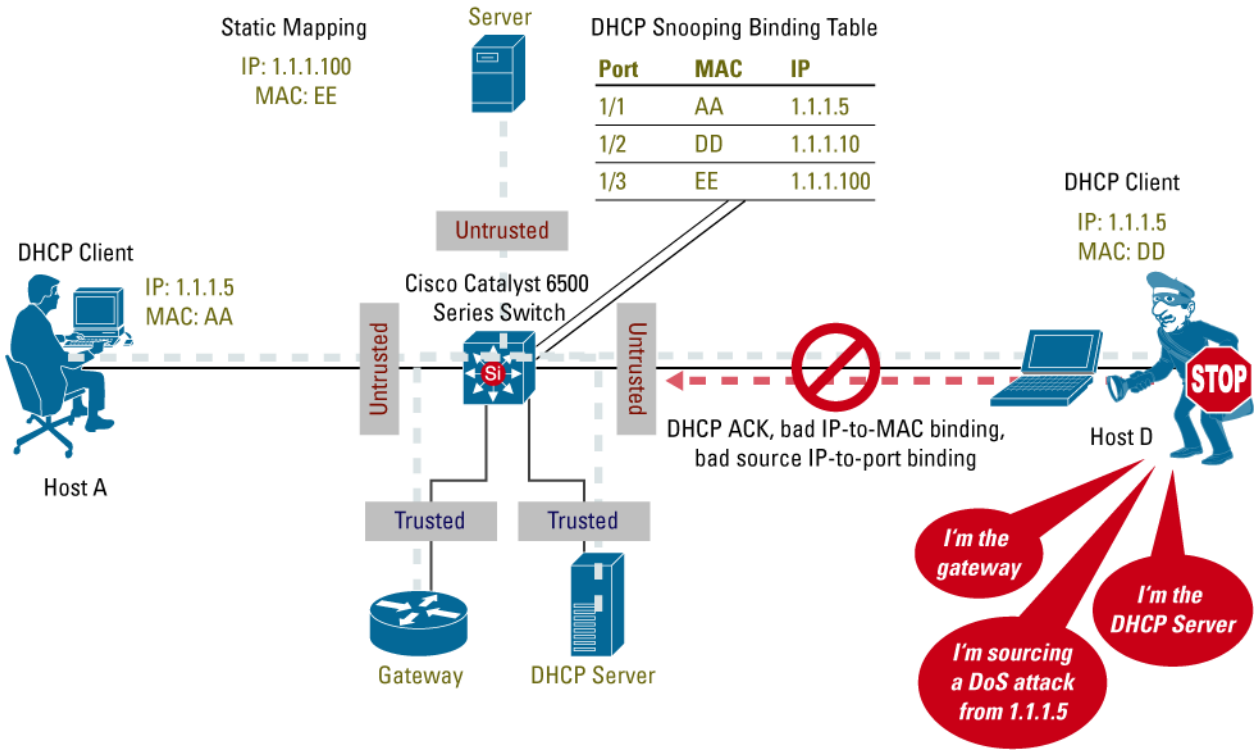

DHCP Snooping verifies DHCP transactions and protects against rogue DHCP servers. DHCP Snooping uses the concept of trusted and untrusted ports to filter the DHCP packets that are received by the switch. Servers and relay ports should be configured as trusted, while client ports should be set to untrusted. Any DHCP server-to-client reply message coming from untrusted ports is dropped. This includes messages such as DHCPOFFER, DHCPACK, and DHCPNAK.

Aside from verifying DHCP transactions, DHCP Snooping also forms the basis for other security features such as IP Source Guard and DAI by building a table of authorized IP and MAC addresses per interface that can be checked against to validate traffic coming in from untrusted ports.

Additional DoS protection can also be obtained with DHCP Snooping by rate limiting incoming DHCP requests and enabling the DHCP option 82 to offer port-to-port DHCP isolation.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Dynamic ARP Inspection (DAI) prevents ARP spoofing and man-in-the-middle attacks, where an attacker could send forged ARP packets (for example, gratuitous ARPs) carrying a bogus IP/MAC binding in the payload to a host or to the default gateway. This attack would "poison" the ARP table of the target device. The DAI feature intercepts all ARP requests and replies on the untrusted ports and uses the binding information that is built by the DHCP Snooping feature to verify ARP packets have valid IP-to-MAC bindings. IP hosts that have static addresses assigned and do not use DHCP, such as servers, are supported using static ARP access control lists (ACLs). Violated hosts can be logged, and the attacker port can be error disabled.

Additionally, DAI also helps in the prevention of DoS attacks by limiting the number of incoming ARPs per second on an interface.

IP Source Guard prevents IP spoofing by forwarding only packets that have a source address consistent with the DHCP Snooping table. The IP Source Guard function can be performed at line rate by dynamically maintaining per-port ACLs (PACLs) based on DHCP bindings, so no rate limiting is required. IP Source Guard secures all DHCP-allocated IP addresses, and all other IP addresses may be supported through static bindings.

In Cisco Catalyst OS, the DHCP Snooping and DAI features are software mechanisms supported on all supervisor engines. In Cisco IOS Software, they are supported on Cisco Catalyst 6500 Series Supervisor Engine 2, Supervisor Engine 32 and Supervisor Engine 720. On Supervisor Engine 32 and Supervisor Engine 720 running Catalyst OS, a special rate limiter can also be turned on to control the rate of DHCP traffic and incoming ARP packets bound to the CPU (see Table 1), which provides for another level of DoS protection when combined with the platform-independent software DHCP and ARP rate limiter.

IP Source Guard can only be supported on forwarding engines that have inherent PACL hardware support; therefore it is supported only on the Cisco Catalyst 6500 Series Supervisor Engine 32 and Supervisor Engine 720.

The minimum software release for DHCP Snooping and DAI is Cisco Catalyst OS Release 8.3 and Cisco IOS Software Release 12.2(18)SXE. The minimum software release for IP Source Guard is Cisco Catalyst OS 8.3 (only Catalyst OS support today).

Anti-Spoofing Throughout the Network

Unicast Reverse Path Forwarding (URPF) discards packets that lack a consistent source IP address, such as spoofed IP source addresses created by malicious users to intercept valuable data. This feature uses Cisco Express Forwarding tables to verify that the source addresses and the interfaces on which packets were received are consistent with the forwarding tables on the supervisor engine. Should the packet be received from reverse path routes, the packet is forwarded. If there is no reverse path route on the interface on which the packet was received, the packet fails the URPF check and is discarded.

The Cisco Catalyst 6500 Series Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720 support URPF in hardware. In Layer 3 networks with multiple equal-cost paths, traffic from an individual source can enter the switch on more than one interface. The Supervisor Engine 32 and Supervisor Engine 720 support multiple return-path URPF checks in hardware, with up to two parallel paths for all prefixes in the routing table, and up to six parallel paths for prefixes reached through any four user-configurable RPF interface groups. However, URPF should be applied with caution on Supervisor Engine 2 because Supervisor Engine 2 allows only one single return-path check in hardware and when URPF is enabled on a Supervisor Engine 2, its Cisco Express Forwarding table capacity is halved.

URPF can be deployed throughout the campus at hardware-based performance rates. Aside from being an efficient anti-spoofing solution, URPF can also be used effectively in combination with static routes pointing to Null0 adjacencies to drop traffic from specific hosts or networks. Null0 Forwarding Information Base (FIB) adjacencies can in turn be distributed automatically with internal BGP (iBGP) to quickly drop packets from infected hosts (also known as remotely triggered blackhole filtering).

More information on remotely triggered blackhole filtering is available at: http://www.cisco.com/warp/public/732/Tech/security/docs/blackhole.pdf

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Access Control Lists

ACLs determine whether particular packets should be forwarded or dropped by the switch. More generally, security ACLs can be used to protect against source address spoofing or to restrict network access to only legitimate sources, networks, and applications. For example, ACLs should be used to deny private address space at the ingress of the Internet and perform some filtering in the campus such that packets can only originate from customer-assigned addresses. ACLs should also be used to deny unused multicast addresses, to prevent multicast DoS attacks. Another interesting example is that of MAC ACLs which could be used to deny packets with invalid IP versions.

ACLs can be used to react to DoS attacks: ACLs can be an efficient mechanism for dropping the DoS packets prior to reaching their intended target. When the switch is used in conjunction with a Cisco intrusion detection system (IDS) module, ACLs can also be installed dynamically as a response to the detection of the attack by the sensing engine.

The Cisco Catalyst 6500 Series supports the following types of security ACLs:

- *Router ACLs (RACLs)* provide access control for routed traffic. RACLs are applied to a VLAN interface or routed port.
- *VLAN ACLs (VACLs)* provide access control for bridged and routed traffic. VACLs are applied to VLANs.
- *Port-Based ACLs (PACLs)* provide access control for bridged and routed traffic. PACLs are applied to a physical Layer 2 switch port.

All Cisco Catalyst 6500 Series supervisor engines support RACLs and VACLs. However, PACLs are only supported on the Catalyst 6500 Series Supervisor Engine 32 and Supervisor Engine 720 (any Cisco Catalyst 6500 Series Policy Feature Card 3 [PFC3] type is supported in Catalyst OS only as of Catalyst OS 8.3). The PFC3B and PFC3BXL provide for enhanced hardware ACL capacity, and support for hardware access control entries (ACEs) statistics. When access lists are in use, it is highly recommended to double check the hardware ACL resources utilization because exceeding the hardware capacity could cause software switching.

More information on Cisco Catalyst 6500 Series ACLs is available at: http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/65acl_wp.pdf

Managing Resource Saturation

One of the most important factors in DoS mitigation is restricting access to network resources so that one user or device cannot deplete all the hardware resources and forwarding capacity. To do so, an administrator should ensure that only legitimate traffic can traverse the network, and in the right proportions. Mechanisms to restrict access to network resources include the previously mentioned anti-spoofing mechanisms. Other very effective mechanisms limit the amount of hardware resources in use by different sources. The following section emphasizes QoS and Port Security as ways to limit bandwidth usage and CAM usage per source on the switch.

Quality of Service

QoS policers can be used to limit the amount of traffic allowed through a switch on a per-source, per-flow, or per-port basis. In this manner, QoS can continue to minimize data path congestion and protect high-priority, mission-critical data, and time-sensitive applications such as voice and video traffic, should a DoS attack be initiated against the switch.

Another integral part of QoS is priority queuing and edge re-marking. Traffic should be marked at the network edge in bulk or using QoS policies such that high-priority traffic and low-priority traffic can be treated differently at queuing structures throughout the network. One application is “scavenger” class queuing, where excess traffic is marked to a low priority such that upstream network devices can prioritize critical traffic more effectively; there is no penalty assigned to legitimate traffic flows that exceed a certain threshold aside from re-marking. Only sustained, abnormal streams generated simultaneously by one or multiple hosts are subject to aggressive dropping. When possible, it is preferable to mark down out-of-profile traffic to a lower-priority threshold in the same queue, so that the possibility of out-of-sequence traffic is avoided. Figure 2 shows how excess traffic is classified on downstream devices as scavenger-class traffic, and how that traffic is then assigned to the lowest-priority queue as opposed to regular, gold, and voice traffic which go in different queues and are thus not affected by scavenger-class traffic.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Figure 2. Scavenger-Class Queuing on the Cisco Catalyst 6500 Series

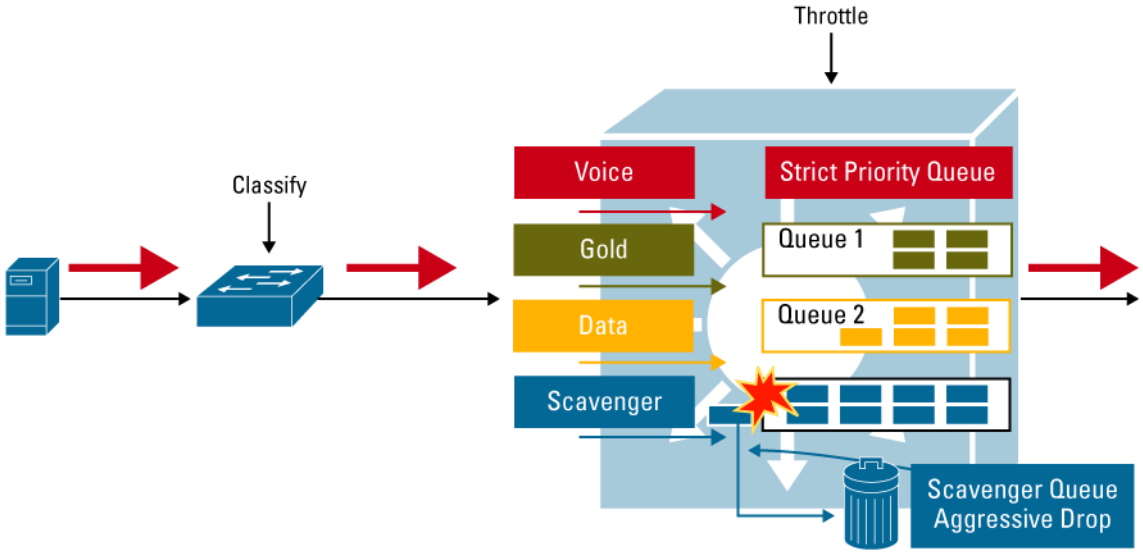

All Cisco Catalyst 6500 Series supervisor engines support QoS policers. However, granular user-based rate limiting (UBRL), extended QoS ACL hardware capacity, and egress QoS aggregate policers are only supported on the Supervisor 32 and Supervisor 720. Queuing and buffer capacity is a property of the line card on which ports reside and is not a property of the supervisor engine itself.

The following paper provides in-depth queuing and buffer information for the Cisco Catalyst 6500 Series:

http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/prodlit/buffe_wp.pdf

Port Security

Port Security mitigates against Layer 2 MAC address table overflow attacks (also known as CAM overflow), which exhaust switch hardware CAM tables by bombarding the switch with random MAC addresses, so that new host MACs are flooded in the network, thus potentially slowing network performance and increasing CPU load on clients and hosts. Port Security can limit the number of MAC addresses learned on a particular port. Using this feature, hosts cannot overload the CAM tables with more than the configured amount of MAC addresses for their port.

CONTROL PLANE AND MANAGEMENT PLANE PROTECTION

The vast majority of traffic generally travels through the router via the data plane; however, the switch processor and the route processor must manage certain packets. These packets will be referred to as control-plane packets for the remainder of the document.

The switch processor and route processor are critical for system operations. To protect the switch's control plane effectively, it is first important to profile the CPU traffic to better understand which types of packets should be allowed to reach the CPU and how critical each of these packet types are. Packets bound to the CPU include usual control-plane and management-plane traffic. Examples of such traffic include:

- Routing protocol packets—Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS), etc.
- First Hop Redundancy Protocol packets—Hot Standby Router Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), Virtual Router Redundancy Protocol (VRRP)
- Multicast control packets—Internet Group Management Protocol (IGMP), Protocol Independent Multicast (PIM), etc.
- Remote access and management traffic—Simple Network Management Protocol (SNMP), Network Time Protocol (NTP), Telnet, Secure Shell (SSH) Protocol, Trivial File Transfer Protocol (TFTP), etc.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

- Monitoring and troubleshooting traffic—Internet Control Message Protocol (ICMP), Traceroute, etc.
- Address Resolution Protocol (ARP)
- Layer 2 protocols—Spanning Tree Protocol, Cisco Discovery Protocol, VLAN Trunking Protocol (VTP), Link Aggregation Control Protocol (LACP), Port Aggregation Protocol (PaGP), Unidirectional Link Detection Protocol (UDLD), 802.1x, etc.

Some data-plane traffic may have to be processed in software as well. This type of traffic is referred to as data-plane "punt" traffic. Examples of software-processed data-plane packets include:

- Packets with IP options
- Packets with time-to-live (TTL) field equal to 1
- Packets for which destination prefix cannot be found in the routing table, also referred to as "FIB-miss"
- Packets that require ACL logging
- Packets that cannot be switched in hardware because a non-hardware-supported feature is applied to the packet
- Packets that are not classified by the hardware (AppleTalk, Internetwork Packet Exchange [IPX] in the Cisco Catalyst 6500 Series Supervisor Engine 32 and Supervisor Engine 720, etc.

A DoS attack targeting the Cisco Catalyst 6500 Series, which can be perpetrated either inadvertently or maliciously, typically involves high rates of traffic destined to the switch or route processor itself. This can result in the following symptoms:

- High CPU utilization on the route processor or switch processor
- Route Flaps due to loss of routing protocol updates or keepalives
- Unstable Layer 2 topology
- Slow or unresponsive interactive sessions with the command-line interface (CLI)
- Processor resource exhaustion such as memory and buffers
- Indiscriminate drops of incoming packets

Several features address the need to protect the control plane, ultimately helping to ensure network stability, reachability, and packet delivery. Figure 3 depicts a Cisco Catalyst 6500 Series, its two Multilayer Switching Feature Card (MSFC) processors, and the mechanisms available on the Catalyst 6500 Series to protect the switch processor and route processor against DoS attacks. Specifically, the Cisco Catalyst 6500 Series supports a layered defense with control-plane policing (CoPP) and special-case route processor CPU rate limiters. The CoPP feature applies to traffic going to the route processor control-plane interface. CoPP is applied in hardware on a per-forwarding-engine basis (Cisco Catalyst 6500 Series Policy Feature Card [PFC] and Distributed Forwarding Card [DFC]). The special-case CPU rate limiters are platform-dependant rate limiters applied in hardware to traffic going to the switch processor or route processor.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Figure 3. Control-Plane and Management-Plane DoS Protection with Control Plane Policing and Special-Case Rate Limiters

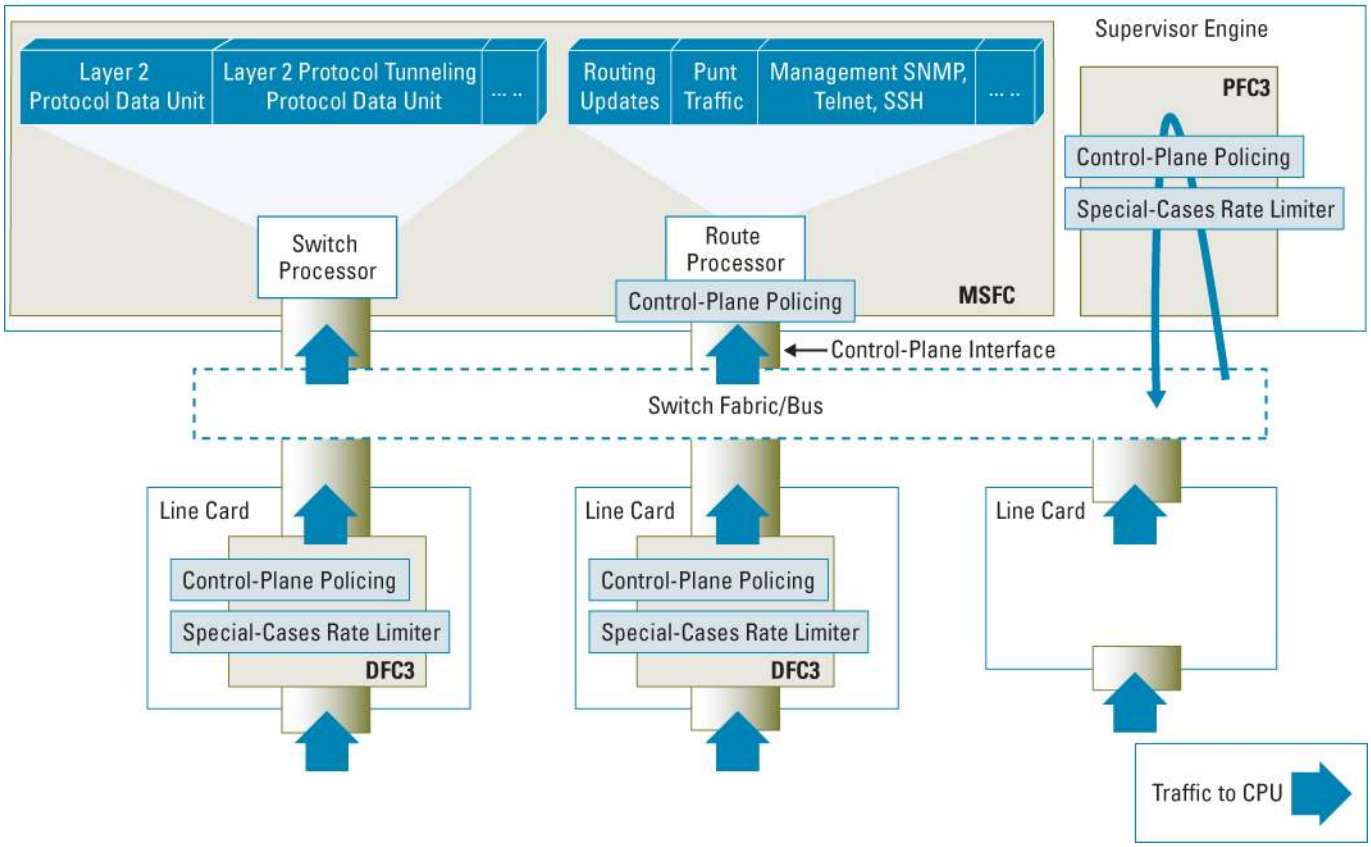

Control-Plane Policing

Control-plane policing (CoPP) allows filtering and rate-limiting of traffic sent to the route processor. This CoPP capability is achieved by using existing QoS policers and applying these to a new interface, the "control-plane" interface. This interface is attached to the route processor (see the control-plane interface in Figure 3). As a result, a control-plane policy protects traffic inbound to the route processor CPU (CoPP only affects input packets, not output packets), and it can thus prevent DoS traffic from congesting the route processor CPU.

CoPP configured policies depend heavily on the customer environment and where the switch is used in this environment. For example, an enterprise access-layer switch and an enterprise core switch may run different protocols and the expected CPU load for a given common protocol will be different in these two environments. The following methodology can be used to determine the right CoPP policies for a given switch:

1. Determine the classification scheme for your network: enumerate the known types of traffic that access the route processor and divide them into categories (classes). Examples of categories include an Exterior Gateway Protocol (EGP) class, Interior Gateway Protocol (IGP) class, management class, reporting class, monitoring class, critical application class, undesirable class, and default class.
2. Classify traffic going to the route processor CPU using ACLs. For each category identified in step 1, different types of traffic can be further categorized using granular access control entries.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

3. Review Identified traffic, adjust Classification, and Apply liberal CoPP policies for each class of traffic. It is essential to apply a corresponding policing action for each class, because the Cisco Catalyst 6500 Series will ignore a class that does not have a corresponding policing action. If the traffic in a given class should not be rate limited, configure a transmit policing conform-action with a high rate and a policing exceed-action of drop (for example "police 31500000 conform-action transmit exceed-action drop"). Alternatively, both conform-action and exceed-action could be set to transmit, but doing so will allocate a default policer as opposed to a dedicated policer with its own hardware counters.
4. Narrow the ACL permit statements to only allow known authorized source addresses
5. Refine CoPP policies based on CoPP and software ACE counters. On the Cisco Catalyst 6500 Series, the administrator can use the following commands to collect these statistics: **`show policy-map control-plane [input class <class_name>], show mls qos ip, show access-list`**.

The Cisco Catalyst 6500 Series supports CoPP on the Supervisor Engine 32 and Supervisor Engine 720 in hardware starting with Cisco IOS Software Release 12.2(18)SXD1. CoPP is actually applied at two different levels on the Cisco Catalyst 6500 Series. The first level is the hardware-based forwarding engine mitigation, and the second level is the software CoPP. Forwarding engines are programmed with the same global CoPP policy even though they each police traffic independently, so the route processor CPU could ultimately be presented N times the configured traffic rate, where N denotes the number of forwarding engines (active PFCs and DFCs) present in a Cisco Catalyst 6500 Series chassis. In Figure 3, after each forwarding engine has independently mitigated a line-rate attack in hardware, CoPP is enforced in software at interrupt level to make sure that only the exact rate configured in the control-plane policy is processed by the route processor. This should be taken into account when configuring a control-plane policer.

The administrator should be aware that CoPP on the Cisco Catalyst 6500 Series is not enforced in hardware unless Multilayer Switching (MLS) QoS is enabled globally: if the "mls qos" global command-line interface (CLI) is not entered, CoPP will only work in software. CoPP is supported in hardware for unicast IPv4 and IPv6 traffic. It is not supported in hardware for multicast and broadcast traffic. However, CoPP software protection should still be used to mitigate multicast and broadcast DoS attacks targeted at the route processor CPU. Use software CoPP in conjunction with other hardware-based mechanisms such as multicast CPU rate limiters, access-list, and traffic storm control to provide robust protection against line-rate multicast and broadcast DoS attacks (see next section). Note also that hardware and software CoPP does not apply to ARP traffic. An ARP policing mechanism (see Routing Protocol and ARP Policing Mechanism section) should be applied instead.

The following is a configuration example of CoPP for the reporting class of traffic such as ICMP. It is first essential to use the "mls qos" CLI to allow hardware CoPP DoS mitigation. Access-list and class-maps should then be defined to match the ICMP traffic. The following CoPP policy-map limits reporting traffic to 100 Kbps. After the policy-map is applied to the control-plane interface with the **`service-policy input`** command, ICMP traffic to the route processor is limited to 100 Kbps in hardware.

```
Router(config)# mls qos

Router(config)# access-list 101 permit icmp any any

Router(config)# class-map reporting
Router(config-cmap)# match access-group 101

Router(config)# policy-map control-plane-policy
Router(config-pmap)# class reporting
Router(config-pmap-c)# police 100000 conform-action transmit exceed-action drop
```

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

```
Router(config)# control-plane
Router(config-cp)# service-policy input control-plane-policy
```

More platform-independent CoPP information can be found at:
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6642/prod_white_paper0900aecd804fa16a.shtml

Built-In "Special-Case" CPU Rate Limiters

In addition to CoPP, the Cisco Catalyst 6500 Series Supervisor Engine 32 and Supervisor Engine 720 support platform-specific, hardware-based rate limiters for special networking scenarios resembling DoS attacks. These hardware CPU rate limiters are called "special-case" rate limiters because they cover a specific predefined set of IPv4, IPv6, unicast, and multicast DoS scenarios. These DoS scenarios identify special cases where traffic needs to be processed by the switch processor or route processor CPU. Examples include multicast traffic for which a destination prefix cannot be found in the routing table, dropped traffic that needs to be processed by the CPU to send an ICMP unreachable back to the source, and special packet types that cannot be identified with an access list.

The special-case rate limiters do not provide the same level of granularity as CoPP and are thus especially useful for cases where hardware CoPP cannot be used to classify particular types of traffic. Such special packet types include packets with TTL equal to 1, packets that fail the maximum transmission unit (MTU) check, packets with IP Options, and IP packets with errors. Other examples of DoS scenarios not covered by CoPP include CPU protection against line-rate attacks using multicast packets, and switch processor CPU protection. CoPP and special-case rate limiters should be used in conjunction. However, the administrator should be aware that the special-case rate limiters will override the hardware CoPP policy for packets matching the rate limiters criteria. That is, as shown in Figure 4, if traffic matches a special-case rate limiter, it is never compared against the hardware CoPP policy. It will only be compared against the software CoPP policy. Therefore, it is strongly recommended to disable the "Cisco Express Forwarding Receive" rate limiter when using CoPP (see explanation for "Cisco Express Forwarding Receive" in Table 1). The Cisco Express Forwarding Receive rate limiter is disabled by default and can be easily disabled by issuing the `"no mls rate-limit unicast cef receive"` command.

Figure 4. Control-Plane Policing and Special-Case Rate Limiters Interaction: Special-Case Rate Limiters Override CoPP in Hardware

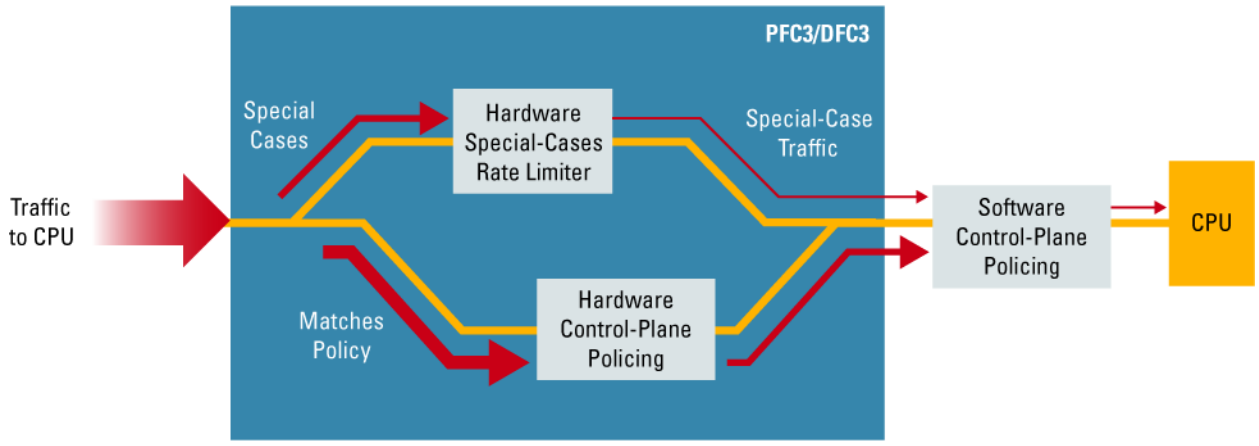

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

The Cisco Catalyst 6500 Series Supervisor Engine 32 and Supervisor Engine 720 forwarding engines provide 10 hardware registers to be used for special-case rate limiters. Eight of these registers are present in the Layer 3 forwarding engine and two of these registers are present in the Layer 2 forwarding engine. The registers are assigned on a first-come, first-serve basis and some rate limiters share one register. Should all registers be used, the only means to configure another special-case rate limiter is to free one register. It is recommended to use all ten special-case rate limiter hardware resources available. These are supported in all available Catalyst OS and Cisco IOS Software releases for the Supervisor Engine 720 and Supervisor Engine 32. However, some rate limiters have been added over time. An exhaustive list of special-case rate limiters can be obtained by issuing the **`show mls rate-limit`** command.

The following output shows all rate limiters available as of Cisco IOS Software Release 12.2(18)SXE1 (with default settings):

```
Router#sh mls rate-limit
 Sharing Codes: S - static, D – dynamic
 Codes dynamic sharing: H - owner (head) of the group, g - guest of the group
   Rate Limiter Type     Status       Packets/s   Burst  Sharing
 --------------------   ----------   ---------   -----  -------
         MCAST NON RPF   Off                -       -     -
        MCAST DFLT ADJ   On            100000     100  Not sharing
      MCAST DIRECT CON   Off                -       -     -
        ACL BRIDGED IN   Off                -       -     -
       ACL BRIDGED OUT   Off                -       -     -
           IP FEATURES   Off                -       -     -
          ACL VACL LOG   On              2000       1  Not sharing
           CEF RECEIVE   Off                -       -     -
             CEF GLEAN   Off                -       -     -
      MCAST PARTIAL SC   On            100000     100  Not sharing
        IP RPF FAILURE   On               100      10  Group:0 S
           TTL FAILURE   Off                -       -     -
 ICMP UNREAC. NO-ROUTE   On               100      10  Group:0 S
 ICMP UNREAC. ACL-DROP   On               100      10  Group:0 S
         ICMP REDIRECT   Off                -       -     -
           MTU FAILURE   Off                -       -     -
       MCAST IP OPTION   Off                -       -     -
       UCAST IP OPTION   Off                -       -     -
           LAYER_2 PDU   Off                -       -     -
            LAYER_2 PT   Off                -       -     -
             IP ERRORS   On                 100      10  Group:0 S
           CAPTURE PKT   Off                -       -     -
            MCAST IGMP   Off                -       -     -
 MCAST IPv6 DIRECT CON   Off                -       -     -
 MCAST IPv6 ROUTE CNTL   Off                -       -     -
 MCAST IPv6 *G M BRIDG   Off                -       -     -
  MCAST IPv6 SG BRIDGE   Off                -       -     -
  MCAST IPv6 DFLT DROP   Off                -       -     -
 MCAST IPv6 SECOND. DR   Off                -       -     -
```

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

```
  MCAST IPv6 *G BRIDGE   Off                  -       -     -
        MCAST IPv6 MLD   Off                  -       -     -
```

Tables 1–4 provide a list of hardware-based rate limiters available on the Cisco Catalyst 6500 Series Supervisor Engine 32 and Supervisor Engine 720. This list may evolve over time as new cases are identified and it does not include IPv6 multicast rate limiters for simplicity reasons. The tables break down features into the following categories: unicast (Table 1), multicast (Table 2), all (Table 3), and Layer 2 (Table 4) as is the case in the special-case rate limiter Cisco IOS CLI. The following example shows the special-case rate limiter CLI structure. It also shows how to configure the TTL failure rate limiter to limit TTL 1 packets going to the route processor CPU at 10 packets per second (pps). Just like hardware CoPP, the Layer 3 special-case rate limiters are applied on a per-forwarding engine basis. So packets failing the TTL check will be rate limited at 10 pps per forwarding engine and the aggregate rate presented to the CPU could thus be more than 10 pps. However, note that Layer 2 special-case rate limiters are applied globally.

```
Router(config)#mls rate-limit ?
  all        Rate Limiting for both Unicast and Multicast packets
  layer2     layer2 protocol cases
  multicast  Rate limiting for Multicast packets
  unicast    Rate limiting for Unicast packets
Router(config)#mls rate-limit all ttl-failure 10
```

The “unicast” rate limiter category cannot rate limit multicast traffic and vice-versa. The “all” category affects both unicast and multicast packets. The “layer2” rate limiter category uses Layer 2 hardware rate limiter resources (as opposed to other categories which share eight Layer 3 rate limiter registers). Layer 2 rate limiters do not work in the truncated fabric-switching mode: that is, they do not work in a system that has a mix of fabric-enabled and traditional, non-fabric-enabled line cards or if the system has classic line cards and redundant Cisco Catalyst 6500 Series Supervisor Engine 720s.

Tables 1–4 should help you understand which special-case rate limiter is right for your network. Highlighted special-case rate limiters should be used in most environments and you should not use those that do not apply to your environment. For example, if VACL log is not used in your network, do not enable the VACL log special-case rate limiter (Instead, you should disable it since it is enabled by default). This would waste a hardware rate limiter and would not serve any purpose.

Table 1. Unicast Special-Case Rate Limiters

Rate Limiter	Default	Description
Ingress ACL Bridged Packets	OFF	Limits packets punted to the route processor CPU because of an Ingress/Egress ACL bridge result. Ingress/Egress ACL rate limiters can be used independently. However, if both rate limiters are turned on, they must share the same value and are limited in aggregate. Examples of ACL bridged packets include packets hitting ACEs with the “log” keyword, packets requiring special ACL features, and unsupported hardware packet types such as IPX and AppleTalk. When a CoPP policy is present, it is better to disable the Ingress ACL bridged rate limiter because these may overlap.
Egress ACL Bridged Packets		
Cisco Express Forwarding Receive cases	OFF	Rate limits all packets that contain any route processor IP address as the destination address. It is better to use CoPP than to use this rate limiter; do not use both mechanisms in conjunction.
Cisco Express Forwarding Glean Cases	OFF	Limits traffic that is punted to the CPU when no ARP entry exists for the destination host, and the CPU thus needs to ARP for a next hop. Note that this does not affect ARP traffic but just traffic that requires address resolution.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Rate Limiter	Default	Description
IP Features	OFF	Limits traffic requiring special software feature processing. Software features include Authentication proxy, IPSec, and Inspection in the Native implementation. Note that the Cisco Catalyst OS feature rate limiter is used to rate limit dot1x, DHCP, and ARP Inspection traffic. Do not use this rate limiter unless you are using one of these features.
ICMP Redirect	OFF	Limits traffic requiring generation of ICMP redirect messages. ICMP redirect packets are sent back to the originating hosts to advertise optimal routes.
ICMP Unreachable—no route ICMP Unreachable—ACL drop IP Errors IP RPF Failures	ON 100 pps	Limits traffic requiring generation of ICMP unreachable messages due to FIB miss, ACL drop, or RPF failure. The IP Error rate limiter limits bad IP traffic with invalid length or checksums that needs to be sent to the route processor for further processing. All four rate limiters share a single hardware rate limiter and should any of these rate limiters be enabled, all the rate limiters in this group will share the same value and state (ON/ON/ON/ON).
IP Options	OFF	Packets with IP Options are sent to the CPU for further processing. This special-case rate limiter limits unicast packets with IP Options punted to the CPU. This option is only supported on the Cisco Catalyst 6500 Series policy feature cards 3B and 3BXL (PFC3B, PFC3BXL).
VACL Log	ON 2000 pps	Limits packets punted to the CPU because of VLAN ACL logging. VACLs are processed in hardware, but the logging function is done by the route processor. This rate limiter is enabled by default. Do not use this rate limiter unless VACL Log is configured.
Capture Packet	OFF	Limits packets punted to the CPU because of Optimized ACL Logging (OAL).

Table 2. Multicast Special-Case Rate Limiters

Rate Limiter	Default	Description
Multicast Partial-SC	ON 100000 pps	Some multicast flows can be partially software switched if special processing is required. It is desirable to rate limit these flows destined to the Multilayer Switching Feature Card (MSFC). Note that this rate limiter uses a special register that is not accounted for in the available ten hardware registers and it is applied globally, not on a per-forwarding-engine basis.
Multicast Default Adjacency	ON 100000 pps	Limits traffic requiring special software processing because of an FIB miss (for example, if multicast traffic does not match an entry in the hardware mroute table).
Multicast Non-RPF	OFF	Same as unicast RPF Failure rate limiter, but applies to multicast traffic.
Multicast IP Options	OFF	Same as unicast IP Options rate limiter, but applies to multicast traffic. This option is only supported on the Cisco Catalyst 6500 Series PFC3B and PFC3BXL.

Table 3. Unicast and Multicast (“All”) Special-Case Rate Limiters

Rate Limiter	Default	Description
TTL Failure	OFF	Limits packets punted to the route processor CPU because of a time-to-live (TTL) check failure. Do not use this rate limiter in conjunction with Layer 2 multicast in a system with the Cisco Catalyst 6500 Series PFC3A. Doing so would break multicast bridging.
MTU Failure	OFF	Limits packets that are punted to the route processor CPU because of to maximum transmission unit (MTU) failure. Do not use this rate limiter in conjunction with Layer 2 multicast in a system with the Cisco Catalyst 6500 Series PFC3A.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Table 4. Layer 2 Special-Case Rate Limiters

Rate Limiter	Default	Description
L2 PDU	OFF	Limits Layer 2 control protocol data unit (PDU) packets destined for the switch processor. That includes bridge protocol data units (BPDUs), VTP, UDLD, PagP, and LACP. Being too aggressive with the rate limiter could have adverse effects on the Layer 2 network stability. This rate limiter is not supported when the switch is in the truncated fabric-switching mode*.
L2 Protocol Tunneling	OFF	Limits Layer 2 protocol tunneled PDUs destined for the switch processor (Cisco Discovery Protocol, Spanning Tree Protocol, and VTP). This rate limiter is not supported when the switch is in the truncated fabric-switching mode (See footnote).
Multicast IGMP	OFF	Limits IGMP control messages sent to the CPU for IGMP snooping. This rate limiter should be used when IGMP Snooping is enabled: a switch with IGMP Snooping enabled will listen for IGMP messages to optimize the flow of multicast traffic at Layer 2. This rate limiter is not supported when the switch is in the truncated fabric-switching (See footnote).

* Mix of fabric enabled and classic/non-fabric enabled linecards or if the system has classic linecards and redundant Cisco Catalyst 6500 Series Supervisor Engine 720s.

Selective Packet Discard

Selective Packet Discard (SPD) is a software mechanism to manage the process-level input queues on the route processor. The goal of SPD is to provide priority to routing protocol packets and other critical control packets (high-precedence packets with IP precedence set to 6 or 7) during periods of process-level queue congestion. SPD provides special headroom in the process-level queue for these important packets. SPD is enabled by default on the Cisco Catalyst 6500 Series. The default process-level queue size is 74 and the default headroom size that accommodates special high-priority control-plane packets is 100 on the Catalyst 6500 Series. Where software CoPP operates at interrupt level to prioritize important traffic, SPD adds another level of software prioritization at the process level.

For more information on Selective Packet Discard, please visit
http://www.cisco.com/en/US/products/hw/routers/ps167/products_tech_note09186a008012fb87.shtml

ACLs and QoS

Access-list and QoS features are covered in detail in the data plane DoS protection sections. These features should be used for DoS mitigation cases that cannot be managed by CoPP and CPU rate limiters. For example, a MAC ACL can be used to drop non-IP traffic.

Routing Protocol and ARP Policing Mechanisms

Malicious users may try to overwhelm the route processor CPU with control packets such as routing protocol or ARP packets. These special control packets can be hardware rate limited using a Cisco Catalyst 6500 Series specific routing protocol and ARP policing mechanism configurable with the **mls qos protocol** command. The routing protocols supported include RIP, BGP, LDP, OSPF, IS-IS, IGRP, and EIGRP. For example, the following rate limits ARP packets in hardware at 32,000 bps.

```
Router(config)#mls qos protocol arp police 32000
```

Note, however, that although this throttling mechanism effectively protects the route processor CPU against such attacks as line-rate ARP attacks, this mechanism does not only police routing protocols and ARP packets to the switch. It also polices traffic through the box and is less granular than CoPP.

This policing mechanism shares the root configuration with a policing-avoidance mechanism. The policing-avoidance mechanism lets the aforementioned routing protocol and ARP packets flow through the network when these particular packets reach a QoS policer. This function can be configured using the **mls qos protocol <protocol> pass-through** command.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Traffic Storm Control

A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. The traffic storm control feature prevents LAN ports from being disrupted by a broadcast, multicast, or unicast traffic storm on physical interfaces from either network configuration mistakes, or from malicious users issuing a DoS attack. Traffic storm control (also called broadcast suppression) monitors incoming traffic levels over a 1-second traffic storm control interval and, during the interval, compares the traffic level with the traffic storm control level configured. The traffic storm control level is a percentage of the total available bandwidth of the port. For example, configuring broadcast suppression on interface Gigabit 1/1 to a level of 20 percent will limit all broadcast control-plane traffic such as ARP requests, DHCP requests, and other broadcast traffic to a combined amount of 200 Mbps on interface Gigabit 1/1. This is achieved using the following command:

```
Router(config)#int GigabitEthernet 1/1
Router(config-if)# storm-control broadcast level 20
```

The Cisco Catalyst 6500 Series applies traffic storm control at the port ASIC level. Thus, unicast, multicast, and broadcast suppression is a property of the line card itself, not the supervisor engine. Each port has a single traffic storm control level that is used for all types of traffic (unicast, multicast, and broadcast). Though broadcast suppression does not differentiate between good and bad broadcast traffic, it can be configured to combat effectively against line-rate broadcast DoS attacks.

Directed Broadcast

IP directed broadcast traffic is traffic directed to an IP address that specifies “all hosts” on a given subnet. A single copy of a directed broadcast address is routed to the target network where it is eventually broadcasted. IP directed broadcast packets are dropped by default on the Cisco Catalyst 6500 Series, so that DoS attacks cannot exploit this vulnerability. For example, “smurf” attacks exploit the use of ICMP directed broadcast packets. In this attack, the perpetrator sends an IP ping (or “echo my message back to me”) request to a broadcast address within the receiving site. The ping packet is broadcast to all hosts within the receiving site's local network. The packet contains a “spoofed” source address, which is the intended address of the recipient of this DoS attack. Each host that receives the ping will reply to the spoofed source address. The result will be lots of ping replies flooding back to the innocent, spoofed host. If the flood is great enough, the spoofed host will no longer be able to receive or distinguish real traffic.

If needed, IP directed broadcast can be enabled on a particular VLAN or physical interface with the interface-level **ip directed-broadcast** command. When enabled, IP directed-broadcast packets are managed at the process level by the route processor in the default mode. To enable the hardware switching of the IP directed broadcast, use the **mls ip directed-broadcast** command. ACLs should also be deployed to provide further security. The ACLs can be configured to only allow the expected sources of the directed broadcast.

Cisco Express Forwarding

The Cisco Catalyst 6500 Series Supervisor Engine 2, Supervisor Engine 32, and Supervisor 720 use the Cisco Express Forwarding technology to help ensure that hardware routing tables are populated based on topology information before any user traffic flows through the switch. This architecture helps ensure that data packets are processed in hardware.

The Cisco Express Forwarding-based architecture is superior to flow- or cache-based architectures where the first few packets of all new flows have to be forwarded to the CPU until a hardware cache entry is created. This is especially true when dealing with certain worms. A worm is a self-contained program that can propagate itself through systems or networks. One spreading strategy is based on random scanning, where the worm operates by selecting random IP addresses to find vulnerable hosts to infect. On flow-based systems, scanning random address spaces is CPU-intensive and could easily result in flow cache overflow. When the flow cache is full, all new flows have to be software-switched. This is in contrast with Cisco Express Forwarding switching, where routing decisions are topology-based and thus do not depend on the number of existing switched flows.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

DOS ATTACK DETECTION

While mitigation mechanisms are needed, it is also important to make sure that data can be collected to identify these attacks and that mechanisms are in place to trigger appropriate actions when detecting anomalous behavior. NetFlow, SNMP, Capture, and SPAN are amongst the most important DoS attack detection mechanisms.

NetFlow

NetFlow is one of the most efficient ways to detect DoS attacks. It is a technology that provides visibility into the traffic traversing a given switch. The NetFlow data can provide flow information such as source and destination IP address, protocol, source and destination port, and incoming interface. It also provides statistics on the number of packets seen for that given flow. NetFlow could thus be used to detect DoS and DDoS attacks, trace back attacks, and to characterize the switch's traffic profile. NetFlow records can be sent periodically to NetFlow collectors which can store and analyze the NetFlow data.

The Cisco Catalyst 6500 Series supports NetFlow in hardware. The Cisco Catalyst 6500 Series Policy Feature Card 3A and 3B (PFC3A and PFC3B) can keep track of up to 128,000 flows at a given time in hardware. The PFC3BXL can scale to up to 256,000 recorded flows. On the Cisco Catalyst 6500 Series, the hardware NetFlow table can also be used for feature hardware acceleration.

SNMP

It is important to monitor the critical switch resources to make sure that they are not overwhelmed, because this could indicate a DoS attack. For example, the CPU utilization could be monitored and SNMP traps could be sent to a network management server as an early warning sign of a possible DoS attack. Other critical resources to watch include the memory utilization, the link utilization, drop counters, ternary content addressable memory (TCAM) resource utilization, the total number of flows traversing the switch, and syslog messages. These resources could be monitored and retrieved on a central management station through SNMP, and SNMP traps or other proper action should be taken when abnormal behaviors are seen. The Cisco Catalyst 6500 Series supports SNMP Version 1 (SNMPv1), SNMPv2, and SNMPv3 (authentication and encryption).

VACL Capture and Switched Port Analyzer

Should a possible DoS attack be detected, it is possible to capture complete traffic flows and send these packets to a sniffer or a special appliance for complete packet payload analysis. The two features providing this port mirroring ability are Switched Port Analyzer (SPAN) and VLAN ACL (VACL) Capture. The SPAN function provides the capability to send a copy of all traffic traversing a given physical port or VLAN to a sniffer. Different types of SPAN exist: SPAN is local to the switch, Remote SPAN (RSPAN) sends a copy of the traffic to a remote Layer 2 destination, and Encapsulated Remote SPAN (ERSPAN) sends a copy of the traffic to a remote Layer 3 destination using generic routing encapsulation (GRE). VACL Capture works locally on a given switch and is applied to a VLAN or a set of VLANs. One advantage of using VACL capture over SPAN is that ACEs can be used to specify which exact traffic should be sent to a sniffer. VACL capture presents only relevant traffic to the sniffer.

SPECIALIZED DENIAL OF SERVICE PROTECTION

This section will highlight some characteristics of the main services modules that provide advanced DoS protection. Services modules can offer the same level of security as separate security appliances with the added benefit of providing an integrated solution that takes advantage of the underlying Cisco Catalyst 6500 Series infrastructure.

Firewall Services Module

The Cisco Catalyst 6500 Series Firewall Services Module (FWSM) is a high-speed, application-aware firewall that provides stateful inspection and advanced DoS protection with features such as Flood Guard, FragGuard, DNS Guard, and TCP Intercept. The FWSM can be deployed in both transparent (Layer 2) or routed (Layer 3) mode and can be virtualized into multiple logical firewalls, with each security context having its own security policies and administration rules.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Cisco Traffic Anomaly Detector Module and Cisco Guard Module

The Cisco Guard and Traffic Anomaly Detector modules are used to monitor traffic flows and mitigate DDoS attacks that can flood networks and servers with unneeded traffic. The Cisco Traffic Anomaly Detector passively monitors the network traffic looking for abnormal behavior that could indicate DDoS attacks. When an attack is identified, the Cisco Traffic Anomaly Detector alerts the Cisco Guard, providing information on which destination is being attacked. All the traffic destined to that destination, both legitimate and malicious, is dynamically diverted to the Cisco Guard and is then subject to multiple levels of analysis and counter measures resulting in the removal of malicious traffic and the forwarding of legitimate traffic to the destination.

Intrusion Detection System Services Module

The Cisco Catalyst 6500 Series Intrusion Detection System Services Module (IDSM) monitors the network traffic and compares the packet payload against well-known attack signatures. Should an attack be detected, the IDSM can instruct the switch to take proper action against these attack packets such as dropping the DoS packets with an access list. The Cisco Catalyst 6500 Series IDSM-2 is part of the Cisco Intrusion Prevention System, providing the ability to analyze, identify, classify, and stop threats including worms, spyware and adware, network viruses, and application abuse.

Network Analysis Module

The Cisco Network Analysis Module (NAM) for the Catalyst 6500 Series gives network managers visibility into the network by providing integrated traffic-monitoring services. Such services include capture capability, switch health parameters, port-level, application-level, hosts, and conversations traffic statistics, and network-based services such as QoS and voice over IP (VoIP). The Cisco NAM can be used to detect DoS attacks or to profile an attack after it has been detected.

© 2005 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

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 International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

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, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • 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 • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R)

205309.BB_ETMG_CC_7.05

Printed in the USA