Safe and Sound:
High Availability and Security for Cisco Catalyst 4500 Series Switches

prepared for Cisco Systems
October 2004

CONTENTS

ILLUSTRATIONS

Executive Summary

High availability ranks among the top network infrastructure requirements – more so than price. There's good reason for this kind of thinking: Highly available networks prevent losses in productivity and revenue. In fact, high-availability networking is the foundation of business resilience.

A recent study by Infonetics Research makes clear the importance of resilient networks. When asked to name their top requirements for Internet infrastructure, network managers rated high availability ahead of nearly all other factors[1]. Figure 1 below presents results from the Infonetics study.

Figure 1: Key Factors for Network Infrastructure

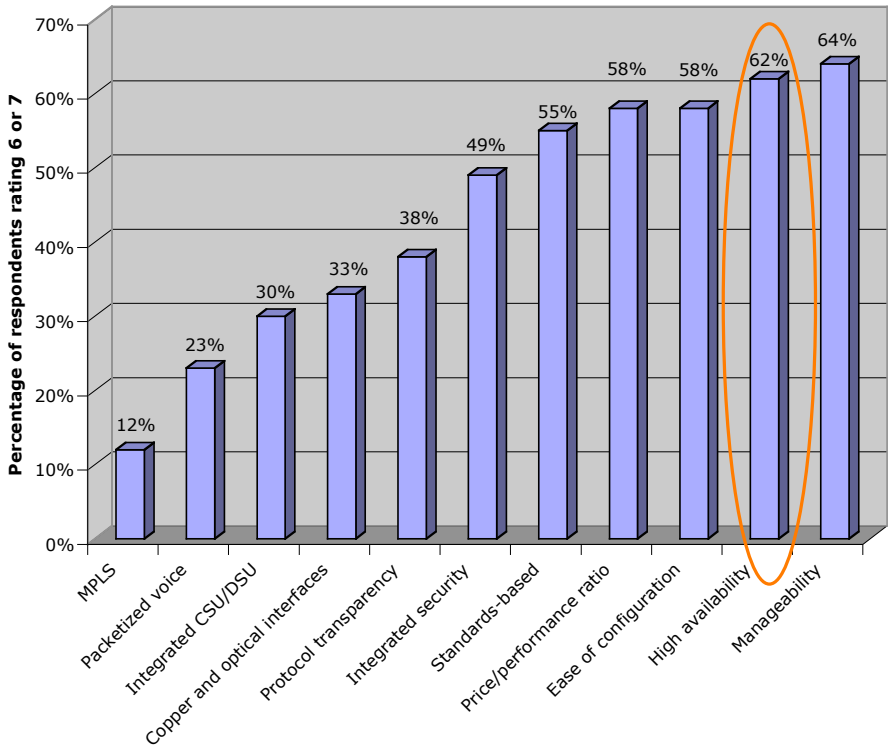

Cisco Systems is addressing the need for resilient network infrastructure by adding many new high availability and security features to Cisco IOS software for Catalyst 4500 series switches. The new software enhances reliability and security for enterprise and service provider edge customers. The new Catalyst 4500 software enhances reliability and security while continuing to deliver wire-rate performance in edge deployments[2].

---

[1] Infonetics Research, *User Plans for WAN and Internet Access*, US/Canada, 2003.

[2] The Supervisor Engines II-Plus and IV for the Catalyst 4507R are rated at 48 million packets per second (mpps) and the Supervisor Engine V for the Catalyst 4507R/4510R is rated at 72 mpps.

Cisco commissioned Opus One, an independent networking consultancy, to assess the effectiveness of Cisco's new resiliency and security mechanisms. In addition to conducting performance tests, Opus One launched various network-based attacks against the Catalyst 4500 to test the new security mechanisms. The tests were intended to verify that security policies remained intact before, during, and after the failure of an active Supervisor card.

Among the key findings of Opus One's tests:

- In the event of a Supervisor Engine failure, Cisco's Stateful Switchover (SSO) redundancy mechanism reduces recovery time from approximately 60 seconds to less than 50 milliseconds (0.05 seconds), a 1,200-fold improvement.

- With SSO, there is no perceived degradation in VOIP or video application performance in the event of a Supervisor Engine failure.

- SSO is equally effective in reducing failover time for unicast and multicast traffic.

- Cisco's Dynamic ARP Inspection security mechanism blocks potential attackers from performing man-in-the-middle (MITM) attacks. MITM attacks can redirect and capture users' passwords, IP phone calls, and other sensitive traffic.

- DHCP snooping secures DHCP transactions by locking out rogue DHCP servers and thwarts denial-of-service attacks by rate-limiting DHCP packets.

- Cisco's IP Source Guard security mechanism automatically prevents potential attackers from using spoofed source IP addresses.

This report is organized as follows. This executive summary introduces the tests performed. The "SSO failover" section describes the performance tests Opus One conducted to measure recovery times. The section "Securing the Wiring Closet" introduces three new security mechanisms and discusses tests for each: DHCP Snooping, IP Source Guard, and Dynamic ARP Inspection. Then the section "Security and Failover Mechanisms Working Together" presents results from a combined test involving SSO and all three security mechanisms working together while the Catalyst 4500 handles a heavy load of voice, video, and data traffic, including both unicast and multicast flows.

SSO Failover

With a growing number of enterprises using delay-sensitive applications such as voice and video, minimizing downtime has become more critical than ever. With these time-sensitive applications, even very brief outages can adversely affect performance or even lead to session loss.

With Cisco's Stateful Switchover (SSO), Catalyst 4500 series switches can continue forwarding traffic despite a hardware or software failure in a Supervisor Engine card.

SSO works by synchronizing layer-2 forwarding tables between redundant Supervisor cards in the same chassis. If one card fails, the other continues to forward traffic.

Cisco claims SSO failover time is less than 1 second, a dramatic reduction from the failover times of approximately 60 seconds using other resiliency mechanisms.

To validate this claim, Cisco asked Opus One to compare failover times using SSO with those using Redundant Route Processor (RPR), an existing method of transferring control between redundant Supervisor cards.

Figure 2 below illustrates the test bed we used for the failover tests. We attached a Spirent SmartBits traffic generator/analyzer to a Catalyst 4507R switch equipped with redundant Supervisor Engine IV cards. The switch also housed copper gigabit Ethernet line cards, to which we attached 8 SmartBits test interfaces.

Figure 2: The SSO Test Bed

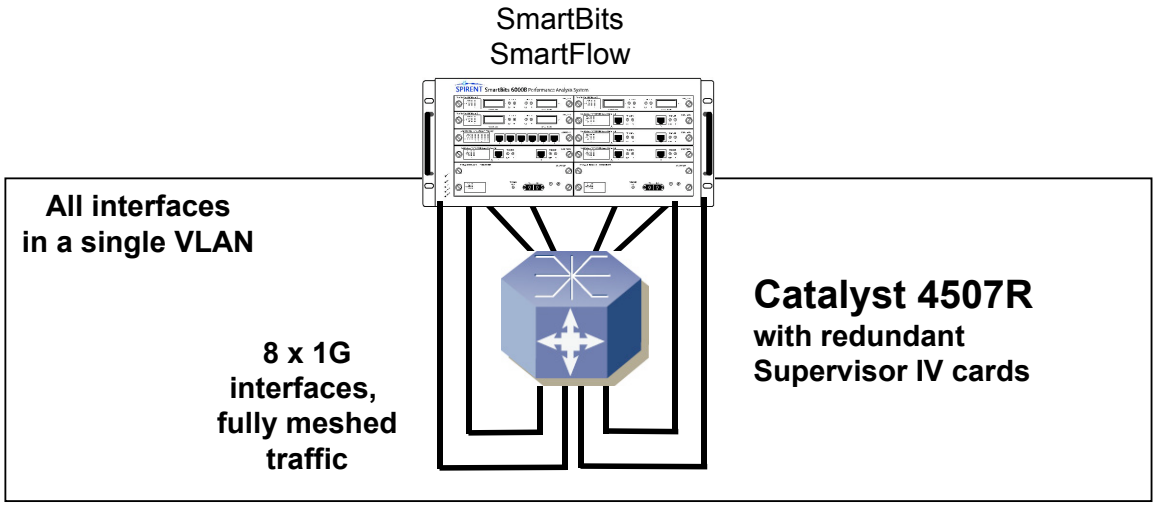

We configured the SmartFlow application for the SmartBits to offer minimum-sized IP packets[3] to all 8 interfaces at line rate in a fully meshed pattern (meaning every SmartBits interface generated traffic destined to every other interface). Short packets, heavy loads, and fully meshed traffic patterns are among the most stressful conditions a switch must handle.

We began with a baseline test in which we offered traffic for 100 seconds without forcing a Supervisor card failure. The goal of this baseline test was simply to verify the Catalyst 4500 forwarded all packets with zero loss, which it did.

Then we configured the Catalyst 4507R to use Redundant Route Processor Plus (RPR), a resiliency method known to deliver failover times in the neighborhood of 60 seconds. About 10 seconds into our 100-second test duration, we physically removed the active Supervisor card, forcing a failover to the standby Supervisor. Based on packet loss statistics, we observed a failover time of 62.98 seconds using RPR.

Next, we configured the Catalyst 4507R to use SSO instead of RPR and conducted three trials of the failover test. Each time, we physically removed the active Supervisor card from the switch approximately 10 seconds into the 100-second test duration. Once the test was complete, we calculated failover time using packet loss measurements from the SmartBits.

Over three trials, the average failover time using SSO was 49.884 milliseconds, or approximately 1/20 of a second.

That represents just 0.08 percent of the RPR failover time, or better than a 1,200-fold improvement in recovery time.

Figure 3 below summarizes results from the SSO failover tests.

[3] In this case, the packets were 68 bytes in length, including Ethernet header, VLAN tag, and Ethernet CRC.

Figure 3: SSO Failover Times on the Catalyst 4500

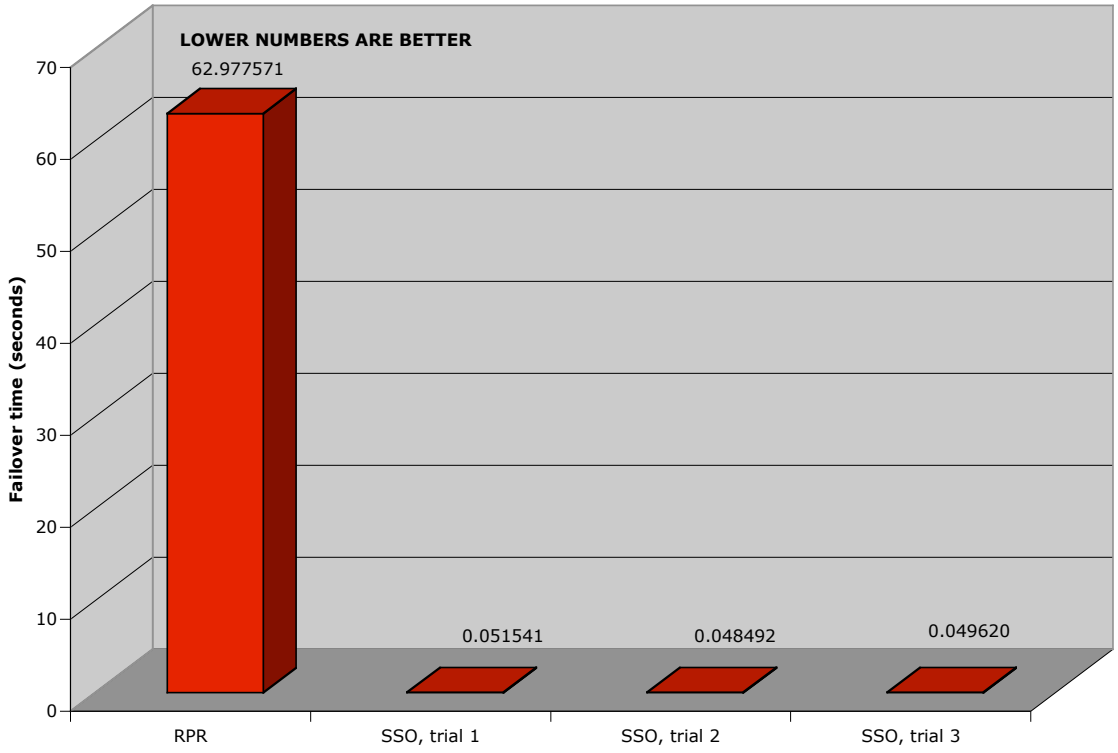

SSO vastly reduces failover time in the event of a Supervisor failure. Even at gigabit speeds under highly stressful test conditions – short packets, heavy load, and fully meshed traffic patterns – SSO greatly reduces recovery times for Catalyst 4500 series switches.

Note also that SSO failover test results are consistent across multiple trials. Since the Catalysts synchronize layer-2 forwarding state across Supervisor IV cards, there is no additional time needed to populate a new layer-2 forwarding table after the failure of an active Supervisor card.

Securing the Wiring Closet

While component failure such as loss of a Supervisor card can affect network uptime, an even more serious threat is downtime caused by various forms of attack. This is an often overlooked aspect of high availability: Even if the *network* is available thanks to mechanisms like SSO, the *services* delivered by that network may remain vulnerable due to denial-of-service attacks, man-in-the-middle-attacks, and other threats.

Cisco IOS protects the network with numerous security mechanisms for Catalyst 4500 series switches. For this project, we tested three such mechanisms: DHCP Snooping, IP Source Guard, and Dynamic ARP Inspection.[4]

Figure 4 below shows the test bed we used to validate the effectiveness of these mechanisms. A Catalyst 4507R switch is connected to a Catalyst 6500 switch. The Catalyst 4507R is configured in "layer-2+ mode," meaning that although it performs conventional layer-2 switching it also inspects upper-layer fields in the packets it receives. The Catalyst 6500 is configured as a layer-3 router.

Figure 4: Catalyst 4500 Security Test Bed

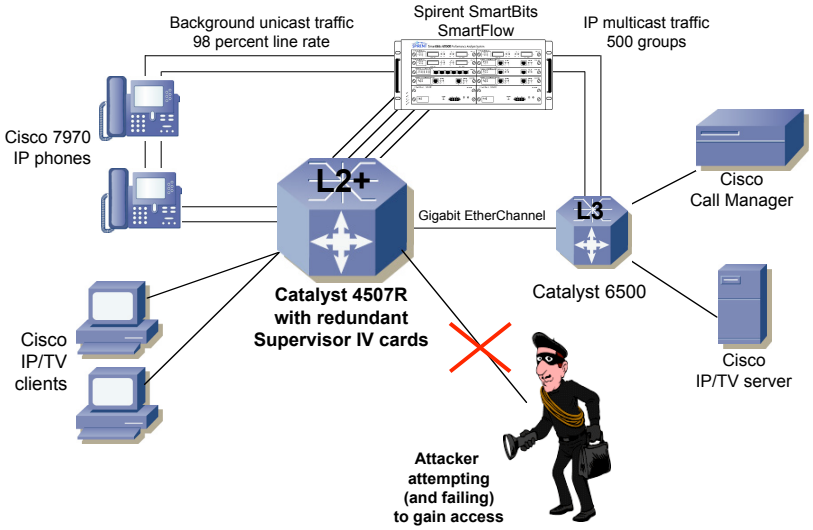

Note the presence of both unicast and multicast traffic flowing through the Catalyst 4507R. For the unicast traffic, Cisco 7970 series IP phones conduct voice calls while a

[4] For more information, see "Catalyst 4500 Security Features Best Practices for Supervisors," available at http://www.cisco.com/en/US/products/hw/switches/ps4324/products_white_paper09186a00801faa79.shtml .

Spirent SmartBits simultaneously loads the network at 98 percent of line rate. A Cisco Call Manager attached to the Catalyst 6500 handles call setup and teardown functions.

For the IP multicast traffic, a Cisco IP/TV server attached to the Catalyst 6500 generates streaming video feeds to IP/TV clients attached tot he Catalyst 4507R. At the same time, the SmartBits transmits IP multicast traffic to 500 IGMPv2 groups, again utilizing 98 percent of each interface's capacity.

We chose heavy unicast and multicast loads to demonstrate two things: First, the security mechanisms we tested continued to operate regardless of network utilization. Second, security mechanisms continued to operate during and after a loss of a Supervisor card because they work together with SSO to replicate information about authorized hosts across redundant Supervisor cards.

We validated that each security mechanism thwarted actual attacks both during and after a Supervisor card failure. We now discuss each security mechanism in turn.

DHCP Snooping

The Dynamic Host Configuration Protocol (DHCP) greatly simplifies IP address management, but it is inherently insecure[5]. Because DHCP lacks authentication and encryption capabilities, a rogue DHCP server in an enterprise network may hand out unauthorized IP addresses and default gateways. In a metro Ethernet setting, users may unwittingly deploy cable or DSL routers with unauthorized DHCP services enabled. And attackers in any setting may swamp legitimate DHCP servers by flooding the network with DHCP requests.

The Catalyst 4500 secures the DHCP protocol from all these threats. By "snooping" on DHCP traffic in much the same way IGMP snooping monitors multicast traffic, the Catalyst 4500 will drop unauthorized replies to DHCP requests. DHCP Snooping eliminates the problem of rogue or unauthorized DHCP servers. Further, DHCP Snooping allows rate-limiting of DHCP traffic, thus preventing an attacker from flooding the network with DHCP traffic.

The Catalyst 4500 also can employ DHCP option 82 to insert information about itself in request packets destined for the DHCP server. Using both DHCP Snooping and option 82 data, the Catalyst 4500 builds a table that binds each switch port to authorized IP and MAC addresses as well as VLAN and DHCP lease information. Because it maintains state on DHCP bindings, the Catalyst 4500 ensures that only legitimate DHCP packets are forwarded.

DHCP Snooping works not only on individual switch ports but also on EtherChannels, 802.1Q trunks, and private VLANs.

[5] RFC 2131 describes DHCP for IPv4, while RFC 3315 describes DHCP for IPv6.

To validate the correct operation of DHCP Snooping, we attached a rogue DHCP server to the Catalyst 4507R. The rogue server was configured to hand out incorrect addresses and other parameters. This failed; the switch forwarded DHCP responses only from the legitimate DHCP server and not the rogue server.

Figure 5 below shows an excerpt from a DHCP Snooping table. Note that the table binds an IP address and other information to a switch port. Thus, a rogue server cannot attempt to supply the same address on a different switch port.

Figure 5: DHCP Snooping Binding Table

```
Switch# show ip dhcp snooping binding
MacAddress     IP Address   Lease (seconds)   Type        VLAN      Interface
-----------    ----------  ----------------  -----       -----     ------------
011.92ba.19ec  11.1.1.8    1600              dynamic     11        GigabitEthernet4/8
```

IP Source Guard

DHCP Snooping secures the switch from rogue DHCP responses, but it does not prevent an attacker from spoofing a legitimate user's IP and MAC address. That is where IP Source Guard comes in: It dynamically secures all IP addresses, regardless of whether they are allocated through dynamically or statically.

IP Source Guard requires DHCP Snooping and uses its binding table to populate a separate source IP binding table. With this feature, a legitimate IP address is bound to a specific switch port. If an attacker attempts to spoof the same address on a different port, the Catalyst 4500 will drop the spoofed packets.

IP Source Guard works not only on individual switch ports but also on EtherChannels, 802.1Q trunks, and private VLANs.

In our test, we attempted to generate packets from a spoofed IP address. With IP Source Guard disabled, the Catalyst 4507R forwarded such packets; however, the Catalyst 4507R blocked the traffic once we enabled IP Source Guard.

Figure 6 below shows a sample entry from the IP Source Guard binding table. Since Catalyst 4500 series switches bind IP addresses (and other information) with switch ports, no spoofing of source addresses is possible.

Figure 6: IP Source Guard Table Entry

```
ip source binding 0001.0003.0034 vlan 10 10.1.3.34 interface Gi7/4
```

Dynamic ARP Inspection

The address resolution protocol (ARP) allows the use of "gratuitous ARP," in which a host sends an ARP response message containing its IP and MAC addresses without first being prompted by an ARP request.[6] An attacker can use gratuitous ARP messages to poison a switch's ARP cache. As a result of this "man-in-the-middle attack" (MITM), traffic to and from a legitimate user's IP address will be redirected through the attacker's machine, allowing the attacker to capture cleartext passwords (commonly used in email), entire VOIP phone calls, and any and all other traffic. As our test results demonstrated, host operating systems typically are not aware of this man-in-the-middle attack.

Figure 7 below illustrates the danger posed by an MITM attack. By sending a gratuitous ARP to an unprotected switch, an attacker can cause traffic to be redirected, capturing sensitive data such as passwords.

Figure 7: MITM Attack in Progress

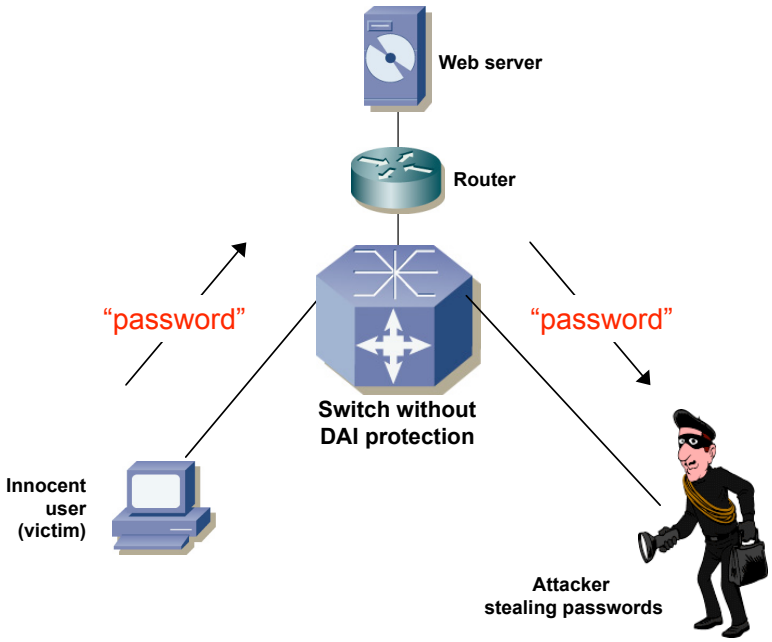

Catalyst 4500 series switches protect against gratuitous ARP attacks through the use of Dynamic ARP Inspection (DAI), which intercepts and examines every ARP packet in

[6] RFC 826 describes ARP.

every VLAN. If an attacker sends a gratuitous ARP message, a Catalyst 4500 series switch configured with DAI will not forward the message.

DAI works not only on individual switch ports but also on EtherChannels and 802.1Q trunks. Since VLANs may span multiple switches and since ARP is a different protocol from DHCP, DAI maintains a separate table of IP-MAC address bindings, distinct from the DHCP Snooping table.

To demonstrate the effectiveness of DAI, we began with the mechanism disabled to show an attack at work. First, an authorized host received the IP address 11.1.1.9 via DHCP. Then we offered a gratuitous ARP message from a rogue host using the same MAC address as the authorized host, as shown in Figure 8 below.

Figure 8: Successful Gratuitous ARP Attack

```
Switch#show ip arp vlan 11 | include 5ee6
Internet  11.1.1.9                0   0040.ca16.5ee6  ARPA   Vlan11
Internet  11.1.1.6                0   0040.ca16.5ee6  ARPA   Vlan11
```

Note that both the legitimate client's IP address (11.1.1.9) and the attacker's IP address (11.1.1.6) use the same MAC address. As a result, any traffic exchanged by the legitimate client will be copied to the rogue client.

We then went on to exploit the now-poisoned ARP cache by capturing passwords sent by the legitimate user. Figure 9 below shows the output produced by dsniff, a widely available security auditing tool.

Figure 9: Capturing Cleartext Passwords

```
08/31/04 10:30:12 tcp 11.1.1.6.1566 -> 11.1.1.1.80 (http)
GET / HTTP/1.1
Host: 11.1.1.1
Authorization: Basic cm9vdDpoYWNrbTM= [root:mypassword]

-----------------
08/31/04 10:30:24 tcp 11.1.1.6.1567 -> 11.1.1.1.80 (http)
GET / HTTP/1.1
Host: 11.1.1.1
Authorization: Basic cm9vdDpnYWdhZ2E= [root:otherpass]

-----------------
08/31/04 10:30:36 tcp 11.1.1.6.1569 -> 11.1.1.1.80 (http)
GET / HTTP/1.1
Host: 11.1.1.1
Authorization: Basic cm9vdDpnYWdhZ2E= [root:H4ckm3]
```

In this case, the attacker watched as a legitimate user (logged in under the root account) attempted to use three passwords to gain access to a Web site: "mypassword," "otherpass," and "H4ckm3." In the same manner, the attacker can record calls from Cisco IP phones as well as traffic from virtually any other application.

After clearing the ARP cache on the Catalyst 4507R, we enabled DAI and attempted the attack once again. This time the attack failed: The only entry with the MAC address in question was that of the legitimate user, as shown in Figure 10 below.

Figure 10: Blocking a Gratuitous ARP Attack

```
Switch#show ip arp vlan 11 | include 5ee6
Internet  11.1.1.9               0   0040.ca16.5ee6  ARPA   Vlan11
```

Since the attacker could not gain access to the network through gratuitous ARP, no further monitoring of legitimate users' traffic was possible.

Security and Failover Mechanisms Working Together

After verifying that all security mechanisms worked properly, we forced a system failure by removing the active Supervisor card, fan tray, and one power supply. The goal of this test was not only to measure recovery time, but also to verify that the Catalyst 4507R remained invulnerable to attack both during and after the component failures.

Before offering data traffic or forcing a failover, we started IP phone and IP/TV sessions. Once these were active, we offered unicast and multicast traffic at 98 percent of gigabit Ethernet line, thus heavily loading the network.[7]

We also configured the Catalyst 4507R to use DHCP Snooping, IP Source Guard, and Dynamic ARP Inspection, and one additional security mechanism: Port security, the mechanism whereby Catalyst 4500 series switches learn the MAC addresses of attached hosts and maintain state across reboots or switchovers. So-called sticky port security prevents MAC address spoofing at all times. Figure 11 below illustrates the logical overlay of all the security mechanisms working together.

Figure 11: Security Mechanisms Working Together

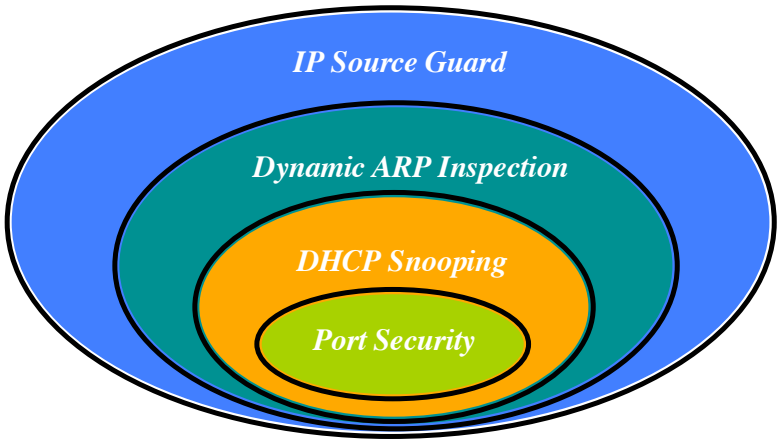

[7] With the benefit of hindsight, we probably could have raised the offered load from the SmartBits even higher, to 99 percent utilization, and still had some bandwidth left over for voice and video traffic. However, since this was not a forwarding rate test the difference between 98 percent and 99 percent utilization is not meaningful.

Once these mechanisms were in place, we launched custom-developed attack scripts to run continuously throughout the test duration.

After generating a heavy load of unicast and multicast traffic from the SmartBits, we then forced a failover by removing the active Supervisor and other system components.

We made several observations in this test:

- Video traffic was not perceptibly affected by the failover
- Voice traffic (music on hold) was not perceptibly affected by the failover
- Man-in-the-middle and IP spoofing attacks failed before, during, and after the sub-second switchover; in other words, the network was never vulnerable to these attacks

Opus One observed average failover times of 24 milliseconds (0.024 seconds) for unicast and multicast traffic alike. The similarity between unicast and multicast failover times demonstrates that there is no extra performance cost to handling multicast, even when the network is heavily utilized and hundreds of multicast groups are involved.

Figure 12 below shows failover times for three trials.

Figure 12: Failover Times With System Under Attack

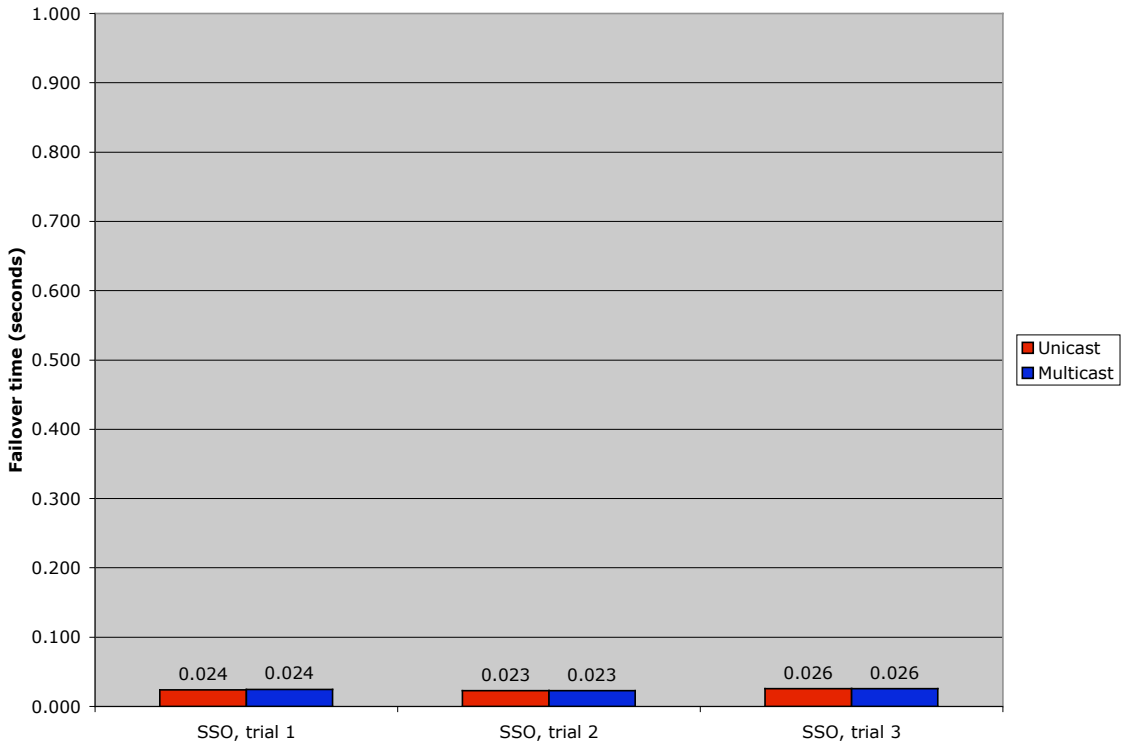

Conclusion

These tests have demonstrated new levels of availability and security in IOS running on Catalyst 4500 series switches. In these tests, Opus One has validated Cisco's claims of vastly improved network resilience and resistance to attack.

To recap the major findings from these tests:

- SSO failover time of approximately 50 *milliseconds* or less for both unicast and multicast traffic, compared with approximately 60 *seconds* with RPR
- No compromise in security during or after active Supervisor card failure in the face of DHCP spoofing attacks
- No compromise in security during or after active Supervisor card failure in the face of source IP spoofing attacks
- No compromise in security during or after active Supervisor card failure in the face of gratuitous ARP attacks
- No perceptible effect on VOIP or video application performance during or after Supervisor card failure

With these new redundancy and security mechanisms in place, Cisco Systems has significantly improved performance and security in two areas often cited as critical to network managers: High availability and security.

Acknowledgements

Opus One gratefully acknowledges the support of Spirent Communications, which supplied engineering assistance for this project. Spirent test engineers Mark Hall, Gary Hansen, and Brooks Hickman assisted with configuration of Spirent's SmartFlow and SmartVOIP/QOS test applications.

About Opus One®

Opus One® is a consulting and information technology firm based in Tucson, AZ. Founded in 1989, Opus One's corporate goal is to help our clients make the best use of information technology. We focus on efficient and effective solutions in the areas of data networking, electronic mail, and security. For more information, see http://opus1.com or contact us at:

Opus One
1404 East Lind Road
Tucson, AZ 85719
+1-520-324-0494