Always Up:
High Availability Features for Cisco Catalyst 6500/Cisco 7600 Switches and Routers

prepared for Cisco Systems
September 2004

CONTENTS

ILLUSTRATIONS

Executive Summary

High availability ranks among the top network infrastructure requirements – more so than security, standards support, performance, or even price. There's good reason for this kind of thinking: High availability features increase uptime and prevent losses in productivity and revenue.

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

Figure 1: Key Factors for Network Infrastructure

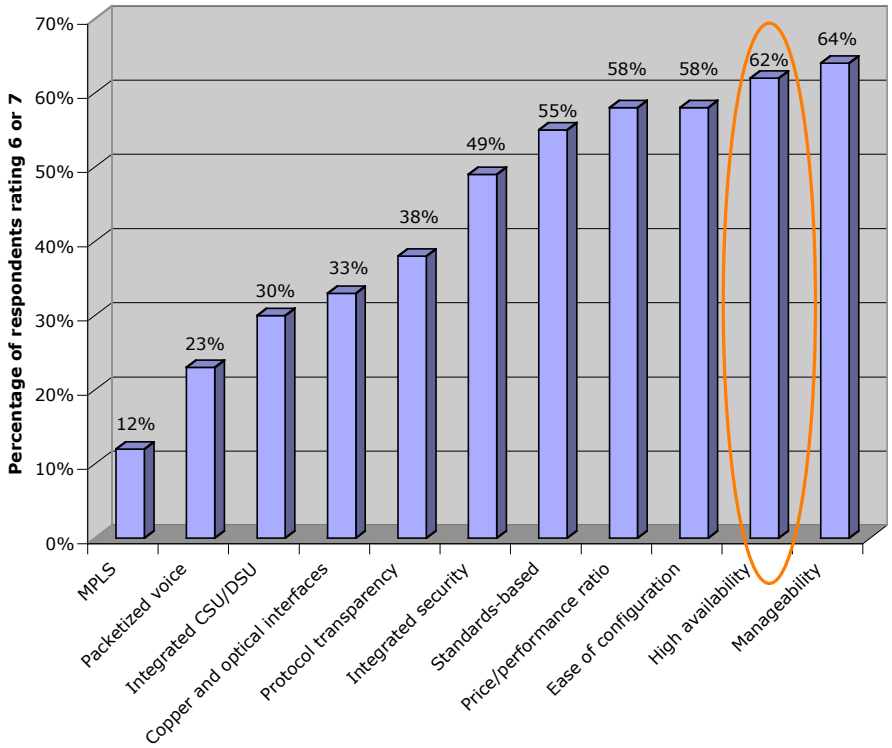

Cisco Systems is addressing the requirement for resilient network infrastructure by adding several new features to its Cisco Catalyst 6500 series switches and Cisco 7600 series routers – Gateway Load Balancing Protocol (GLBP), Non-Stop Forwarding (NSF), and Stateful Switchover (SSO). These features ensure greater uptime with no loss in functionality of existing switch or router features.

Cisco commissioned Opus One, an independent networking consultancy, to conduct performance tests measuring the effectiveness of Cisco's new resiliency mechanisms.

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

Opus One not only tested each resiliency mechanism, but also applied many of the factors at work in large enterprise settings: Unicast and multicast traffic; voice over IP traffic; Policy Based Routing; QoS enforcement; attacks using spoofed IP addresses; and very large access control lists. In addition to the resiliency tests, Opus One tested Cisco's new 10GBase-CX4 interfaces, a cost-effective new standard for running 10-gigabit Ethernet over copper.

Among the key findings of Opus One's tests:

- NSF/SSO provides zero packet loss on any of 4 million flows despite the loss of a Supervisor Engine card and 10,000 OSPF routes when line cards are equipped with Distributed Forwarding Card (DFC) modules

- NSF/SSO provides zero packet loss on any of 4 million flows despite the loss of a Supervisor Engine card and 10,000 BGP routes when line cards are equipped with Distributed Forwarding Card (DFC) modules

- No loss in functionality during or after Supervisor Engine failure for any of the following features: Policy Based Routing, access control lists, rate limiting, and Unicast Reverse Path Forwarding (uRPF, which protects against the use of spoofed IP addresses in DoS attacks)

- Thanks to enhanced wiring-closet device resilience provided by Cisco's new Gateway Load-Balancing Protocol (GLBP), first-hop router or switch recovery of 2.01 seconds or less

- Perfect load balancing across protected VLANs and subnets using GLBP, making full use of two uplinks to each wiring closet and doubling capacity compared with VRRP

- NSF/SSO failover times are virtually identical with unicast and multicast traffic, even when 10,000 s,g mroutes are involved

- Minimal degradation of voice over IP audio quality during Supervisor Engine failover

- NSF/SSO protects upper-layer session state through tight integration with other services modules for Cisco Catalyst switches

- NSF/SSO delivers high availability to wireless as well as wired clients through tight integration with the new Wireless LAN Services Module (WLSM) for Cisco Catalyst 6500 series switches

- Line-rate throughput for the new 10GBase-CX4 interfaces

These results underscore the ability of Cisco Catalyst 6500 series switches and Cisco 7600 series routers to deliver near-perfect uptime, despite the loss of a Supervisor Engine card.

This report is organized as follows. An introduction describes the various high availability mechanisms tested. Then we move on to discuss test bed configuration, procedures and results from tests of GLBP, NSF/SSO with OSPF, NSF/SSO with BGP, and NSF/SSO with IP multicast traffic.

Introduction

Our tests focused on three of Cisco's resiliency features for Cisco Catalyst 6500 series switches and Cisco 7600 series routers: the Gateway Load Balancing Protocol (GLBP), Non-Stop Forwarding (NSF), and Stateful Switchover (SSO). We also benchmarked the performance of new 10Gbase-CX interfaces, which give the Cisco Catalyst 6500 and Cisco 7600 10-gigabit-Ethernet-over-copper capability.

The Gateway Load Balancing Protocol is a patent-pending evolution of Cisco's Hot Standby Router Protocol (HSRP). With first-hop router redundancy protocols such as the Virtual Router Redundancy Protocol (VRRP) or Cisco's Hot Standby Routing Protocol (HSRP), only a single "active forwarder" is permitted per protected subnet/VLAN[2]. In addition, VRRP permits only one of the two uplinks from each wiring closet to be active; the other is held in standby mode and cannot be used to carry traffic.

GLBP, in contrast, allows the use of both redundant uplinks during normal operation. This allows both GLBP routers to be "active forwarders" simultaneously. With GLBP, both GLBP routers are active in the routed topology. The rest of the network will see equal-cost paths to the protected subnet, and traffic to that subnet is load-balanced across the two routers. In the reverse direction, a patent-pending method load-balances traffic from end-stations between the two GLBP routers. With GLBP, failover times are configurable.

The net result: GLBP doubles available bandwidth while allowing users to deploy a single subnet in the wiring closet.

GLBP can be said to be an "active-active" protocol, while VRRP is an "active-passive" protocol. VRRP supports a single active uplink from the wiring closet at any one time.

GLBP, in contrast, makes use of both uplinks during normal operation. Further, it balances the load across uplinks. Our test results confirmed that GLBP distributes loads evenly across links. In fact, the load was so evenly distributed in our tests that interface counters on each of two Cisco Catalyst switches running GLBP matched to the packet.

[2] RFC 3768 describes VRRP, while RFC 2281 describes HSRP.

Figure 2 below compares forwarding paths for VRRP (on the left) and GLBP (on the right.)

Figure 2: Comparing VRRP and GLBP

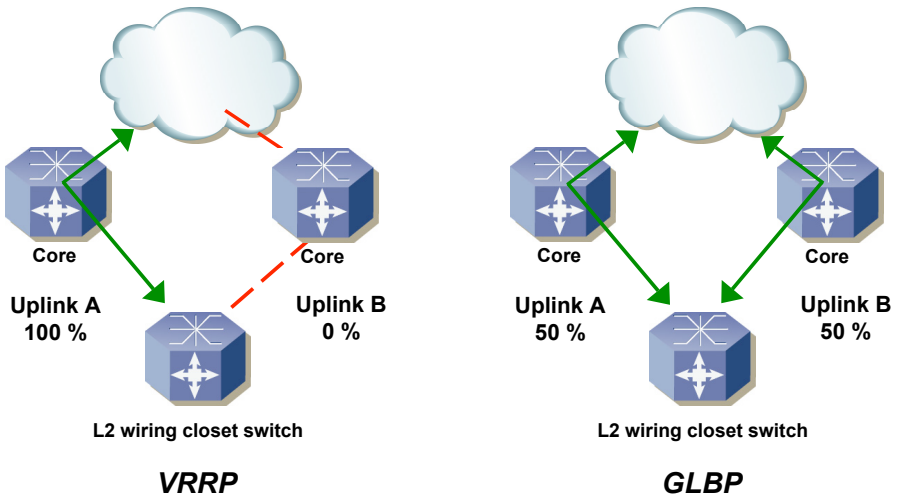

GLBP also enhances routing resiliency. If one GLBP router fails, another is instantly able to forward traffic to/from the core network since its routing adjacencies are already established. This is not the case with VRRP.

Non-Stop Forwarding (NSF) makes use of the industry-standard graceful restart mechanisms developed by the IETF. It preserves layer-3 forwarding state during the loss and restart of a routing session, as might occur due to the failure of a Supervisor card.

Without NSF, reconvergence after loss of a routing session may take tens of seconds or even minutes. For example, the OSPF routing protocol's default timer values require 40 seconds to pass before a router will declare a routing session to be dead. Then a new routing session must be re-established, followed by a potentially lengthy exchange of routing updates.

Our tests show that NSF can reduce this interval to 2 seconds or less for packets centrally switched by the failed Supervisor Engine card, or zero loss if NSF/SSO is used in conjunction with line cards equipped with Cisco's Distributed Forwarding Card (DFC) modules.

Cisco's NSF works with EIGRP, BGP, OSPF, and IS-IS. We used OSPF and BGP in these enterprise-focused tests.

Stateful Switchover (SSO) is Cisco's method of preserving layer-2 forwarding state despite the failure of a Supervisor Engine card. SSO synchronizes layer-2 forwarding tables and spanning tree topology state between redundant Supervisor cards in the same chassis. This ensures forwarding will continue even after the loss of an active Supervisor card, and that no spanning tree topology change will be triggered by the failover to the standby Supervisor.

GLBP Baseline Tests

The Gateway Load Balancing Protocol feature of Cisco IOS provides both fault tolerance and load-sharing, something we demonstrated in tests involving multiple failure scenarios. As noted in the introduction, GLBP improves on existing redundancy technologies like Virtual Router Redundancy Protocol (VRRP) by providing "active-active" rather than "active-standby" availability of redundant routers.

Figure 3 below illustrates the test bed used in the GLBP baseline tests. Four Cisco Catalyst 6500 switches – designated A, B, C, and D – are interconnected with 10-gigabit Ethernet circuits.[3] While we used Cisco Catalyst switches for this project, the same features are available on Cisco 7600 series routers.

Figure 3: The GLBP Test Bed

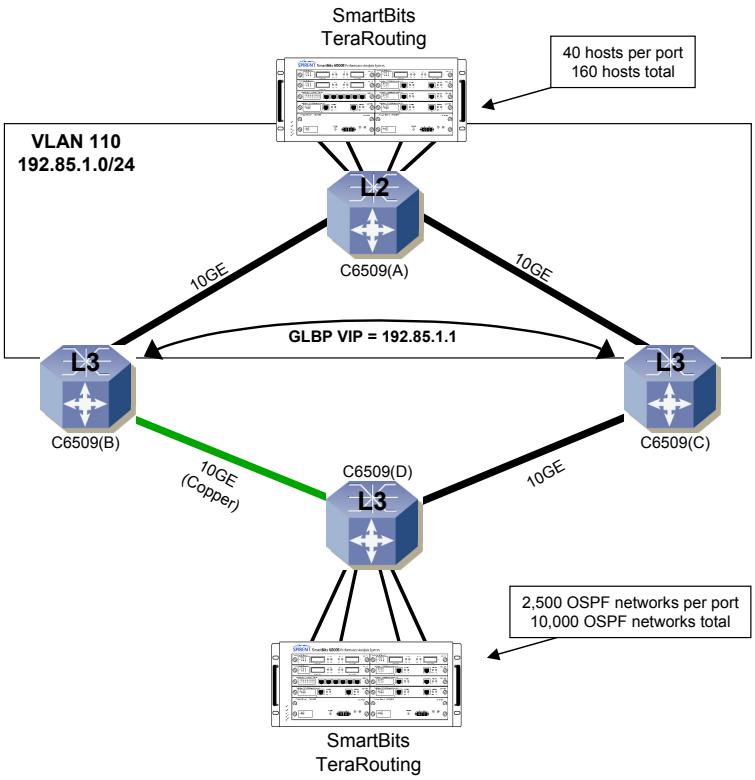

[3] We used Cisco Catalyst 6500 series switches for these tests, but all test results in this document apply equally to Cisco 7600 series routers. Any references to Cisco Catalyst switches in text cover the Cisco 7600 series routers as well.

Switch A represents a layer-2 wiring-closet device. Behind it, a SmartBits traffic analyzer/generator offers traffic from 40 emulated hosts on each of four switch ports, for a total of 160 emulated hosts. The interfaces linking Switch A with Switches B and C share a common VLAN ID.

Switches B and C represent redundant layer-3 devices at the core of the network. These two GLBP-enabled routers share a single virtual IP address used by end-stations (emulated by the SmartBits) as their default gateway. By responding to end-station ARP requests with alternating MAC addresses representing Switch B or C, GLBP directs end-stations to use one or the other GLBP router as their default gateway. In this way, traffic from the end-stations is balanced evenly across the A-B and A-C links. This virtual IP address is in the same VLAN and IP subnet as the end-stations being protected by GLBP.

Switch D represents another layer-3 core device with a large number of networks behind it. A SmartBits attached to Switch D establishes OSPF adjacencies and advertises 2,500 networks behind each of four interfaces, for a total of 10,000 networks.

We offered test traffic to four ports on Switch A, destined to all 10,000 networks beyond Switch D, at a rate of 1 million packets per second. At that rate, each dropped packet represents 1 microsecond of failover time.

We ran this test multiple times: First as a baseline case with no failure to verify that GLBP load-balanced traffic as claimed, and then with separate failover test cases involving a link failure and failures of the Supervisor 720 card in Switch B and the Supervisor 2 card in Switch C.

By testing both Supervisor 720 and Supervisor 2 scenarios, we covered the major portion of Cisco's installed base of users. This validated the functionality of GLBP in either environment, or indeed in a hybrid network as used in these tests.

In the no-failure baseline, we verified that the system under test could forward to all ports at 1 million packets per second with zero loss. This test also determined that GLBP balanced the load across the A-B and A-C links.

We verified load balancing using the Cisco Catalyst 6500 port counters, which showed uniform distribution of packets across the two paths. We then verified the accuracy of the Cisco Catalyst port counters by comparing them with SmartBits transmit and receive counters. All the counters matched: Load balancing was perfect across the A-B and A-C links.

Next, we offered the same traffic and tested the effects of link failure. Approximately 30 seconds into the 60-second test, we physically disconnected the A-C link, forcing GLBP to redirect all traffic onto the A-B link.

GLBP worked correctly here: All traffic arrived at the destination ports with zero loss despite the loss of the A-C link. Since ample bandwidth existed on the A-B link to carry

traffic redirected from the A-C link, zero loss was the expected result. We noted that there was no routing protocol convergence needed on Switch B, allowing traffic to be forwarded with no delay.

In the next test case, we forced a Supervisor card failure by removing the active Supervisor 720 card from Switch B approximately 30 seconds into the test. This removal forced GLBP to redirect traffic onto the A-C link and through Switch C. In three trials, the failover took an average of 1.2 seconds. This test result represents the time needed for flows to be redirected and switched through Switch C.

We then repeated the test while removing the active Supervisor 2 card from Switch C, thus forcing the system to redirect traffic via Switch B. This time, the failover took an average of 2.0 seconds over three trials.

Table 1 below summarizes results from the GLBP failover tests.

Table 1: GLBP Failover Tests

Test case	Failover time (seconds)
GLBP, Supervisor 720 card failure in Switch B	1.207804
GLBP, Supervisor 2 card failure in Switch C	2.016601

NSF/SSO Testing: A Complex Configuration

Large-scale enterprise networks are anything but simple, and we used an accordingly complex setup in our NSF/SSO tests. The test bed configuration modeled many aspects of large-scale production networks, involving not only multiple OSPF areas or BGP autonomous systems, but also many other factors that can affect network performance.

The features simultaneously active in this test included all of the following:

Policy-Based Routing (PBR). It is often desirable for administrative or technical reasons to override OSPF or BGP shortest-path calculations and force some traffic to use "high-cost" links. For this event, the Cisco Catalyst 6500 switches were configured to enforce PBR on a subset of test traffic. The expected result was that the switches would continue to enforce policy during and after a failover, sending specified traffic – and only specified traffic – over a high-cost link.

Access Control Lists (ACLs). We used a 10,000-line access control list in this test. That is considerably larger than the ACLs in use at even many large organizations, and thus considerably more stressful a test case.

The 9,999th rule required routers to discard all test traffic with a particular destination IP address and TCP port number. Placing this rule at the bottom of the list forced the switches to compare every packet to nearly 10,000 entries before making a forwarding decision.

The expected result was that the switches would drop all traffic matching the 9,999th rule during and after failover, demonstrating that ACLs remain in effect at all times. To simplify results tracking, we configured the Spirent SmartBits traffic generator/analyzer to send all traffic matching the 9999th rule to a single destination interface. Therefore, we expected one SmartBits interface to receive zero packets before, during, and after a failover.

Unicast Reverse Path Forwarding (uRPF). Denial-of-service and other attacks commonly originate from spoofed IP addresses. Tracing spoofed addresses to their actual origin can be quite difficult. If a packet with a spoofed address originates from a directly attached subnet, ACLs can help by blocking traffic not originating from that subnet. However, such ACLs do nothing to stop forged packets sourced from one or more hops away.

Cisco's hardware-based Unicast Reverse Path Forwarding (uRPF) feature compares the source address of every received packet with the known shortest-path route back to the source subnet. If the packet's source interface does not represent the shortest known path back to the source subnet, uRPF discards the packet – thus blocking the attack and saving the network from possible meltdown.

We tested uRPF by generating packets with spoofed source addresses. These were not just any spoofed addresses; we used legitimate source addresses from other subnets in the test network, forcing the Cisco Catalyst switches to distinguish between legitimate and spoofed traffic.

The expected result was that uRPF would block all packets offered with spoofed addresses, both during and after a failover. In this case, packets with spoofed addresses represented 20 percent of all traffic offered to one of the switches; thus, we expected loss of exactly 20 percent on each destination interface for traffic received from this switch. Further, we expected the system to forward traffic from legitimate sources (packets from the same source network as the spoofed traffic), demonstrating that uRPF distinguishes between legitimate and spoofed addresses.

QoS Enforcement. Opus One examined the ability of the Cisco Catalyst switches to classify traffic, apply a rate limiter, and forward packets to the appropriate priority queues during and after a Supervisor card failure.

We offered all test traffic with a diff-serv codepoint (DSCP) value of 63, the highest possible priority. For a subset of test traffic, the switches were configured to enforce a very low forwarding rate of 2,000 packets per second. Rather than simply dropping packets above that rate, the rate limiter was configured to "mark down" any traffic exceeding 2,000 pps with a lower DSCP value of 40.

The expected result was that the Cisco Catalysts would forward all traffic to the expected priority queues, and correctly re-mark DSCPs for any traffic exceeding the rate limiter's maximum setting.

This re-marking capability is an important safeguard for devices enforcing QoS policies; it prevents errant or maliciously mismarked packets from using all available bandwidth by claiming high-priority status. By default, Cisco Catalyst 6500 series switches re-mark all ingress traffic with a "normal" DSCP value of 0 unless otherwise configured.

IP Telephony. In addition to the various other conditions, we offered voice-over-IP (VoIP) traffic to determine whether the switches would protect audio quality by ensuring low, consistent delay during and after the failure of a Supervisor card.

Using the SmartVoIP/QoS application for the SmartBits, we offered G.711-encoded VoIP traffic and used the application's Perceptual Speech Quality Measurement (PSQM) scores to determine audio quality. We offered voice traffic both with and without a Supervisor card failover. The expected result was minimal variation in PSQM scores across the two test cases.

We mixed this alphabet soup of factors – PBR, ACLs, uRPF, QoS, and VoIP – together, introducing all features simultaneously in our tests of with NSF/SSO failover with OSPF and BGP.

OSPF NSF/SSO Failover

We tested the effects of NSF/SSO in OSPF networks using a complex configuration modeling many aspects of large-scale enterprise networks. The test included not only multiple OSPF areas but also numerous other factors – including Policy Based Routing, a 10,000-line access control list, Unicast Reverse Path Forwarding, QoS enforcements, and VoIP traffic.

Figure 4 below illustrates the test bed used in the OSPF NSF/SSO event. OSPF Area 0 represents the core switches in a data center. Area 1 represents internal subnets at a local site, while Area 4 represents subnets at remote sites.

Figure 4: The OSPF NSF/SSO Test Bed

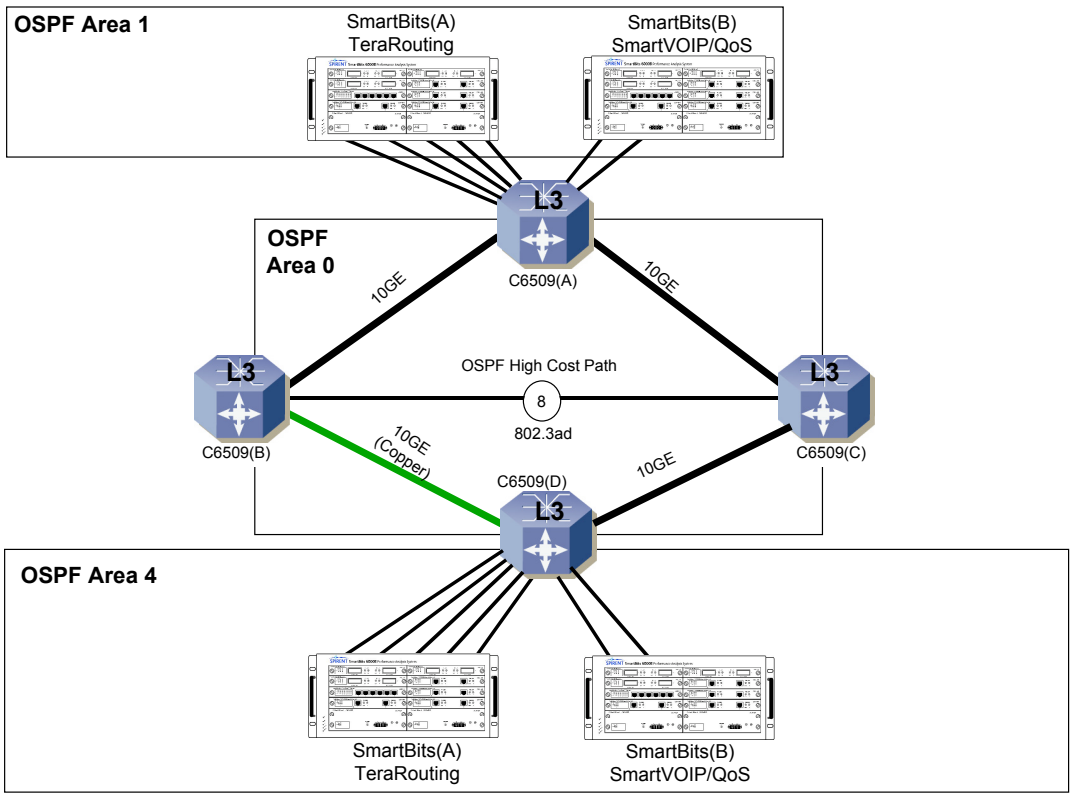

The test bed differed from that used in the GLBP baseline tests in these respects:

- All switches are configured as layer-3 routers running OSPF with graceful restart extensions.

- The TeraRouting application for SmartBits advertises 1,000 networks behind each of 10 interfaces, for a total of 10,000 routes advertised. The test traffic represents

200 hosts on each of these 10,000 networks, for a total of 2 million layer-3 flows on the test bed.

- Switch A has seven edge interfaces. On five of these, the TeraRouting application for SmartBits advertises 1,000 networks behind each interface. On the other two interfaces, we generate voice traffic using the SmartVoIP/QoS application.

- Switches B and C communicate with Switch A using router interfaces and OSPF.

  Also, Switches B and C are interconnected using an IEEE 802.3ad link aggregation group (LAG) consisting of eight interfaces on each switch. The devices treat this LAG as an OSPF path with a higher cost than the other 10-gigabit Ethernet interswitch links.

- The Switch D configuration uses seven edge router interfaces. On five of the interfaces, the TeraRouting application brings up OSPF adjacencies and advertises 1,000 routes behind each interface. SmartVoIP/QoS offers voice traffic on two additional interfaces concurrently with the routed data traffic.

To ensure all traffic flowed through the device under test (either Switch B or C in the diagram), we adjusted OSPF cost metrics in Switch A and D to force the traffic through the appropriate switch.

For each test, we began with a baseline case with no failure to verify there was zero packet loss under normal conditions. Then we repeated the test and removed an active Supervisor card with traffic running, forcing failover to a redundant Supervisor card in the same switch. We ran the baseline and failover cases for both Switches B and C.

In tests of Supervisor 720 failover on Switch B, we offered 64-byte packets at a rate of 1 million pps. Thus, each dropped packet represented 1 microsecond of failover time.

When we removed the active Supervisor card from Switch B, it took an average of 1.4 seconds (over three trials) for the secondary Supervisor card to become active. Some loss was expected in this test, since the loss of an active Supervisor 720 card also means the loss of the data plane (switch fabric) over which packets are forwarded.

In tests of Supervisor 2 failover on switch C, we used 256-byte packets and a rate of 333,333 pps. At this rate, each dropped packet represented 3 microseconds of failover time.

The change in packet sizes did not affect the test results. These were tests of resilience, not measurements of forwarding performance. As such, forwarding rate and packet length are not major factors in gathering accurate results.

In the Supervisor 2 failover test, we observed packet loss equivalent to about 0.4 seconds of failover time (again averaged over three trials). As in the Supervisor 720 failover test,

some loss was expected here as well, even though Cisco Catalysts equipped with Supervisor 2 cards use a separate module for the switch fabric. The loss occurs because the line cards we used were not equipped with the Distributed Forwarding Card (DFC) daughtercards. As a result, packet headers are inspected by the active Supervisor Engine, which makes a centralized forwarding decision on behalf of the line cards. Thus, loss of a Supervisor card leads to momentary loss of forwarding capability in a no-DFC configuration.

It is possible to achieve zero loss with Supervisor 2 cards when line cards are equipped with DFC daughtercards. The DFCs place an independent forwarding engine on the line cards, obviating the need for Supervisor lookups.

Further, Cisco claims that there is zero packet loss with the Supervisor 720 in situations where the switch fabric is not in the data path – in other words, where the ingress and egress for a flow are on the same side of the switch fabric. However, time constraints prevented us from verifying this claim.

To demonstrate zero loss when DFCs are present, we reran the tests using three Cisco Catalysts with DFC-equipped 10-gigabit Ethernet line cards. The OSPF configuration was essentially the same as before, with each of 10 SmartBits interfaces advertising 1,000 routes. Also, as before, we offered 256-byte packets at a rate of 333,333 pps, so that each dropped packet represented 3 microseconds of failover time.

This time, there was zero loss in the NSF/SSO failover case. This result demonstrates that NSF/SSO resiliency, when used in conjunction with DFC-equipped line cards, will result in no downtime for end-users, even during a Supervisor Engine card failure.

Separately, we used each of three Cisco resiliency mechanisms in tests with DFC-equipped cards to show the relative efficiency of each mechanism:

- Route processor redundancy plus (RPR+)
- Stateful switchover (SSO)
- NSF/SSO

Table 2 below shows the failover times with the various redundancy mechanisms configured on the Cisco Catalyst switches. Note that NSF/SSO represents a massive improvement over SSO and RPR+.

Table 2: Failover With Various Cisco Redundancy Methods

Test case	Failover time (seconds)
RPR+ with Supervisor 2 and DFC-equipped line cards	94.1320059
SSO with Supervisor 2 and DFC-equipped line cards	10.0485618
NSF/SSO with Supervisor 2 and DFC-equipped line cards	0

Table 3 below summarizes OSPF NSF/SSO failover times for tests of Switches B and C, tested both with and without DFC-equipped line cards.

Table 3: OSPF NSF/SSO Failover Times

Test case	Failover time (seconds)
Switch B with Supervisor 720 average failover time	1.407337
Switch C with Supervisor 2 average failover time	0.432234
Switch C with Supervisor 2 and DFC-equipped line cards, average failover time	0.000000

Although failover times were the primary metrics in this test, these were not the only measurements. We also sought to determine whether the system would continue to enforce the various other conditions related to QoS enforcement, policy-based routing, access control lists, unicast reverse path forwarding, and voice quality.

To verify QoS enforcement, we captured packets on egress interfaces and examined their DSCPs. In all cases, the switches "marked down" the DSCP of out-of-contract packets from the initial value of 63 to a value of 40.

To show the effects of the other traffic conditions, we configured the TeraRouting test application to classify packets into five groups. Figure 5 below shows results of this classification for one baseline test and three trials of the Supervisor 720 failover test.

Figure 5: OSPF NSF/SSO Traffic Classification, Supervisor 720

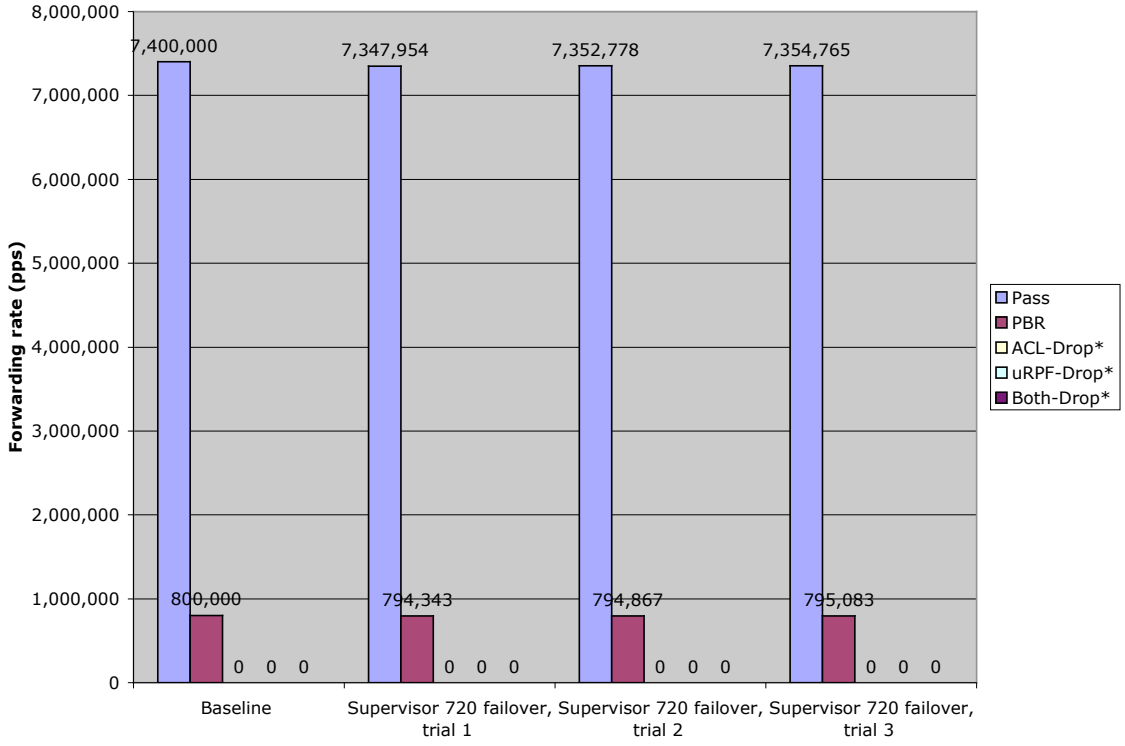

*A forwarding rate of 0 is a perfect result for these traffic classes.

The "Pass" group was test traffic we expected the switches to forward. We also expected the system to forward packets belonging to the "PBR" group, but only across the high-cost link between Switches B and C (something we verified by examining the switches' packet counters, and comparing them with those of the TeraRouting application).

We expected the system under test to drop all traffic belonging to the final three groups: The "ACL-Drop" group represented test traffic matching the 9,999th "deny" rule of a 10,000-entry ACL. The "uRPF-Drop group" represented packets with a spoofed source IP address. The "Both-Drop" group represented packets that matched both the ACL and uRPF conditions. As indicated in Figure 5, a forwarding rate of 0 was a perfect result for these final three groups.

A major goal of adding all these additional features was to verify that the NSF/SSO-enabled switches would always enforce the various rules, even during and after failover. This is important in demonstrating that the system remains protected from attack, even during and after failover. Also note that the forwarding rates are quite consistent across multiple trials.

Figure 6 below presents results of the failover test for the active Supervisor 2 card in Switch C.

Figure 6: OSPF NSF/SSO Traffic Classification, Supervisor 2

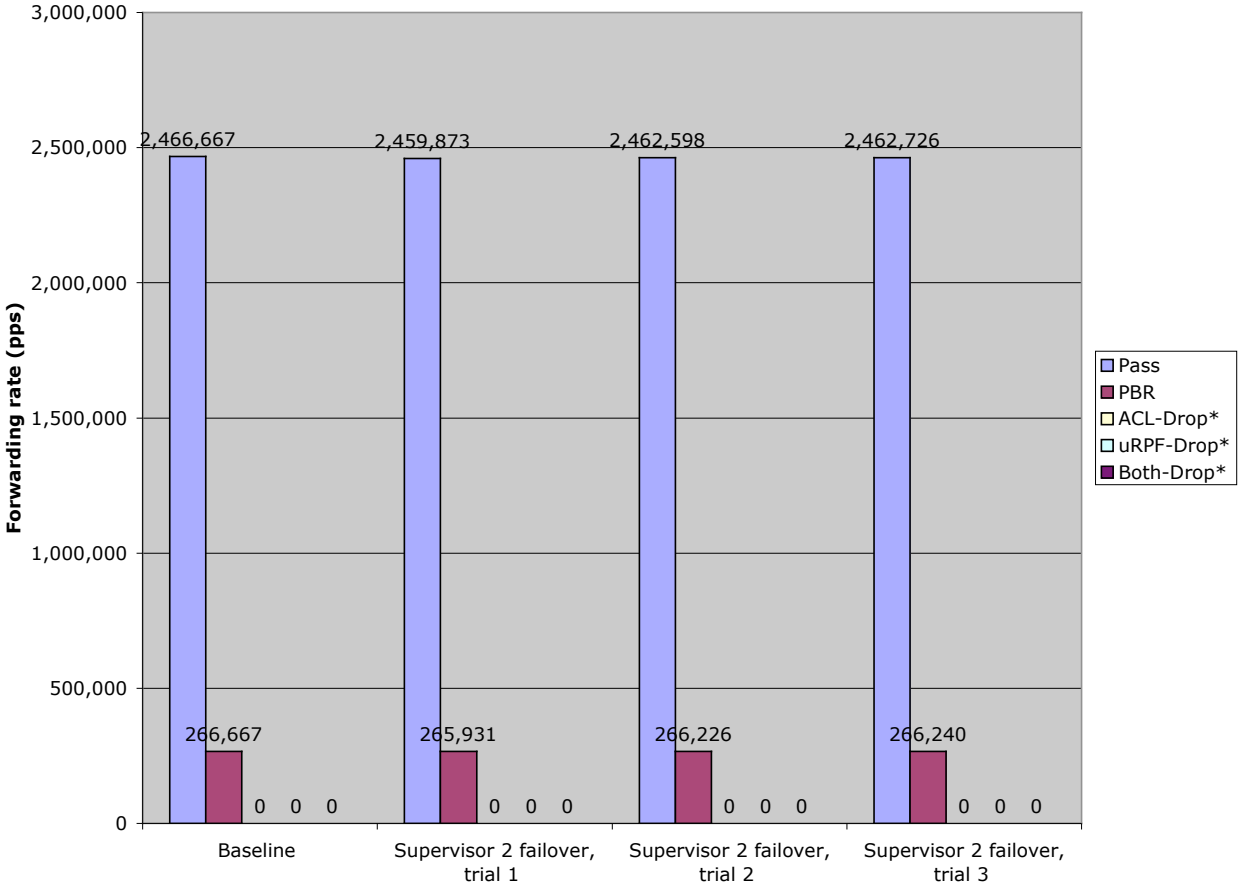

*A forwarding rate of 0 is a perfect result for these traffic classes.

Once again, the switches enforced all the various conditions, both during and after a failover.

Tables 4 and 5 below present the OSPF NSF/SSO traffic classification results in tabular form.

Table 4: OSPF NSF/SSO Traffic Classification, Supervisor 720

Group	Baseline (pps)	Supervisor 720 failover, trial 1 (pps)	Supervisor 720 failover, trial 2 (pps)	Supervisor 720 failover, trial 3 (pps)	Failover average (pps)
Pass	7,400,000	7,347,954	7,352,778	7,354,765	7,351,832
PBR	800,000	794,343	794,867	795,083	794,764
ACL-Drop*	0	0	0	0	0
uRPF-Drop*	0	0	0	0	0
Both-Drop*	0	0	0	0	0

Table 5: OSPF NSF/SSO Traffic Classification, Supervisor 2

Group	Baseline (pps)	Supervisor 2 failover, trial 1 (pps)	Supervisor 2 failover, trial 2 (pps)	Supervisor 2 failover, trial 3 (pps)	Failover average (pps)
Pass	2,466,667	2,459,873	2,462,598	2,462,726	2,461,732
PBR	266,667	265,931	266,226	266,240	266,132
ACL-Drop*	0	0	0	0	0
uRPF-Drop*	0	0	0	0	0
Both-Drop*	0	0	0	0	0

*A forwarding rate of 0 is a perfect result for these traffic classes.

We also measured voice call quality during the failover event. We offered G.711-encoded voice traffic concurrently with routed data traffic, and used Perceptual Speech Quality Measurement (PSQM) scoring to assess audio quality. As described in ITU recommendation P.861, PSQM scoring predicts the subjective quality of speech without requiring subjective testing.[4]

PSQM scoring works on a sliding scale. In Spirent's SmartVoIP/QoS application, the "best" possible PSQM score for a G.711 codec is 0.4, meaning a jury would rate an audio sample as having the highest audio quality. The "worst" score in SmartVoIP/QoS is 6.5, meaning a jury would consider audio to be unintelligible.

In a baseline test with no failover, we recorded a PSQM score of 0.4, the best possible with a G.711 codec. With the failure of an active Supervisor 720 card, the PSQM scores rose to anywhere from 1.8 to 2.0 – higher than the baseline score, to be sure, but still nowhere near the 6.5 "worst" score. These slightly elevated scores are due to the momentary loss of the data path (the switch fabric) integrated into the Supervisor 720 card.

[4] Spirent Communications. "Voice Over IP." 2001. Available at http://www.spirentcom.com/documents/100.pdf?wt=2&az-c=dc.

There was less degradation in the Supervisor 2 failover tests, where PSQM scores ranged from 0.9 to 1.2. Again, these scores are nowhere near the levels where users would consider voice signals to be unintelligible. The lower PSQM scores reflect the fact that the Supervisor 2 and switch fabric are located on separate cards.

SmartVoIP/QoS also records latency and jitter measurements. Note that both metrics remained low and consistent across all trials.

Table 6 below presents VoIP traffic measurements taken during the OSPF NSF/SSO tests.

Table 7: OSPF NSF/SSO VoIP Traffic Handling

Test case	Average PSQM	Average latency across 3 switch hops (usec)	Average jitter across 3 switch hops (usec)
Baseline	0.4	53.0	0.4
Supervisor 720 failover, trial 1	2.0	52.7	0.3
Supervisor 720 failover, trial 2	1.9	52.4	0.4
Supervisor 720 failover, trial 3	1.8	52.6	0.4
Supervisor 2 failover, trial 1	1.2	54.7	2.6
Supervisor 2 failover, trial 2	0.9	54.4	2.5
Supervisor 2 failover, trial 3	0.9	54.4	2.1

BGP NSF/SSO Failover

We reran the NSF/SSO tests using BGP to test the effects of NSF/SSO on this routing protocol, the most popular method of connecting different domains on public networks.

Figure 7 below shows the test bed for the BGP NSF/SSO failover tests. The physical test bed was identical to that of the OSPF NSF/SSO event. In terms of BGP configuration, Switch A and D resided in their own autonomous systems (ASs), while switches B and C shared a common AS (and also shared a common OSPF area for administrative traffic). This models the scenario in which data center devices (Switches B and C here) exchange traffic with multiple external domains over BGP.

Figure 7: The BGP NSF/SSO Test Bed

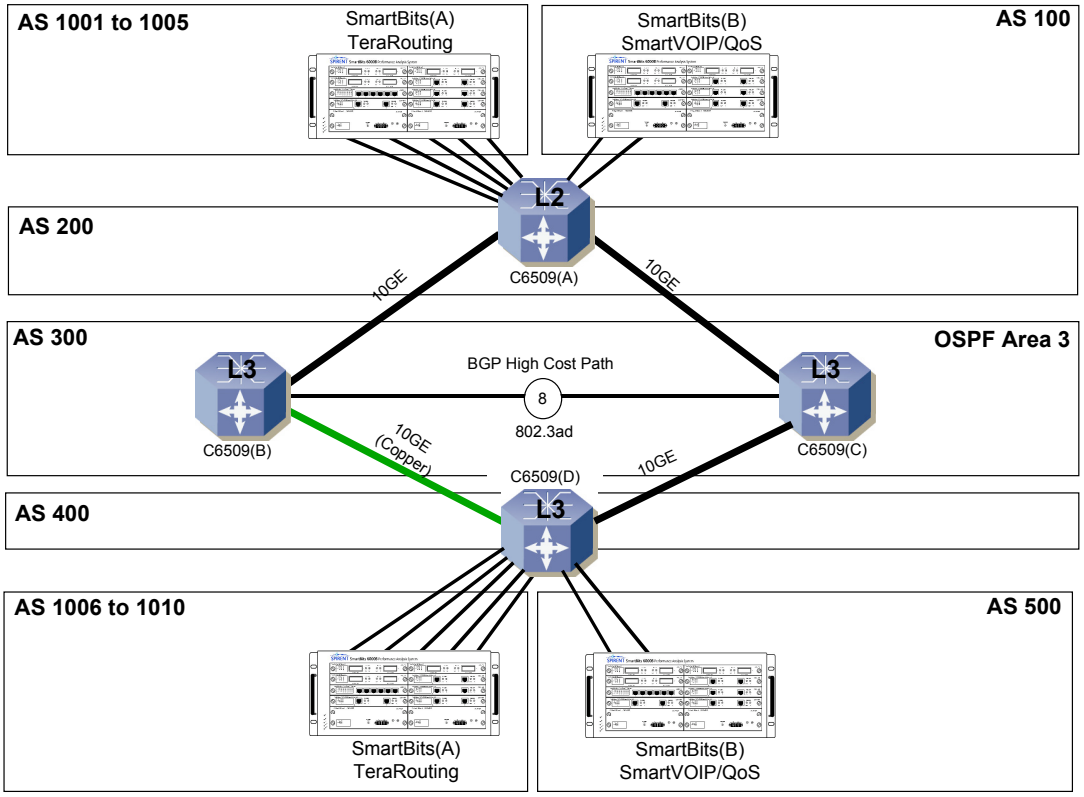

As in the OSPF tests, we measured the failover times for redundant Supervisor 720 modules (in Switch B) and Supervisor 2 modules (in Switch C).

In the Supervisor 720 failover tests, it took approximately 1.2 seconds to fail over from an active to a standby Supervisor card, averaged over three trials. Here again, some loss is expected, since removing an active Supervisor 720 card also removes the switch fabric over which packets are forwarded.

In the Supervisor 2 failover tests, the transition from primary to secondary Supervisor card took about 0.4 seconds (again, averaged over three trials). This is roughly equivalent to the result from the OSPF failover tests. Even though there were separate switch fabric cards in Switch C, some loss was still expected: As in the OSPF tests, the line cards were not equipped with Distributed Forwarding Cards (DFCs) and thus packet headers were sent to the active Supervisor card for a centralized forwarding decision to be made. This configuration represents a worst-case scenario.

It is possible to achieve zero loss with Supervisor 2 cards when line cards are equipped with DFC daughtercards. The DFCs add an independent switching engine to the card and obviate the need for Supervisor lookups.

To demonstrate zero loss when DFCs are present, we reran the BGP failover tests using three Cisco Catalysts with DFC-equipped 10-gigabit Ethernet line cards. The BGP configuration was essentially the same as before, with each of 10 SmartBits interfaces advertising routes to 1,000 networks.

This time, there was zero loss in the NSF/SSO failover case. This result demonstrates that NSF/SSO resiliency, when used in conjunction with DFC-equipped line cards, will result in no downtime for end-users, even during a Supervisor Engine card failure.

Table 6 below summarizes Supervisor failover times for tests of Switches B and C, tested both with and without DFC-equipped line cards.

Table 6: BGP NSF/SSO Failover Times

Test case	Failover time (seconds)
Switch B with Supervisor 720 average failover time	1.206398
Switch C with Supervisor 2 average failover time	0.429090
Switch C with Supervisor 2 and DFC-equipped line cards, average failover time	0.000000

We also tracked various classes of traffic in the BGP NSF/SSO tests to verify the switches enforced various policy rules in place. Once again, policy enforcement was flawless: The Cisco Catalysts adhered to all the various policy rules before, during, and after failover.

In the area of QoS enforcement, we verified with captured traffic that the switches "marked down" out-of-contract packets with a lower DSCP value. This was true for both the Supervisor 720 and Supervisor 2 failover test cases.

We also verified PBR worked correctly by examining the switches' port counters and comparing these with those of the TeraRouting test application.

For enforcement of various other policy rules – including PBR, ACLs, and uRPF – the switches again performed as expected. With BGP running, traffic that should have been dropped was dropped, and traffic that should have been forwarded (except for the small amount of loss during failover) was sent to the correct destinations.

Figures 8 and 9 below summarize the results of traffic classification during the BGP NSF/SSO failover tests.

Figure 8: BGP NSF/SSO Traffic Classification, Supervisor 720

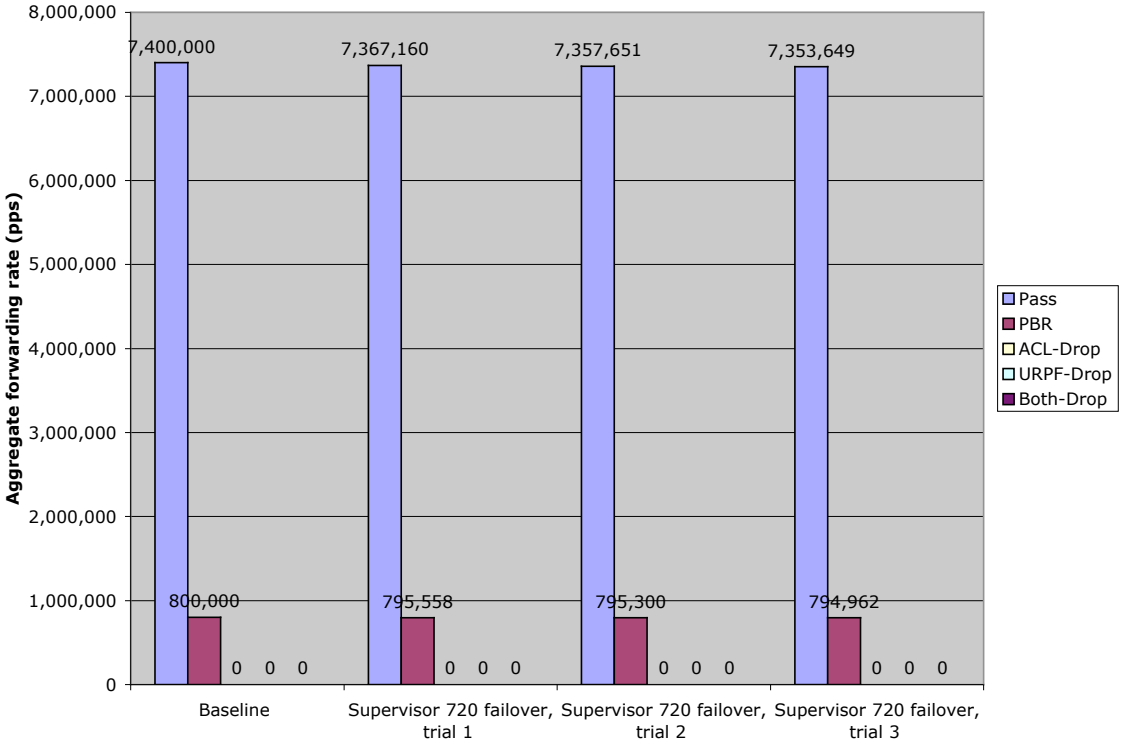

Figure 9: BGP NSF/SSO Traffic Classification, Supervisor 2

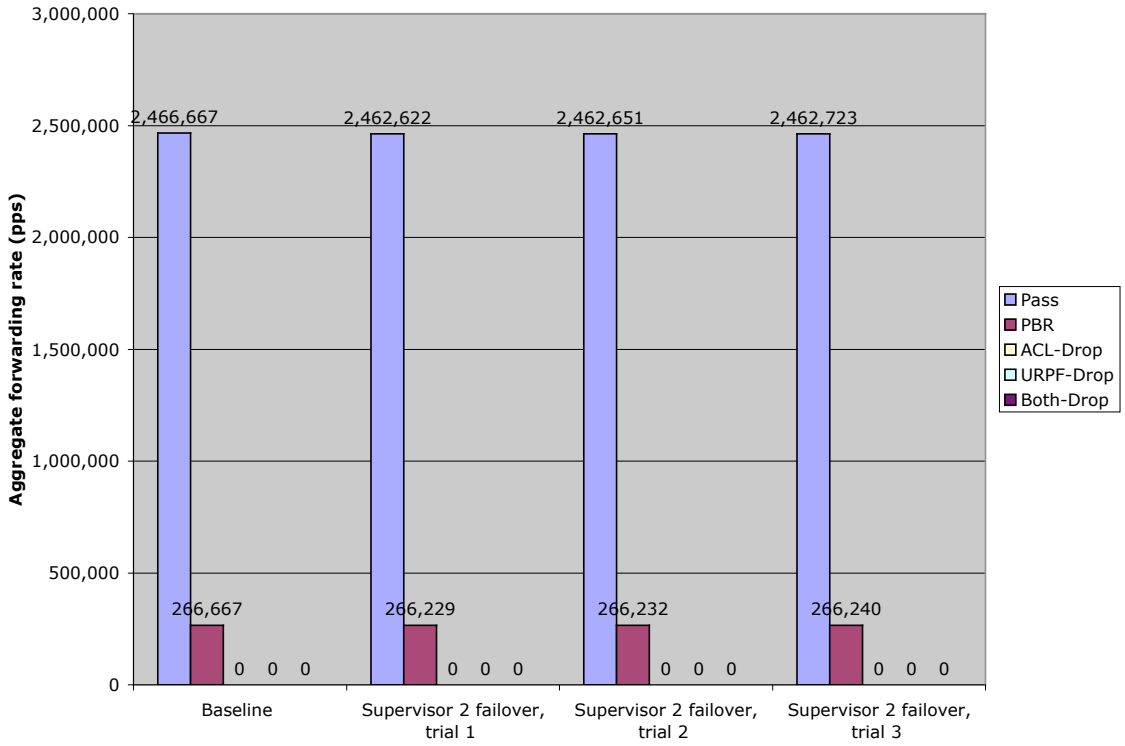

As in the OSPF tests, note that results are very consistent across multiple failover trials. More importantly, in no case did a failover cause traffic to be misrouted.

Tables 7 and 8 below present the BGP NSF/SSO traffic classification results in tabular form.

Table 7: BGP NSF/SSO Traffic Classification, Supervisor 720

Group	Baseline	Supervisor 720 failover, trial 1 (pps)	Supervisor 720 failover, trial 2 (pps)	Supervisor 720 failover, trial 3 (pps)	Failover average (pps)
Pass	7,400,000	7,367,160	7,357,651	7,353,649	7,359,487
PBR	800,000	795,558	795,300	794,962	795,273
ACL-Drop	0	0	0	0	0
uRPF-Drop	0	0	0	0	0
Both-Drop	0	0	0	0	0

Table 8: BGP NSF/SSO Traffic Classification, Supervisor 2

Group	Baseline	Supervisor 2 failover, trial 1 (pps)	Supervisor 2 failover, trial 2 (pps)	Supervisor 2 failover, trial 3 (pps)	Failover average (pps)
Pass	2,466,667	2,462,622	2,462,651	2,462,723	2,462,665
PBR	266,667	266,229	266,232	266,240	266,233
ACL-Drop	0	0	0	0	0
uRPF-Drop	0	0	0	0	0
Both-Drop	0	0	0	0	0

When it came to handling VoIP traffic, the switches running NSF/SSO did even better with BGP than with OSPF.

In a baseline test with no failover, we recorded a PSQM score of 0.4, the best possible score with a G.711 codec. With the failure of an active Supervisor 720 card, the PSQM scores rose to anywhere from 1.5 to 1.8 – higher than the baseline score, to be sure, but still nowhere near the 6.5 "worst" score.

Voice tests in the Supervisor 2 failover case with BGP produced invalid results. Note that PSQM scores in the failover trials were virtually identical to those from the baseline test, strongly suggesting voice traffic was not subject to failover. The most likely explanation is that we forgot to reset the OSPF metrics from a previous test, thereby causing voice traffic to be forwarded through the other switch and not the DUT. We present all results below, but only for the sake of completeness.

Table 9 below summarizes results the VoIP measurements taken during the BGP NSF/SSO tests.

Table 9: BGP NSF/SSO VoIP Traffic Handling

Test case	Average PSQM	Average latency (usec)	Average jitter (usec)
Supervisor 720 baseline	0.4	49.0	0.4
Supervisor 720 failover, trial 1	1.5	49.2	0.3
Supervisor 720 failover, trial 2	1.7	48.9	0.4
Supervisor 720 failover, trial 3	1.8	49.1	0.4
Supervisor 2 baseline	0.4	49.0	0.4
Supervisor 2 failover, trial 1[5]	0.4	49.5	0.3
Supervisor 2 failover, trial 2[6]	0.4	49.9	0.2
Supervisor 2 failover, trial 3[7]	0.4	49.4	0.5

[5] Invalid result; see comments in text.
[6] Invalid result; see comments in text.
[7] Invalid result; see comments in text.

Multicast Multilayer Switching NSF/SSO Failover

Just as NSF/SSO increases network availability for unicast and broadcast traffic, Cisco Multicast Multilayer Switching NSF/SSO (MMLS/NSF/SSO) provides resiliency for multicast traffic. Protection of multicast routing state is especially important given that multicast applications, almost by definition, are intended to serve large numbers of users.

Cisco asked Opus One to measure failover time for a network handling a mix of multicast and unicast traffic, including concurrent multicast and VoIP sessions.

Figure 10 below shows the test bed for the MMLS/NSF/SSO failover event. One SmartBits interface attached to Switch A emulates 10 multicast sources, each sending traffic to 1,000 groups. This created a total of 10,000 s,g mroutes in the system under test, thus forcing it to cope with a large amount of multicast routing table state.[8]

Figure 10: MMLS/NSF/SSO Failover Test Bed

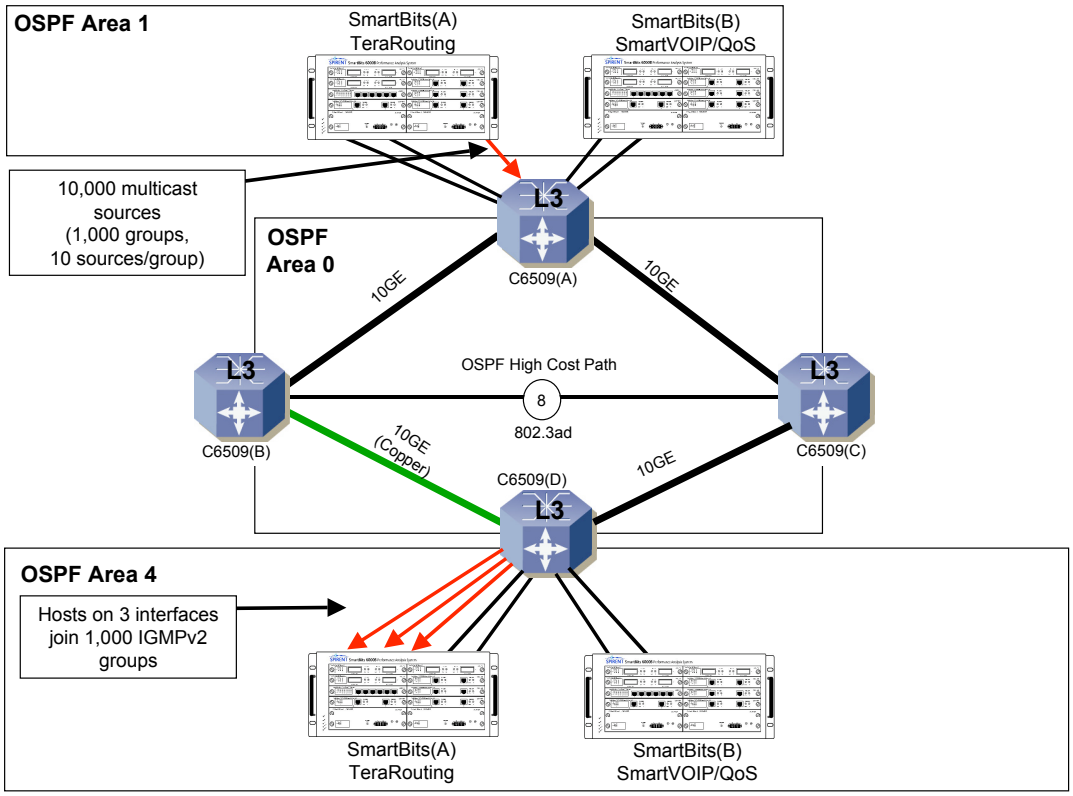

[8]We intended to use 10,000 groups, but a problem in the TeraRouting test application prevented us from using that configuration. The substitute configuration – 1,000 multicast groups, each with 10 sources – still results in 10,000 s,g mroute entries in the Cisco Catalyst 6500's multicast routing tables.

We used Protocol Independent Multicast-Sparse Mode (PIM-SM) as our multicast routing protocol, running on all switches.

To create a single point of failure and a single multicast path through the test network, we used a combination of OSPF route metrics in Switches A and D to force all unicast traffic via the device under test. Since PIM-SM uses the underlying unicast routing protocol to calculate multicast routes, this configuration also forced all multicast traffic to use the device under test.

We also configured a single PIM-SM Rendezvous Point (RP) for the whole test network. The RP was configured on the device under test, so that a failure of the online Supervisor card would also affect this critical function.

Further, in Switch D, the last-hop router, we configured the PIM-SM shortest-path-tree (SPT) threshold so that shortest-path mroutes would be created, adding further multicast state to the device under test.

This test also involved unicast traffic routed over OSPF. Four SmartBits interfaces – two attached to Switches A and D – each brought up OSPF adjacencies and advertised 2,500 external routes, for a total of 10,000 unicast routes. We configured the SmartBits to emulate 200 hosts on each of these external networks. This made for a total of 4 million unique unicast flows.

Voice traffic was also present, as in previous tests. We configured Spirent's SmartVoIP/QoS application to generate G.711-encoded voice calls on each of four interfaces – two each attached to Switches A and D.

As in earlier tests, Switches B and C were linked with an 802.3ad link aggregation group (LAG) configured as a high-cost OSPF path. We then configured Policy-Based Routing (PBR) on the device under test and forwarded a selected subset of the traffic across this link as in previous tests of OSPF and BGP NSF/SSO failover.

Switches B and C were equipped with redundant Supervisor cards. Switch B used dual Supervisor Engine 720 cards, while Switch C used dual Supervisor Engine 2 cards.

The test procedure for this event was similar to that of previous tests: We first brought up all the necessary routing protocols and through the use of OSPF cost metrics on Switches A and D, forced all traffic to flow through the device under test – either Switch B or C. Once we had verified the routing path, we began to transmit both multicast and unicast test traffic. Then we ran a baseline test with no Supervisor card failure to verify the system could forward all traffic with zero loss.

For the failover tests, we physically removed one of the Supervisor cards from either Switch B or C approximately 30 seconds after we began offering data and voice traffic.

Since we offered traffic at a rate of 1 million pps per interface, each dropped packet represented 1 microsecond of failover time.

Figure 11 below presents results from the MMLS/NSF/SSO failover test. Over three trials, multicast and unicast failover times ranged from approximately 1.4 to 1.6 seconds.

Figure 11: MMLS/NSF/SSO Failover for Supervisor 720

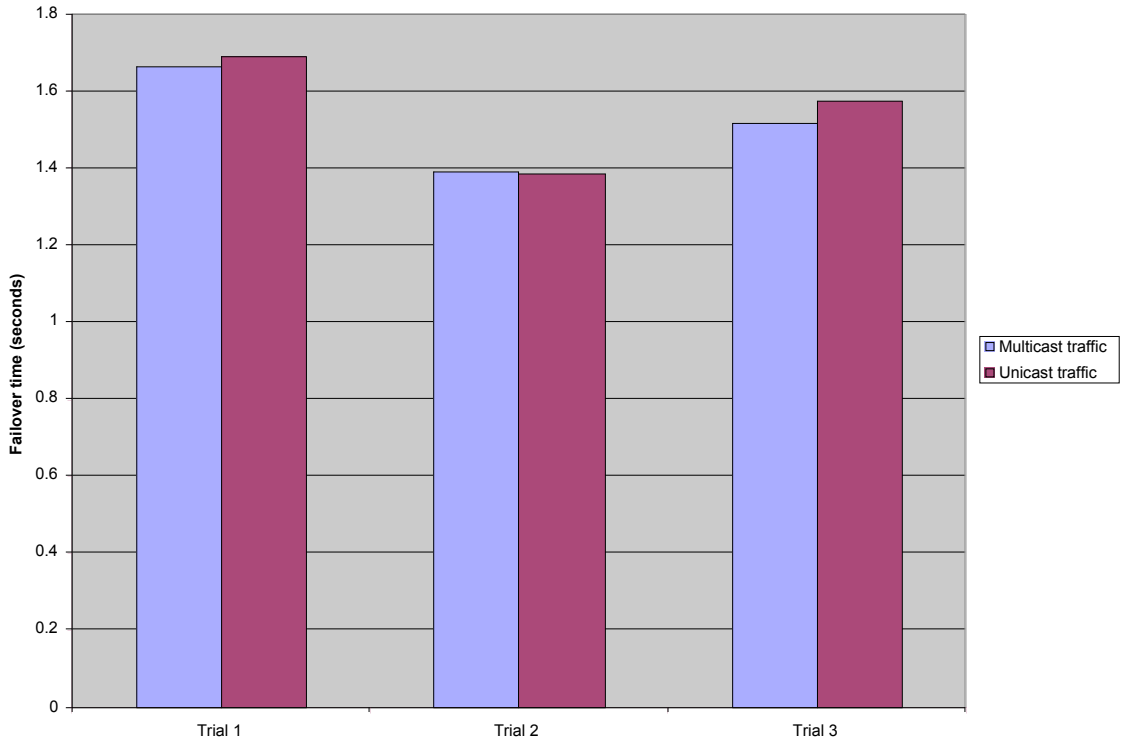

There is no extra performance penalty to protecting multicast traffic with MMLS/NSF/SSO resiliency. The test results show this in two ways: First, failover times are nearly identical for multicast and unicast traffic. Second, failover times in this event are roughly equivalent to those from earlier OSPF and BGP tests.

The IOS version available at test time in August 2004 supported multicast resilience on the Supervisor 720 card, but not on the Supervisor 2 card. Accordingly, we tested IP multicast NSF/SSO failover only on Switch B, equipped with redundant Supervisor 720 cards.

Table 10 below summarizes multicast and unicast results in tabular form.

Table 10: MMLS/NSF/SSO Failover Times

Test case	Multicast failover (seconds)	Unicast failover (seconds)
Baseline	NA	NA
Trial 1	1.662343	1.689121
Trial 2	1.390877	1.385399
Trial 3	1.515533	1.572533

As with the unicast tests, we also measured voice call quality while offering a mix of multicast and unicast data traffic. With the failure of an active Supervisor 720 card, PSQM scores ranged from 1.8 to 2.0 – higher than the baseline score of 0.4 but well below the "worst" possible score of 6.5.

SmartVoIP/QoS also records latency and jitter measurements. Both latency and jitter remained low and consistent across all trials.

Table 11 below presents VoIP traffic measurements taken during the IP multicast NSF/SSO tests.

Table 11: MMLS/NSF/SSO VoIP Traffic Handling

	Average PSQM	Average latency (usec)	Average jitter (usec)
Baseline	0.4	52.3	0.4
Supervisor 720 failover, trial 1	2.0	52.4	0.4
Supervisor 720 failover, trial 2	1.8	52.5	0.4
Supervisor 720 failover, trial 3	1.9	52.6	0.6

NSF/SSO Protection for Upper-Layer Services

NSF/SSO not only provides resiliency at the network layer, but also works in concert with the various services modules available for Cisco Catalyst 6500 series switches. The result of this tight integration between NSF/SSO and the various services modules is higher availability for applications as well as network infrastructure.

Although this report focuses primarily on NSF/SSO and MMLS/NSF/SSO at the network layer, we also sought to determine what effect, if any, the loss of a Supervisor card would have on upper-layer connection state for three services modules: The Firewall Services Module (FWSM), the Content Switching Module (CSM), and the SSL Services Modules.

We tested the interaction of upper-layer services and NSF/SSO by establishing approximately 900,000 HTTP sessions using the Spirent Avalanche and Reflector test instruments. Then, as in earlier tests discussed in this report, we physically removed the active Supervisor card from the switch under test.

In the Switch B Supervisor failover tests, a Cisco Catalyst 6509 maintained 900,003 concurrent HTTP sessions, with no loss of client TCP connections. In the Switch C failover tests, a Cisco Catalyst 6509 maintained 900,008 concurrent sessions, again with no loss in connectivity for clients.

In both tests, traffic continued to flow uninterrupted through the same FWSM, CSM, and SSL Services Modules on each switch. In no case was traffic failed over to the integrated services modules on the other switch, nor did the services modules lose connection state.

Although NSF/SSO preserves L2 and L3 forwarding in the event of a Supervisor card failure, note that it does not preserve upper-layer connection state. For example, NSF/SSO will not save TCP connection state if a single (non-redundant) integrated services module fails. For this reason, Cisco recommends the use of NSF/SSO in conjunction with redundant services modules for maximum reliability: The services modules provide high availability for layer 4-7 connections, while NSF/SSO adds resiliency with non-stop forwarding at layers 2 and 3.

Table 12 below summarizes results from the NSF/SSO failover tests with HTTP traffic. Entries in *green italic type* show the number of concurrent connections after loss of a Supervisor module.

Table 12: NSF/SSO Supervisor Failover With Long-Lived HTTP Sessions

Elapsed time (seconds)	Established TCP connections, loss of Switch B Supervisor 720	Established TCP connections, loss of Switch C Supervisor 2 module
4	900,003	900,008
8	900,003	900,008
12	900,003	900,008
16	900,003	900,008
20	900,003	900,008
24	900,003	900,008
28	900,003	900,008
32 (post-failover)	*900,003*	*900,008*
36 (post-failover)	*900,003*	*900,008*
40 (post-failover)	*900,003*	*900,008*
44 (post-failover)	*900,003*	*900,008*
48 (post-failover)	*900,003*	*900,008*
52 (post-failover)	*900,003*	*900,008*
56 (post-failover)	*900,003*	*900,008*
60 (post-failover)	*900,003*	*900,008*

In tests involving both HTTP and HTTPS, we observed behavior similar to that of other events involving SSL: A small loss in connectivity during the failover, followed by modest gain in concurrent connections.

Table 13 below summarizes results from the NSF/SSO failover tests. Entries in *green italic type* show the number of concurrent connections after the loss of a Supervisor module.

Table 13: NSF/SSO Supervisor Failover With Long-Lived HTTP and HTTPS Sessions

Elapsed time (seconds)	Established TCP connections, loss of Switch B Supervisor 720	Established TCP connections, loss of Switch C Supervisor 2 module
4	200,000	200,000
8	200,000	200,000
12	200,000	200,000
16	200,000	200,000
20	200,000	200,000
24	200,000	200,000
28	200,000	200,000
32 (post-failover)	*200,017*	*200,015*
36 (post-failover)	*200,021*	*200,019*
40 (post-failover)	*200,021*	*200,019*
44 (post-failover)	*200,021*	*200,019*
48 (post-failover)	*200,021*	*200,019*
52 (post-failover)	*200,021*	*200,019*
56 (post-failover)	*200,021*	*200,019*
60 (post-failover)	*200,021*	*200,019*

NSF/SSO Failover for Wireless LAN Traffic

With the fast-growing adoption of wireless LAN infrastructure in enterprise networks, ensuring high availability for wireless clients has taken on new significance. The recently introduced Wireless Services Module (WLSM) for the Cisco Catalyst 6500 takes full advantage of NSF/SSO, ensuring the same highly available network infrastructure for wired and wireless end-stations alike.

The WLSM occupies a single slot in Cisco Catalyst 6500 series switches. No separate infrastructure is needed for WLAN management, thus protecting investment in existing equipment. The WLSM on the Cisco Catalyst 6500 serves as central ingress for all wireless traffic, both on the control and data planes, not only managing traffic from wireless clients but also ensuring high availability through the use of NSF/SSO.

Although this report focuses primarily on NSF/SSO and related services, the WLSM also delivers the same level of security and management services through tight integration with other Cisco Catalyst services modules such as those for firewall and intrusion detection. A separate report in this series examines the WLSM's security and performance in more detail.

Cisco asked Opus One to determine failover times for wireless traffic forwarded through a Cisco Catalyst 6500 switch equipped with a WLSM and redundant Supervisor Engine 720 cards.

Figure 12 below shows the test bed for the wireless NSF/SSO failover event. Virtual clients attached to a pair of Cisco Aironet 1100 access points which, in turn, attached to a Cisco Catalyst 6500 switch. The switch housed two Supervisor 720 cards, one in Active mode and the other in Standby mode.

Figure 12: WLSM with NSF/SSO Failover Test Bed

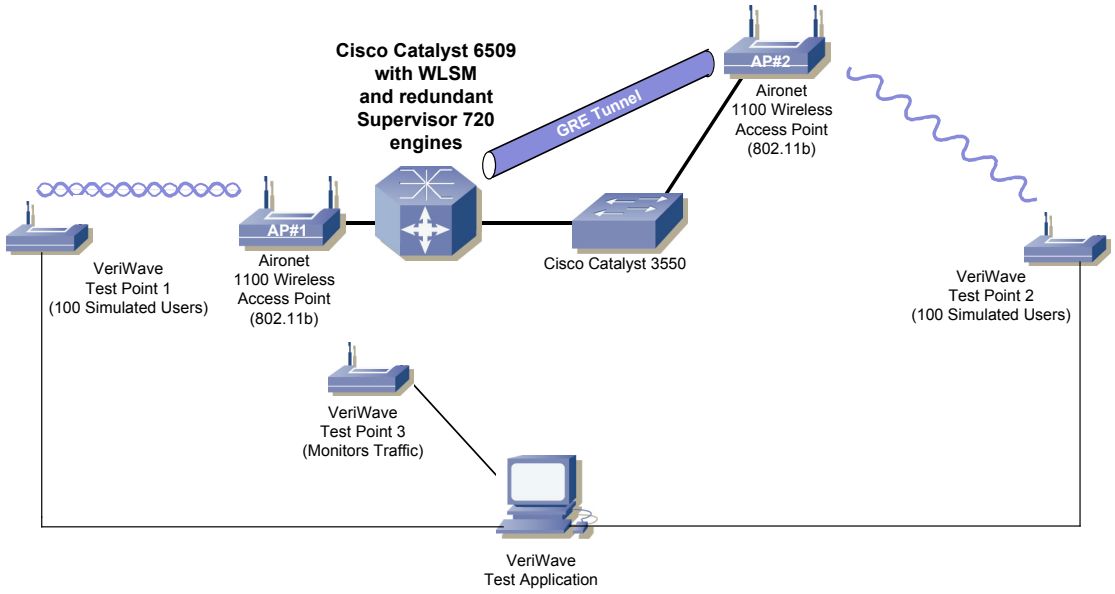

To determine failover time, we used the VeriWave TestPoints test instrument to emulate 200 clients transmitting and receiving data at an aggregate rate of 1,000 packets per second.

At that rate, each dropped packet was equivalent to 1 millisecond of failover time. (We began with baseline tests both with and without the WLSM active to verify that the system dropped zero data packets without a Supervisor failure.)

We then removed the active Supervisor card while offering traffic, thus forcing a failover to the standby Supervisor card, and measured packet loss. We repeated the failover test three times.

We expected some packet loss in these tests. The Supervisor 720 not only establishes and maintains routing protocol and spanning-tree state, but also integrates a 720-Gbit/s switch fabric. Since the loss of an active Supervisor card also means the temporary loss of the switch fabric, some packet loss is inevitable.

Figure 13 below presents results from the NSF/SSO failover tests for the WLSM. The average failover time was 1.776 seconds across three trials. As noted, this loss is a result of the temporary removal of the data path (the switch fabric) in the active Supervisor 720 card.

Without NSF/SSO, the normal routing protocol convergence mechanisms would apply. In the case of OSPF with default timer values, the "router dead" interval does not occur until 40 seconds (the period for the loss of three "hello" messages), followed by additional time to restart the routing session and reconverge the network. This last step may itself last many minutes, depending on the size of the network. With NSF/SSO, failover time is less than 2 seconds.

Figure 13: WLSM Failover With NSF/SSO

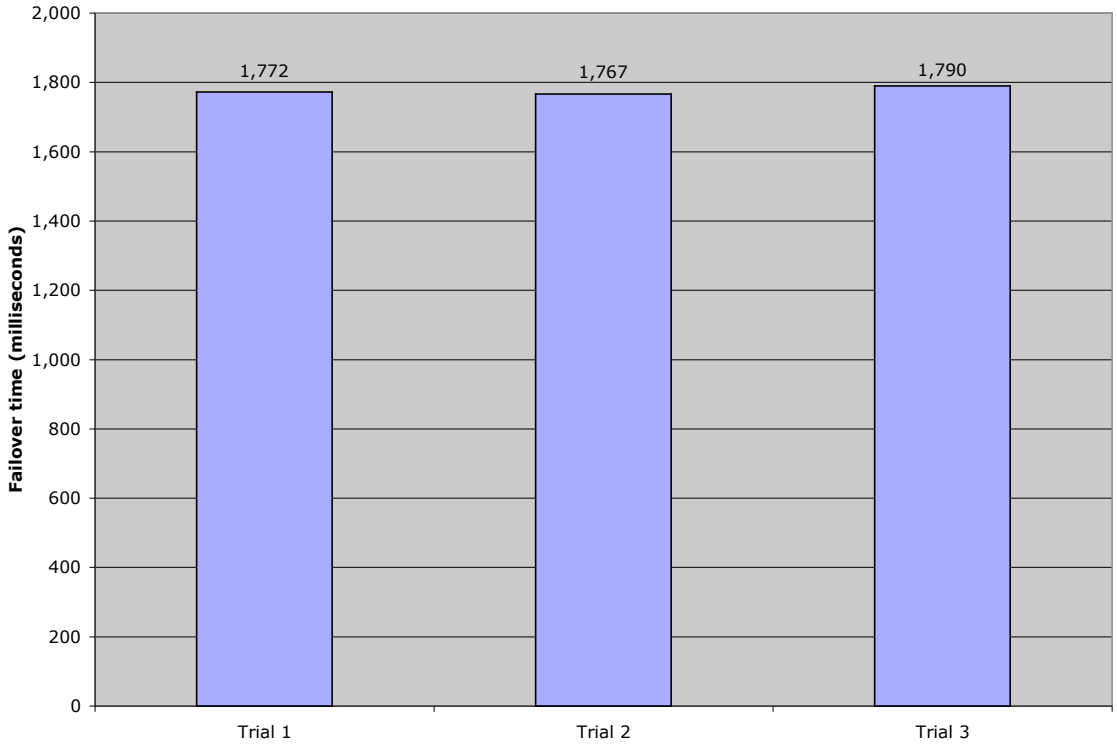

Although we did not use this configuration in our tests, it is possible to achieve zero loss upon failure of a Supervisor card by using line cards equipped with Distributed Forwarding Card (DFC) modules. DFC modules include a local switching engine and switch fabric, obviating the need to go through a central fabric for traffic passing between ports on the same card. Unfortunately, DFC-equipped line cards were not available for inclusion in our tests. Even in this worst-case scenario without DFCs on the line cards, disruption was minimal.

The failover tests demonstrate that loss of a layer-3 routing session introduces only a brief loss in forwarding capabilities.

10GBase-CX4 Throughput

Cisco supplied Cisco Catalyst 6500 line cards with interfaces supporting 10GBase-CX4, the new standard for 10 gigabit Ethernet over copper cabling. As described in IEEE standard 802.3ak, CX4 uses twinaxial copper cabling and relatively inexpensive transceivers to link 10 gigabit devices at distances of up to 50 feet[9]. CX4 offers excellent price/performance for data-center applications such as the switch interconnections used in these tests.

Cisco asked Opus One to validate its claim of line-rate throughput for CX4 interfaces in the Cisco Catalyst 6500.

Working with Spirent's SmartFlow application, we configured the SmartBits traffic generator/analyzer to offer test traffic to a pair of CX4 interfaces. Our test traffic consisted entirely of 64-byte packets, the shortest allowed in Ethernet and thus the most stressful possible load. We used bidirectional traffic to fully load the interfaces.

Over a 60-second duration, the CX4 interfaces forwarded all traffic at line rate – 14,880,952 pps in each direction, or nearly 30 million pps total – with zero loss. Further, the switches delivered all packets in the sequence we offered them; the switches did not reorder frames in flight.

The test results validate Cisco's claim of line-rate performance for CX4.

Table 13 below summarizes results from the 10GBase-CX4 performance tests.

Table 13: 10GBase-CX4 Performance

Total packets transmitted (60 seconds)	1,785,714,228
Total packets received (60 seconds)	1,785,714,228
Total packets received in sequence	1,785,714,228
Bidirectional throughput (pps)	29,761,904
One-way throughput (pps)	14,880,952

[9] IEEE 802.3ak documents are available at http://www.ieee802.org/3/ak.

Conclusion

These tests demonstrate Cisco's new resiliency mechanisms working in a wide variety of settings. In all cases, these mechanisms greatly reduce or eliminate downtime, both in the control plane and in the data plane.

To recap the major findings of these tests:

- Global Load Balancing Protocol (GLBP) reduces failover times from tens of seconds to user-defined values of 2 seconds or less, and doubles available bandwidth during normal operation compared with VRRP

- Non-Stop Forwarding/Stateful Switchover (NSF/SSO) preserves routing state for OSPF and BGP sessions despite the loss of a Supervisor card, ensuring that data traffic continues to be forwarded with little or no disruption

- NSF/SSO delivers zero packet loss when used in conjunction with line cards equipped with Distributed Forwarding Card (DFC) modules, even when handling millions of flows

- Networks remain protected from attack after the loss of a Supervisor card, since NSF/SSO ensures that all security mechanisms continue to work

- NSF/SSO is effective in protecting IP multicast traffic, even when 10,000 s,g mroutes are involved

- NSF/SSO protects upper-layer session state through tight integration with other services modules for Cisco Catalyst switches

- NSF/SSO delivers high availability to wireless as well as wired clients through tight integration with the new Wireless LAN Services Module (WLSM) for Cisco Catalyst 6500 series switches

- The new 10GBase-CX4 line card delivers 10-gigabit Ethernet line-rate throughput over copper cabling

Although these tests covered many different technologies and resiliency mechanisms, the common thread among them all is *high availability*. With GLBP and NSF/SSO deployed on Cisco Catalyst 6500 series switches and Cisco 7600 series routers, network managers can now boost uptime to unprecedented levels.

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 TeraRouting, SmartFlow, and SmartVoIP/QoS test applications. Thanks also to VeriWave, which supplied its WaveTest system for assessing failover performance for wireless clients.

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