



**Competitive Performance Testing Results
Carrier Class Ethernet Services Routers**

DR110209

**Cisco Aggregation Services Router 9000
Alcatel-Lucent 7750 Services Router**

February 2011

Contents

1.0 Executive Summary	3
2.0 Overview	4
Key Findings and Conclusions	5
3.0 Test Setup & Topology	6
Test Equipment Configuration – IXIA	7
4.0 Test Scale Configuration.....	8
5.0 System High Availability Tests.....	10
5.1 Soft Failover.....	11
5.2 Hard Failover	14
6.0 Path Failure Convergence	17
7.0 BGP Multi-Path Convergence Test.....	20
7.1 Single PE Failure Test.....	22
7.2 Multiple PE Failure Test.....	25
8.0 FIB Convergence Speed Test.....	29
8.1 Sequential FIB Convergence.....	30
8.2 Random FIB Convergence	33
9.0 Fabric Priority Test.....	37
9.1 Fabric Priority on Alcatel-Lucent 7750 SR with Two Priorities.....	38
9.2 Fabric Priority on Cisco ASR 9000 with Two Priorities	41
9.3 Fabric Priority on Cisco ASR 9000 with Three Priorities	43
10.0 Multicast.....	46
10.1 Multicast Throughput Test	47
10.2 Multicast Convergence	53
10.3 Multicast Line Card Redundancy Test.....	56
10.4 Multicast Fabric/Route Processor Redundancy and Resiliency	58

1.0 Executive Summary

The Cisco ASR 9000 is Cisco's Unified Converged Edge Router that is designed to offer customers a unique service and application-level intelligence that focuses on optimized video delivery and mobile aggregation. Additionally, it is equipped with the support of a full set of service activation and provisioning systems that are designed to simplify and enhance the operational and deployment aspects of service-delivery networks.

This architectural validation provided strong conclusions that the Cisco ASR 9000 router outperformed the Alcatel-Lucent 7750 SR in the areas of Scale, Quality of Service, Convergence, Network Traffic Prioritization, Multicast for Video Distribution and Video Streaming services and above all High Availability and Resiliency.

The resiliency and high availability shown by the Cisco ASR 9000 was extremely impressive. In both hard and soft failover scenarios, the product experienced minimal to zero packet loss for services. We noted in repeated testing that the ASR 9000 maintained strict fabric priority and never dropped a single packet of high priority traffic during periods of congestion. On the other hand, the Alcatel-Lucent 7750 SR showed severe instability with unpredictable behavior upon a failure. We noticed active services were interrupted for significant periods of time while convergence times took longer as well. For example, the Forwarding Table convergence (FIB) on the Cisco ASR 9000 was 16 times faster than the Alcatel-Lucent 7750 at a very moderate scale.

The performance of the Cisco ASR 9000, regardless of scale was very stable and predictable. The architecture inherently enables service providers the ability to offer guaranteed service level agreements to customers. The Cisco ASR 9000 consistently maintained high priority traffic at any time during periods of congestion whereas the Alcatel-Lucent 7750 SR was not able to do that.

When it came to Video services such Video Distribution and High Speed Video Streaming, we found the Cisco ASR 9000 was optimally designed to scale a very large number of sessions whereby new multicast streams (content) was added dynamically without any service impact to existing streams. Unfortunately, the same was not true for the Alcatel-Lucent 7750 SR. The 7750 SR experienced service interruptions for existing streams as new multicast streams were added. This proved to be a concern that could be an issue for service providers who are looking to offer large scale high availability video services.

Finally, we were impressed with some of Cisco's innovative features and architecture that could be a value add to service providers looking for a true carrier-class platform for next generation services. We observed a feature known as Border Gateway Protocol (BGP) Prefix Independent Convergence (PIC) which significantly reduced convergence times, up to 750 times faster, in the event of multiple path failures. This could be a huge performance enhancement for service providers offering IP and Packet services. We were also impressed with the unique options such as line cards which feature both long and short pins to detect that they are being removed and in turn pro-actively failover the services gracefully and seamlessly.

The Cisco ASR 9000 is a state of the art carrier class services edge router that can exceed the expectations of service providers who are looking for a next generation platform that can deliver Video, Mobile and Carrier Ethernet services.

Rob Smithers
CEO
Miercom

2.0 Overview

The Cisco ASR 9000 Service Aggregation Router facilitates the evolution of Carrier Ethernet networks with its modular and fully distributed architecture combined with the exceptional scale (6.4Tbps per system), comprehensive system redundancy and full complement of network resiliency schemes. The purpose-built high-density Ethernet line cards are equipped with a flexible programming infrastructure, robust Hierarchical Quality of Service (H-QoS), advanced Ethernet services, wide ranging security mechanisms, and integrated Synchronous Ethernet capabilities. The distributed implementation on the ASR 9000 extends to the control plane and provides improved scalability.

The tests in this report are intended to be reproducible for customers who wish to recreate them with the appropriate test and measurement equipment. Please contact reviews@miercom.com for additional details on the configurations applied to the System Under Test and test tools used in this evaluation. Miercom recommends customers conduct their own needs analysis study and test specifically for the expected environment for product deployment before making a selection.


Alcatel-Lucent was advised of this testing in accordance with the Miercom fair testing policy and was afforded the opportunity to participate. Although they were not an active participant, they are welcome to demonstrate their product and its capabilities to Miercom in a similar test environment.

Key Findings and Conclusions

- Cisco ASR 9000 achieved minimal to zero packet loss for all services across a number of network resilience tests, encompassing node, link and service convergence scenarios. The Alcatel-Lucent 7750 SR demonstrated unpredictable and severe packet loss under moderate service scale. We saw extended periods of packet loss on the Alcatel system during convergence testing.
- During link failures, the Cisco ASR 9000 converged up to 93.7% times faster than the Alcatel-Lucent 7750 SR for random FIB convergence. The Cisco system protected services and reduced service outages to a minimum under all service scales tested.
- The fabric-based multicast replication of the Cisco ASR 9000 architecture, proved to be more scalable and resilient than the multicast implementation of the Alcatel-Lucent 7750. The Cisco ASR 9000 protected all high priority traffic, IP Video, Video on Demand (VOD) and other multicast traffic at all times, regardless of the type of failure whereas the Alcatel-Lucent 7750 SR failed to do the same.
- Cisco ASR 9000 hardware architecture adhered to two levels of strict fabric priority and protected high priority traffic under all congestion conditions tested. The Alcatel-Lucent 7750 SR failed to maintain a single level of strict fabric priority traffic under the same conditions.

3.0 Test Setup & Topology

Figure 1: Test Bed Diagram



This topology was used for all tests, unless otherwise noted in the test topology section of the individual test.

The Cisco ASR 9000 Aggregation Services Router (IOS XR version asr9k v4.0.0) included: Two route switch processor cards (ASR9K-RSP-8G), one active and one stand-by and quantity eight 8x 10GE line cards.

Alcatel-Lucent 7750 Services Router (version 8.0.R4) Chassis, Mode D included: two routing engines (Sfm3-12), 7 port line cards: 6 x iom3-xp modular, 1 x imm5-10gb-xp-xfp.

The Cisco CRS-1 was used to provide load balancing and scaling. It provided connectivity between the SUT and IXIA test equipment, and was configured as a simple Core (P) node. Four link bundles were configured between the CRS-1 and the SUT where each bundle had 2 x 10GE interfaces, supplying a total 8 x 10GE interfaces. Each link is protected with multiple options for complete link redundancy.

In addition to providing ECMP paths to the SUT representing a real network redundancy model, the availability of this Core node allowed for scaling the simulated services ports without requiring additional NNI ports on the SUT.


The CRS-1 had 18 x 10GE ports connected to the test equipment allowing the use of the link bundles to load balance for the different tests and services. The following diagram illustrates the physical topology for the CRS.

The test equipment was used to create a simulated topology that combines different types of Layer 2 and Layer 3 services (or multi-service edge) towards the SUT.

24 x 10GE ports on the SUT were configured as customer facing UNI interfaces to handle the mixed services with the distribution of services per 10GE interface.

Note: This test was conducted with the most up-to-date and Generally Available (GA) software release from both vendors.

Figure 2: Ixia Test Bed Set Up



Connectivity between SUT and CRS1 (P) Router including Redundancy Models.

Test Equipment Configuration – IXIA

The test systems used in this evaluation were two IXIA-XM12 chassis with IXOS 5.70.600.build 9; IxExplorer 5.60.301.30; IX Network Protocols 5.60.140.6; TCL v 8.4.14.0. The two separate test beds (one for each SUT) running the IxNetworks application were used to simulate the UNI (Customer Edge) and NNI (Provider Edge) ports and services.

4.0 Test Scale Configuration

24 x 10GE interfaces on the SUT were configured as customer facing User-to-Network Interfaces (UNI) or Customer Edge (CE) interfaces to handle the mixed services with the distribution of services per 10GE interface. Each service interface is configured through a single VLAN, using IEEE 802.1Q, and corresponds to a unique Service as follows:

VPLS (Business VPN)

84x2 VPLS service interfaces per 10G using LDP signaling
Total $24 * (84+84) = 4032$ **VPLS based interfaces**
84x2 VPLS service interfaces per 10G using BGP signaling
Total $24 * (84+84) = 4032$ **VPLS based interfaces**

EoMPLS (Business VPN)

167 x 2 EoMPLS service interfaces per 10G (using IETF Martini draft)
Total $24 * (167+167) = 8016$ **EoMPLS based service interfaces**

L3VPN (Business VPN)

120 VPNv4 service interfaces per 10G
Total $24 * 120 = 2880$ **L3VPN interfaces**

Each L3VPN VLAN interface was configured with eBGP (CE-PE protocol) and enabled with BFD (@ 350 milliseconds interval) and enabled with Netflow Version 5 (@ 1:1000 sampling-rate).

6VPE (Business VPN)

20 VPNv4 service interfaces per 10G
Total $24 * 20 = 480$ **6vPE interfaces**

Multicast VPN (Business VPN)

Multicast VPN services using Default-MDT and Data-MDT services (IETF Rosen draft) for corporate/business streaming video type or high-speed multicast services.

100 multicast Virtual Routing and Forwarding (VRF) service interfaces per UNI port
Total $24 * 100 = 2400$ **Multicast VPN interfaces (OIL)**

Video Broadcast and Video On Demand (VoD)

Shared single L3 interface per 10G for L3 VoD and L3 multicast for Broadcast TV
100 multicast groups per VOD interfaces per UNI port
Total $100 * 24 = 2400$ **interfaces (OIL)**

High Speed Internet (HSI) / Voice over IP (VoIP) / EoMPLS

Shared Single L2 UNI per 10G for Pseudowire-based backhaul of HSI and VoIP traffic to BRAS/BNG

The test equipment simulated 27 separate Remote PEs (Provider Equipment) connected through the CRS P (Core) node simulating all services.

Each EoMPLS service interface was configured to connect point-to-point to one simulated PE.

Each Virtual Private LAN Service (VPLS) interface represented a unique VPLS domain connected to 3 remote simulated PEs. The SUT had a total of 8064 VPLS service interfaces, 4032 were BGP signaled VPLS domains and 4032 LDP signaled VPLS domains.

Each IP-VPNv4 (L3VPN) service interface represents a unique VRF, which was connected to one of four simulated PEs with an equal distribution. Each simulated PE had 720 such interfaces, which give a total of 2880 L3VPN instances.

Each 6VPE (IP-VPNv6) service interface represents a unique VRF, which was connected to one of the four simulated PEs. Each simulated PE has 20 such interfaces, which give a total of 480 6VPE instances.

Furthermore, 4 additional simulated PEs were created to advertise the same prefixes and provide Layer 3 VPN BGP multipath for the SUT and allow fast-convergence during failover situations.

L3 based Multicast: A simulated 10GE PE interface is used to feed IP based multicast traffic into the SUT. This test port is configured to run PIM SSM on the source side; all UNI interfaces are running IGMPv3.

L3 based VOD Service: Shared VOD VLAN with L3 IP multicast service interface VOD from Remote PE simulated Video Server Farm and Backend.

High Speed Internet and VoIP services: A shared VLAN is configured on each UNI interface, bidirectional symmetric traffic is generated directly towards the test equipment simulating the Broadband Remote Access Server (BRAS) on the remote side.

Note: We discovered that the Alcatel-Lucent 7750 SR could not run more than 85 VLANs per port while running Bi-directional Forwarding Detection (BFD). We observed unstable behavior by the platform above this level. We also observed that at one-tenth (1/10th) of the scale configuration, the Alcatel-Lucent 7750 SR began showing predictable results for scale configuration test cases.

The ASR 9000 was also tested with double its services scale configuration at 64K L2+L3VPN Service instances and 48 UNI (CE-facing) interfaces, and produced similar and predictable results.

5.0 System High Availability Tests

Service providers offer continuous network operations as a basic requirement for all applications. Residential customers require access to data, voice, and video services at all times. Enterprise business customers depend on 24-hour network operations with guaranteed service-level agreements (SLAs) for mission-critical applications. Mobile phone subscribers expect to make calls and access data services at all times. Both voice and video services must be delivered in an uninterrupted manner and in turn provide a high quality of experience to end users. As a result, minimizing the impact of component failover is critical to meet such strict service level requirements.

In the following tests, we measured the impact on active services as a result of a failover scenario involving the Route Processor/ Switch Fabric components of the Cisco ASR 9000 and the Alcatel-Lucent 7750 SR routers.

On the Alcatel-Lucent 7750 SR, control plane resiliency across redundant Route Processors was ensured through Non Stop Routing (NSR) mechanisms, while the Cisco ASR 9000 adopted a combination of NSR and Graceful Restart (GR).

Two types of high availability tests were conducted; i.e. soft (initiating the failover from a command line interface) and hard failover were conducted.

5.1 Soft Failover

Test Objective

The goal of this test was to measure the impact on active services as a result of a soft failover on the Route Processor (RP) component.

A controlled soft failover is essential for any In-Service Software Upgrade (ISSU). The hardware and software architecture of a resilient node should ensure that a soft (controlled) RP failover happens without incurring any loss of service.

Test Description

The test was performed using the service configuration and scale described in Section 4 running bidirectional traffic.

We initiated a failover using CLI commands and measured the impact on traffic. Failover using the CLI only failed the RP and did not interrupt the switch-fabric forwarding. Please note that the controlled soft failover test is the best case scenario as it allows the system to guarantee the synchronization of the Standby RP before the switchover.

We allowed the test to run for several minutes beyond the failover time to ensure full system recovery which included complete resynchronization between the new active and standby RP.

Test Results

During the test, the Alcatel-Lucent 7750 SR experienced severe outages on most services ranging anywhere from 91 seconds to 280 seconds depending on the service type. Severe packet loss was observed and this occurred around 8 minutes after the RP switchover took place.


MAC flooding was detected for some of the Layer 2 services, which was a result of the flushing of Layer 2 tables and slow MAC re-learning afterwards.

The Cisco ASR 9000 gracefully recovered from the RP failover with no traffic interruption across all services, showing impressive stability and performance.

Overall the Alcatel-Lucent 7750 SR took 245% longer (22 minutes) to fully restore the system after the failover whereas the Cisco ASR 9000 took only a total of 9 minutes. A system is considered fully restored once it has been brought back to a state where both RP cards are fully operational. Restoration time is an essential parameter in failover scenarios, since it directly affects the time that needs to elapse between two subsequent failovers, which in turn influences the duration of a maintenance window for software upgrades.


Figure 3 below shows results of a soft failover for the Alcatel-Lucent 7750 SR and Figure 4 shows the results for the Cisco ASR 9000. Note that the Alcatel-Lucent 7750 SR chart displays traffic loss at the beginning of the outage, which happens few minutes after the failover was triggered.

Figure 3: Alcatel-Lucent 7750 SR – Soft Failover (CLI Initiated)



Duration of outage in the case of soft failover (CLI initiated) for all services except native multicast and VoD, displayed by service. Outage by service ranged anywhere from 91 seconds to 280 seconds depending on service type. Flooding occurs for the BGP-VPLS service in Customer Edge (CE) to Provider Edge (PE) direction. Full system restoration was not achieved until 22 minutes after failover was initiated.

Figure 4: Cisco ASR 9000 – Soft Failover (CLI initiated)



No disruption of services was experienced by the Cisco ASR 9000 during a soft failover (CLI initiated) of the RP.

5.2 Hard Failover

Test Objective

The goal of this test was to measure the impact on active services of a hard failover on the Route Processor / Switch Fabric which may result from a hardware malfunction or any other unplanned failure.

While such occurrences are not desired in any Service Provider environment, a resilient system architecture must be able to address such failures by ensuring minimal to zero packet loss.

Test Description

The test was performed by using the service configuration and scale discussed in Section 4, running bidirectional traffic.

We simulated a RP failure by pulling the active RP card and measuring the impact on traffic.


Similarly to test 5.1, we allowed this test to run for several minutes beyond the failover time to ensure full system recovery.

Test Results

The Alcatel-Lucent 7750 SR showed traffic disruption for most services at the time the RP card was pulled, due to the loss of switch fabric. Additionally, around 8 minutes into the failover, it experienced a massive outage as described earlier in Section 5.1 during the soft failover scenario. This validates the instability and unpredictable behavior demonstrated by the Alcatel-Lucent 7750 SR.


MAC flooding was detected for some of the Layer 2 services, as a result of flushing of Layer 2 tables and slow MAC re-learning afterwards (represented in Figures 5 and 6 below).

Figure 5: Alcatel-Lucent 7750 SR – Hard Failover



The inset shows the disruption of services when the RP was physically removed from the Alcatel-Lucent 7750SR.


Figure 6: Alcatel-Lucent 7750 SR – Recovery of Services after Hard Failover



The inset section shows the loss of all services that occurred as the Alcatel-Lucent 7750 SR attempted to rebuild route tables following the hard failover. All traffic was recovered after 14 minutes following the physical removal of the RP.

The Cisco ASR 9000 was able to demonstrate zero packet loss across all services. Please note one unique aspect of the Cisco ASR 9000 architecture was its ability to preserve all packets that were in transit through the system during the RP removal. The data plane operation was not impacted in anyway. The Cisco ASR 9000 maintained all services operational without any issue, as shown in [Figure 7](#) on the following page.

Figure 7: Cisco ASR 9000 – Hard Failover



Services were not impacted on the Cisco ASR 9000 as a result of the RP physical failover.

6.0 Path Failure Convergence

To minimize the impact of a link failure, a network is frequently designed with multiple load-sharing paths. Link resiliency, deployed using either link bundling or Equal Cost Multiple Paths (ECMP) or both, is the next logical step toward the self healing and fast converging networks that today's service providers demand.

Test Objective

The goal of this test was to measure convergence time for the different services as a result of core link failures when multiple equal cost paths exist.


Test Description

We performed this test using the service configuration and scale discussed in Section 4 running bidirectional traffic. Throughput on these link-bundle interfaces was kept at less than 50% so as to avoid oversubscription of the active bundles after a failover.

As seen in *Figure 8* below, the physical topology included a total of four ECMP connections between the System Under Test (SUT) and an upstream Core (P) router. Each path consisted of a pair of links bundled together to form a single logical forwarding entity, commonly referred to as a bundle interface. By using bundles instead of physical links, the test was able to evaluate the efficiency of the hashing algorithm used for routing re-convergence towards the remaining alternate paths, as well as that of the algorithm used for the selection of the bundle member ports. Such efficiency can be easily assessed by observing the outages experienced by traffic traveling in the Customer Edge (CE) to Provider Edge (PE) direction, which is exiting the system over the redundant paths.

Path failures were simulated by physically pulling out the corresponding Network to Network Interface (NNI) linecard, which caused the loss of two out of the four bundles (4 links), and resulted in an Interior Gateway Protocol (IGP) path recalculation.


Figure 8: Test Topology



Test Results

Removal of the NNI card on the Alcatel-Lucent 7750 SR caused traffic flowing from the CE to PE to drop by approximately 50% across most services and for several seconds before recovering. CE to PE traffic convergence is directly influenced by the SUT. Multicast traffic (Native and VoD types) experienced a 7.7 second outage. Such a service interruption is considered lengthy and disruptive for real-time video traffic. Traffic flowing in PE to CE direction had outages ranging from 297 milliseconds to 462 milliseconds across all the service types as shown in [Figure 9](#) below.


Figure 9: Alcatel-Lucent 7750 SR Path Failure Line Card Removal



The graph shows the duration of outage for services on the Alcatel-Lucent 7750 SR as result of a Core facing line card removal. All CE -> PE traffic experienced losses of 4-5 seconds, while PE ->CE traffic saw losses of 297-462 milliseconds.

On the Cisco ASR 9000, we observed minor traffic losses of up to 566 milliseconds as shown in [Figure 10](#) on the following page. Most services experienced outages of less than 500 milliseconds.

Figure 10: Cisco ASR 9000 Path Failure Line Card Removal




The minimal disruption of services experienced on the Cisco ASR 9000 as a result of a Core facing line card removal is displayed. Most services experienced loss of less than 500 milliseconds.

7.0 BGP Multi-Path Convergence Test

The BGP protocol and BGP-based services such as MPLS Virtual Private Networks (VPNs) has been largely adopted by Service Providers and large enterprises, mainly because BGP provides the ability to learn network path intelligence and to offer additional levels of path convergence and resiliency among redundant Edge Routers or Interior Border Gateway Protocol (iBGP peers). With this feature, the network can learn routes (or prefixes) with the best and alternative paths, which in turn provide the ability to load-share routes and further reduce the impact of the outages during failover to the next best path. Therefore the terminology of fast-convergence or Prefix Independent Convergence (PIC), and the predictability of convergence times regardless of route table size or BGP service type (address-families) are of importance to Service Providers and their customers.

Figure 11: Test Topology used for BGP Multipath Convergence Test



Note that unidirectional traffic from the SUT to the emulated PE was used in this test.

This test measures BGP Multipath (which utilizes BGP Best Path Selection Algorithm) convergence time for IP-VPNv4 and IP-VPNv6 addresses. We used two next-hop paths for all the BGP prefixes advertised from remote Multi Path (MP)-BGP Provider Edge (PE) neighbors emulated by the IXIA test equipment (PE80 - PE87).

The CRS-1 Provider (P) node was configured with 1 VLAN interface for each of the PE nodes (VLAN's 80-87 corresponding to PEs 80-87), thereby ensuring path failures could be emulated for each individual PE (and its advertised prefixes) by simply shutting down the respective interface.

PE80 -PE83 are used as the four primary PEs, advertising unique IP-VPNv4 and IP-VPNv6 prefixes. The same prefix distribution was also used across PE84 - PE87, thereby creating a multi-path forwarding scenario at the SUT. For correct multipath forwarding operations, the FIB table will indicate the next two hop entries for each prefix, and load-balance across the two available PEs. An Interior Gateway Protocol (IGP) link failure on the remote PE link will then result in a forwarding path convergence. The SUT will have to go through its best-path BGP selection algorithm to re-converge the failed prefixes and its corresponding traffic.

A total of 720 IP-VPNv4 and 120 IP-VPNv6 instances (VRFs) per remote PE pair and a total of 2880 IP-VPNv4 and 480 IP-VPNv6 (VRFs) across all 4 PE pairs was emulated. The prefix distribution used in this test is summarized in [Table 1](#). What is relevant for this test is the convergence time for the prefixes learned from one remote PE and four remote PEs where the number of prefixes ratio is also linearly increased by 1:4. The failover time was measured for a single PE and multiple (all four) PE failure cases, for both IP-VPNv4 and IP-VPNv6 prefixes as outlined in the test cases below.

Table 1: Prefix Distribution Summary

Number of Prefixes	Local per VRF	Remote per VRF	1 Remote PE Pair	4 Remote PE Pairs	Total CE + PE
IP-VPNv4	135	135	97,200	388,800	777,600
IP-VPNv6	10	135	16,200	64,800	69,600

7.1 Single PE Failure Test

Test Objective

To measure the convergence time of BGP/VPN prefixes (IP-VPNv4 and IP-VPNv6 address families) during next best path selection resulting from a single path failure to a remote PE in a multi-path environment. This test provides a measure of how dependent convergence is to the number of prefixes withdrawn; in other words, whether the convergence time is a function of the number of prefixes.

Test Description

In this test case, we simulated the failure of a single PE neighbor or path, by shutting down the corresponding interface (VLAN 80) on the Core (P) router; i.e., the CRS-1. This in turn should trigger BGP to converge all the prefixes and its corresponding traffic from the failed path (VLAN 80) to the redundant (or multi) path (VLAN 84), as also shown in the test topology above in [Figure 9](#) respectively.


Test Results

The Alcatel-Lucent 7750 SR experienced losses ranging from hundreds of milliseconds to several seconds depending on the number of prefixes and address family, as shown in [Table 2](#) and the graph in [Figure 12](#) on the following page.

Table 2: Alcatel-Lucent 7750 SR Test Results for Single PE Failure

Traffic Type	Loss	Prefixes
IP-VPNv4	3.72 seconds	97,200
IP-VPNv6	416 milliseconds	16,200

Figure 12: Alcatel-Lucent 7750 SR Single PE Failover



VPNv4 traffic experienced a loss of 3.72 seconds, while VPNv6 traffic experienced a loss of 416 milliseconds when a single PE is failed over to its redundant peer.


The Cisco ASR 9000 showed minimal loss in traffic as shown in Table 3 below and the graph in [Figure 13](#) on the following page. The loss amount remained the same between VPNv4 and VPNv6 prefixes.

Table 3: Cisco ASR 9000 Test Results for Single PE Failure

Traffic Type	Loss	Prefixes
IP-VPNv4	16 milliseconds	97,200
IP-VPNv6	16 milliseconds	16,200

Cisco ASR 9000 demonstrated the ability to re-converge 26 times faster than the Alcatel-Lucent 7750 SR for VPNv6 prefixes, and 232.5 times faster for VPNv4 prefixes.

Figure 13: Cisco ASR 9000 Single PE Failover



Minimal impact is seen on the Cisco ASR 9000 when a single PE is failed. BGP Prefix Independent Convergence (PIC) provides instant convergence in the event of a failure.

7.2 Multiple PE Failure Test

Test Objective

To measure the convergence time of BGP/VPN prefixes (IP-VPNv4 & IP-VPNv6 address families) during next best path selection resulting from a single path failure to a remote PE in a multi-path environment. This test also provides a measure of how dependent convergence is to the number of prefixes withdrawn; that is, whether the convergence time is a function of the number of prefixes.

Test Description

In this test case, we simulated the failure of four PE neighbors, by shutting down the corresponding interfaces (VLAN 80-83) on the Core (P) router; i.e., the CRS-1. This triggers BGP to converge all the prefixes and its corresponding traffic from the failed paths (VLAN 80-83) to the redundant or multipath (VLAN 84-87), as also shown in the test topology above in [Figure 9](#).

Test Results


The Alcatel-Lucent 7750 SR showed outages in the order of several seconds, reaching almost 60 seconds for VPNv4 traffic. As noted before, convergence time increased as the number of prefixes and address family increased, thereby increasing the overall service impact time.

The following results in [Table 4](#) and the graph in [Figure 14](#) on the following page show the exact loss for each type of traffic.

Table 4: Alcatel-Lucent 7750 SR Test Results for Multiple PE Failure

Traffic Type	Loss	Prefixes
IP-VPNv4	54 secs	388,800
IP-VPNv6	3 secs	64,800

Figure 14: Alcatel-Lucent 7750 SR Multiple PE Failover




The magnified areas show the effect of failing multiple PEs on the Alcatel-Lucent 7750 SR. All traffic is impacted and for a longer period since the FIB table is larger. VPNv4 traffic took 54 seconds to recover, while 6vPE traffic recovered in 3 seconds.

The Cisco ASR 9000 again showed an insignificant loss in traffic across all services as shown in [Table 5](#) and the graph in [Figure 15](#) on the following page.

Table 5: Cisco ASR 9000 Test Results for Multiple PE Failure

Traffic Type	Loss	Prefixes
IP-VPNv4	79 milliseconds	388,800
IP-VPNv6	72 milliseconds	64,800

Figure 15: Cisco ASR 9000 Multiple PE Failover



Cisco ASR 9000 experienced minimal disruption to VPNv4 and VPNv6 services, even when all PEs were failed, demonstrating superior performance of Prefix Independent Convergence (PIC) for large FIB tables.

The Cisco ASR 9000 re-converge time was comparable across address families and number of prefixes, and overall 41.6 times faster than the Alcatel-Lucent 7750 SR for VPNv6 prefixes and 683.5 times faster for VPNv4 prefixes

Cisco ASR 9000 achieved less than 100 milliseconds outages consistently, regardless of the number of prefixes, address families, or injected failures. Alcatel-Lucent 7750 SR convergence time (ranging from hundreds of milliseconds to almost a minute) seemed to linearly increase with the number of prefixes and seemed to lack any PIC capability.

Table 6 summarizes both test results for the BGP/MPLS convergence.

Table 6: BGP Multipath Convergence Test Summary

Convergence Time	1 PE Fail/VPNv4	1 PE Fail/VPNv6	4 PE Fail/VPNv4	4 PE Fail/VPNv6
Cisco ASR 9000	16 milliseconds	16 milliseconds	79 milliseconds	72 milliseconds
ALU 7750 SR	3.7seconds	400 milliseconds	54 seconds	3 seconds


Note: We found the Cisco IOS-XR Prefix Independent Convergence (PIC) feature to be a powerful and innovative feature. It significantly reduced convergence times and packet loss for BGP/VPN networks. PIC re-engineers how forwarding information is represented in the data plane (FIB) to make re-convergence faster and simultaneous for share-fate prefixes regardless of their number. The use of a hierarchical FIB architecture across BGP and IGP entries results in faster convergence times regardless of table size. Furthermore, it enables the system to scale better by utilizing less memory and CPU resources. This is a tool that network engineers can use as they prepare to re-design or expand their networks which are experiencing explosive growth in IP and packet traffic.

8.0 FIB Convergence Speed Test

Network changes happen continuously. It is important for today's networks to be able to react and process route changes quickly in response to any network event. During a link failure and resulting convergence, packet loss due to packets routed to the wrong interface lead to service disruptions and SLAs may be affected. Therefore it is essential for a router to learn and apply any routing changes as quickly as possible and in a predictable amount of time to meet SLAs.

This test was designed to verify the Forward Information Base (FIB) convergence speed of the SUT. Different systems may use varying types of compression or lookup trees on the data plane that may cause different performance for any FIB insert or remove operation. *Figure 16* shows the test topology used.

Figure 16: FIB Convergence Test Setup



8.1 Sequential FIB Convergence

Test Objective

The objective here is to test the SUT's ability to react and process route changes quickly in response to a given network event. In this test, routes are injected in a sequential and controlled manner.

Test Description

The test was conducted by injecting BGP routes but the results are applicable to any routing protocol. The test was run with a sequential FIB table first and then again using a random FIB table (see Test 8.2) to show any possible difference in convergence speed. We measured the time it takes for the routes to be actually installed and used in the FIB on the line card. The sequential test uses the best case routing table scenario. All routes are in sequence, with no gaps and all the same network masks. Approximately 10% of the announced routes are converged.

Testing began with 585,000 routes spread between three egress ports. After the initial convergence we sent test traffic to 65,000 of the prefixes reachable through the first egress port (Out 1). The test then announced more specific subnets for these 65,000 routes on the second egress port (Out 2) and measured the convergence time until the egress traffic completely moved to the second egress port. This was repeated one more time on the third egress port (Out 3) before additional, more specific routes are withdrawn again and all test traffic is back to the very first egress port (Out 1).


Test Results

The test showed both platforms converged in a short amount of time. The Alcatel-Lucent 7750 SR exhibited a slightly faster convergence for the sequential FIB per second convergence test as shown in [Table 7](#):

Table 7: Alcatel-Lucent 7750 SR Time to Converge 65,000 Sequential BGP Routes

Test Run 1	Test Run 2	Average Routes/Second
Announce	Announce	
11 seconds	11 seconds	5909
14 seconds	14 seconds	4643
Average FIB entries learned per second:		5276
Withdrawal	Withdrawal	
12 seconds	12 seconds	5417
8 seconds	10 seconds	7313
Average FIB entries withdrawn per second:		6365

Figure 17: Sequential FIB Convergence – Alcatel-Lucent 7750 SR




Rapid convergence of sequential BGP routes for the Alcatel-Lucent 7750 SR. Total traffic rate was 6.5 Mpps (100 pps per prefix).

Conversely, the Cisco ASR 9000 demonstrated a slightly longer sequential FIB per second convergence, as shown in [Table 8](#).

Table 8: Cisco ASR 9000 Time to Converge 65,000 Sequential BGP Routes

Test Run 1	Test Run 2	Average Routes/Second
Announce	Announce	
16 secs	17 secs	3944
23 secs	20 secs	3038
Average FIB entries learned per second:		3491
Withdrawal	Withdrawal	
15 secs	16 secs	4079
15 secs	19 secs	3623
Average FIB entries withdrawn per second:		3851

Figure 18: Sequential FIB Convergence – Cisco ASR 9000



Rapid convergence of sequential BGP routes for the Cisco ASR 9000. Total traffic rate was 6.5 Mpps (100 pps per prefix).

Note: Measurements for this test was done in one second intervals, so there may be an error of up to two seconds in the convergence time.

8.2 Random FIB Convergence

Test Objective

The objective here is to test the SUT's ability to react and process route changes quickly in response to a given network event. In this test we will inject routes in a random manner.

Test Description

The random FIB test uses the worst case routing table scenario. All routes are randomized, with gaps, overlapping networks and different network masks. The same number of routes, as with the sequential test, is used and converged.


Test Results

The Cisco ASR 9000 converged the routes in a similar amount of time as the sequential test conducted in Section 8.1. However the Alcatel-Lucent 7750 SR took much longer for the convergence, as compared to the sequential FIB convergence time, as seen in [Table 9](#).

Table 9: Alcatel-Lucent 7750 SR Times to Converge Random FIBs

Test Run 1	Test Run 2	Average Routes/second
Announce	Announce	
224 secs	240 secs	280
257 secs	257 secs	253
Average FIB entries learned per second:		267
Withdrawal	Withdrawal	
241 secs	236 secs	272.5
162 secs	160 secs	403.5
Average FIB entries withdrawn per second:		338

Figure 19: Random FIB Convergence – Alcatel-Lucent 7750 SR



The Alcatel-Lucent 7750 SR shows much slower convergence when random BGP routes are announced. Total traffic rate is 6.5 Mpps.


Note: The Alcatel-Lucent 7750 SR Management Console provided statistics that are basic and incomplete. For example, we suspected that the CPU on the line card was the bottleneck in this test. We were unable to check and confirm the line card CPU utilization in the management console. The main CPU on the Route Processor was not being stressed at the time. Additionally, verifying the FIB table on the line card frequently only gave the error message “CLI resources busy - try again later” during the convergence. However we were able to continue with the test and 585,000 routes were announced successfully (reported as 59% FIB Current Occupation).

The Cisco ASR 9000 took the same amount of time for the random routes convergence as the sequential route convergence. [Table 10](#) on the following page summarizes the random route convergence times.

Table 10: Cisco ASR 9000 Time to Converge Random FIBs:

Test Run 1	Test Run 2	Average Routes/Second
Announce	Announce	
19 secs	16 secs	3742
17 secs	16 secs	3943
Average FIB entries learned per second:		3843
Withdrawal	Withdrawal	
14secs	14 secs	4643
14 secs	14 secs	4643
Average FIB entries withdrawn per second:		4643

Figure 20: Random FIB Convergence Cisco ASR 9000



Convergence is shown when random BGP routes are announced for the Cisco ASR 9000. The convergence time was consistent for both sequential and random tests. Total traffic rate was 6.5 Mpps.

This test confirmed that even though Alcatel-Lucent 7750 SR converged slightly faster in the sequential test, the Cisco ASR 9000 demonstrated consistent convergence speed during both tests. The consistent fast convergence times of the Cisco ASR 9000 allows network engineers to build a predictable network with guaranteed SLAs. The Alcatel-Lucent 7750 SR showed unpredictable convergence times down to less than 300 routes per second for announcements or withdraws – 14 times slower than the Cisco ASR 9000. While the convergence speed on the Route Processor was fast, routes were very slowly installed on the line cards. This caused severe discrepancy between the forwarding tables as shown on the RP and those used on the line cards.

9.0 Fabric Priority Test

Strict adherence to assigned network traffic priorities is a must for enterprise and carrier class networks. High priority traffic needs to be protected at the expense of lower priority traffic, especially when a network becomes congested. Most networks are over-provisioned for Best Effort (BE) traffic, so congestion is part of regular network operations. Video and Voice applications are sensitive to traffic loss and are the highest priority services in a network. There can be no packet loss for these services. Differentiation between high priority traffic is required since video and voice have different SLAs, latency and jitter tolerance.


Test Objective

Our objective here was to test the SUT's ability to prioritize and protect traffic levels defined in the system. We verified Fabric Priority Levels and validated strict fabric priority adherence.

Test Description

The topology used is shown in Figure 21 below.

Figure 21: Fabric Priority Test Setup



Traffic with different priorities was sent from different ingress line cards to one single egress line card. Exact topology will differ between the two systems under test, since the test is to oversubscribe the fabric. The Alcatel-Lucent 7750 SR with SFM3-500G fabric and IOM3 line cards supports 50Gbps per slot, while the Cisco ASR 9000 supports 80Gbps per slot.


No routing protocol is involved and all traffic is sent from directly connected Layer 3 interfaces to other directly connected interfaces. All traffic is IPv4 unicast with 1,500-byte packets.

9.1 Fabric Priority on Alcatel-Lucent 7750 SR with Two Priorities

Test Description

Alcatel-Lucent 7750 SR claims two fabric priorities (Priority and Best Effort) and supports 50 Gbps per line card with IOM3 and CP/SFM3 fabric cards. This test verifies the fabric for strict priority.

Figure 22: Alcatel-Lucent 7750 SR Test Topology



As shown in Figure 22, traffic was sent from 2 line cards, each with 4 x 10Gbps interfaces, to a single Egress card with 8x10Gbps Interfaces. Ingress traffic from one line card was marked Best Effort (BE) while the traffic from the other line card was marked High Priority (HP). We gradually increased, in 0.25Gbps increments, traffic from both ingress line cards to the same egress card until we exceeded the egress line card fabric throughput (50Gbps) and verified which traffic was maintained versus dropped.

Under normal circumstances, if the SUT supports strict fabric priority, when the egress linecard reaches the oversubscription point, only the BE traffic will be dropped and 100% of HP traffic should continue to egress the system.


Test Results

Alcatel-Lucent 7750 SR traffic ramped up without loss until it reached 52Gbps on the egress card. When the egress line card capacity was exceeded, both BE and HP traffic were dropped, as shown in [Table 11](#) and [Figure 23](#).

Table 11: Alcatel-Lucent 7750 SR Priority Traffic Test Results Table

Total attempted Egress Traffic (Gbps)	Attempted Egress Traffic per QoS Class (Gbps)	Best Efforts Traffic Loss %	High Priority Traffic Loss %
36	18	0	0
40	20	0	0
44	22	0	0
48	24	0	0
50	25	0	0
52	26	3.85	0
54	27	10.00	0.78
56	28	13.53	4.32
58	29	16.51	7.62
60	30	19.30	10.71
62	31	21.90	13.59
64	32	24.40	16.29
66	33	26.63	18.82
68	34	28.79	21.21
70	35	30.82	23.46
72	36	32.75	25.59
74	37	34.56	27.60
76	38	36.29	29.51
78	39	37.92	31.31
80	40	38.66	32.14

Figure 23: Alcatel-Lucent 7750 SR Fabric Priority with Two Priorities



Alcatel-Lucent 7750 SR starts dropping High Priority traffic in addition to Best Effort traffic when the egress fabric begins to become oversubscribed. The rate of losses continues to increase as the fabric becomes oversubscribed.


The Alcatel-Lucent 7750 SR loses packets for both BE and HP traffic starting at 52Gbps total traffic. There is no protection of HP traffic in such conditions, indicating that Alcatel-Lucent 7750 SR Fabric Priority mechanisms cease to function as soon as egress fabric is over-subscribed (between 50Gbps and 52Gbps total traffic). Above 52Gbps, the additional traffic is equally dropped between both classes.

9.2 Fabric Priority on Cisco ASR 9000 with Two Priorities

Test Description

Cisco ASR 9000 supports up to three fabric priorities (Priority1, Priority2, and Best Effort) and supports 80Gbps per line card on the Cisco ASR 9000. This test verified the fabric's ability to handle High Priority and Best Effort traffic similar to the test conducted in 9.1, except the traffic levels were changed to accommodate the 80Gbps throughput per line card capacity instead of 50Gbps supported by the Alcatel-Lucent 7750 SR.

Figure 24: Cisco ASR 9000 Test Topology



As shown in *Figure 24*, traffic was sent from two line cards, each supporting 8x10Gbps interfaces to a single egress line card with 8x10Gbps interfaces. The ingress traffic from one line card was marked Best Effort (BE) while the traffic from the other ingress line card was marked High Priority (HP or P1). We gradually increased traffic from both ingress line cards to the same egress card, up to and exceeding the egress line card throughput 80Gbps and verified traffic behavior.

Under normal circumstances, if the SUT supports strict fabric priority, then only the BE traffic will be dropped in the event of congestion. We expect the SUT to properly handle traffic based on priorities assigned and to protect High Priority traffic at the expense of BE traffic.


Test Results

At 40Gbps (per class) traffic throughput, the Cisco ASR 9000 showed 2.5% BE traffic packet loss and 0% HP traffic loss, as displayed in *Table 12* on the following page. As the traffic level increased, the BE traffic loss increased as well, but the HP traffic was always protected up to the maximum rate at which time all the BE traffic was dropped. The results in *Table 12* and *Figure 25* on the next page show how the Cisco ASR 9000 maintained HP traffic at all times.

Table 12: Cisco ASR 9000 Priority Traffic Test Results

Total attempted Egress Traffic (Gbps)	Attempted Egress Traffic per QoS Class (Gbps)	Best Efforts Traffic Loss (%)	High Priority Traffic Loss (%)
72	36	0	0
76	38	0	0
80	40	2.50	0
84	42	12.00	0
88	44	21.00	0
92	46	28.37	0
96	48	35.52	0
100	50	42.10	0
120	60	68.42	0
140	70	87.22	0
156	78	98.78	0
160	80	99.99	0

Figure 25: Cisco ASR 9000 Fabric Priority with Two Priorities



Cisco ASR 9000 protects High Priority traffic (HP) during oversubscription of up to 80Gbps per slot, and dropped only the BE traffic as expected under congested conditions.


When oversubscribed, the Cisco ASR 9000 ports begin to experience loss at 40Gbps (80Gbps total) as expected. There is no HP traffic loss experienced throughout the test. HP traffic is protected at all traffic rates, up to 160Gbps total traffic (80Gbps HP traffic – full egress bandwidth). This is critical as service providers offer a variety of services which include real-time video services combined with Internet access.

9.3 Fabric Priority on Cisco ASR 9000 with Three Priorities

Description

The Cisco ASR 9000 supports 3 strict priority levels on the fabric, which helps to distinguish and guarantee different high priority traffic (e.g. Voice and Video). This test verified the three strict priority levels (High Priority1, High Priority2, and Best Effort).

Figure 26: Cisco ASR 9000 Test Topology for Three Priority Levels



As shown in [Figure 26](#), traffic was sent from three line cards, each supporting 8x10Gbps interfaces to a single egress card with 8x10Gbps interfaces. Ingress traffic from one line card was marked Best Effort (BE) while ingress traffic from the other two line cards were marked HP1 and HP2 respectively. Traffic was gradually increased from all ingress line cards to the same egress card, until we exceeded the egress line card throughput 80Gbps and verified traffic behavior.

Under normal circumstances, if the SUT supports strict fabric priority, then only the BE traffic (first) followed by HP2 traffic (next) would get discarded as a result of congestion, ensuring that HP1 traffic is maintained at all times.


Test Results

As shown in the results in [Table 13](#) and [Figure 27](#) on the following pages, we see the Cisco ASR 9000 did protect High Priority traffic against any loss. Traffic marked with ToS 0 (BE traffic) was the first to experience any loss, followed by traffic marked ToS 5 (HP 2), ensuring all traffic marked ToS 7 (HP1) was protected at all times.

Table 13: Cisco ASR 9000 with Three Fabric Priority Levels Test Results

Total attempted Egress Traffic (Gbps)	Attempted Egress Traffic per QoS Class (Gbps)	Loss		
		ToS 0 (BE)	ToS 5 (HP2)	ToS 7 (HP1)
72	24	0.00	0.00	0.00
84	28	18.02	0.00	0.00
96	32	53.28	0.00	0.00
108	36	80.70	0.00	0.00
120	40	100	2.60	0.00
132	44	100	20.57	0.00
144	48	100	35.52	0.00
156	52	100	48.18	0.00
168	56	100	59.02	0.00
180	60	100	68.42	0.00
192	64	100	76.64	0.00
204	68	100	83.90	0.00
216	72	100	90.35	0.00
228	76	100	96.12	0.00
240	80	100	99.99	0.00

Figure 27: Cisco ASR 9000 Fabric Priority with Three Priorities



With three traffic priority levels, Best Effort traffic is dropped first, while ToS 5 and ToS 7 traffic is protected. As oversubscription conditions become more severe, ToS 5 traffic begins to be dropped while ToS 7 is still protected. Up to 80Gbps top priority traffic is protected, demonstrating strict fabric priority.

The Cisco ASR 9000 exhibits stability in its ability to prioritize traffic as seen in these tests. The three levels of fabric priority on the Cisco ASR 9000 allow service providers to guarantee different levels of traffic with oversubscribed Best Effort (BE) traffic. The two strict fabric priorities allow for different protected traffic as needed today for voice and video.

Alcatel-Lucent 7750 SR did not demonstrate the ability to support strict priorities across the fabric, showing that the current architecture cannot guarantee any high priority traffic in a congested network.

10.0 Multicast

The demand for next-generation Broadcast TV and Packetized Digital Video distribution has been growing, leading to 3D video distribution and bringing resolution requirements as high as 4,000p. These are bandwidth demanding applications and require anywhere from 40Mbps up to 6Gbps of Multicast throughput delivery per channel. High resiliency and availability is also essential to offer an optimal Video Quality of Experience (QoE), and become very important parameters to be considered on a Provider Edge (PE) type platform.

The tests in this section provide insight into Multicast throughput, convergence, and resiliency on the Cisco ASR 9000 and the Alcatel-Lucent 7750 SR.

The Cisco ASR 9000 was built from the ground up with a focus on large-scale and resilient Multicast distribution. Cisco innovation and engineering to the highest degree have been employed in the design of software, line cards, and system architectures to meet the robustness and high-speed requirements of multicast-based applications for today and in the future. The Alcatel-Lucent 7750 SR showed far more limitations and instability in the delivery, convergence and resiliency of high-throughput multicast traffic as shown in the following test cases.

10.1 Multicast Throughput Test

Test Objective

To measure the maximum multicast throughput that a linecard on the SUT can achieve. Evaluate the packet loss and latency in the delivery of the traffic.

Test Description

The test was designed to use a 1:1 correlation between ingress and egress ports involved in the multicast traffic distribution. This implies that all multicast flows entering a port in the system (source port) will only be forwarded out of one other port in the system (receiver port). There is no multicast replication involved (as shown in [Figure 28](#)).


The source and receiver ports connect directly to the IXIA tester ports, which will generate and receive multicast traffic up to line rate port speed (10Gbps) or maximum forwarding rate without dropping any packets (No Drop Rate). The multicast routing protocols utilized in this test are dynamic IGMPv3 (Internet Group Management Protocol version 3) and PIM-SSM (Protocol Independent Multicast Source Specific Multicast). The Alcatel-Lucent 7750 SR required an additional feature, called Ingress Multicast Path Manager (IMPM), to be enabled on its ingress line cards to support multicast throughput rates greater than 2Gbps.

The first part of the test only uses a single multicast stream or (S, G), sourced to a single receiver, to measure the maximum throughput supported for a single multicast stream.

The second part of the test, each receiver port will receive (or join via dynamic IGMPv3) 125 multicast streams for a total of 500 (S,G) unique groups across the 4x10GE ports on an egress line card. Traffic rate will be adjusted to reach a state in which no packet-loss or instability is observed over an extended period of time.

Test Topology

Figure 28: Multicast Throughput Test Topology



Test Results

A range of throughput tests were conducted from a minimum value of 80Mbps (single multicast stream used) to maximum of 40Gbps (500 streams used), across 4x10GE ports per line card. We removed traffic until we found a stable level for each unit without any loss or latency.

Baseline Throughput Test (One Multicast Channel up to Max Port Speed)

The Alcatel-Lucent 7750 SR was only able to support up to 2Gbps of throughput for a single multicast stream and would drop the stream (“blackhole,” not forward its traffic) if its throughput exceeded this threshold, as seen in [Figure 29](#) on the following page. This seems to be a fabric limitation on the Alcatel-Lucent 7750 SR.


Maximum Throughput Test (500 Multicast Channels at 80Mbps per Multicast Channel)

The Alcatel-Lucent 7750 SR was not able to achieve line rate throughput at 40Gbps between two IOM3-XP line cards. A total of 500 multicast streams (or groups) were used at 80Mbps per group, and the Alcatel-Lucent 7750 SR system was constantly dropping (“blackhole”) almost all multicast streams as seen in [Figure 30](#) on page 50. Frequent syslog messages also flooded the router console with “channel blackhole” messages. It appears that the Alcatel-Lucent 7750 SR was not able to guarantee or protect any multicast channel delivery at high throughput.

Alcatel-Lucent 7750 SR continued to be very unstable even as low as 24Gbps with minimal packet loss being observed as seen in [Figure 31](#) on page 51.


The Cisco ASR 9000 performed all tests without traffic loss or latency (see [Figure 32](#) on page 52).

Figure 29: Alcatel-Lucent 7750SR Multicast Throughput




Each fabric channel on the Alcatel-Lucent 7750 SR is limited to 2Gbps. Traffic sent at a rate of 2,100 Mbps was blackholed, reducing throughput to zero.

Figure 30: Alcatel-Lucent 7750SR Multicast Throughput at 40Gbps




With 500 multicast streams at a traffic rate of 80Mbps per stream, Alcatel-Lucent 7750 SR “blackholed” many channels and packet loss was observed.

Figure 31: Alcatel-Lucent 7750SR Multicast Throughput at 24Gbps showing Minimal Loss



The maximum multicast rate that the Alcatel-Lucent 7750 SR could sustain without packet loss was 48Mbps on each of the 500 streams for a total of 24Gbps. At this rate, random channels were still being “blackholed,” as seen in the enlarged portions above.

Figure 32: Cisco ASR 9000 Multicast Throughput at 40Gbps



Cisco ASR 9000 was able to sustain a rate of 80Mbps for each of 500 multicast streams simultaneously without incurring any packet loss, proving that it can provide line rate throughput of video multicast traffic.

10.2 Multicast Convergence

Test Objective

Measure the time taken by multicast streams (or groups) being broadcasted simultaneously to converge or get forwarded across the system from the ingress port to its corresponding egress port. Also measure the impact of packet loss on existing multicast streams, on different ports, as new streams were being added. The best throughput from the previous test (10.1) is used as the benchmark for this multicast convergence test.

Test Description

Multicast traffic was received from the tester on the 4x10GE source ports sequentially. Since the best throughput achieved for the Alcatel-Lucent 7750 SR was ~24Gbps across 4x10GE ports per IOM3-XP line card, 6Gbps of traffic at 125 multicast groups (S, G) per 10GE port was forwarded sequentially. In the case of the Cisco ASR 9000, the best throughput used was line rate throughput at 40Gbps per line card or 10Gbps per port.

The same test topology was used as in the previous section (10.1).


Test Results

As seen from the results in [Figure 33](#), every time a new port receives multicast traffic on the Alcatel-Lucent 7750 SR, it affects existing multicast traffic across all existing ports with an average traffic-loss of 11 seconds.

The Cisco ASR 9000 did not experience any traffic loss when new multicast traffic was enabled nor did it impact traffic already flowing on existing ports ([Figure 34](#)).


This test was run for an extended period of time (1 hour) after convergence of all traffic and we noticed the Alcatel-Lucent 7750 SR continued to drop (“blackhole”) traffic and experienced an intermittent 5 milliseconds of packet loss across every multicast stream. The Cisco ASR 9000 did not exhibit any packet loss and was consistent and predictable in its behavior.

Figure 33: Alcatel Multicast Convergence



Alcatel-Lucent 7750 SR experienced packet loss initially when new multicast (S, G) streams are received on the first port. As subsequent ports also receive traffic, it affects traffic delivery of already received multicast streams. Every channel, each of which represents a broadcast video customer, is affected. Furthermore, the IMPM seems to continue black-holing traffic (as seen in middle dip of graph), trying to accommodate new streams with existing streams across the fabric-planes.

Figure 34: Cisco ASR 9000 Multicast Convergence



No packet loss was seen on the Cisco ASR 9000 during this test.

10.3 Multicast Line Card Redundancy Test

Test Objective


The goal is to measure multicast redundancy and resiliency across different source ports on two different ingress line cards. In order to provide path redundancy for downstream multicast traffic, for the same multicast (S,G) traffic streams, dual ports across different line cards are necessary.

Test Description

Two ingress line cards were used to provide redundant 500 (S,G) streams to a common receiver line card. The receiver line card ports, in turn, will only select and forward one of the two 500 (S,G) streams. We then failed one of the two ingress line cards and observed the convergence time for all the multicast streams. The test ran at a fixed duration of 300 seconds. The test was run on the Alcatel-Lucent 7750 SR at a lower, more stable rate of 24Gbps on both the ingress line cards, and run at 40Gbps line rate on the Cisco ASR 9000.

Test Topology

Figure 35: Multicast Line Card Redundancy Test Topology




Test Results

When one of the two ingress line cards was physically removed from the system, we observed that the Alcatel-Lucent 7750 SR experienced approximately 7.3 seconds of loss.

A loss of approximately 180 milliseconds was observed on the Cisco ASR 9000.


It was expected that some loss would be experienced from physically removing the line card due to multicast convergence, but the Cisco ASR 9000 converged 40 times faster than the Alcatel-Lucent 7750 SR.

Figure 36: Alcatel-Lucent 7750SR Redundant Ingress Multicast Line Card Resiliency



Alcatel-Lucent 7750 SR experienced approximately 7.3 seconds of packet-loss during multicast path convergence.

Figure 37: Cisco ASR 9000 Redundant Ingress Multicast Line Card Resiliency



The Cisco ASR 9000 experienced approximately 180 milliseconds of packet loss during multicast path convergence.

10.4 Multicast Fabric/Route Processor Redundancy and Resiliency

Test Objective

To measure the multicast traffic-loss during the redundant Route Processor/Switch Fabric card failover.

Test Description

The multicast traffic rate used on the Alcatel-Lucent 7750 SR was set to 24Gbps for stability purposes while the Cisco ASR9000 was set to 40Gbps line rate across all 4x 10GE ports for all 500 S, G streams cumulative.


Two tests were performed:

- 1) The card hosting the primary or active RP will be failed/removed.
- 2) The card hosting the standby RP will be failed/removed.

Note: The RP operates in active/standby mode, and the Switch-Fabric operates in active/active mode for both the Alcatel-Lucent 7750 SR and Cisco ASR 9000.

Test Topology

Figure 38: Multicast Fabric/RP Redundancy and Resiliency Test Topology




Test Results

When the primary/active RP card was pulled to initiate the failover, the Alcatel-Lucent 7750 SR experienced approximately 10 seconds of packet loss. The following graph indicates how the system restored traffic on the standby/secondary RP (see [Figure 39](#) on the next page). Also when we reinserted the card, the Alcatel-Lucent 7750 SR experienced 2 seconds of packet loss.


Upon pulling the standby card, the Alcatel-Lucent 7750 SR experienced an average of 0.5 seconds loss across all channels. Some channels did not experience any loss, while others experienced as much as 6.5 seconds of packet loss. (See [Figure 40](#) on the next page.)

Figure 39: Alcatel-Lucent 7750SR Multicast Primary RP Failover



Alcatel-Lucent 7750 SR experienced 10 seconds of packet loss when the primary RP was failed by physical removal. This test was performed at a traffic rate of 24Gbps.


Figure 40: Alcatel-Lucent 7750SR Multicast Standby RP Failover



When the Standby RP was removed from the Alcatel-Lucent 7750 SR, multicast streams across all ports experienced an average of 0.5 seconds packet loss, while some channels experienced as much as 6.5 seconds packet loss.


When the active RP card was pulled to initiate the failover on the CiscoASR 9000, it experienced no loss and traffic continued without interruption(see [Figure 41](#)). With the reinsertion of the card, the Cisco ASR 9000 had no loss and no change in traffic flow as well. When the standby RP card was pulled on the Cisco ASR 9000, we observed no change in traffic pattern or loss(see [Figure 42](#)).

Figure 41: Cisco ASR 9000 Multicast Primary RP Failover



No loss was observed on the Cisco ASR 9000 when the primary RP was failed over by physical removal.

Figure 42: Cisco ASR 9000 Multicast Standby RP Failover



Standby RP failure experienced no packet loss on Cisco ASR 9000.

The Cisco ASR 9000 was able to failover all multicast traffic without loss or impact to users, despite physical removal and reinsertion of cards. The Alcatel-Lucent 7750 SR experienced loss during both the process of removal and reinsertion of cards. Traffic on the Alcatel-Lucent 7750 SR was impacted to varying degrees in a random fashion with some channels experiencing heavy loss and others none.