TESTED AT ISOCORE INTERNETWORKING LAB

Isocore Summary Report

Cisco ASR1000 – Aggregation Service Router Product Evaluation Report

Interface Scaling for IPv4 and IPv6
Multicast Replication
High-Availability for Forwarding
Route Scalability

ISOCORE Internetworking Lab
ISOCORE Technical Document
Reference: ITD: 13005
Version (v1.0): 2/17/08

12359 Sunrise Valley Drive, Suite 100
Reston, VA 20190
Phone: 703.860.1777
Fax: 703.860.1778
www.isocore.com

For more information about the testing capabilities of the Internetworking lab, please email testing@isocore.com.

About Isocore
Isocore provides technology validation, certification and product evaluation services in emerging and next generation Internet and wireless technologies. Isocore is leading validation and interoperability of novel technologies including Carrier Ethernet, IPv6, IP Optical Integration, wireless backhauling and Layer 2/3 Virtual Private Networks (VPNs) and currently focuses on IPTV service deployment architecture validation and design. Major router and switch vendors, Service Providers, and test equipment suppliers participate in Isocore activities. Isocore has major offices in the USA (the Washington DC area), Europe (Paris, France) and Asia (Tokyo, Japan).

1 EXECUTIVE SUMMARY

This report presents the highlights of an independent product evaluation conducted by Isocore of Cisco ASR 1000 series Aggregation Service Router (ASR) based on a request by Cisco. The testing primarily focused on evaluating the ability of the Cisco ASR platform to support a large number of layer 3 interfaces. In addition, the testing included verification of the Cisco ASR High-Availability implementation which ensures minimal or no traffic loss during a forwarding or route processor failure, BGP route scalability, and multicast replication to a large number of groups. The testing performed was in accordance with Isocore's stringent requirements for validating aggregation routers for their ability to support large number of interfaces, subscribers, and maintaining hitless data forwarding.

Testing was conducted at Isocore's state-of-the-art facilities in the Washington Metro area in a completely independent environment. The testbed was created from scratch and comprised of two Cisco ASR 1006 platforms. For this test evaluation, Isocore used the Ixia XM2 test solution.

Key Findings:

During the tests, the Cisco ASR1000 series router demonstrated that it meets the interface scalability as per specifications. The key findings of the evaluation include:

1. ASR1006 successfully a added large number of virtual IPv4/ IPv6 layer 3 interfaces and maintained zero traffic loss traffic while forwarding across all interfaces during long test run hours
2. ASR1006 successfully demonstrated its high-availability feature set guaranteeing zero packet loss during a primary route processor (RP) failure and sub milli-second traffic loss while primary forwarding processor (FP) switchover.
3. ASR1006 successfully supported data replication to a large number of multicast groups
4. Considering the pre-production equipment under test, no abnormal behavior observed during the scaled configurations testing
5. ASR1006 successfully allowed configuration of a large number of queues, reverse path forwarding (RPF), and access list configuration on each of layer 3 interfaces configured on the platform. This further confirms the ASR's ability to filter and control traffic through policies for an enterprise or service provider network
6. ASR1006 successfully demonstrated that its routing information base (RIB) can hold a large number of IPv6 and IPv4 routes

Based on the observed results and experience gained during the evaluation process Isocore believes that the Cisco ASR 1006 platform successfully demonstrated dual-stack IPv4 and IPv6 interface scalability, ability to support large number of access control lists (ACLs), multicast data replication, high-availability, and a scaled existence of Quality of Service for both IPv4 and IPv6 unicast and multicast simultaneously. The consistency of observations made and stability demonstrated by the system during the test establishes the readiness of the ASR 1006 platform to be deployed for tested services in a real–world environment.

Table of Contents

List of Figures

2 TEST OVERVIEW AND SUMMARY

As part its independent edge router evaluation program, Isocore tested Cisco ASR 1000 series routers for their ability to be deployed as either an edge router, for WAN aggregation, or as a BGP router-reflector. During the evaluation, Isocore focused primarily on these three aspects which required testing the ASR1000 for utmost interface scalability, application of QoS, RPF and ACLs across all configured interfaces. In addition, verification of high-availability feature set and multicast replication capabilities was carried out.

Table 1 provides summary observations made during the entire test effort. The IPv4/ IPv6 scaling observations are presented separately for simplicity, however a dual-stack scaling test was carried out with all virtual interfaces configured for both IPv4 and IPv6 concurrently. The test for IPv4 and IPv6 could have been carried out separately, but one of the key benchmarks of the testing was also to verify whether the Cisco ASR router would be able to handle large configurations while handling QoS and ACLs for both IPv4 and IPv6 traffic flows.

Table 1: Summary of Test Results

Verified Capabilities	Observations/Measured Results
IPv4 Interface Scaling	• **16,000 Layer 3 IPv4 Virtual Interfaces** o With 1000 VLANs/Port • **16,000 VLANs** • **Each Interface configured with 2000 rules entries** o 1000 entries for IPv4 ingress per ACL o 1000 entries for IPv4 egress per ACL • **Unicast RPF (uRPF) applied to all 16,000 Virtual Interfaces** • **80,048 Queues configured across 16,000 interfaces** o Child Policy – 2 queues/interface o Parent Policy – 3 queues/interface • **Traffic Forwarding verification** o 32,000 bi-directional flows created o No traffic loss observed o Each interface sending and receiving IPv4 traffic
IPv6 Interface Scaling	• **16,000 Layer 3 IPv6 Virtual Interfaces** o With 1000 VLANs/Port • **16,000 VLANs** • **Each Interface configured with 2000 ACE entries** o 1000 entries for IPv4 ingress per ACL o 1000 entries for IPv4 egress per ACL • **Unicast RPF (uRPF) applied to all 16,000 Virtual Interfaces** • **80,048 Queues configured across 16,000 interfaces** o Child Policy – 2 queues/ interface o Parent Policy – 3 queues/ interface • **Traffic Forwarding verification** o 32,000 bi-directional flows created o No traffic loss observed o Each interface sending and receiving IPv6 traffic
Route Scaling	• **1.1 Million IPv4 and IPv6 Routes** o 1 Million Unique IPv4 BGP routes o 100,000 IPv6 OSPFv3 Routes

High-Availability Verification	• **Non-stop Forwarding - RP switchover** o No service loss observed when primary RP was failed o The test utilized bi-directional flows across 10G interfaces in two frame sizes – 64 byte and 1518 byte • **Non-stop Forwarding - FP switchover** o Worst measured switchover time of 0.02ms for bi-directional flows with 64 byte frames o Worst measured switchover time of 0.03ms for bi-directional flows with 1518 byte frames
IPv6 Multicast Replication	• 1000 groups replication verified • MLD Joins for 1000 groups on 10 virtual egress ports • Worst latency measurements of 77.07 µs for 68 byte frames

The results observed during the independent evaluation confirm that Cisco ASR series routers offer a high degree of resiliency, high-availability, and significant dual-stack interface scalability for supporting a large number of IPv4 and IPv6 subscribers at the edge of carrier-class networks. To summarize, even though only a sub-set of capabilities were verified, the observed test results establish the Cisco ASR 1006 as an optimal WAN aggregation, route-reflector or an edge router which can handle varying bandwidth requirements, hardware redundancy, offer enhanced hierarchical QoS functionality, BGP route scalability, multicast support, and hitless forwarding plane functionality.

2.1 CISCO ASR 1000 SERIES AGGREGATION ROUTER OVERVIEW

Cisco ASR1000 series routers are available in 3 different form factors: Cisco ASR 1002, 1004, and 1006. The Isocore testbed utilized only Cisco ASR 1006 for the entire evaluation. Other variants were not offered for the testing. The tested chassis supports 10Gbps throughput (egress) performance for both IPv4 and IPv6 flows, and is equipped with redundant RPs (Route Processors) and FPs (Forwarding Processors). Cisco ASR1006 is the only variant in the ASR family that offers hardware redundancy for RP and FPs. This is the primary reason for using ASR 1006 for this test series. Cisco ASR1006 also offers software redundancy, but verification of this capability was out of scope of this test.

The Cisco ASR1006 forwarding plane capability is based on its embedded services processor (ESP) engine which forms the central FP. Depending upon the application and performance, different types of FPs could be used. For this test Isocore tested the Cisco ASR1006 with the ESP10 type of forwarding processor. Cisco ASR offers distributed control plane architecture with central processing unit (CPU) presence on both RP and FPs. The CPUs run IOS XE software consisting of a Linux-based Kernel. The Cisco IOS process is responsible for the control plane processing including processing network configuration and CLI, running routing protocols, computing routes, managing interfaces and tunnels, and session setup.

Both Cisco ASR1006 routers used for this testing were loaded with the same version of Cisco IOS XE software - 12.2(33) XN. The primary device under test (DUT) – Cisco ASR1006 chassis was equipped with dual RPs, dual FPs (ESP10), and 2 SIP10 housing one 10GigE SPA, one 10-port 1GigE SPA, and two 5-port 1GigE SPAs.

2.2 TEST ENVIRONMENT AND SETUP

The evaluation was conducted by Isocore at its Internetworking Lab. The entire testbed was designed and created by Isocore engineers. Figure 1 provides an overview of the test network setup that was used for the interface and route scalability evaluation. In the setup used, the two Cisco ASR 1006 routers were provisioned as service provider edge routers. The primary DUT was configured with multiple gigabit Ethernet interfaces and a couple of 10 gigabit Ethernet interfaces. This configuration gave the flexibility to conduct the evaluation in a system under test (SUT) and DUT environments. Both nodes were connected to Ixia XM2 platform via gigabit Ethernet interfaces. Ixia XM2 chassis was equipped with XMV16 module cards for simulating the access connections. Considering the scale, a large number of access ports were used in the test.

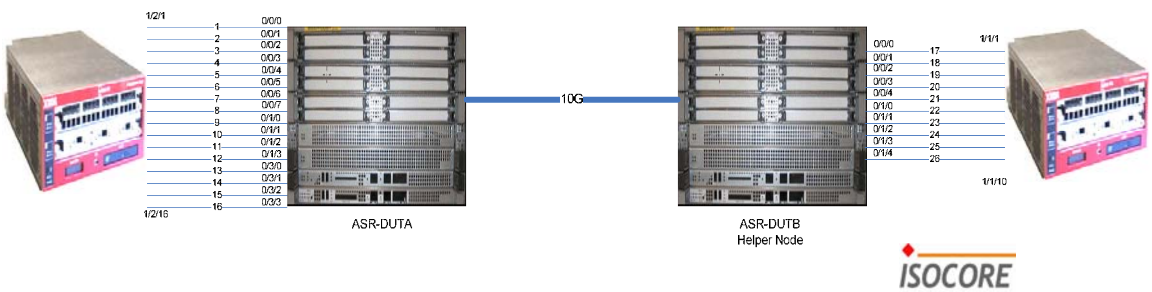

Figure 1: Cisco ASR 1006 Interface and route scaling Verification Test Bed

Two setups were used during the entire test, the testbed shown in Figure 1 was used primarily for interface scaling test while the testbed shown in figure 2 was used for multicast replication and High-availability testing. Both setups were configured from the ground-up and no canned configurations were used during the testing.

Figure 2: Isocore Cisco 1006 Verification testbed for HA and multicast replication test

Isocore selected the Ixia XM2 test solution for this testing considering the port scalability of the XMV cards used in the setup. The Ixia XM2 two-slot chassis offered the flexibility of either using up to 32 1-gigabit Ethernet ports or a combination of 16 1-gigabit Ethernet and 3 10-gigabit Ethernet ports. IxNetwork was used for protocol emulation, traffic generation and analysis. The virtual LAN interfaces were configured to emulate Ethernet service subscribers across 16 1-gigabit Ethernet ports. No triple-play type of traffic was configured during this test, only basic IPv4 and IPv6 flows were created for data plane verification. For multicast replication test Tcl scripts were used to configure Ixia to send MLD joins on 10 VLANs for 1,000 groups on each VLAN across 10-gigabit Ethernet interfaces.

3 TEST DETAILS AND OBSERVATIONS

This section provides the overview of the tests that were conducted as part of this test effort. The tests were designed to verify the scalability of the Cisco ASR 1000 series routers and evaluate the ability of the platforms to adapt as a WAN aggregation, edge router and BGP route-reflector in any real-world network. Even though ASR can be positioned in other network segments of a network (such as providing managed services real-time voice and video services), the specific tests necessary to verify these functionalities were not included in this test but may be verified in near future. Isocore testing of ASR 1000 evaluated the following areas –

a. Verify the scalability of layer 3 IPv4 and IPv6 interfaces along with QoS and Access list applications
b. Benchmark the switchover times for both route processors and forwarding processors thus confirming the high-availability capabilities of the platform
c. Verify Multicast replication to a large number of groups.
d. Test Co-existence of large number of IPv4 routes along with IPv6 routes

The following sections provide brief overview of test methodologies, objective, and results of each of these tests.

3.1 IPV4 AND IPV6 INTERFACE SCALING

The primary objective of this test was to evaluate the scalability of layer 3 interfaces on the Cisco ASR1006 router shown as DUT A in Figure 1. The idea was not only to scale the number of layer 3 dual-stacked interfaces, but also to verify the bi-directional forwarding across scaled configured interfaces. The test also required configuring hierarchical QoS on all virtual interfaces with ingress and egress ACL rules. In addition, Unicast Reverse Path Forwarding (uRPF) needed to be configured on all the interfaces. The requirement of QoS was included to verify whether Cisco ASR1006 as a WAN aggregation router can support hierarchical queue structure at maximum scale on all interfaces. The test however did not require verifying the functionality of the QoS, but only to verify the QoS configuration at a scale.

For the test configuration, 16 1-gigabit Ethernet interfaces on Cisco ASR1006 were connected to Ixia, with each interface configured for 1000 Virtual interfaces (VLAN) making a total of 16,000 configured IPv4 and IPv6 interfaces. High-availability redundancy was also configured. However, the verification of hitless switchover was out of scope of this particular test and was verified in depth separately. All 16,000 interfaces were configured as dual-stack interfaces with IPv6 and IPv4 routing configured concurrently. On each interface, 2 ACLs with 2000 rules for IPv4 were applied, one for ingress traffic and other for the outgoing traffic. Also, on each interface, 2 ACLs with the same configuration as IPv4 was applied for IPv6 making a total of 4 ALCs with 4000 rules on each configured 16,000 virtual interfaces. Additionally, uRPF (unicast RPF) configuration was added to test the ability of Cisco ASR1006 to block or limit the malicious traffic when deployed in a real enterprise environment. For QoS configuration, each subscriber policy used 5 queues: 2 queues defining the child policy and 3 queues defining the parent policy. Each main interface policy used 3 queues. Figure 3 shows the calculation of the total number of total queues configured[1] on the system.

The QoS queues configured on Cisco ASR1006 –
Number of queues per virtual interface = 5
Total number of virtual interfaces configured = 16,000
Total number of queues across 16,000 VLANs – 16,000 * 5 = 80,000
Number of main interfaces used in the test = 16
Number of main interface policies = 3
Total number of queues configured on the system = (16,000*5) + (16*3) = 80,048 queues

Figure 3: Total number of queues configured on the Cisco ASR1006 – DUT A

For data plane verification the primary requirement was to show that subscribers attached to each virtual interface send and receive traffic. To achieve this, 16 ports on Ixia were configured as

[1] The QoS configuration was only used as background; no shaping, and policing or QoS verification was performed to test the behavior of QoS implementation of Cisco ASR1000.

remote access end points with dynamic ARP resolution. Traffic flows were created from 16,000 end points towards 8 access ports connected to Cisco ASR1006 DUTB (refer figure 1) and vice-versa. This setup offered 32,000 unique flows based on unique destination and source pairs. Latency was also measured across 32,000 flows and the test was run overnight to see if any packet drop was observed. Similar traffic streams were configured for IPv6 subscribers and traffic was verified. However, static neighbor discovery was configured for IPv6. Even though all 16,000 interfaces were configured for dual-stack IPv4 and IPv6 addressing, the traffic verification was performed separately for IPv4 and IPv6 to measure latency of IPv6 flows and IPv4 flows separately.

Table 1 summarizes the performance and scaling results of the IPv4 and IPv6 interface scaling test.

Table 1: Results of the IPv4/IPv6 Scaling Test

Dual-Stack Interface Scaling: IPv4/ IPv6
IPv4 configured VLAN Scaling (Per DUTA) • 16,000 Layer 3 IPv4 Virtual Interfaces with 1000 VLANs/Port • Each Interface configured with 2000 rules entries o 1000 entries for IPv4 ingress per ACL o 1000 entries for IPv4 egress per ACL • Unicast RPF (uRPF) applied to all 16,000 Virtual Interfaces • Traffic Forwarding verification o Traffic run duration 14 hrs and 59 minutes o 32,000 bi-directional flows created o No traffic loss observed o Worst observed Latency across 32,000 flows = 65.846 µs o Frame size of 68 bytes o Each interface sending and receiving IPv4 traffic o Average throughput across FP = 9,110,320pps **IPv6 configured VLAN Scaling (Per DUTA)** • 16,000 Layer 3 IPv6 Virtual Interfaces with 1000 VLANs/Port • Each Interface configured with 2000 rules entries o 1000 entries for IPv6 ingress per ACL o 1000 entries for IPv6 egress per ACL • Unicast RPF (uRPF) applied to all 16,000 Virtual Interfaces • Traffic Forwarding verification o Traffic run duration 16hrs and 16minutes o 32,000 bi-directional flows created o No traffic loss observed o Worst observed Latency across 32,000 flows = 65.651 µs o Frame size of 90 bytes[2] o Each interface sending and receiving IPv4 traffic o Average throughput across FP = 7,307,157 pps **QoS configuration Across 16,000 Virtual Dual-Stack Interfaces** • 80048 Queues configured across 16,000 interfaces o Child Policy – 2 queues/interface o Parent Policy – 3 queues/interface o 3 main interface policies • **Overall Results** o No traffic loss observed either in IPv4 or IPv6 forwarding test

[2] Due to the instrumentation limitation, the IPv6 frame size was not lowered further.

- Both tests were run for overnight to monitor the abnormal behavior, if any. During the overnight runs no issues were observed
- No signs of stress observed on CLI performance during the overnight runs

Throughout the test the Cisco ASR1006 proved to be stable considering large dual-stack configuration with a large number of ACLs, and queues configured on the system. Also, the overnight run added a further level of credence to the stability of the platform and its suitability deployment in a real world environment.

3.2 NON-STOP FORWARDING VERIFICATION: SOFTWARE-INDUCED SWITCHOVER

The objective of this test was to verify the high-availability implementation of Cisco ASR1000 series routers. Cisco ASR1006, which was used in the testbed, was configured with dual RPs and dual ESPs (FP10). This was necessary to verify the hardware redundancy. The test required taking the active RP and FP out of service using CLI commands and monitoring the impact on the traffic forwarding. For this testing, the setup shown in Figure 2 was used. The Cisco ASR1006 was configured with an IPv6 address with static neighbor pointing towards Ixia. IPv6 streams were configured and the test was repeated at two frame sizes (1518 bytes and 64 bytes). Two scenarios were verified as part of this test:

Case 1: Monitor the impact on the traffic flows while switching over forwarding processor.
Case 2: Monitor the impact on the traffic flows while switching over the route processor.

For case 1, initially the state of both FPs were identified. One of the FPs (FP0) was in active/ primary mode and other FP (FP1) was in standby mode. Following the identification, traffic was initiated from Ixia to verify the baseline forwarding performance. Once the traffic reached steady state, the primary FP0 was forced to stop/ reload using CLI commands on the Cisco ASR. This forced the FP1 to switchover to the new active/ primary FP. Following the successful transition, the traffic was stopped and the traffic loss was noted. The observed results of this test are provided in the Table 2.

Table 2: HA FP Switchover Results

	Frame Size = 64 bytes		Frame Size = 1518 bytes	
	Port 0/1/0	Port 1/0/0	Port 0/1/0	Port 1/0/0
Packets TX	1,380,329,883	1,380,315,779	28,159,827	28,159,546
Packets RX	1,380,315,638	1,380,329,741	28,159,536	28,159,816
Packet-Rate (pps)	7,440,474	7,440,474	405,844	405,844
Frames Lost	141	142	10	11
Loss Interval (in ms)	0.02	0.02	0.02	0.03

The total time loss (in ms) was calculated as follows:

Total time loss (in ms) = (packet loss/ packets per second) * 1000

Table 3: HA RP Switchover Results

	Frame Size = 64 bytes		Frame Size = 1518 bytes	
	Port 0/1/0	Port 1/0/0	Port 0/1/0	Port 1/0/0
Packets TX	366,493,672	366,489,928	50,220,297	50,219,796
Packets RX	366,489,928	366,493,672	50,219,796	50,220,297
Packet-Rate (pps)	7,440,474	7,440,474	405,844	405,844
Frames Lost	0	0	0	0
Loss Interval (in ms)	0	0	0	0

For case 2, the same methodology as case 1 was followed, but this time active/ primary RP was failed over while the traffic was being forwarded. Table 3 provides the results of this test.

The tests confirmed that there was no traffic loss when switching over the route processors but there was a sub milli-second hit when the forwarding processor was switched over. Both tests involved software-induced switchovers.

3.3 MULTICAST REPLICATION VERIFICATION

The objective of this test was to verify the IPv6 multicast replication capability of the Cisco ASR1006 platform. The test setup used for this verification is shown in figure 2.

The test adopted a simple methodology in which the Cisco ASR1006 (DUT) platform was configured for 10 VLANs on one of the two ports to act as an IPv6 multicast egress interface. Ingress interface is configured with an IPv6 address as well. The test was automated to distinguish between the actual received tests packets from the control packets. The automated script that was used for the test configured the Ixia platform to send the MLD joins on the 10 configured VLANs for 1000 groups on each VLAN. Following the configuration multicast traffic script was used which employed a binary search algorithm to find the no drop rate and the latency for the multicast traffic sent across the DUT.

Table 4 presents the summarized results observed during this test.

Pkt Size (incl. Tag)	Total Pkt send	Total Pkt rcvd	Pkt per sec	Avg. Latency (ns)	Max Latency (ns)	Result
68	852,270,000	852,270,000	710,227	60,162	69,320	Pass
68	1,065,340,000	1,065,340,000	887,784	62,331	72,020	Pass
68	1,171,870,000	1,171,870,000	976,562	64,720	77,700	Pass

The results confirmed that the Cisco ASR1006 successfully demonstrated multicast replication for IPv6 SSM (source-specific multicast) groups.

3.4 ROUTE SCALABILITY VERIFICATION

The objective of this test was to verify whether Cisco ASR 1000 series BGP routing information base (RIB) could store 1 million unique entries while maintaining a large number of IPv6 routes. The test setup used for this test is shown in figure 1. However, only a pair of 1-giagbit Ethernet ports was used for the test. The methodology required to verify the installed routes using the CLI show commands. Only control plane scale was verified during the test and no forwarding verification was verified during this test.

The Cisco ASR 1006 platform successfully demonstrated its ability to scale to 1 million unique IPv4 routes across single eBGP session in addition to 100,000 unique IPv6 routes learnt via OSPFv3 peering set up with an Ixia XM2 chassis using IxNetwork emulation software. The BGP routes advertised from Ixia resembled the real Internet routing table. The routing protocol test validated that Cisco ASR 1000 series routers can store a large number of both IPv4 and IPv6 routes and the total scale that was verified with a combined routing table size of 1.1 million IPv4 and IPv6 routes. IPv4 routes were learnt via eBGP Peering, and IPv6 routes were learnt via OSPFv3 adjacency.

4 CONCLUSION

Cisco ASR1000 series routers are a part of Cisco next generation networking portfolio and caters to both service provider and enterprise markets. The features evaluated as part of this testing are applicable to these edge applications and WAN aggregation only.

Isocore, through this evaluation, was able to successfully verify the dual-stack interface scalability along with Cisco ASR1000 series routers ability to support large QoS and ACL configurations. One of the key characteristics of this testing was the duration of each test. The data verification for interface scaling tests was carried out for more than 12 hours through overnight runs. These longevity tests add another level of stability acceptance that ties to the consistency of the platform behavior and readiness to the operational deployment of the platform. During these long test run times no abnormal behavior or performance degradation was observed, further cementing high-availability capability of the Cisco ASR 1000 series routers. Also, the successful verification non-stop forwarding feature set of the ASR1006 clearly demonstrated the separation of data and control planes through no drop forwarding of user traffic when a router processor is failed. However, it does take a service hit of a less than 50 microseconds when a primary forwarding processor is failed. The failures that were induced to verify the capability were only software-triggered.

Multicast replication is very important, especially if a platform is positioned for multi-play services and visual networking. The Cisco ASR1006 in the Isocore test passed the multicast replication test while keeping the latency to an acceptable value. Finally, route scalability was also verified for the control plane only, and the Cisco ASR1006 comfortably scaled to the observed results for routing table capacity.

Isocore concludes through observed results that the Cisco ASR 1006 platform offers a combination interface scalability, high-availability, and apt multicast traffic handling capability along with large route storage capacity. These areas must necessarily be supported by all platforms that are being positioned as a WAN aggregation, edge or a route-reflecting router in any carrier-class network.

5 LIST OF ACRONYMS

TE	Traffic Engineering
SPA	Shared Port Adaptor
SIP	SPA Interface Processor
RP	Route Processor
FP	Forwarding Processor
ESP	Embedded Service Processor
RIB	Routing Information Base
FIB	Forwarding Information Base
GE	Gigabit Ethernet
SP	Service Provider
IPv6	Internet Protocol version 6
IPv4	Internet Protocol version 4
OSPF	Open Shortest Path First
BGP	Border Gateway Protocol
ASR	Aggregation Service Router