



Isocore Technical Report

Validation of Cisco ASR 1000

**BGP Route-Reflector Control Plane Scaling and Convergence
Layer 3 VPN Provider Edge Scale and Performance**

**ISOCORE Internetworking Lab
ISOCORE Technical Document
Reference: ITD: 13029
Version (v1.1): 3/20/09**



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 Isocore Internetworking lab, please email testing@isocore.com.

This document should not be reproduced as a certification or a validation reference from Isocore. Permission is granted to Cisco to reproduce this document in its current state and form.

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 key findings of an independent product evaluation executed by Isocore for the Cisco ASR 1000 series Aggregation Service Router (ASR). In February 2009, Isocore was commissioned by Cisco Systems to evaluate the ASR1000 latest route processor (ASR1000-RP2) for validating its route-reflector control plane scalability and convergence, and verifying layer-3 VPN provider-edge (PE) scalability and performance.

For this test, Isocore set forth future scaling and performance requirements for a Route-Reflector (RR) and a PE router role in a service provider network. The event involved testing of ASR1000 - focusing on the latest route-processor (RP2) for its ability to support a very large dedicated route-reflector deployment. The test scope additionally included verification of a scaled MPLS/BGP based layer-3 VPN PE deployment with significant number of customer edge (CE) devices connected through Virtual LAN (VLAN) sub-interfaces within the virtual forwarding and routing (VRF) instances, with QoS, security access lists (ACL) and Unicast Reverse Path Forwarding (uRPF) enabled on each VLAN sub-interface.

The testing also exercised the forwarding capability of ASR1000 embedded services processor (ASR1000-ESP20) engine. ESP20 forms the central forwarding processor (FP) for ASR1000.

Key Findings:

During the tests, the Cisco ASR1000 series router exhibited that it meets or exceeds the scalability requirements set forth for this test. The testing was split into two parts: the first part focusing on BGP RR convergence and scalability testing, while the second part involving the layer-3 VPN PE performance and scalability aspects. The key findings of the evaluation include:

1. All BGP Route-Reflector, and PE layer-3 VPN scalability and performance tests were successfully executed
2. Dedicated BGP RR Test Results
 - a. ASR1000 successfully demonstrated Route-Reflector route scalability for all BGP address-families –
 - i. 20 Million unique IPv4 routes with 1000 RR-clients, and 20,000 routes per client
 - ii. 20 Million unique VPNv4 routes (spread across 8000 route-targets) with 200 RR-clients, and 100,000 routes per client
 - b. ASR1000 successfully scaled to 8000 RR-client sessions with a combination of IPv4/VPNv4/IPv6/VPNv6 address families and routes
 - c. Demonstrated Route-Reflector control plane convergence of 75 seconds with 1000 RR-clients and 1 Million IPv4 routes
3. Layer-3 VPN Test Results
 - a. ASR1000 impressively showcased the scalability and performance of a Provider-edge router that can handle a sizeable routing table, VRF instances and eBGP PE-CE sessions
 - i. Scaled to 4,000 VRFs with 1 M unique VPNv4 routes and 300,000 IPv4 routes with 2 multi-paths
 - ii. 8,000 Virtual LAN interfaces
 - iii. Scaled to 8,000 PE-CE eBGP sessions
 - iv. Security ACLs, uRPF and ingress QOS policy configured on all 8,000 VLAN sub-interfaces with 1000 unique QOS policies, 4000 unique access list and 84,000 access list entries
 - v. No-drop data plane verification performed successfully to all advertised routes

Table of Contents

1 EXECUTIVE SUMMARY.....	3
2 TEST OVERVIEW AND SUMMARY.....	5
2.1 CISCO ASR 1004 –DEVICE UNDER TEST	6
2.2 TEST ENVIRONMENT AND SETUP	6
3 TEST DETAILS AND OBSERVATIONS.....	7
3.1 CONTROL PLANE SCALING AND CONVERGENCE FOR ROUTE-REFLECTOR APPLICATION.....	7
3.2 PROVIDER-EDGE VERIFICATION TO DEMONSTRATE VRF SCALING	10
4 CONCLUSION.....	12
LIST OF ACRONYMS	13

List of Figures

Figure 1: Cisco ASR 1004 Control Plane and provider-edge testing	7
Figure 2: Cisco ASR1004 RP2 Route-Reflector Convergence and Scaling Test Setup.....	9
Figure 3: Cisco ASR1004 RP2/ESP20 Provider Edge Test Setup	11

2 TEST OVERVIEW AND SUMMARY

For this evaluation, a test setup was prepared to offer large number of peering RR-clients, CE end points, route-reflectors, provider (P) and remote PE routers for validating the capabilities of ASR1000 series router. Cisco ASR1000 was tested for its ability to be deployed as a dedicated BGP router-reflector and a layer-3 VPN PE router. During the evaluation, the tests focused primarily on these two aspects, which required testing the ASR1000 for utmost BGP control-plane scalability, and support for large number of layer-3 Virtual Private Networks (VPNs). In addition, verification of forwarding processor was performed to ensure the handling of data forwarding at different levels of class-of-service while limiting malicious traffic coming from any enterprise network.

Table 1 provides summary observations made during the entire test effort. The testing was split into two parts, first focusing on RR control plane scaling and convergence, and the second one focused on the multi-service capabilities of the ASR1000 – primarily targeting the Layer 3 PE scale for VRFs and the number eBGP sessions.

Table 1: Summary of Overall Results

Verified Capabilities	Observations/Measured Results
RR Control Plane Scaling Verification	<ul style="list-style-type: none"> • IPv4 address family scaling – Routes reflected <ul style="list-style-type: none"> ○ Total of 20 Million unique IPv4 Routes learnt from 1000 clients with 20,000 routes/client ○ Clients simulated by Ixia IxNetwork ○ Tested with dedicated BGP route-reflector feature for IPv4 that prevents the routes learned from BGP to be installed in the routing table or forwarding table. This feature reduces memory usage and CPU utilization for dedicated RR application purposes • VPNv4 address family scaling – Routes reflected <ul style="list-style-type: none"> ○ Total of 20 Million unique VPNv4 Routes learnt from 200 clients with 100,000 routes/client ○ Clients simulated by Ixia IxNetwork • BGP RR-Client session scalability <ul style="list-style-type: none"> ○ 8,000 sessions with 590,000 unique routes ○ Advertised routes consisted of IPv4, VPNv4, IPv6, and VPNv6 address families ○ 250,000 unique IPv4 routes with 4 paths for each prefix ○ 300,000 unique VPNv4 routes with 8,000 route-targets ○ 20,000 unique IPv6 routes ○ 20,000 unique VPNv6 routes ○ All routes advertised, and learnt by Ixia IxNetwork • No Traffic forwarding verification performed
RR Control Plane Convergence Verification	<ul style="list-style-type: none"> • BGP Control Plane Convergence with 1000 peers, and 1 Million unique routes <ul style="list-style-type: none"> ○ IPv4 route convergence - 75 seconds ○ VPN4 route convergence - 221 seconds ○ IPv6 route convergence - 194 seconds ○ VPNv6 route convergence - 293 seconds

PE Service scaling and Performance Verification	<ul style="list-style-type: none">• 4,000 Layer-3 VPNs with 1M unique VPNv4 routes<ul style="list-style-type: none">○ 4,000 VRFs configured on ASR1000 PE○ iBGP peering with 2 emulated Router-reflectors while simulating routes learnt from 200 remote PEs.○ Total 8,000 CE connected to ASR1000 PE through 8,000 VLAN sub-interfaces○ Total of 8,000 PE-CE eBGP sessions○ Security ACLs, uRPFs, ingress QoS policy configured on all 8,000 VLAN sub-interfaces with total of 1000 unique policy-maps, 4000 unique access lists○ No traffic loss observed in the overnight run to all destinations advertised for the directions of local CEs to remote PEs and vice versa
--	--

The results observed during the independent evaluation confirm that Cisco ASR series routers offer a high degree of scalability and great convergence with large number of peers/clients while excelling at the service scale and forwarding performance. It is also confirmed through this exercise that ASR offers features that are critical for preventing malicious traffic on to the carrier network. To summarize, even though only targeted feature-sets of multi-service ASR were verified, the observed test results establish the Cisco ASR 1000 as one of the leading route-reflectors or edge routers which can handle large number of services, and a very high BGP scale. These are two key requirements of a multi-service provider-edge, and a route-reflector respectively.

2.1 CISCO ASR 1004 –DEVICE UNDER TEST

Cisco ASR1000 series routers are available in 3 different form factors: Cisco ASR 1002, 1004, and 1006. For this test, Cisco offered 2- ASR1004 with two configurations of (RP2) processors, one with 16GB of physical memory and the other with 8GB of physical memory. RP2 processors are considered second-generation route-processors for ASR1000, which runs 64-bit IOS-XE operating system. Two setups were used for convenience, one used for RR testing, and other used for PE service scaling testing. Both set-ups used one Cisco ASR1004, and one Ixia XM2 multi-service test chassis.

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. Cisco ASR1000 Series also offers software redundancy, but verification of this capability was out of scope of this test.

First setup used one Cisco ASR1004 with a single ASR1000-RP2, ASR1000-ESP20, SPA-5X1GE-V2 and ASR1000-SIP10. The second setup utilized one Cisco ASR1004 with a single ASR1000-RP2, ASR1000-ESP20, SPA-8X1GE-V2, SPA-1X10GE-L-V2 and two ASR1000-SIP10. Both setups used derivatives of experimental version of ASR release 3 image for RP2 since the test was conducted before the first customer shipping date for RP2.

2.2 TEST ENVIRONMENT AND SETUP

Figure 1 provides an overview of the test network setup that was used for both RR and PE tests. In the setup used, the sole DUT was configured as RR for the control plane testing, and as multi-

service PE for the provider-edge testing. The DUT was connected to XM2 chassis through both Gigabit Ethernet (GE) and 10-Gigabit Ethernet (10GE) interfaces. Ixia XM2 chassis used for this test was equipped with LSM1000XMV16 GE module, and LSM10GXM8X 10GE module with a total of 16 GE and 8 10GE ports. For emulating the entire test network, we chose to use IxNetwork 5.30 SP3 Patch 1.



Figure 1: Cisco ASR 1004 Control Plane and provider-edge testing

IxNetwork was used for protocol emulation (including RR clients, CE, RR, and remote PEs), traffic generation and analysis. The virtual LAN interfaces were configured to emulate CE end points within the VRFs of ASR1004. No triple-play type of traffic was configured during this test, however both unlabeled and labeled IPv4 traffic was generated. IPv4 traffic utilized 4 different class of services.

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 validate the scalability of ASR1000 RP2 control plane for RR application, and verifying the performance of ASR1000 in a provider-edge role. For the effort specific tests were created to assess the convergence of ASR1000 by varying number of routes from different BGP address families and number of RR-clients peering with the ASR1000. The intention was to investigate the stability of the ASR1000 as focal point for iBGP sessions in large-scale Autonomous System (AS) deployments. Similar to the RR control plane testing, verification of AS1000 as a PE role followed similar methodology. In addition to having a scaled BGP/MPLS layer-3 VPN control plane configuration, a combination of multiple embedded services was also enabled within each VLAN sub-interfaces. These test placed emphasis on following areas:

- a. Verify the control plane capability of the ASR1000 in supporting a considerable number of routes, and BGP client sessions
- b. Benchmark the control plane convergence of IPv4/VPNv4 and IPv6/VPNv6 routes with large number of BGP peers
- c. Verify the ASR1000 RP2/ESP20 layer-3 BGP/MPLS VPNs, PE-CE eBGP session scaling and forwarding performance with scaled QoS ingress policies, security ACLs, and uRPF

For all above tests, emphasis was laid on reliability, repeatability, and sustainability of configured scale on the ASR1000. To ensure that the test results are reliable, at least two iterations were carried for all control plane tests. For forwarding performance tests, the traffic was passed for extended period of time to confirm no-loss packet forwarding.

3.1 CONTROL PLANE SCALING AND CONVERGENCE FOR ROUTE-REFLECTOR APPLICATION

This section provides the details of the ASR1000 control plane scaling and BGP route convergence tests. It also presents the BGP session scalability results as well. Typically, in a large Internet Service Provider (ISP) networks, the internal BGP mesh can become

unmanageable with the increase in the number of BGP speakers. BGP offers two alternatives to solve this scaling problem, one through the configuration of confederations, and the other through the use of route reflectors. Route-reflector is an iBGP speaker that advertises routes learnt from each of its iBGP peers to other iBGP peers. This characteristic eliminates the need for a full iBGP mesh, and the routes are reflected among the non-meshed iBGP speakers. This testing was structured to verify the ability of the ASR1000 to be deployed as a Route-reflector in a scaled network environment. The level of stress a RR is subjected to in any network is dependent on three factors. These include:

- Total number of routes learnt from/reflected to RR-clients
- Number of peering clients or BGP neighbors, and
- The time elapsed while converging learnt routes

Also, these key parameters impact the number of RRs need to be deployed for a given number of clients and routes, and the faster network convergence in case of any unplanned failures or network transients. Considering these as primary concerns for any service provider, the testing focused on evaluating ASR1000 on these three fronts.

This evaluation only focused on RR control plane benchmarking and no data forwarding tests were carried out since dedicated RR is normally never in the forwarding path of BGP learnt routes. However, to ensure that ASR1000 reflected all the learnt routes to its peers/clients, Ixia IxNetwork was used to confirm and validate that each emulated route-reflector client does receive the expected number of routes. Figure 2 shows the detailed setup used to evaluate the ASR1000 route-reflector route/session scalability and convergence. The primary DUT was ASR1004. DUT was connected to Ixia XM2 test system through 4-GE interfaces. These four interfaces were configured in dual-stack (IPv4/IPv6) mode. On each of the Ixia port, ISIS session is used to emulate an IGP cloud, which consisted of large number of iBGP nodes (depending upon the test case, the number of RR clients in each cloud was set to 50, 250 or 500 with a total of 200, 1000 and 2000 peers). ASR1000 was configured to peer with all emulated nodes through the ISIS adjacencies created on each of the four ports. In all convergence and route-scalability tests, BGP peer-groups (with each peer-group containing up to 1000 RR clients) were configured and interface maximum transmission unit was set to 9000 bytes.

Same test setup was used for BGP route scalability test for IPv4/ VPNv4 routes and for BGP route convergence tests for all address families (IPv4, IPv6, VPNv4, and VPNv6) except for BGP session scalability test.

At least two iterations were carried out for convergence tests and if the values for observed times were not consistent additional iterations were carried out. Table 2 summarizes the results from control plane convergence tests.

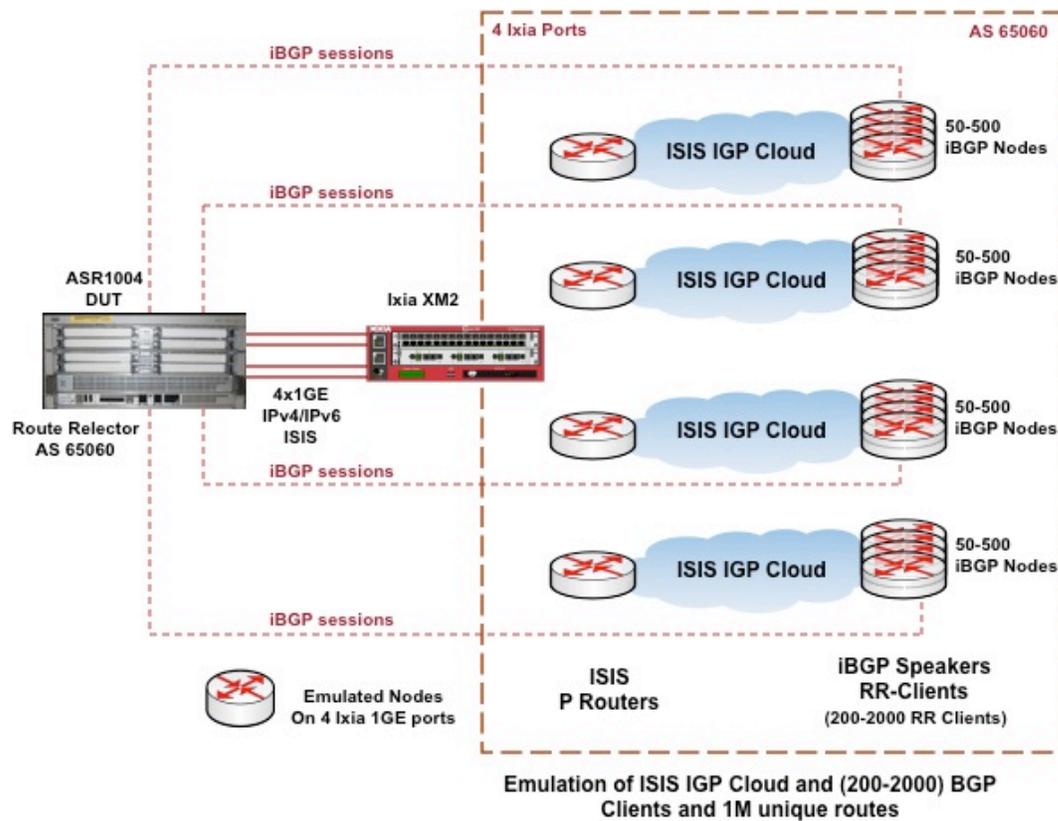


Figure 2: Cisco ASR1004 RP2 Route-Reflector Convergence and Scaling Test Setup

The convergence time for the purpose of this test was defined as - the total elapsed time when all sessions were enabled on the tester until the CPU utilization of the DUT falls to steady state and the tester receives all the advertised routes for each port (which is equal to the multiplication of total routes advertised from all emulated clients and the number of total client/ sessions per tester port). The convergence time presented in table 2 below was measured using this consideration.

Table 2: Control Plane Convergence Results for RR tests

No. of RR-clients/ BGP sessions	Total Number of routes learnt per client	Total Number of Routes learnt from all RR-clients	Address family (For VPN routes, 8,000 Route targets were used)	Convergence Time (seconds)
1000	1000	1Million	IPv4	75
1000	1000	1Million	VPNv4	221
1000	1000	1Million	IPv6	194
1000	1000	1Million	VPNv6	293
2000	500	1Million	IPv4	138
2000	500	1Million	VPNv4	394
2000	500	1Million	IPv6	284
2000	500	1Million	VPNv6	551

For BGP RR session scalability test, additional four Ixia XM ports were added to the setup shown in Figure 2 with a total 8-GE ports being utilized on the ASR1004 DUT. Also, the profile of the routes advertised included all address families mimicking a carrier-class network environment

with major percentage of routes advertised were from IPv4 and VPNv4 address families. IPv6 and VPNv6 routes were also included in the overall mix of routes advertised. Table 3 presents the findings from and route/session scalability tests.

Table 3: Results of the RR control plane route and session scale

RR Control Plane Route and Session Scaling
<ul style="list-style-type: none">• IPv4 Route Scaling<ul style="list-style-type: none">○ 1,000 BGP RR-clients○ 20,000 routes learnt per client/peer○ Total of 20Million unique routes learnt○ Tested on ASR1000-RP2 with 16GB of memory• VPNv4 Route Scaling<ul style="list-style-type: none">○ 200 BGP RR-clients○ 100,000 routes learnt per client/peer○ Total of 20Million unique routes learnt○ Tested on ASR1000-RP2 with 16GB of memory• RR scaling with maximum number of client sessions<ul style="list-style-type: none">○ 8,000 BGP RR-clients○ Total of 590,000 routes learnt○ Route profile included 250,000 unique IPv4 routes with 4 paths for each prefix, 300,000 unique VPNv4 routes, 20,000 routes - each for IPv6 and VPNv6 address families

3.2 PROVIDER-EDGE VERIFICATION TO DEMONSTRATE VRF SCALING

This section presents the results of the evaluation of provider-edge testing performed. In this test the intention was to validate the scalability of Cisco ASR1000 implementation to support scaled Layer-3 MPLS based VPN services while exercising the forwarding plane by hitting all advertised routes with traffic flows. The total number of routes advertised to the ASR1000 for this test was limited by the testbed restrictions to generate very large number of traffic flows.

In addition, verifying the functionality of ingress QoS policies, security ACLs, and uRPF filters that are enabled on all VLAN sub-interfaces was also carried out.

This part of the test focused both on the PE control plane scaling for layer-3 VPN services and forwarding plane verification. Figure 3 shows the detailed setup used for PE testing. As compared to the test setup shown in Figure 2, this contained an additional 10GE core-facing link to emulate the core provider-network with two core routers, two route reflectors, and 200 remote PE devices. Each emulated route-reflector peered with 100 remote PE devices. The P1, and P2 in Figure 3 emulated the core nodes connected to the ASR (DUT) through the Ixia XM2 10GE physical interface. DUT establishes ISIS adjacency with P1 and P2 emulating the IGP cloud, which provides provides the reachability between the RRs, the PEs and the DUT. The IGP connectivity is required for the DUT to establish iBGP sessions to the route-reflectors (RR1, RR2) and to validate nexthop address in the advertised VPNv4 routes from the remote PEs. The LDP sessions between the DUT and the core nodes provides label switched path from the DUT to the remote PE devices for transporting L3VPN traffic. DUT learns identical 300,000 IPv4 and unique 400,000 VPNv4 routes from each of the route-reflectors, for a total of 300,000 IPv4 routes (2 paths per route) and 800,000 VPNv4 routes. The nexthop in the advertised VPNv4 routes is set to emulated PEs. Each emulated remote PE advertises 100 routes per route-target with 40 route-targets per PE. For forwarding plane verification, traffic flows were created to mimic MPLS VPN labeled traffic to routes advertised from the access side.

On the access side, 1,000 virtual interfaces were provisioned on each Ixia XM 1GE port with a total of 8,000 VLAN sub-interfaces on 8 ports to emulate 8,000 CE devices, with 2 CE devices

4 CONCLUSION

The evaluation of ASR1000 series routers confirms the ability of device under test to significantly scale beyond the current industry requirements for any dedicated route-reflector deployment that requires a large number of iBGP sessions, total routes reflected to its client, and total number of routes learnt from all peering clients. The verified fast convergence of routes in a route-reflector application also established that the Cisco ASR1000 would recover quickly from any unplanned network failures with minimum down time. The number of iterations carried out during the evaluation had definitely put stress on the DUT, but during the entire event no signs of performance degradation were observed. The test results also demonstrated the resolute forwarding plane, which was extensively verified during the PE testing, in which Cisco ASR1000 comfortably supported large number of VRFs and PE-CE eBGP sessions. To ensure the reliability of the DUT, the forwarding tests were run for extended duration to constantly monitor for traffic drops, or unexpected jump in the memory or control processor utilization.

The observed results concludes that the Cisco ASR 1000 offers a combination of both control plane scalability for route-reflector applications and PE scalability for supporting large number of enterprise customers. With the consistency of results obtained and the stability of the platform during the entire duration of the testing, it is justifiable to establish Cisco ASR1000 as one of the leading options for a provider edge or a route-reflector router in a carrier-class network.

LIST OF ACRONYMS

RR	Route-Reflector
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
BGP	Border Gateway Protocol
ASR	Aggregation Service Router