Cisco 7600/Catalyst 6500:
Scaling Up for MPLS VPN Service with 10 Gigabit Ethernet

By David Newman
prepared for Cisco Systems
January 2004

CONTENTS

ILLUSTRATIONS

Executive Summary

MPLS VPNs offer service providers many benefits, but come with one big catch: Cost. Devices targeted for deployment in MPLS VPNs typically carry seven-figure price tags, and even then may not scale to deliver services to hundreds or thousands of customers.

Cisco Systems has introduced new components for its Cisco 7600 and Catalyst 6500 devices that offer vastly improved price/performance for MPLS VPN services. The components include the Supervisor 720 management module with the Policy Feature Card PFC3-BXL daughtercard, and the DFC3-BXL distributed switching engine daughtercard for each of the system's line cards.

Cisco commissioned Opus One, an independent benchmarking and network design consultancy, to assess the performance of the Cisco 7600/Catalyst 6500 equipped with the PFC3-BXL and DFC3-BXL switching engines. Cisco supplied the Cisco 7600/Catalyst 6500 with its full complement of 32 10-gigabit Ethernet interfaces.

Although the entire system as tested carries a US list price of less than $500,000, its performance was more in line with far more expensive devices. Among the highlights:

- Routing table capacity of 1 million IPv4 routes, and the ability to forward traffic with zero packet loss to all routes. One million routes is approximately eight times larger than the full Internet routing table

- Routing table capacity of 500,000 IPv6 routes, and the ability to forward traffic with zero packet loss to all routes

- RFC 2547bis - Support for 800,000 routes learned in an MPLS VPN environment

- RFC 2547bis - Support for more than 1,500 MPLS VPNs (VRFs)

- Zero packet loss on stable paths during massive route flaps

- Ability to function as a P (provider) router as well as a PE (provider edge) router

- Low and consistent latency across all tests

If anything, these scalability numbers are conservative. Although CPU and memory utilization were not formal test metrics, our observations of these measurements suggest sufficient memory and CPU cycles exist for the Cisco 7600/Catalyst 6500 to exercise additional software features.

Table 1 on the next page summarizes US list pricing for the device under test. The entire system as tested carried a US list price of less than $500,000.

Table 1: Cisco 7600/Catalyst 6500 US List Pricing

Product ID	Description	List Price	Qty.	Subtotal
WS-C6509	6509 Chassis	$ 9,500	1	$ 9,500
WS-C6K-9SLOT-FAN2	Fan2 (for 6509)	$ 495	1	$ 495
WS-CAC-4000W-US	4000W Power Supply	$ 5,000	1	$ 5,000
WS-SUP720-3BXL	Sup720 with PFC3BXL	$ 40,000	1	$ 40,000
WS-X6704-10GE	4-port 10GE line card	$ 20,000	8	$ 160,000
WS-F6700-3BXL	DFC3BXL daughtercard	$ 15,000	8	$ 120,000
XENPAK-10GB-LR	XENPAKS-LR	$ 4,000	32	$ 128,000
Total List Price				**$ 462,995**

Introduction

Layer 3 MPLS VPNs deliver routed IP services over a service provider's MPLS core network. This gives customers the ability to connect among different sites over what appears to be a private routed IP network. From the service provider's perspective, all traffic shares a common infrastructure. This lets service providers deliver routed IP services to multiple customers from a single router at the provider edge, while significantly reducing the amount of routing state required in the core network.

Layer 3 MPLS VPNs use the RFC 2547*bis* draft specification from the Internet Engineering Task Force (IETF).[1]

We measured the suitability to task of the Cisco 7600/Catalyst 6500[2] as a layer-3 MPLS VPN device using seven different tests:

- IPv4 BGPv4 RIB/FIB capacity
- IPv6 BGP4+ RIB/FIB capacity
- PE/CE scalability for small and midsized customers
- PE/CE scalability for large customers
- PE/CE scalability for global customers
- RFC 2547*bis* MPLS VPN route flap handling
- P router scalability

For all these tests, we fully loaded a Cisco 7600/Catalyst 6500 with 32 10-gigabit Ethernet interfaces, the maximum complement for this chassis, each equipped with the DFC3-BXL distributed switching engine daughtercard. We also equipped the device under test with Supervisor720 module with PFC3-BXL daughtercard.

[1] E. Rosen and Y. Rekhter, "BGP/MPLS IP VPNs." Internet-Draft.
[2] Cisco's 7600 router and Catalyst 6500 switch are functionally identical. Both run the same IOS routing code, and both accept the same management modules and line cards.

We used the SmartBits traffic generator/analyzer system from Spirent Communications as the test instrument for this project. To generate routes and test traffic, we used version 3.1 of Spirent's TeraRouting Tester (TRT) test application.

IPv4 BGP4 RIB/FIB capacity

Routing scalability is a bedrock requirement for all devices in a service provider's network. This is especially true with RFC 2547*bis* VPNs, where a single device at the edge of the service provider network may hold routing tables for numerous customers.

A primary goal of Opus One's routing scalability tests was to validate Cisco's claim that the PFC3-BXL daughtercard supports up to 1 million IPv4 routes. For these tests, we used BGP (Border Gateway Protocol) since it's the most commonly used method to exchange dynamic routing information in service provider networks.

A routing table with 1 million entries is approximately eight times larger than a full BGP table today. In other words, 1 million represents eight times more networks than today's entire public Internet.

It's important to note that routing table size alone is an overly simple measure of routing scalability. In this era of inexpensive hard disks, it's relatively easy to build a system that accepts an arbitrarily large number of routes – but actually forwards traffic to only a fraction of those routes. In practice, routing *and* forwarding to a large number of networks are both critical measures of a router's scalability.

Tests of routing table size alone measure a device IPv6's RIB (routing information base), while tests of route forwarding capacity measure the device's FIB (forwarding information base). The tests in this report simultaneously measured RIB and FIB capacity.

For this and all other tests, our test bed comprised of a single Cisco 7600/Catalyst 6500 equipped with 32 10-Gigabit Ethernet interfaces. Using the Spirent SmartBits test instrument, we established E-BGP (external BGP) sessions with each of the 32 interfaces of the Cisco 7600/Catalyst 6500. Once all sessions were established, we advertised 1,000,000 routes, or 31,250 routes per interface.

To verify all routes were usable, the SmartBits then offered traffic destined to all 1 million networks. The traffic we used consisted exclusively of 64-byte Ethernet frames; this is the shortest length in Ethernet, and thus the most stressful possible setting. Further, we offered traffic in a fully meshed pattern, in which traffic offered to each interface is destined for networks on all other interfaces; again, this traffic pattern places the greatest possible stress on the Cisco 7600/Catalyst 6500, passing the majority of traffic across the backplane and switch fabric of the device.

The results validate the ability of the Cisco 7600/Catalyst 6500 to support 1 million routes using IPv4. With an aggregate offered load of approximately 313 million packets per second (equivalent to 66.5 percent of line rate with 64-byte frames), the Cisco device forwarded traffic to all 1 million routes without dropping a single packet. (Cisco claims that the aggregate forwarding rates of the distributed switching engines (DFC3BXL) remain the same regardless of packet length. This means that with larger packet sizes, forwarding rates will be closer to line rate than those achieved with 64-byte frames. Unfortunately, time constraints prevented us from verifying this claim.)

Further, average latency was just 11.2 microseconds. This is orders of magnitude below the level where delay might degrade application performance.

Table 2 below summarizes findings from the IPv4 routing scalability tests.

Table 2: IPv4 BGP RIB/FIB Capacity

32-port full mesh, 64-byte frames, 66.5 percent offered load	
Routes learned	1,000,000
Throughput (aggregate packets per second)	312,991,903
Average latency (microseconds)	11.2

IPv6 BGP4+ RIB/FIB capacity

IPv6 is already in widespread use in networks in the Asia/Pacific region, and its use in North America is likely to grow rapidly. In late 2003 the U.S. Defense Department mandated IPv6 support in equipment it buys; this order is likely to be followed by other government agencies, in turn creating demand for IPv6 support in service provider networks. Meanwhile, the IETF has defined a set of BGP extensions to support IPv6 routing colloquially known as BGP4+[3].

All this activity naturally raises questions about IPv6 scalability in Cisco routers and switches. Our goal here was to validate Cisco's claim of support for 500,000 IPv6 routes. As before, we also verified the device's ability to forward traffic to all destinations with zero loss.

The test bed was essentially identical to that used in the IPv4 routing scalability tests. The only notable exception (other than the obvious use of IPv6 routing prefixes) was our use of 58-byte IPv6 packets, which in turn resulted in 76-byte Ethernet frames. This was a requirement imposed by the TeraRouting Testing (TRT) software running on the SmartBits test instrument; 76 bytes is the shortest Ethernet frame TRT will generate when the frames carry IPv6 packets.

As in the IPv4 tests, the Cisco 7600/Catalyst 6500 forwarded traffic in a fully meshed pattern without dropping a packet. With 500,000 unique routes learned, we offered traffic at 43 percent of line rate. The Cisco 7600/Catalyst 6500 forwarded traffic at an aggregate rate of approximately 179 million frames per second with zero loss.

Table 3 below summarizes findings from the IPv4 routing scalability tests.

Table 3: IPv6 BGP4+ RIB/FIB Capacity

32-port full mesh, 76-byte frames, 43 percent offered load	
Routes learned	500,000
Throughput (aggregate packets per second)	178,740,529
Average latency (microseconds)	10.8

[3] The actual specifications are in two IETF RFCs (requests for comment):
T. Bates, Y. Rekhter, et al. "Multiprotocol Extensions for BGP-4." RFC 2858. http://www.ietf.org/rfc/rfc2858.txt
P. Marques and F. Dupont. "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing." RFC 2545. http://www.ietf.org/rfc/rfc2545.txt

RFC 2547*bis* MPLS VPN PE/CE Scalability

While highly scalable IP routing is important, it is far from the only requirement when selecting equipment to provide RFC 2547*bis* MPLS VPN services. With MPLS, large-scale IP routing exists mainly at the edge of the service provider's network[4]. MPLS is a "layer two and a half" technology that requires many other functions:

- **Maintain a large number of virtual routing and forwarding (VRF) instances.** A key attraction of MPLS VPNs is the ability to provision service to many customers from a single edge device. As the number of customers grows, so too does the number of VRF instances each edge device must support.

- **Ensure privacy of customer traffic.** Many customers use private IP addresses as defined in RFC 1918 to extend address space or simplify administration. Since service provider devices may carry multiple customers' traffic, the potential exists for address overlapping. For example, a provider edge router may receive two packets from two customers, both destined to the "same" address. The ability to keep different customers' traffic separate is a key requirement.

- **Set up MPLS VPNs using a potentially large number of label-switched paths (LSPs).** MPLS-capable devices use a signaling protocol like LDP (label distribution protocol) to set up LSPs, a connection-oriented link-layer tunnel, through the service provider's network. The size of the service provider's network, the number of customers, and the number of customer sites are all factors that can affect the number of LSPs an MPLS-capable device must support.

- **Map customer routes to LSPs and vice-versa.** While MPLS VPNs eliminate some of the complexity of IP routing from the core of the service provider's network, the requirement for path selection remains. As the number of customers and network paths grows, so too do the numbers of FECs (forwarding equivalency classes) and label bindings an MPLS-capable device must track.

Cisco asked Opus One to determine the scalability of the Cisco 7600/Catalyst 6500 when used as a PE (provider edge) router. We tested in three common configurations:

- **VPNs for small and midsized organizations:** In this scenario, we set up 1,504 VRF instances on a Cisco 7600/Catalyst 6500, each supporting 100 BGP routes.

- **VPNs for large organizations:** In this scenario, we set up 1,024 VRF instances, each of which used a routing table with 700 BGP entries.

- **VPNs for global organizations:** In this scenario, we set up 80 VRF instances, each of which used a routing table with 10,000 BGP entries.

---

[4] Many service providers use IP routing within their own networks, even when provisioning MPLS services. The size of the service providers' IP network is actually unrelated to the number of MPLS VPNs in use. Typically, however, the bulk of IP routing in an RFC 2547*bis* MPLS VPN network is done at the edge of the service provider's network on provider edge devices (PE routers).

For each scenario, we measured device throughput and latency. We also compared results across the three different test setups to determine whether the number of VRF instances or routing table size had any impact on either throughput or latency. We also used the same private addressing scheme for all customer networks.

Figure 1 below shows the logical layout of the test bed (the one with VPNs for small and midsized organizations). On the left side of the figure, the SmartBits test instrument emulates CE (customer edge) routers. In this configuration, the SmartBits emulated 94 CE routers per interface. These CE routers attached to each of 16 customer-facing 10-gigabit Ethernet interfaces of the Cisco device, for a total of 1,504 VRF instances. The CE and PE routers exchange routing updates using E-BGP (external BGP).

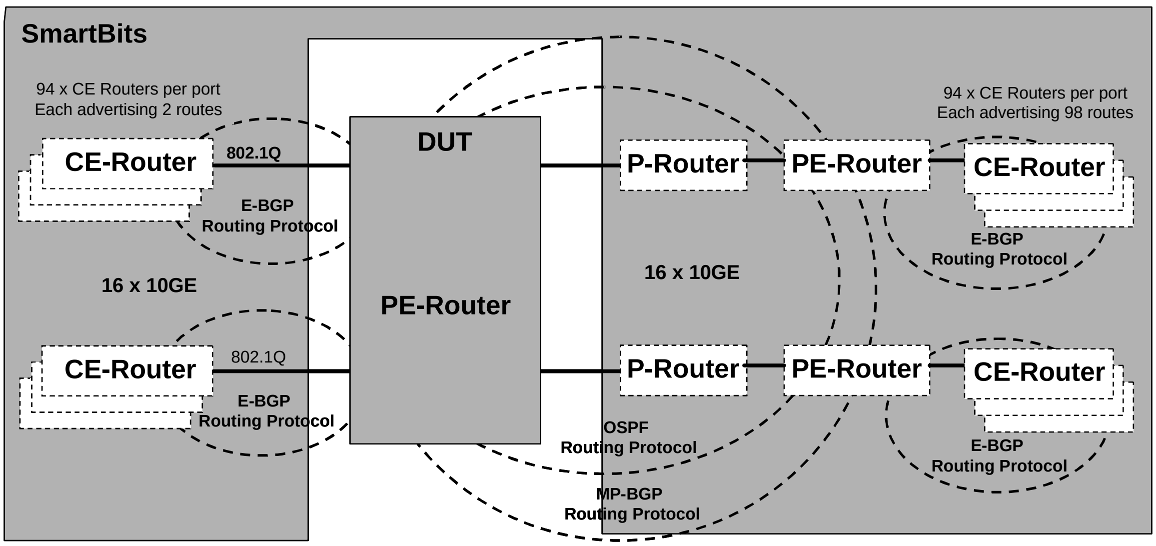

Figure 1: The CE/PE Scalability Test Bed

The service provider and remote customer sites are on the right side of the figure. Here, the Cisco device connects with one or more P (provider) routers, which in turn connect with PE routers and ultimately with CE routers at customer sites. (The SmartBits emulates the P, PE, and remote CE devices.)

In this case, we used 16 interfaces of the network-facing side of the Cisco device; however, the actual number of remote P and PE routers is unimportant. Within the service provider's network, the routers use OSPF to distribute internal topology information[5] and multiprotocol BGP (MP-BGP) to exchange customer BGP information across the service provider's MPLS "cloud."

[5] An interoperability issue with the SmartBits TRT test application prevented us from using I-BGP (internal BGP) to distribute routing data inside the service provider's network. Cisco routers and switches commonly use both OSPF and I-BGP in production networks.

For all the CE/PE tests, we offered traffic in a "port-pair" topology, meaning the first CE router on the left side of Figure 1 exchanged traffic with the first CE router on the right side of the figure. This is a configuration necessitated by a limitation in the TRT test software of the SmartBits and not a limitation of the Cisco device. A port pair topology is inherently less stressful than a partial mesh topology, in which all CE routers on the left side exchange traffic with all CE routers on the right side. Previous tests of the Cisco 7600/Catalyst 6500 involving partial and full mesh topologies suggests the device can handle such loads without loss; however, we were unable to verify that for this configuration.

VPNs for Small and Mid-sized Organizations

In the first CE/PE scalability test, we loaded the Cisco 7600/Catalyst 6500 with more than 1,500 VRF instances, with each handling 100 routes. A key objective of this scenario was to validate Cisco's assertion that service providers can provision MPLS VPNs to more than 1,500 customers using a single Cisco 7600/Catalyst 6500 as a PE router.

Once all the BGP and MPLS routing information had been propagated, we configured the SmartBits to offer data between all networks within each customer's VPN. Although we used a port-pair topology rather than a mesh, our test was actually far more stressful than any production network would experience, on two counts.

First, we used minimum-length packets[6], since these place the greatest processing burden on the device under test.

Second, we used a test duration of 300 seconds. This is five times greater than the 60-second duration recommended in RFC 2544, the IETF's router testing methodology[7], and thus far more likely to increase latency and even induce congestion on some devices. In practice, no network handles bursts of all minimum-length packets lasting for 300 seconds. If a device were prone to packet loss or elevated latency, our test would discover it.

Our results verified Cisco's claim to support more than 1,500 VRF instances on one device. With 1,504 total VRF instances and 100 routes on each, the Cisco 7600/Catalyst 6500 forwarded traffic at an aggregate rate of more than 200 million packets per second with zero loss. This load equates to an offered load of 45 percent of line rate.

Average latency was 10.9 microseconds, a level comparable to that in the BGP-only tests. This suggests there is no delay penalty in moving from IP routed networks to MPLS VPNs. Here again, delay is far below the point where application performance might suffer.

---

[6] Since MPLS inserts and removes labels, the actual frame size varied depending on traffic direction. In the CE->PE direction, we offered 68-byte frames. In the PE->P direction, we offered 64-byte frames.

[7] S. Bradner and J. McQuaid. "Benchmarking Methodology for Network Interconnect Devices," RFC 2544. http://www.ietf.org/rfc/rfc2544.txt

Table 4 below summarizes results from the CE/PE scalability tests for small to midsized organizations.

Table 4: VPNs for Small and Midsized Organizations

Parameter	Result
Number of VRF instances	1,504
Routes per instance	100
Total routes	150,400
Throughput (aggregate packets per second)	204,324,618
Average latency (microseconds)	10.9

VPNs for Large Organizations

Our large MPLS VPN test bed involved 1,024 VRF instances with 700 routes in each. An organization with 700 IP subnets might more accurately be called “very large”; the number of routing entries in many large enterprises is in the low to mid hundreds of networks.

The test bed setup was identical in most respects to that of the previous test. The major changes here were to scale up the number of VRF instances and routes per instance. Since both numbers increased, the total number of routes learned by the Cisco device rose as well, to more than 700,000 entries.

Here again, the Cisco 7600/Catalyst 6500 forwarded more than 200 million packets per second with zero loss. As in the previous test, the offered load equates to 45 percent of line rate[8].

Average latency was 10.9 microseconds, identical to the previous test. Again, this amount of delay is not enough to impact application performance.

Table 5 below summarizes results from the CE/PE scalability tests for large organizations.

---

[8] Sharp-eyed readers may notice a slight difference between throughput in this test and the previous one. Most of the difference is attributable to an algorithm the TRT test application uses to calculate traffic loads. With different numbers of routes in the two tests, TRT offered slightly different numbers of packets to the Cisco device. In addition, the 10-gigabit Ethernet standard, IEEE 802.3ae, allows traffic rates to vary by plus or minus 1,500 packets per second because of clocking differences between interfaces.

Table 5: VPNs for Large Organizations

Parameter	Result
Number of VRF instances	1,024
Routes per instance	700
Total routes	716,800
Throughput (aggregate packets per second)	204,307,401
Average latency (microseconds)	10.9

VPNs for Global Organizations

The key question of the final CE/PE scalability test was to determine whether the Cisco 7600/Catalyst 6500 would handle a relatively small number of VRF instances, each handling a very large number of routes.

In this case we used 80 VRF instances, with each instance building a routing table with 10,000 entries. A few organizations actually do maintain routing tables this large, but they are very few in number: Just 13 such organizations collectively would form a network as large as the entire public Internet.

The traffic patterns we used were identical to those in the previous two tests: CE routers on either side of the test bed exchanged traffic across the Cisco 7600/Catalyst 6500 device. As before, our test traffic consisted entirely of minimum-size packets.

Even with this maximum-size configuration, the Cisco 7600/Catalyst 6500 turned in results much like those in the smaller tests. Aggregate throughput again exceeded 200 million packets per second, and average latency was 10.9 microseconds, virtually unchanged from the previous tests.

Table 6 below summarizes results of tests from the CE/PE scalability tests for global organizations.

Table 6: VPNs for Global Organizations

Parameter	Result
Number of VRF instances	80
Routes per instance	10,000
Total routes	800,000
Throughput (aggregate packets per second)	204,313,042
Average latency (microseconds)	10.9

Comparing MPLS VPN Scenarios

Throughput and latency were remarkably consistent across all three CE/PE scalability scenarios. This suggests there is no performance penalty in scaling up either the number of VRF instances or the number of routing table entries in each instance.

Figure 2 below summarizes throughput results across all three CE/PE test scenarios. Frame departure rates from the Cisco 7600/Catalyst 6500 varied by just 0.004 percent across the three scenarios.

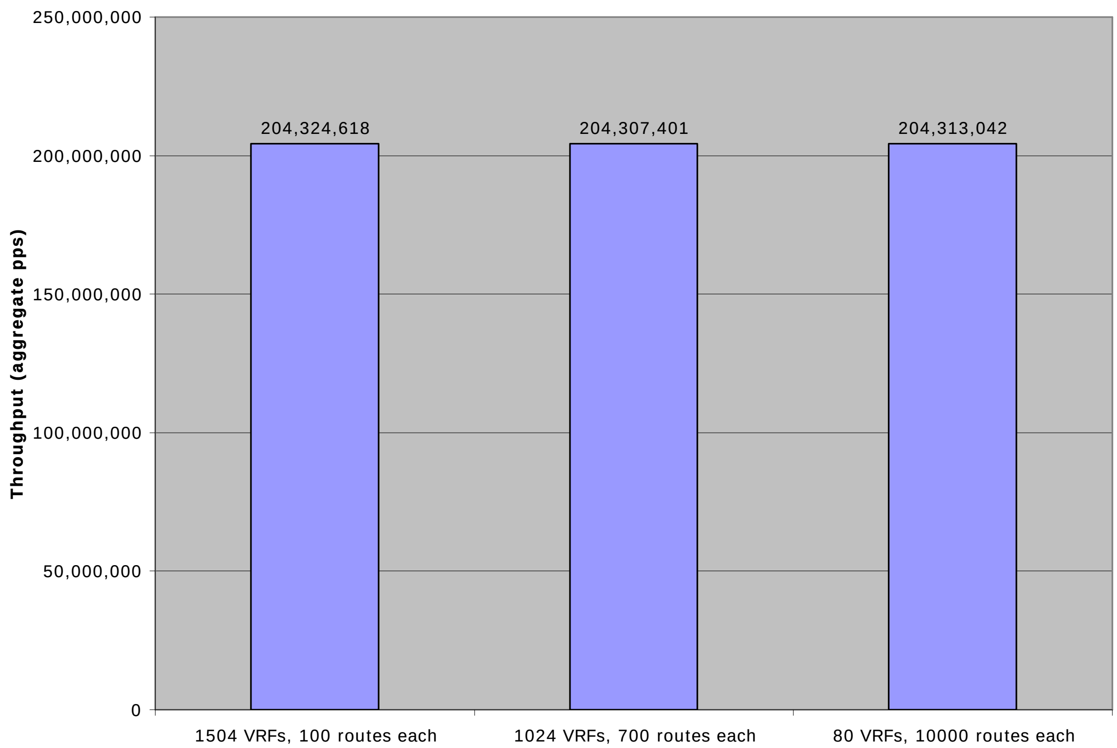

Figure 2: Throughput and CE/PE Scalability

Similarly, average latency was unchanged across all three scenarios. In all, nearly 200 *billion* packets passed through the Cisco 7600/Catalyst 6500 in these three tests, and yet delay for each packet always held to the same average – 10.9 microseconds.

Figure 3 below summarizes latency results across all three CE/PE test scenarios. We present results here in 100-nanosecond increments, which is the timestamp accuracy of the SmartBits test instrument.

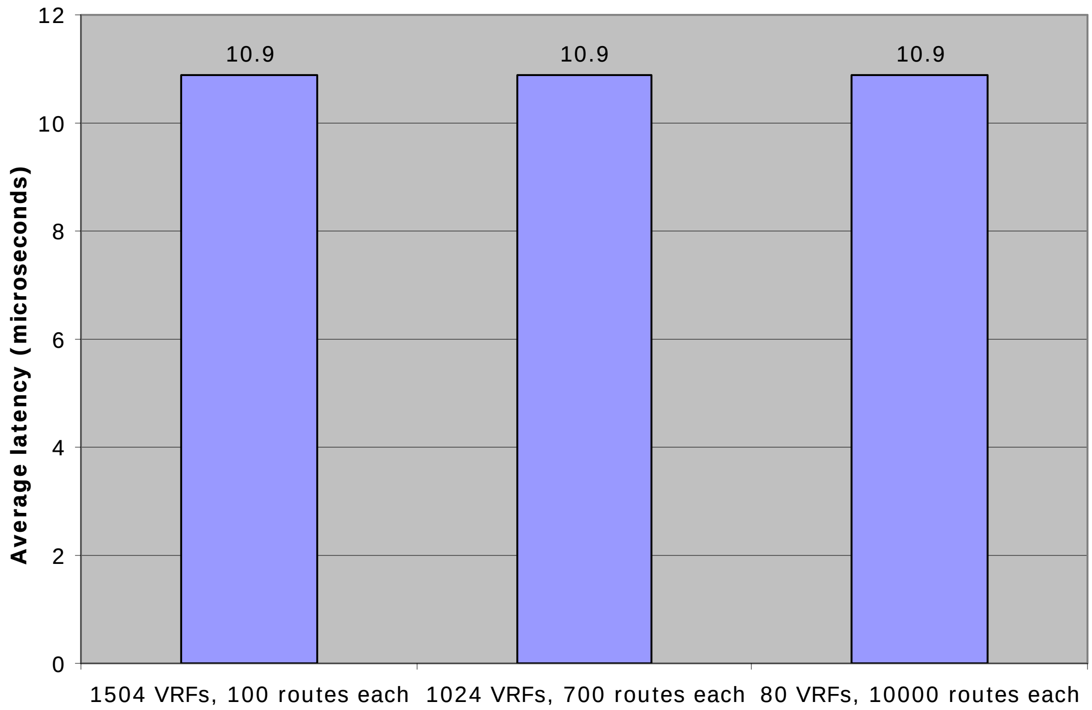

Figure 3: Latency and CE/PE Scalability

RFC 2547*bis* MPLS VPN Route Flap Handling

All tests thus far have used a static routing environment, while in production networks routing information changes second by second. For reasons beyond the administrative or technical control of a given service provider, large numbers of routes may be flapped, or withdrawn and readvertised within a short period.

Flapping places significant strain on IP routers. First among several requirements is protecting traffic on stable paths – packets destined to those routes that have not been flapped. Neither packet loss nor latency should increase on stable paths as a result of flapping. At the same time, routers must quickly process and propagate withdrawal and readvertisement messages for flapped paths.

These challenges are compounded for equipment used to provision RFC 2547*bis* MPLS VPN services. PE routers not only must contend with flapping in the routed IP network but also synchronize any changes with the MPLS network.

To assess the ability of the Cisco 7600/Catalyst 6500 to deal with route flapping, Opus One reran the “global” scenario from the CE/PE tests, the one involving a total of 800,000 routes (80 VRF instances and 10,000 routes per instance). As in the previous test, we offered minimum-length packets at an aggregate rate of about 204 million packets per second.

We configured the SmartBits’ TRT application to withdraw 50,000 E-BGP routes some 60 seconds into the test duration, and readvertise the routes about 100 seconds later. All the flapped routes existed in the CE router-to-PE router (Cisco device) direction, allowing us to compare performance of flapped vs. stable routes.

The 50,000 flapped routes may represent only 6.25 percent of the total 800,000 routes on the entire test bed, but it’s important to put this number in perspective. In interviews with Opus One researchers, network designers at three Tier-1 Internet service providers have estimated the worst-case flapping scenario they’ve handled in production involved only 25,000 routes or fewer. Thus, the severity of the flap event in our tests represents at least twice the worst-case scenario, even for backbone Internet circuits.

Figure 4 below presents a summary of the flapping tests. The most noticeable aspect of this test is that there was no change in forwarding rate for traffic on stable paths. Despite an expected loss of packets on the flapped routes, the Cisco 7600/Catalyst 6500 forwarded traffic on stable paths at a constant aggregate rate of about 102 million packets per second. (Since all flapped routes were in one direction only, this number represents half the total test bed forwarding capacity of 204 million packets per second.)

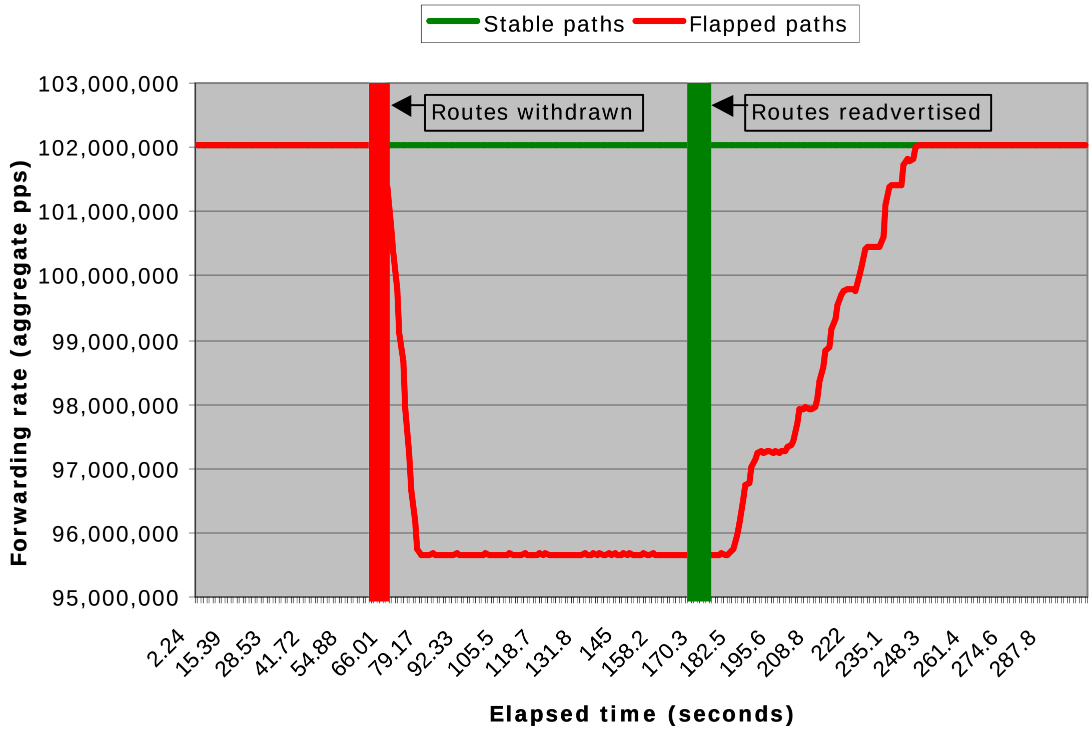

Figure 4: Route Flapping with RFC 2547*bis* VPNs

For traffic on the flapped paths, the aggregate forwarding rate declines in proportion to the number of routes withdrawn – again, about 6.25 percent of the total.[9]

The results also illustrate the time needed for the Cisco 7600/Catalyst 6500 to process routing withdrawals and readvertisements. The shaded areas of the figure represent the time needed for the SmartBits test instrument to send routing updates to the Cisco device – about 5 seconds for withdrawals and 6 seconds for readvertisements. The changes in rates along the flapped paths represent the learning times of the Cisco 7600/Catalyst 6500.

Note that the Cisco device processes withdrawals much more quickly than readvertisements, an unsurprising result given the additional work required to handle a routing advertisement. It took the Cisco 7600/Catalyst 6500 about 16 seconds to scrub all withdrawn routes, compared with a total delay of about 80 seconds to process all readvertised routes and forward traffic at the nominal maximum rate.

[9] Our chart shows forwarding rates on flapped paths dropping by 6.25 percent rather than 100 percent during the flap events. This is because the SmartBits TRT application reports summary traffic statistics per interface, not per route. Since we flapped routes on 6.25 percent of routes on *each* interface, most – but not all – of the traffic continued to be forwarded without loss. Spirent says a forthcoming version of TRT will track groups of routes within each interface.

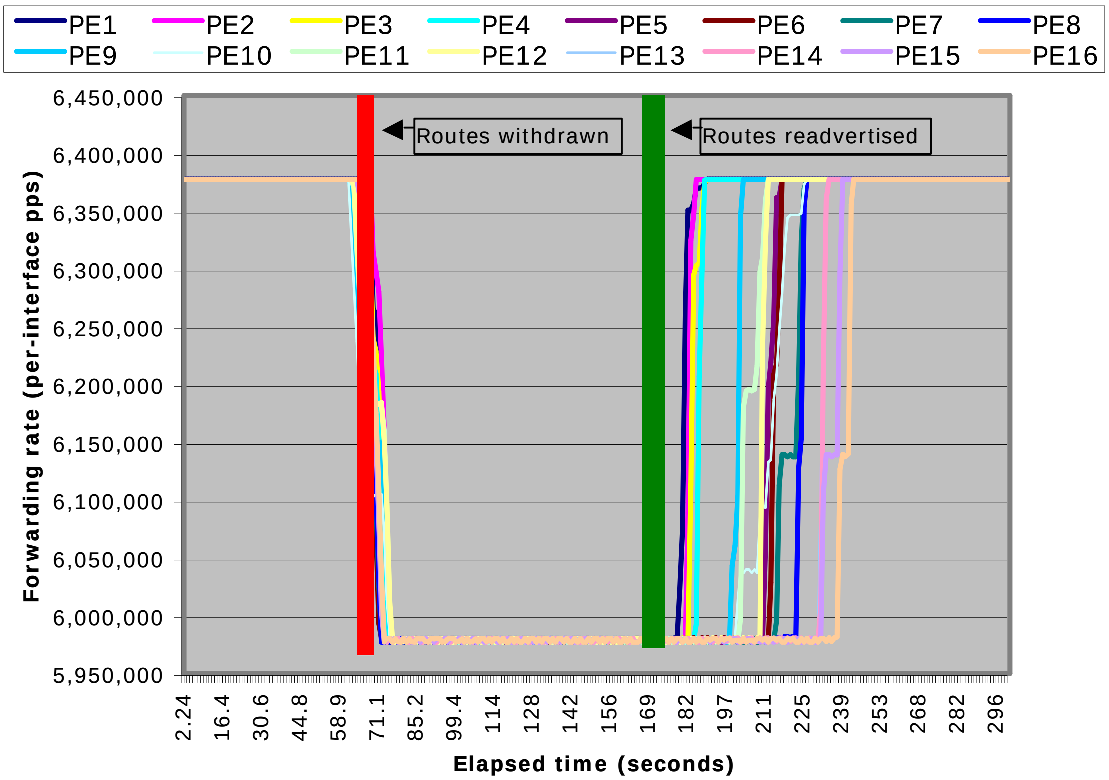

Figure 5: Per-Interface Route Flapping with RFC 2547*bis* VPNs

The key phrase here is "total delay," since the Cisco device relearned some routes faster than others. Figure 5 below presents only the flapped-route results from the same test, this time broken out on a per-interface basis. On the stable paths, forwarding rates were constant throughout the test duration.

When presented this way, we can see that routing reconvergence begins much earlier – in fact, within 20 seconds of receipt of the first readvertisement. It's also notable that the Cisco 7600/Catalyst 6500 processes advertisements in the order received. Reconvergence begins first on the interface belonging to PE 1, followed by PE2, and so on.

P router scalability

Mention the term "P router" to the network architects of many service providers, and the image they'll likely conjure involves an enormous core router with a seven-figure price tag. Mention "Cisco 7600" or "Catalyst 6500," and this time, the likely response will involve terms such as "edge" or "aggregation" or "PE router."

In fact, test results suggest that a Cisco 7600/Catalyst 6500 acting as a P router can provide a cost-effective alternative to larger systems in the core of service provider networks.

To assess the Cisco 7600/Catalyst 6500's suitability to task as a P router, we constructed a test bed in which the Cisco device fielded connections from PE routers on all 32 10-gigabit Ethernet interfaces. Since our test bed involved a single P router connecting multiple PE routers, the Cisco device performed a single label swap on each packet received.

Figure 6 below illustrates the logical P router test bed.

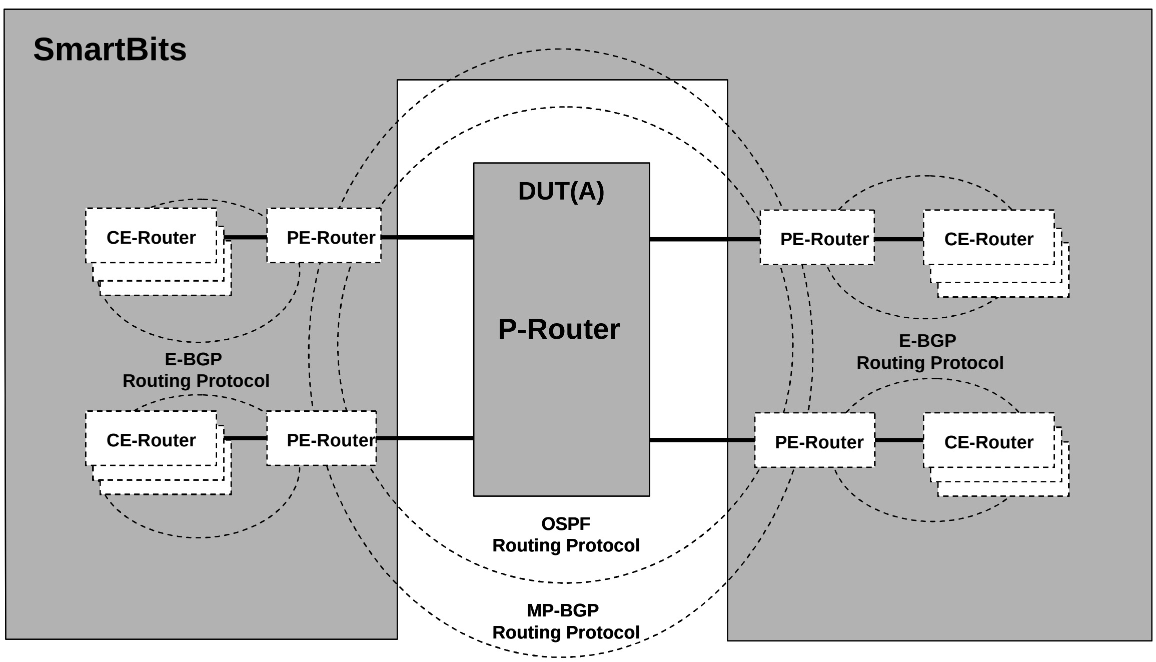

Figure 6: P Router Test Bed

In this test bed, PE routers with attached CE routers (all emulated by the SmartBits) exchanged traffic in a fully meshed pattern across the Cisco 7600/Catalyst 6500 – meaning every interface exchanged traffic with every other interface. Since the test bed involved 32 interfaces[10], this meant the Cisco device handled a fully meshed topology involving 992 networks (one network per CE router, so 32 sources times 31 destinations) and 32 FECs (forwarding equivalency classes). While 992 is a relatively small number compared with earlier tests, the actual number of CEs and routes is unimportant and has no impact on P router performance.

As in other tests, we configured the SmartBits to offer minimum-length[11] packets to the Cisco device. Also as before, the aggregate offered load was about 204 million packets per second (equivalent to around 45 percent of line rate). After performing a single label swap, the Cisco device forwarded all packets to PE and ultimately CE routers.

As in all other tests, the Cisco 7600/Catalyst 6500, when acting as a P router, forwarded all traffic without dropping a single packet. Throughput was an aggregate of 204 million packets per second, identical to the offered load.

Latency averaged 10.9 microseconds, the same measurement as in other tests. This measurement suggests the additional work of a label swap imposes no latency penalty. Further, delay added by the Cisco device is well below the point where of affecting application performance.

Table 7 below summarizes findings from the P router scalability tests.

Table 7: P Router Performance

32-port full mesh, single label swap, 68-byte frames, 45 percent offered load	
Total routes	992
Throughput (aggregate packets per second)	204,322,925
Average latency (microseconds)	10.9

Conclusion

Opus One successfully validated Cisco's claims for the Cisco 7600/Catalyst 6500 for use in MPLS VPN service. When equipped with new PFC3-BXL and DFC-BXL modules, this system can support thousands of MPLS VPN customers, each with substantial routing tables. Further, the Cisco device also can function as a P router. All these functions are available in other equipment, but typically with a much higher price tag. The system we tested carries a US list price of less than $500,000 offering great value for money.

---

[10] We've reduced to PE interface count to 4 in the figure for clarity. The actual number on the test bed was 32 10-gigabit Ethernet interfaces.
[11] Nominally, each PE router would receive a 64-byte Ethernet frame from each CE router, insert a 4-byte label, and forward a 68-byte frame to the P router. Since the SmartBits used virtual CE and PE routers, all traffic appeared on the wire as 68-byte frames.

Acknowledgements

Opus One gratefully acknowledges the support of Spirent Communications, which supplied engineering expertise as well as equipment for this project. Spirent test engineer Mark Hall configured the TeraRouting Tester (TRT) application for many of the tests. In addition, Spirent supplied its SmartBits performance analysis system. We used the SmartBits 6000 chassis and XLW-3721 10-gigabit Ethernet cards for this project. Thanks also to independent consultant Jerry Perser, whose knowledge of 220-volt electrical systems made this test possible.

About the Author

David Newman, an associate researcher at Opus One, has been benchmarking high-end network devices for more than a decade. Newman's comparative tests of routers, switches, firewalls, and IP services appear in trade publications such as *Network World* and *Light Reading*. Newman is the author of RFC 2647 and coauthor of RFC 3511, both on firewall performance measurement. He can be reached at dnewman@opus1.com.