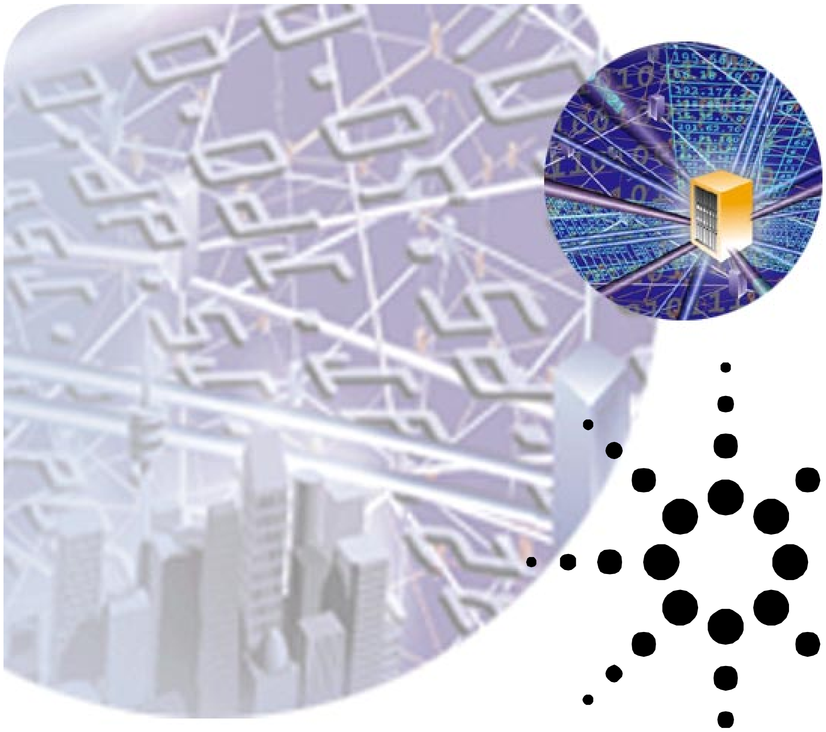

Verifying & Evaluating Gigabit Switch/Routers Performance

with Internet-Scale Simulation

Agilent Technologies RouterTester
Application Note

Introduction

This paper discusses the architectural test challenges within the modern gigabit/terabit router and discusses how *Internet-scale simulation* can be used to effectively measure the performance of these routers.

Agilent Technologies
Innovating the HP Way

Internet Testing of Gigabit & Terabit Routers

Routers with gigabits of switching capacity, such as the Cisco 12000 GSR and the Juniper M40, are now available. Routers with *terabits* of capacity, such as the routers under development at Avici, Nexabit (now part of Lucent Technologies) and Pluris, are under development or in the initial stages of product release.

Measuring the performance of these routers, from their ability to switch millions of packets per second, to their ability to provide multiple levels of Quality of Service (QoS), is necessary in order to verify the reliability and capacity of these devices. A router's true performance must be determined before it is deployed in the network, and verified by service providers so that accurate router requirements, based upon the router's performance can be quantified.

To truly measure the performance of these new routers, a new methodology for testing is required. Complex distributions of source and destination addresses coupled with packet lengths and profiles typical of traffic found within the Internet must be generated into the router under test, and the resultant behavior measured precisely. This new methodology is called *Internet-scale simulation*. Provided by the Agilent Technologies RouterTester, *Internet-scale simulation* subjects routers to traffic similar to that found within the Internet.

Router Architectures and Performance Issues

In this section, we'll briefly examine the architecture of the modern router, while highlighting the key test metrics for gigabit and terabit routers.

Basic Architecture

The modern router consists of a number of improvements to enable it to forward millions of packets per second, from many interfaces, while maintaining Quality of Service guarantees for different types of traffic.

The architecture of a modern router consists of:

- A number of interface cards containing a local CPU, cache of frequently used route table entries, and a number of output queues
- A switching fabric which forwards packets from input interfaces to output interfaces
- A central processor, responsible for management of the router; and
- A routing table (located in the router's memory)

Physical Interfaces

Driven by the demand for network bandwidth, high-speed interfaces are supported by core routers. Typical interfaces and speeds are:

- 622 Mb/s OC-12c/STM-4c and 2.4 Gb/s OC-48c/STM-16c Packet over SONET/SDH
- 155 Mb/s OC-3c/STM-1c and 622 Mb/s OC-12c/STM-4c ATM

Gigabit Ethernet interfaces may also be provided for connectivity to local server clusters or other routers located within the same Point of

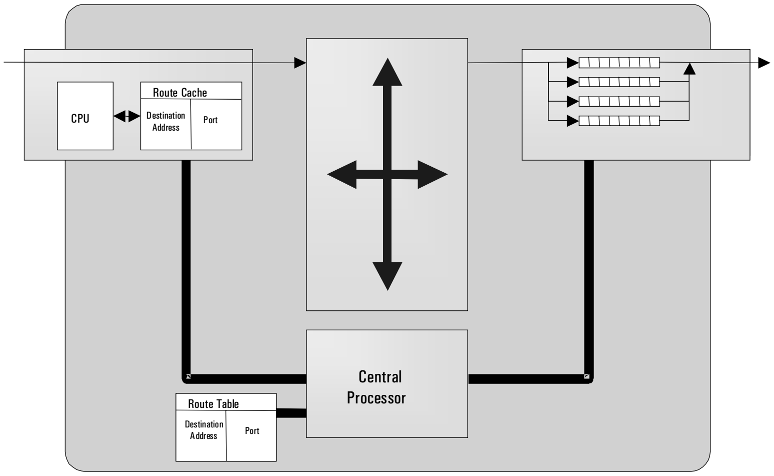

Figure 1: Modern router architecture.

Well-Known Unicast Routing Protocols

Exterior Gateway Protocols

- BGP-4
- EGP (obsolete)

Interior Gateway Protocols

- OSPF
- IS-IS
- RIP
- IGRP (Cisco proprietary)
- EIGRP (Cisco proprietary)

Presence (POP), or perhaps for short reach links, such as those in Metropolitan Area Networks.

Each physical interface typically contains a CPU and memory containing a route cache of frequently used routes. When a packet is received on the interface, the header checksum is verified and the time to live (TTL) field within the packet is decremented. The destination address within the IP packet is then compared to the list of network prefixes contained within the route cache. If a match is found, then the packet is passed into the switching fabric to be forwarded to the appropriate output. If a match is not found, then a request is sent to the central processor to look up the destination address in the complete routing table. The central processor sends back the corresponding output interface, and the packet is passed into the switching fabric.

Performance Metrics

Some of the key metrics that need to be measured are:

- The speed of destination address lookups.
  Clearly, if the destination address within an IP packet is not contained inside the route cache, it will take longer to lookup the destination address and then forward the packet, delaying the packet.
  Routing tables for the Internet backbone currently contain about 65,000 entries - and are growing by about 10,000 entries a year.
- The packet forwarding speed through the physical interface.
  If a physical interface receives many small packets, it may not be able to forward these packets at line rate, since it must process many more packets per second than if it were forwarding only very large packets.

Switching Fabrics

The key to a router's performance is the switching fabric - the component that can take millions of packets per second from many interfaces, and switch these packets to the appropriate output interface.

Switching fabrics have evolved from shared bus, shared memory and crossbar mechanisms into proprietary pipelined and distributed architectures capable of switching packets received at wire speed from many interfaces.

Performance Metrics

The key metric to measure the performance of a switching fabric is to determine the maximum aggregate input rate across as many interfaces as the router can handle. This requires generating traffic at up to line rate into every interface on the router, then measuring the throughput, loss and latency of packets as they traverse the router.

Output Queues

To provide different Quality of Service levels to different types of traffic, traffic with similar qualities is placed into separate queues. A number of different mechanisms can then be used to service these queues. Weighted round robin and weighted fair queuing methods provide a means for providing frequent service to the high priority queues. These mechanisms ensure that high priority, delay and loss sensitive traffic, such as voice, is forwarded ahead of lower priority traffic.

Performance Metrics

Measuring the operation and performance of output queues is extremely subjective. It depends both upon the queue management mechanisms utilized in the router as well as the distribution of traffic that the router is expected to encounter.

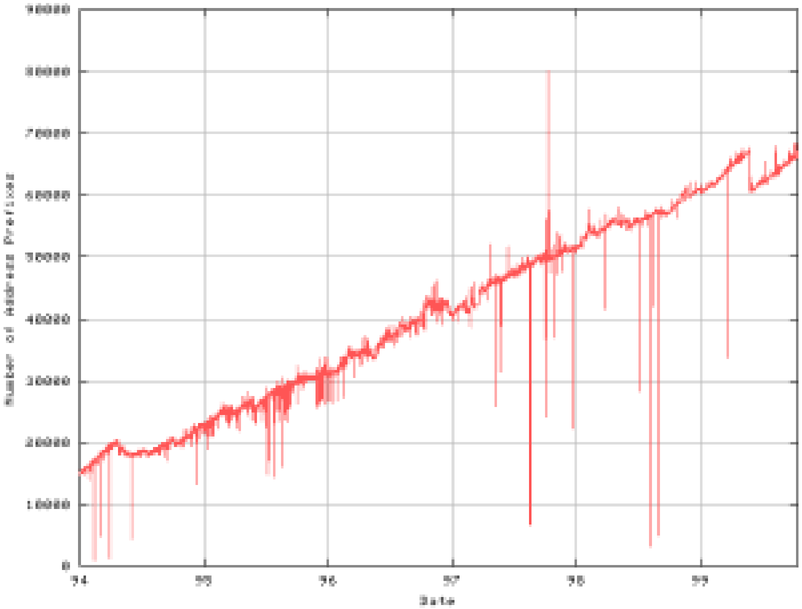

Figure 2: The growth of address prefixes (routes) within the Internet.

The behaviors that must be measured and quantified are:

- Does a router apply the lowest latency and packet loss to the highest priority traffic, especially when there is congestion within the router?
- Does a router maintain an acceptable QoS level for "better than best-effort" traffic?
- Are some traffic types (e.g. voice) adversely impacted when other traffic types (e.g. bursty web data) show increases in load?

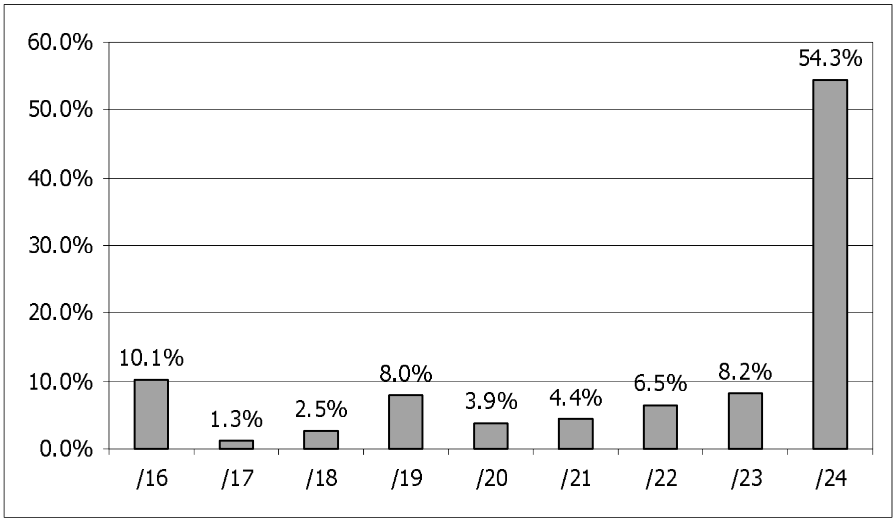

Figure 3: Percentage distribution of network prefix lengths in routing tables within core routers of the Internet.

Realistic Internet-Scale Traffic Performance

So far, we have considered the performance of the individual architectural components. The second aspect to examine when assessing the performance of a gigabit or terabit router is to examine its performance when subjected to realistic traffic.

A router may perform well when it is delivered packets of a consistent size, delivered to a relatively small number of destinations. However, the same router may perform poorly if it is required to forward varying sized frames to a large number of destinations. A broad distribution of destination addresses, packet types and packet lengths characterize realistic traffic.

To fully measure the performance of a router, it must be subjected to traffic that matches realistic Internet traffic as closely as possible.

Routes

As mentioned previously, the routing tables installed on routers within the core of the Internet contain about 65,000 routes. A router must be capable of receiving an IP packet, comparing the destination address to a table containing 65,000 (or more) routes, and then switching the packet to the appropriate destination.

Figure 2 shows the growth of address prefixes (possible routes) visible from one autonomous system (Telstra: AS number 1221). In particular, note that the number of routes increases by approximately 10,000 every year.

Networks connected to the Internet are varied in size. This is reflected in the variance of network prefix lengths within the routing tables of the router. The more different-length prefixes found within the router, the more time it takes to perform a longest length prefix match to find the destination interface corresponding to an IP packet destined for a particular network.

Most network prefix lengths found within the routing tables within the Internet core are between 16 and 24 bits in length (see Figure 3).

Packet Types

The majority of the traffic carried within the Internet is TCP (90% of packets), and the predominant traffic type carried within TCP packets is, of course, web traffic (70% of TCP packets contain HTTP packets). FTP and Telnet account for 5% and 1% of the traffic, respectively. This distribution is shown graphically in Figure 4. Note HTTP, FTP and Telnet protocols are carried within TCP packets.

The type of traffic carried is changing - as the Internet carries more voice traffic, we can expect to see a greater number of short UDP packets containing voice segments.

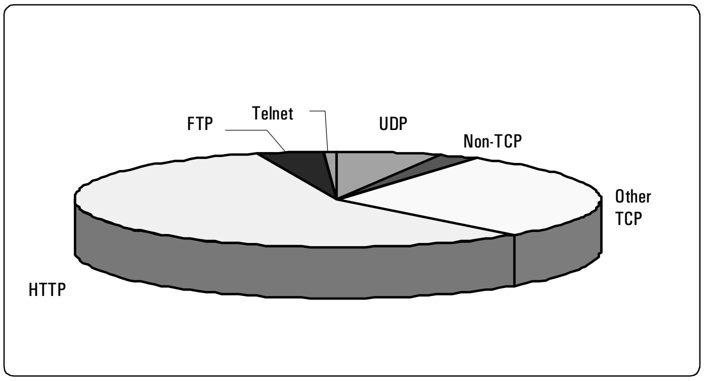

Figure 4: Distribution of packet types within the Internet.

Packet Lengths

As mentioned previously, the lengths of IP packets received by a router can have a significant effect on its performance. IP packet lengths within the Internet are generally modal. 40% of all packets are 40 bytes long (TCP acknowledgments). Other modes occur at 552 and 576 bytes (5% and 6%) representing TCP applications that do not perform maximum transmit unit (MTU) discovery, and 1500 bytes (10%),

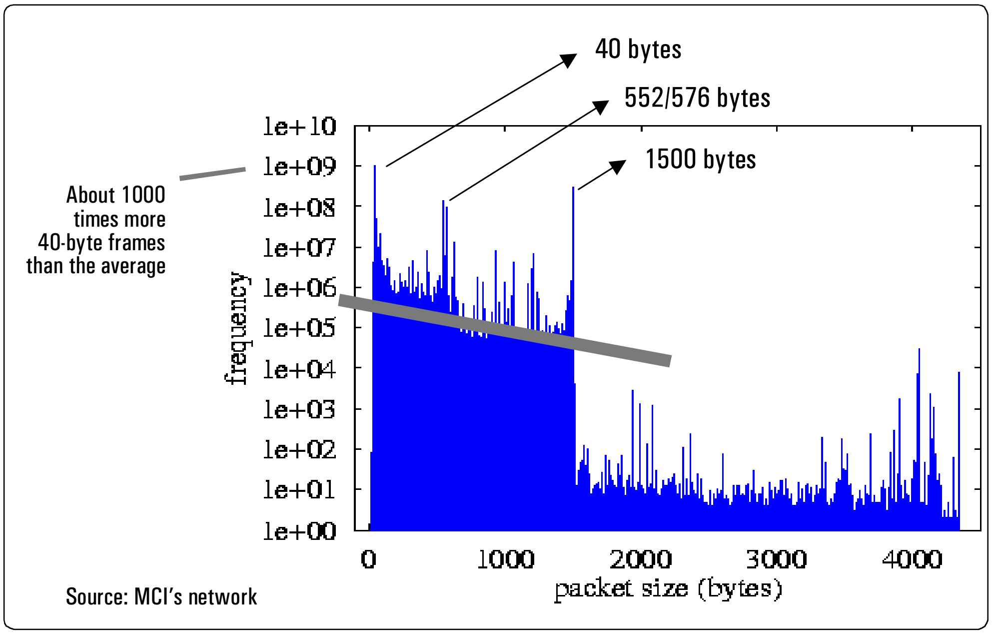

Source: MCI's network

Figure 5: Distribution of packet lengths in a typical network.

corresponding to the maximum segment size for Ethernet.

Burstiness

A typical traffic source within the Internet generally does not transmit data at a constant bandwidth. HTTP and voice traffic, in particular, are characteristically bursty, whereas FTP traffic (which constitutes only 5% of traffic within the Internet) is typically less bursty and more constant.

Realistic traffic thus comprises different packet lengths, different packet types, and different application types transmitted with varying bandwidths over a large number of routes.

Test Challenges for Gigabit and Terabit Routers

Based upon the description of the modern router in the previous section, a number of test challenges come to light:

Challenge #1: Test Router Performance

To verify reliable IP packet delivery, the router's ability to handle extreme traffic loads must be measured. Throughput, latency and loss metrics provide valuable insight into the router's performance capabilities and limitations.

Challenge #2: Test at Internet-Scale

Realistic Internet simulation is essential when testing a gigabit or terabit router. This means creating large aggregates of IP flows over thousands of routes between thousands of networks.

Challenge #3: Verify Quality of Service Guarantees

Queue management techniques are inherently difficult to test. The generation and analysis of multiple IP packet streams with complex traffic parameters is necessary to ascertain the router's ability to effectively manage many different traffic classes.

Understanding the Challenges

Understanding Challenge #1: Test Router Performance

To effectively test the performance of a router, it is necessary to test its performance under stressful, yet realistic conditions. This requires the ability to generate traffic as close to real traffic as possible. Realistic traffic requires the ability to:

- Generate many IP packets as if they came from many networks
- Generate a mix of packet sizes, representing the distribution of packet lengths within the Internet
- Generate traffic with the bandwidth characteristics of realistic traffic

The following sections describe these areas in more detail.

Generate many IP packets as if they came from and are destined to many networks

To generate many IP packets from many networks requires a router to know of the existence of these networks. A routing protocol is thus needed to install a routing table within the router under test. This routing table must be the same size as routing tables typically found within the Internet (65,000 routes and growing). It must also contain a diverse set of network prefix lengths, representative of the large number of differently sized networks within the Internet. The distribution of network prefix lengths shown in Figure 3 should be used.

Next, the test system must generate packets between combinations of these networks. By doing so, traffic with a broad distribution of source and destination addresses is created, representative of the real distribution of addresses within the Internet. A broad distribution of destination addresses stresses the ability of a router to perform the longest prefix match on many different destination addresses.

Generate a mix of packet sizes and packet types, representing the distribution of packet lengths and packet types within the Internet

As mentioned previously, a broad range of packet lengths must be delivered into the router under test. In addition to testing a router with wire-speed packets whose lengths represent both the largest and smallest possible lengths, a distribution of packet lengths representative of the distribution of packet lengths within the Internet must be delivered to the router at wire-speed. Testing a router under these conditions will reveal the ability of a router to forward a broad range of packet lengths, and ensure that the router has not been optimized to forward packets of a particular packet length.

Generate traffic with the bandwidth characteristics of realistic traffic

Traffic sources within the Internet rarely transmit at a constant bandwidth - voice, data and video streams tend to be typically bursty. Consequently, simulated traffic from a number of sources needs to be generated, with characteristics representative of the traffic origins.

Understanding Challenge #2: Test at Internet-Scale

In addition to the generation of realistic traffic (as described previously), *Internet-scale simulation* requires the generation of many IP packets over many interfaces - essentially simulating the entire network to which a router is connected.

For a gigabit or terabit router, *Internet-scale simulation* requires:

- The delivery of packets at wire-speed over high speed interfaces
- The delivery of packets over as many interfaces as possible

Wire-speed traffic generation, coupled with many ports of test capability provides true *Internet-scale simulation* to the router under test.

Understanding Challenge #3: Verify Quality of Service Guarantees

As described earlier, the delivery of QoS within a router is provided through multiple output queues coupled with an algorithm that ensures that high-priority queues are serviced more frequently than low-priority queues.

To test the QoS capabilities of a router, many different streams of IP packets, representing different service types or classes must be generated into the router under test. The conditions that force a router to prioritize outgoing traffic must be created: namely the over-subscription of output ports. Over-subscription forces a router to forward the high priority traffic at the expense of the lower priority traffic.

Each stream must be analyzed in order to measure the QoS metrics: packet throughput, loss, and latency. Analyses of these metrics will reveal the QoS capabilities of the router.

Internet-Scale Simulation: The New Paradigm for Router Performance Testing

More than a marketing term, *Internet-scale simulation* describes a testing philosophy which exposes routers under test to traffic representative of the speed, volume, characteristics and distribution of networks characteristic of the Internet. It is only when routers are exposed to these conditions do they truly exhibit their underlying behavior - whether they have the speed, reliability and performance necessary to operate in Internet-scale networks.

Meeting Challenge #1: Test Router Performance

The Agilent Technologies RouterTester generates realistic traffic, with realistic address distributions, packet lengths and traffic profiles.

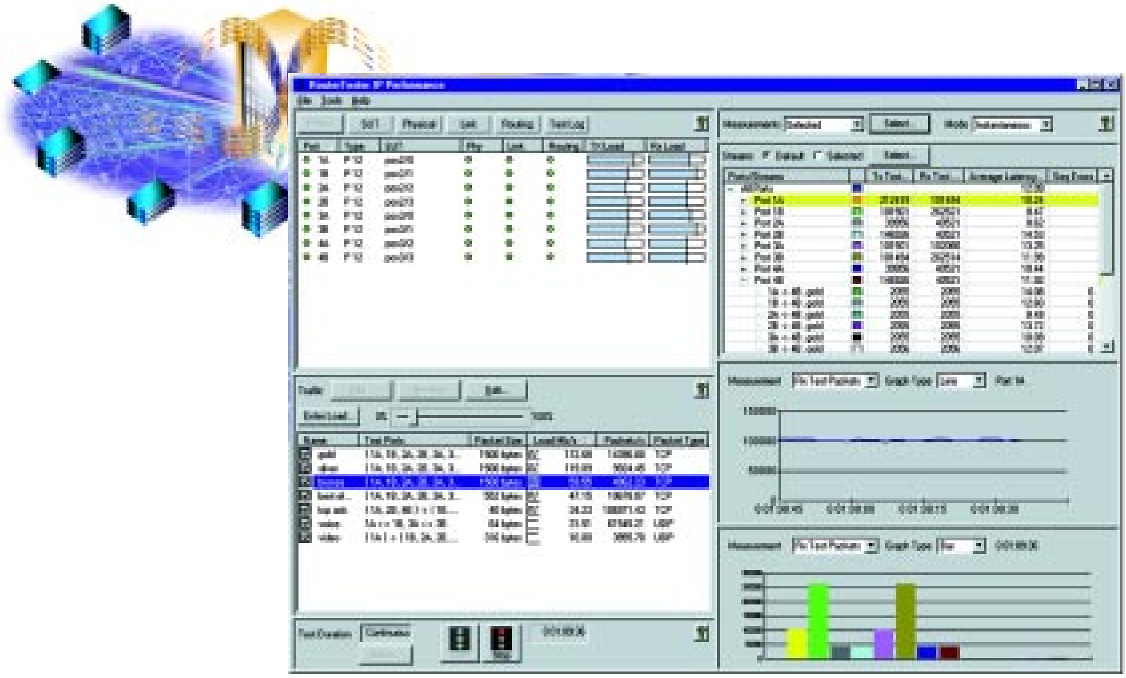

Figure 6: RouterTester generates realistic traffic.

Generate many IP packets as if they came from and are destined to many networks

Realistic traffic cannot be generated by simply incrementing a series of source and destination addresses. This is not representative of the rich distribution of addresses within the Internet. Instead, realistic traffic must be generated as if it came from many different networks of many different sizes. Generating a rich range of destination addresses will fully stress a router's "longest length prefix" match algorithm by forcing it to compare the destination addresses within received IP packets with network prefixes within Internet-scale routing tables (50,000 to 100,000 routes and up).

RouterTester utilizes BGP-4 routing protocol emulation to build up extremely large routing tables within the router under test. Each port on RouterTester can simulate up to 65,536 networks containing any number of networks per CIDR network prefix length. This forces the router under test to install an extremely large routing table. RouterTester then generates traffic from all of these networks, forcing the router to perform exhaustive longest length prefix matches for every destination address.

Figure 7 shows how the route configuration for a single RouterTester test port can be configured to generate a realistic distribution of network prefixes. In this example, there are a total of 20 ports connected to the system under test. To generate a routing table containing 65,000 ports, we generate 3250 (65,000 / 20) network prefixes per port. To create a realistic distribution, we create 10.1% network prefixes with a length of /16 (10.1% x 3250 = 331 network prefixes), 1.3% with a length of /17 (44 network prefixes), and so on, as shown in Figures 3 and 7.

Having generated a realistic routing table within the router under test, RouterTester is then able to generate realistic traffic to these distributed networks.

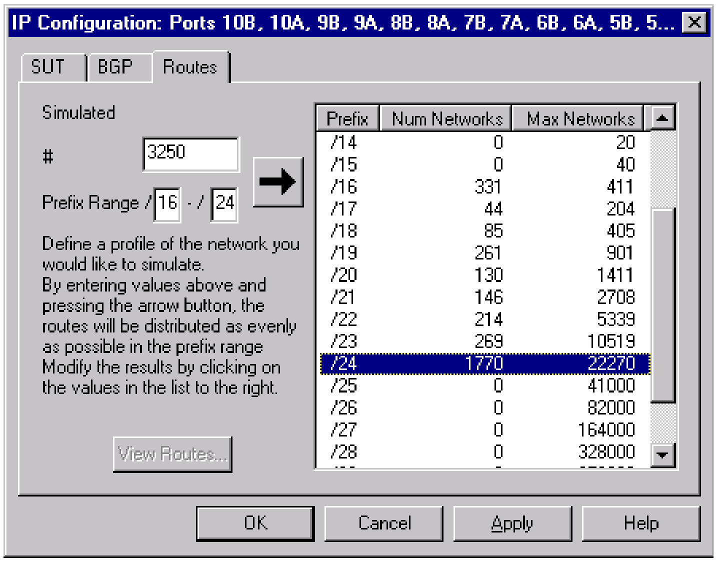

IP Configuration: Ports 10B, 10A, 9B, 9A, 8B, 8A, 7B, 7A, 6B, 6A, 5B, 5...

Simulated

\# 3250

Prefix Range /16 - /24

Define a profile of the network you would like to simulate.
By entering values above and pressing the arrow button, the routes will be distributed as evenly as possible in the prefix range
Modify the results by clicking on the values in the list to the right.

Prefix	Num Networks	Max Networks
/14	0	20
/15	0	40
/16	331	411
/17	44	204
/18	85	405
/19	261	901
/20	130	1411
/21	146	2708
/22	214	5339
/23	269	10519
/24	1770	22270
/25	0	41000
/26	0	82000
/27	0	164000
/28	0	328000

Figure 7: Configuring RouterTester to generate a realistic distribution of network prefixes.

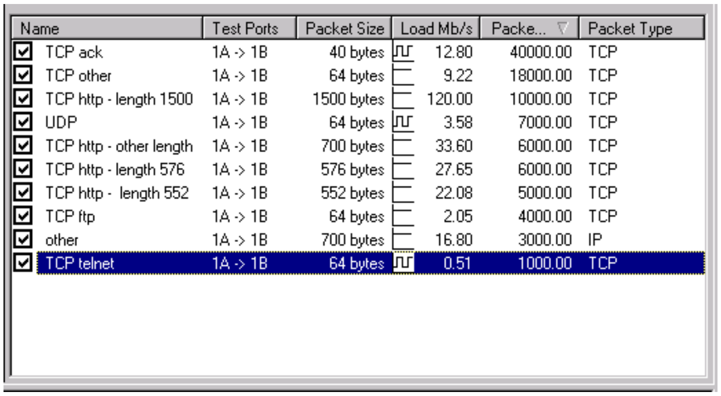

	Name	Test Ports	Packet Size	Load Mb/s	Packe...	Packet Type
☑	TCP ack	1A -> 1B	40 bytes	12.80	40000.00	TCP
☑	TCP other	1A -> 1B	64 bytes	9.22	18000.00	TCP
☑	TCP http - length 1500	1A -> 1B	1500 bytes	120.00	10000.00	TCP
☑	UDP	1A -> 1B	64 bytes	3.58	7000.00	TCP
☑	TCP http - other length	1A -> 1B	700 bytes	33.60	6000.00	TCP
☑	TCP http - length 576	1A -> 1B	576 bytes	27.65	6000.00	TCP
☑	TCP http - length 552	1A -> 1B	552 bytes	22.08	5000.00	TCP
☑	TCP ftp	1A -> 1B	64 bytes	2.05	4000.00	TCP
☑	other	1A -> 1B	700 bytes	16.80	3000.00	IP
☑	TCP telnet	1A -> 1B	64 bytes	0.51	1000.00	TCP

Figure 8: A realistic distribution of packet sizes is generated with RouterTester.

Generate a mix of packet sizes and packet types, representing the distribution of packet lengths and packet types within the Internet

Streams representative of the packet length distributions within the Internet can be quickly set up on RouterTester. A stream can be quickly established with 40 octet IP packets containing TCP acknowledgments, such that the stream carries 40% of the packets on a particular physical link. Similarly, streams can be easily and simultaneously established with 552, 576, and 1500 octet length packets.

Figure 8 shows how RouterTester can be used to generate several traffic streams, with all streams representing the correct distribution of frame types and frame lengths found with realistic Internet traffic. A total of 100,000 packets per second are being generated between two ports of a router. 40% of these packets (40,000 packets per second) are TCP acknowledgments (40 octets in length). 90% of the packets generated are TCP frames, and of the remainder, 7% are UDP packets, and 3% are other types. Of the TCP packets, 70% contain HTTP packets.

Figure 8 also shows the correct distribution of packet lengths: 40% of packets are 40 octets in length, 5% are 552 octets long, 6% are 576 octets long, and 10% are 1500 octets long.

Generate traffic with the bandwidth characteristics of realistic traffic

Different traffic sources have different traffic profiles. Voice traffic is carried within UDP packets and is inherently bursty. FTP file transfers are characterized by a fairly constant bandwidth distribution. HTTP is carried within TCP packets and is both bursty with a 7:1 ratio of small to large packets.

Each type of traffic must be allocated a varying portion of the overall bandwidth available in one particular link. RouterTester applies a different traffic profile to each traffic source and mixes the traffic to create a pipe rich with different types of traffic (see Figure 9).

Figure 8 also shows different traffic profiles applied to different traffic streams. TCP acknowledgments are set as bursty traffic, as are TCP Telnet frames. A bursty profile is also applied to the UDP packets because UDP is used to carry voice over IP.

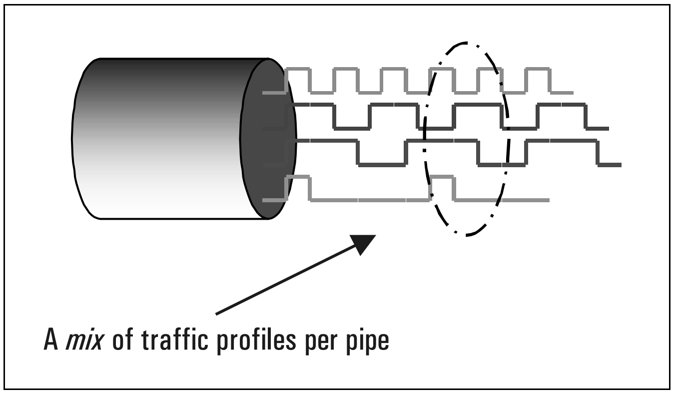

Figure 9: RouterTester mixes traffic streams with different traffic profiles.

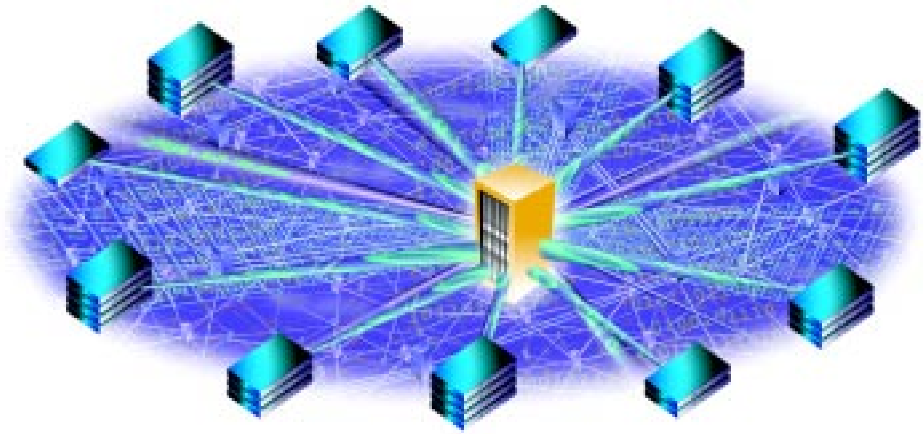

Figure 10: RouterTester simulates the scale of the Internet around the router under test.

Meeting Challenge #2: Test at Internet-Scale

RouterTester scales up to 64 ports of high-speed Packet over SONET/SDH interfaces. Each port uses the BGP-4 routing protocol to advertise up to 65,536 network prefixes (to install a routing table of up to 64 x 65,536 routes), and then generates streams of realistic traffic from these simulated networks.

Meeting Challenge #3: Verify Quality of Service Guarantees

RouterTester can generate multiple streams of IP packets, each representing a different service class, for example:

- gold, silver, bronze, best effort classes (using the "Olympic" model)
- service type (e.g. voice, video, high/low priority data)

These classes are represented by the IP Type of Service field or differentiated services (diffserv) codepoint in the IP header.

Figure 11 shows a number of traffic classes designed to test the QoS capabilities of a router. Each traffic class is distinguished by the differentiated services codepoint appropriate to the service level type. Streams are concurrently analyzed in real time, enabling the comparison of packet loss, latency, and throughput metrics between streams. RouterTester reveals the true Quality of Service capabilities and limitations of the router under test.

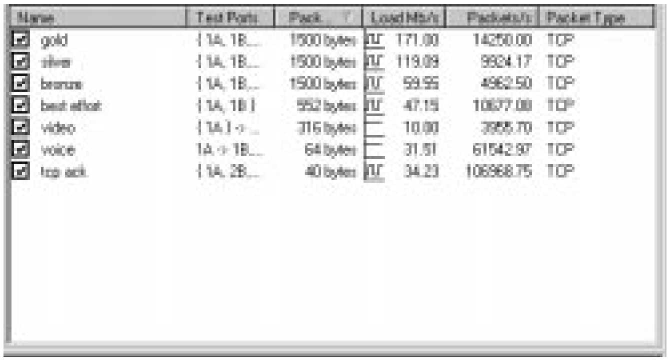

Name	Test Ports	Pack...	Load Mb/s	Packets/s	Packet Type
☑ gold	{ 1A, 1B...	1500 bytes	171.00	14250.00	TCP
☑ silver	{ 1A, 1B...	1500 bytes	119.09	9924.17	TCP
☑ bronze	{ 1A, 1B...	1500 bytes	59.55	4962.50	TCP
☑ best effort	{ 1A, 1B }	552 bytes	47.15	10677.00	TCP
☑ video	{ 1A } -> ...	316 bytes	10.00	3955.70	TCP
☑ voice	1A -> 1B...	64 bytes	31.51	61542.97	TCP
☑ tcp ack	{ 1A, 2B...	40 bytes	34.23	106968.75	TCP

Figure 11: RouterTester generates and analyzes many streams of realistic traffic.

Conclusion

Internet traffic is inherently complex. RouterTester provides the ability to generate traffic representative of that found within the Internet. It achieves this by generating packets with a wide a diverse range of destination addresses, packet sizes, packet types, and traffic profiles.

RouterTester provides realistic *Internet-scale simulation* to effectively simulate the Internet around a gigabit or terabit router. It reveals the performance of these devices within the most realistic possible test scenarios.

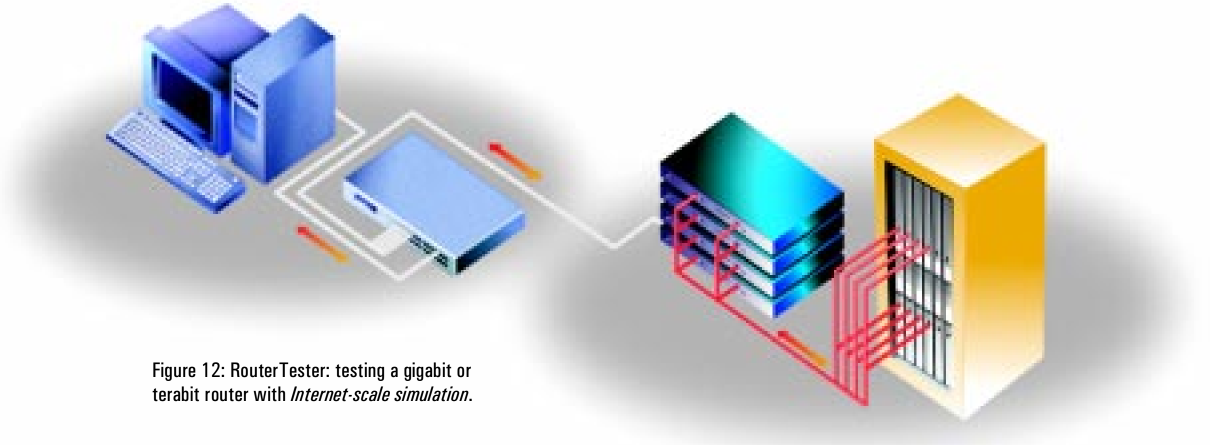

Figure 12: RouterTester: testing a gigabit or terabit router with *Internet-scale simulation*.

Acronyms

AS	Autonomous System
BGP-4	Border Gateway Protocol, Version 4
CIDR	Classless Internet Domain Routing
CPU	Central Processing Unit
EGP	Exterior Gateway Protocol
EIGRP	Enhanced Interior Gateway Routing Protocol (Cisco)
FTP	File Transfer Protocol
HTTP	HyperText Transfer Protocol
IGRP	Interior Gateway Routing Protocol (Cisco)
IP	Internet Protocol
IS-IS	Intermediate System-Intermediate System
MTU	Maximum Transmit Unit
OSPF	Open Shortest Path First
POP	Point of Presence
QoS	Quality of Service
RIP	Routing Information Protocol
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
TCP	Transmission Control Protocol
TTL	Time To Live
UDP	User Datagram Protocol

Agilent RouterTester

RouterTester provides true Internet-scale testing through realistic routing protocol support, multi-stream wire-speed traffic generation and real-time analysis, and multi-port scalability. RouterTester is set to grow as the testing needs of the carrier class router industry evolve to meet the challenges of scale and Quality of Service within the Internet.

www.Agilent.com/comms/RouterTester

United States:

Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:

Agilent Technologies Canada Inc.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
1-877-894-4414

Europe:

Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-9999

Japan:

Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:

Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 267-4245
Fax: (305) 267-4286

Asia Pacific:

Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 2599-7889
Fax: (852) 2506-9233

Australia/New Zealand:

Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

Specifications subject to change.

5968-1233E 03/00 Rev B

Agilent Technologies
Innovating the HP Way