



White Paper
**Characterizing
Layer-4 Forwarding
Performance**
"Moving up the stack"



Agilent N2X



Background and Motivation

Service-oriented networks are a growing focus as service providers look to converge multiple services onto a single network to offer competitive triple-play or multi-play bundles to subscribers with the clear objectives of improving customer retention and increasing the average revenue per user. According to a study by Cox Communications¹, customer churn in two-product households is reduced by 18% and likewise in three-product households by 45%. Furthermore, customers who purchase bundled services spend typically spending twice as much as unbundled customers².

While service bundling is clearly key to retaining customers and driving revenues, converged networks also face tough new challenges in delivering the availability required to meet customers' expectations. For example, according to Network Strategy Partners³, 99.992% availability is sufficient for internet access, while 99.9994% availability is required for voice services based on the "plain old telephone service" implemented with Class 5 switches. Video services have even more stringent requirements, as consumers are reported to be ten times more sensitive to packet loss for video services than for voice services.

Operators of converged networks run the extreme risk of losing subscribers, be they long-term customers, or "converts" won over from the competitors' installed base, if the subscribers' experience does not meet the expectations to which subscribers have grown accustomed. The subscribers' experience extends across the full spectrum of interactions with the service provider, from the moment a subscriber signs up for a particular service, to direct experience with each service, as well as billing and support. At the heart of service delivery is the network itself, which moves content from either server to client, or in some cases, client to client. IPTV, VoD, web, and e-mail are examples of services delivered from servers to clients. VoIP, instant messaging and peer-to-peer, are examples of popular services delivered from clients to clients. Since network resources are finite, underlying network infrastructure must be capable of delivery and scaling a variety of contentious services while enforcing policies and traffic prioritization.

The Challenges of Peer-to-Peer

Peer-to-peer (P2P) traffic in particular presents one of the thorniest challenges to service providers. P2P traffic is reported to consume up to 65% of bandwidth on service providers' networks, but generates relatively little revenue. Moreover, profit-generating services such as VoIP and IPTV are threatened by aggressive bandwidth consumption of P2P traffic.

Furthermore, the problem of P2P traffic is becoming more complicated as file-sharing programs mutate and proliferate. For example, BitTorrent, the dominant P2P application, is evolving to masquerade as other application traffic in order to deliberately fool traffic shaping algorithms¹. Other P2P file-sharing programs such as Share, WinNY, WinMX, and Xunlei Live are emerging in new markets to "compete" with BitTorrent. Clearly, P2P traffic represents a complex and evasive threat to revenues and profits.

Bandwidth and resource contention by these diverse services ultimately means that individual flows must be tracked and prioritized over other flows based on criteria specified in service provider policy.

1. Cox Communications VoIP Whitepaper, May 2004

2. SBC Investor Briefing, July 22, 2004

3. Business Benefits of Modular Software in Carrier Ethernet Routers, Network Strategy Partners, September 2005

Applying Policy—Prioritizing and Tracking in a Stateful Environment

The need for prioritized or policy-based service delivery has pushed the designers of routers and switches to incorporate new strategies to manage these hundreds of thousands, if not millions, of individual flows. One strategy employs subscriber awareness, where the forwarding device organizes flows by subscriber to ensure that the end user receives the committed bandwidth and quality of service (QoS). A second strategy employs service awareness, where the forwarding device organizes flows by service type to ensure the appropriate QoS is applied to each service. Yet a third approach is to blend these various approaches together into a hierarchical scheme, such as applying shaping to flows based on service type at one stage, then applying shaping to a subscriber's bundle of services (for example, all services on a particular VLAN ID) at a second stage.

The need to track flows has likewise pushed the designers of routers and switches to track the state of active connections passing through their devices. While forwarding devices may not be peering directly with a client or server at the other end of a TCP connection, leading-edge forwarding devices are beginning to track the state of each connection as they forward the packets which make up the two flows (one in each direction) in a TCP connection. The problem of tracking flows is further complicated by the nature of TCP itself, which backs off transmission in the face of congestion, then attempts to re-transmit any packets lost in previous transmission attempts. TCP itself exists in many incarnations (Reno, New Reno, Tahoe, etc.) which incorporate various techniques for managing a connection in the presence of congestion and packet loss.

While stateless UDP-based services such as VoIP and IPTV may seem relatively simple to manage compared to stateful TCP-based services, these seemingly independent flows can and do affect one another if shared resources on a forwarding device become exhausted. For example, aggressive P2P traffic may consume bandwidth which should be allocated to IPTV or VoIP, underlining the criticality of testing with both stateless UDP and stateful TCP traffic, over the same interfaces, simultaneously.

If we consider that roughly 80 to 85% of traffic in the Internet is comprised of stateful TCP data, and that several hundred thousand connections will be established through a forwarding device, implementing the above strategies at scale, in a dynamic and complex environment, is clearly a technically daunting task.

Quality of Experience – Beyond Quality of Service

QoS, as presented above, is well established as a technique for prioritizing flows: packets are identified, marked with a DSCP or TOS value, and passed through a network with the objective of meeting a service level agreement that typically specifies bandwidth and latency. However, the consumers of traditional voice and television services do not evaluate these services based on packet loss and delay as laid out in an SLA. Recognition of this fact has given rise to a new category of metrics, broadly referred to as quality of experience (QoE) indicators that more closely estimate the human perception of a service. With the packetization and transport of video services over converged networks, Media Delivery Index (MDI), as per RFC 4445, has emerged as the dominant QoE metric for testing network elements in a video delivery infrastructure. Similarly, with VoIP services, voice quality metrics such as E-Model have emerged to estimate the human experience with packetized voice. Data services are of course native to data networks, but efforts to standardize quality of experience are nevertheless being made by the AppDex forum, which proposes a framework for expressing QoE for data applications.

Moving past basic Traffic Blasting

The nature of the challenges outlined above is making existing forwarding performance test methodologies obsolete. Over the past 8 years, the characterization of forwarding devices has become crystallized around the paradigms outlined in RFCs 2544 (for L3 devices, i.e. routers) and 2889 (for L2 devices, i.e. switches). Both of these documents describe a standardized methodology for finding the non-drop rate of forwarding devices for various frame or packet lengths. Consequent forwarding performance test solutions became very much a reflection of these RFCs, with the addition of enhancements such as support for mixed packet lengths or integrated protocol emulation. While such models and test solutions have been sufficient for simpler networks delivering data on a best-effort basis, they are no longer appropriate in the context of multi-service next generation networks that use intricate prioritization schemes and policies.

As networks evolve to support converged services in the dynamic context of stateful TCP connections, the switching and routing world is waking up to the fact that the basic traffic blasting paradigms in RFCs 2544 and 2889 do not characterize forwarding devices as they are being used in the real world. Furthermore, new architecture, which employs hierarchical queuing, and statefully track individual flows and connections, have created a need to test with realistic traffic.

Realistic traffic must include a mix of services, including those running over UDP, such as IPTV and VoIP, as well as a fully-featured data services running over a fully-featured TCP stack. Credible test solutions should offer integrated stateless traffic (for UDP-based services) and stateful traffic (for TCP-based services) running in the same software application on the same physical port. Emulated TCP traffic should also be realistic in the sense of supporting re-transmissions and dynamic windowing to faithfully re-create actual conditions in the Internet. Further, emulated TCP traffic must also scale in bandwidth in order to achieve at least 85% of line rate (i.e. the percentage commonly found in real Internet traffic), using short packet lengths, over a gigabit Ethernet interface. Finally, realistic application transactions for the most common TCP-based applications, such as BitTorrent, HTTP and FTP, must be simulated within those TCP connections to practically stress application-aware devices⁴.

On the measurement side, the test instrument must provide meaningful measurements for both stateful traffic and stateless traffic to validate that QoS and QoE levels are maintained. In the context of stateful TCP traffic, for example, goodput is a measurement of the useful data transported by TCP, and can be drastically different from measuring throughput at L2 or L3 in a lossy environment. All of this must be wrapped up in an integrated L2-7 solution that can provide accurate and repeatable forwarding characterization of intelligent forwarding devices.

4. "BitTorrent End to End Encryption and Bandwidth Throttling – Part 1", Slyck, February 6, 2006. SBC Investor Briefing, July 22, 2004

This page intentionally left blank.

Sales, Service and Support

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.
2660 Matheson Blvd. E
Mississauga, Ontario
L4V 5M2
1-877-894-4414

Europe:

Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom

07004 666666

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) 269-7500
Fax: (305) 267-4286

Asia Pacific:

Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
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

This information is subject to change without notice.

Printed on recycled paper

© Agilent Technologies, Inc. 2007

Printed in USA June 13, 2007

5989-6994EN

