SmartMetrics

THE KEY TO QoS TESTING

NetCom
SYSTEMS

SmartMetrics

The Key to QoS Testing

NetCom SYSTEMS

P/N 340-1043-001 Rev A, 10/99

Netcom Systems, Inc.
(800) 886-8842 Toll Free
(818) 676-2300 Phone
(818) 881-9154 FAX

Disclaimer
The information contained in this document is the property of Netcom Systems, Inc. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior written permission of Netcom Systems, Inc.

Information contained in this document is subject to change without notice and does not represent a commitment on the part of Netcom Systems, Inc. Netcom Systems, Inc. assumes no responsibility or liability for any errors or inaccuracies that may appear in this document.

Trademarks
- SmartBits and SmartMetrics are trademarks of Netcom Systems, Inc.
- Windows is a registered trademark of Microsoft, Inc.

Table of Contents

Introduction **1**

The Need for QoS **2**

Today's Quality Of Service **3**

Why Network Performance Analysis **5**

The Need For Better Testing **6**

SmartMetrics™ Testing **7**

- Testing All Layers 7
- Per-flow Metrics 9
- Connection Metrics 14
- Application Metrics 16
- Edge Device and Cross Technology Metrics 18
- A Word About Ping 19

Who needs SmartMetrics **20**

The Ideal Equipment for Network Performance Analysis **21**

The Shift to SmartMetrics™ **23**

Introduction

In today’s world, the network is critical. Whether you are a Network Equipment Manufacturer (NEM), a Network Service Provider (NSP), or an Enterprise, your customers are depending on the network for everything from increased productivity to transferring transactions worth millions of dollars per minute. More users, exponentially increasing traffic, support of intranets and extranets, and new applications such as virtual private networking, voice over IP (VoIP), multimedia, and E-commerce have resulted in huge demands for new bandwidth and more efficient ways to use it.

At the same time, networks are becoming much more complex, using multiple technologies (e.g., Ethernet-to-ATM, Frame Relay-to-ATM, Ethernet-to-POS), adding newer mediums (cable modem xDSL), and growing in size to hundreds or even thousands of ports (see figure below). The complexion of the network traffic is also changing. The concept of a converged network, as a way to reduce networking costs has become accepted. So in addition to handling increased traffic, now the network must also contend with the delay-sensitive nature of voice and multimedia. User demands for increased performance and new services are growing faster than the network's ability to support them.

New Quality of Service (QoS) technologies offer the opportunity for NEMs to develop new products, NSPs to offer value-added services, and Enterprises to upgrade their networks, helping all to cost effectively and rapidly respond to these changing requirements.

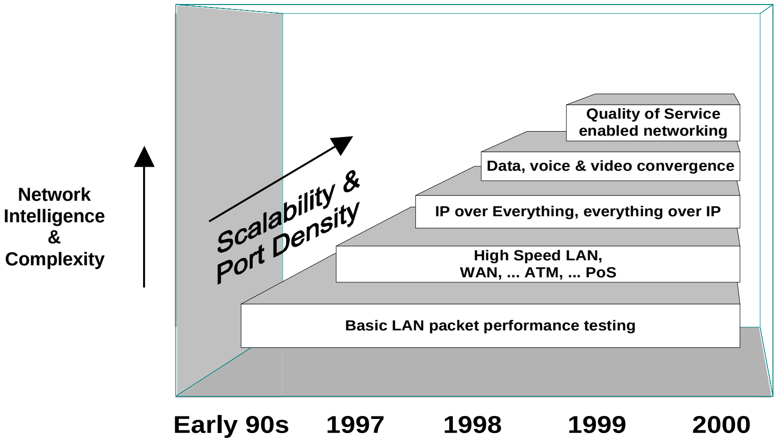

Increasingly complex traffic flows, multiple services and technologies, huge voice, multimedia, and data demands, plus dramatically increased numbers of ports mean that aggressive performance testing is becoming a necessity in all stages of effective network management today -- before the network is built, while the network is being built, and when the network is up and running.

But how do you know which products to choose, how a new application will work and how it will effect existing applications, or whether the network will survive when it is called upon to support huge sudden increases in traffic? Proactive performance testing with new performance analysis systems and testing metrics make it possible to predict which technologies are appropriate, how the network will react to the new demands, when it will fail and precisely which application is the cause. This change is a giant leap forward in verifying network reliability and it's ability to provide Quality of Service (QoS) - before going live.

The Need for QoS

In the midst of all this change, there is one constant: Internet Protocol (IP). IP is the transport of choice for the new network. The convergence of networks, however, is showing the weakness of IP. IP is not able to **guarantee** particular performance characteristics. This weakness has driven the need for IP QoS.

IP is a "best-effort" protocol. Best-effort means that the network does not guarantee how long it will take to send packets, what order they arrive in, or even if the sent data will reach it's destination. Until now, this less-than-perfect performance level has not necessarily been a drawback. For typical network applications like email or file transfer delay is not factor and the integrity of the data is ensured by the Transmission Control Protocol (TCP) layer, which reorders packets and retransmits lost data. In the converged network however, the old rules no longer apply. Today's interactive applications such as E-commerce, voice and video conferencing, and call centers are adversely affected by lost packets, delay, and delay variation (jitter).

On the lightly loaded network, these issues are not typically a problem. When there is ample bandwidth and delay is minimal, jitter is not a factor because the delay is constant and packets are rarely lost. These problems do become critical though when there are bursts of heavy traffic or when the network becomes congested.

One solution to this problem is to design a network that will always have excess bandwidth so that it never gets congested. In most scenarios, this would be neither economically feasible nor desirable.

So how can you create a network that meets your business demands for integrated data, voice, and multimedia? Is it possible to avoid the costs of bandwidth over-subscription? Can an organization get the service it needs by using the Internet?

Positive answers to these questions hinge on the ability of an IP network to ensure that mission-critical and delay-sensitive traffic receives the service needed to maintain the required quality level. Enter IP QoS.

Today’s Quality of Service

What is Quality of Service?

To the user, Quality of Service means quick response, availability, ease-of-use, and trust.

In networking terms, QoS is the ability of all layers of the network to cooperate and provide the performance that each application needs to meet user expectations. QoS also allows network bandwidth to be used more efficiently and thus more economically. The goal is to deliver quality end-to-end service for user applications -- including data, multimedia, and voice.

The key to offering QoS lies in the ability of the network to provide sufficient bandwidth to meet the average network loads, to distinguish between the different applications or classes of traffic, and to allocate the network resources required to deliver the performance characteristics required for that traffic. The nature of the application and/or the importance of the application to your business model, dictates which QoS parameters will be required.

QoS is characterized by a set of network performance metrics including:

- Service availability - ability to gain access to the network.
- Throughput - the rate at which packets are transmitted in a network; measured as average or peak rate.
- Latency or Delay - the time it takes a packet to travel between the transmitting and receiving points on the network – one way or round trip.
- Delay variation or jitter - the variation in time between all the packets in a stream or flow between endpoints.
- Packet loss – the rate at which packets are discarded by the network.
- Packet sequence – the ability of the network to deliver packets in proper sequence.
- Connection availability – the ability of the network to complete the required connections (i.e., TCP or SVCs) with the requested performance characteristics.

There are two main approaches to delivering QoS: resource reservation and prioritization.

Resource reservation (IntServ), such as RSVP, is a signaling protocol which sets up an end-to-end path with specific QoS metrics. If such a path cannot be created, the connection is refused.

Prioritization (DiffServ) classifies each type of traffic according to the specific QoS metrics that it needs. Each classification is mapped into a Per-hop Behavior (PHB) which defines how each node in the network should treat the packet. For example, traffic can be differentiated into real-time (like voice or multimedia) and best-effort (like file transfer or e-mail) traffic. The former traffic type would receive the highest priority through the network as defined by the PHB; the latter would receive lower priority. The nodes in the network use a variety of queuing and congestion management schemes such as Weighted Fair Queuing (WFQ) and Random Early Detection (RED) to give each packet the priority it needs.

Of course, defining certain traffic as "high priority" will not help at all unless all the devices, networking layers, and applications can interpret that priority information and offer each packet the performance it needs. That’s where SmartMetrics™ and proactive network performance testing come in.

QoS may seem subjective, but with the right tools, it is measurable. The best way to guarantee long-term QoS is through rigorous, standardized testing at both the component and system levels under a variety of conditions.

The SmartMetrics layered approach to performance analysis addresses all the necessary measurements that collectively determine the QoS of a network. It does this by creating the testing methodology and implementing new tests to easily and extremely accurately, perform QoS metric analysis.

Why Network Performance Analysis

Because of today's dependency on the network, performance analysis is required to ensure that it will be there when it is needed, that high priority traffic can get through even under the most congested conditions, and that it will not fail under extraordinary loads. Aggressive performance analysis, under all conditions, is the only way to guarantee that a network will successfully meet QoS requirements.

Since no single product or technology can deliver end-to-end QoS, it will most likely require a combination of several. Testing must assess both the individual components and the overall network, under a variety of conditions, in order to be successful. The only way to judge real network performance is to put the whole network or minimally, a representative part of the network, to the test.

Network performance analysis allows one to see how each packet loaded onto a network and run at full wire rate, affects the function of every other part of the network. It is this kind of proactive testing that allows a network operator to plan ahead to eliminate delays and avoid catastrophic failure in the system. This testing addresses the actual user experience to determine how well a network works using "what-if" scenarios – stressing the tested network or network device to its maximum. This type of testing is proactive and is performed on tens, hundreds, or even thousands of ports simultaneously.

Network performance analysis helps the Network Equipment Manufacturer qualify their products for real-life situations, compare performance to expected design goals, and substantially reduce the time to market.

For Network Service Providers and Enterprises, network performance analysis provides the ability to build test scenarios that can emulate the most demanding network environment, allowing comparison between competing equipment, providing a way to measure end-to-end network performance, and ensuring the projected QoS for new applications, before actually connecting users.

Requirements for Network Performance Analysis

- *Must Generate, capture, and analyze traffic at full rate.*
- *Scalable - 100s of ports, 100,000s of streams.*
- *Test layer 2, 3, and 4 networks*
- *Deterministic, repeatable, remotely controllable*
- *End-to-end, edge-to-core, core-to-edge*
- *Test Multi-technology (LAN/WAN/ATM)*
- *Measure new SmartMetrics™*

Network performance analysis should begin in the planning phase before the network is built and must continue to after it becomes operational. A dependable network means productive users and happy customers.

The Need For Better Testing

Before the advent of converged networking, testing was simple. All we needed to know was how many frames/cells per second could pass through our network or could the device switch cells at the designated rate. Scott Bradner of Harvard University devised tests to define the performance of network devices.

These traditional metrics (performance test measurements) were defined in Bradner's RFC 1242 and RFC 1944 (updated to RFC-2544). These RFCs define tests for:

- Throughput - How many frames per second can be forwarded through a device before it starts losing frames.
- Packet Loss - How many packets per second are lost through a device at various throughput loads.
- Latency - What is the delay created by sending a packet through a network device under various load conditions?
- Traffic Bursts - Devices and networks can be very sensitive to bursts of traffic above the normal load. Most devices have memory buffers. This metric examines how a device copes with bursts of traffic.

Traditional performance testing (see figure below) focused on frame/cell/packet measurements on a per-port basis.

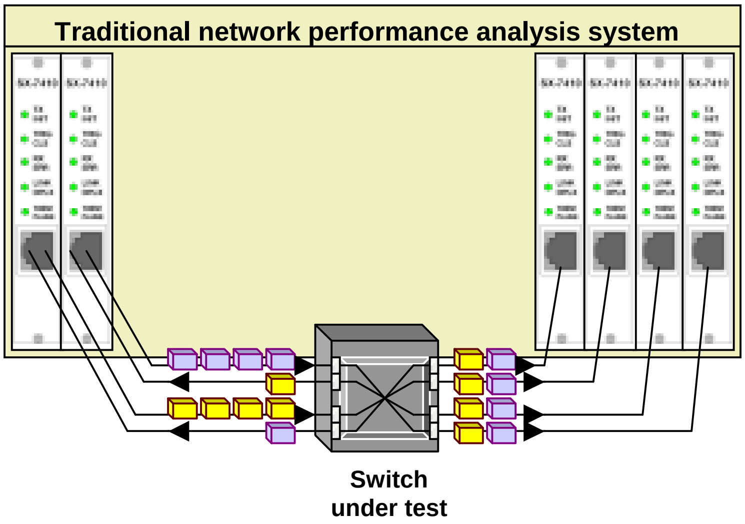

"Packet blasting" - simple per-port traffic loading and measurement

Traditional testing is necessary, but is no longer enough for today's network. It does not measure QoS and the true user experience. In order to do this, the testing must address individual traffic **flows and streams**.

SmartMetrics™ Testing

In today's converged networks, many different applications are multiplexed onto a single device port, combining different class of service traffic, supporting voice, multimedia, data, and multicasting. Each type of traffic could have a distinct set of performance requirements. In order to accurately test such a network the test must take on the attributes of a live network. Testing must emulate many different types of traffic, determine the impact that traffic has on performance, and determine the network's ability to handle lightly to fully loaded traffic conditions.

Today's networking scenario can no longer be properly tested on a per-port basis. To determine the true QoS, testing needs to address:

1) All network layers.

2) Individual IP flows and streams.

3) How new applications will effect the network before it goes "live"

4) The transition between networking technologies.

Testing needs to address the issues created by thousands, even millions, of traffic streams; the complexities of carrying multi-protocol traffic; and the need to meet specific service-level performance requirements.

The SmartMetrics layered approach to performance analysis is designed specifically to address today's complex set of optimization and prioritization methods at all network layers. It is the only true test of QoS. SmartMetrics offers the ability to measure:

- Virtual Local Area Network (VLAN), IP Type of Service (TOS), DiffServ, Multi-protocol label switching (MPLS), voice, and multimedia traffic flows.
- Calls and connections (signaling functions).
- Network applications (management capabilities).
- Transition of the data between different technologies.

SmartMetrics gives you the ability to measure and analyze every aspect of your network, from the performance on each network port, to the performance of millions of IP flows, to the effect of opening and closing thousands of TCP or Multicast sessions.

Testing All Layers

The optimization of traffic can take place at layer 2, layer 3, layer 4, and even up to layer 7 (see the following figure). IEEE 802.1p prioritizes traffic at layer 2 (Datalink). Diffserv and TOS optimize and MPLS and RSVP manage resources at layer 3 (Network). A multitude of new QoS products, such as server load balancers and traffic shaping/access control services, optimize traffic based on criteria including TCP or User Datagram Protocol (UDP) port number at layer 4 (transport). Even higher-layer criteria such as Uniform Resource Locator (URL) or application type can be used.

Network Characteristics Summary

Layer **2** networks	Optimize traffic on MAC addressing, Layer 2 VLANs (IEEE 802.1Q, 802.3ac, port, MAC).
Layer **3** networks	Optimize network performance based on Layer 3 **flows**, traffic management using routing, Layer 3 VLANs (Subnet, Protocol), or by traffic prioritization or reservation**.** May also include optimization features of Layer 2 networks.
Layer **4** networks	Optimize network performance based on TCP/UDP flows and connections. Manage data flow by connection type & security. May also include optimization features of Layer 2 and 3 networks.

Equally important factors for guaranteeing QoS are: the verification of connection capabilities; multicast IP traffic behavior; firewall performance under load; virtual private network (VPN) load handling, efficiency, and leakage under extreme conditions.

> *Since end-to-end network QoS will require the cooperation of several technologies, all of the network components and applications at all layers must be tested.*

SmartMetrics addresses these needs by measuring how well devices and networks optimize, prioritize, and segment traffic using an expanded set of metrics. SmartMetrics measurements fall into the following categories:

- Per-flow metrics.
- Connection metrics.
- Network application metrics.
- Access device and cross-technology metrics.

Per-flow metrics measure the network's ability to efficiently route different priority traffic at all network layers. At layer 2, it tests 802.1Q VLAN and 802.1p where networks optimize based on information placed in the VLAN tag in the Ethernet frame. At layer 3, it verifies the correct operation of the different optimization methods (Diffserv and TOS, VLAN based on subnet or IP address) and reservation protocols including RSVP and even MPLS. At layer 4, SmartMetrics measures prioritization based on UDP or TCP port number simulating (e.g., FTP, TELNET, and HTTP traffic flows).

Connection metrics measure the network's ability to manage sessions. These measurements determine how performance is affected by the network's ability to set up, maintain, and tear down sessions; how many users the it can handle; and the it's breaking point when dealing with extraordinary connection requests. Connection Metrics are available for ATM SVCs at layer 2 and TCP sessions at layer 4.

Network application metrics measure a device's ability to handle multicast traffic on both data and control planes; the performance of firewall/VPNs; router and DiffServ management processes; and the effects of queue depth and processor loading on QoS.

At the higher layers, SmartMetrics tests can generate different classes of traffic, such as those required to test QoS requirements for voice, multimedia, and other critical applications, to verify their operation under all possible network conditions.

Testing all network layers helps you determine how they work or do not work together to provide each application, and ultimately the user, the Quality of Service needed (see figure below).

SmartMetrics – Quick Look

Layer	Concern	Metric
Layer 2 Data Link	- Amount & speed of data between devices. - Delays between packets. - Ability to handle bursts of traffic.	- Throughput - Latency - Packet loss
Layer 3 Network	- Varied network delay within a traffic flow. - All packets in a flow arrive in correct order. - Track every packet of every flow. - VLANs (traffic segmentation).	- Latency variation - Sequence tracking - Throughput, latency, and packet loss - Test by protocol, subnet, services
Layer 4 Transport	- Handling bursts of calls and connections. - Response times for end-station. applications (video, voice, FTP, HTTP, email). - Network fabric performance (firewall, multicast IP, etc.).	- Maximum call rate, - Connection setup rate - Latency variation - Customized to application
Edge and Cross Technology	- End-to-end interoperability, policy mapping. - Performance at technology transition. - Multiple technology support: Cable Modem; xDSL; IP over ATM; ISP access ATM to Frame Relay.	- Latency variation - Sequence tracking - Throughput, latency and packet loss

Per-flow Metrics

The typical network user might be running several applications (i.e. voice, multimedia, data) simultaneously. A large number of users means that numerous applications can be multiplexed over a single port, combining multiple class-of-service traffic or flows (see the following figure). A network can be required to manage thousands or even millions of these flows at any one time. Each flow of traffic could have a distinct set of performance (QoS) requirements, which need to be measured. SmartMetrics Per-flow measurements are the only way to accurately test the QoS capabilities of your network device or network.

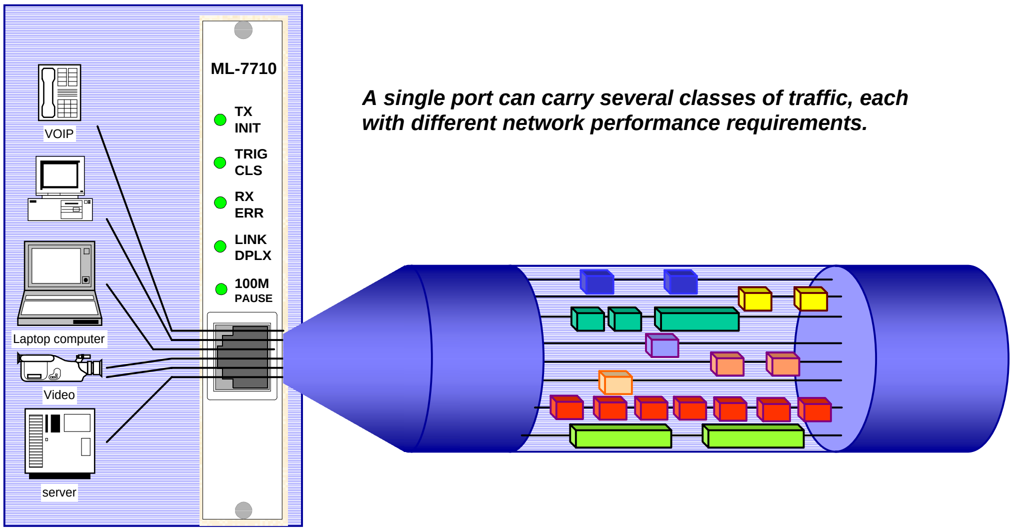

A single port can carry several classes of traffic, each with different network performance requirements.

Per-flow metrics move testing up the stack, beyond the traditional per-port testing. They measure the actual data flow. Quality of Service measurements, made per flow, include:

- Throughput - the rate at which packets traverse a network without loss.
- Latency or Delay - the time it takes a packet to travel between the transmitting and receiving points – one way or round trip. Records the minimum and maximum latency levels and calculates the average latency.
- Delay variation or jitter - the variation in time between all the packets in a stream taking the same route.
- Sequence Tracking - provides precise readings of the number of frames received in sequence, and the number of frames received out of sequence. It also provides frame loss testing on a per-stream basis.
- Latency over Time - for each port, the test records the number of frames received, minimum latency, and maximum latency during a selected time interval. The test also calculates the average latency for each port.
- Latency Distribution – measures the total number of frames received and the number of frames received within each of several variable time intervals.

Per-flow metrics measure:

- *Traffic **optimization** of flows or streams (Layer 3 switching)*
- *Traffic **prioritization** of flows or streams*
  Specifies prioritization of traffic based on policies such as ToS/DiffServ, RSVP, MPLS, or IEEE 802.1 VLAN priority.
- *Traffic **segmentation** of flows or streams (Virtual LANs)*
  Forwarding traffic only to addresses that meet the VLAN criteria.

The SmartMetrics approach defines tests which generate and analyze millions of end-user flows. How does this work? SmartMetrics allows the definition of thousands of **streams** per port. Streams can be set up to represent different types of traffic, having a distinct set of characteristics and performance requirements. A stream specifies the class-of-service, protocol type, VLAN type or priority (IP ToS, DiffServ, IEEE 802.1 VLAN priority, source/destination address, or TCP/UDP port number). In turn, thousands of independently-controlled **flows** are created for each stream, offering the capability to measure millions of flows per port. A flow represents data traveling between a source and destination defined by MAC or IP address, having the performance requirements of the stream. Flows can be set up to move from any source address to any destination address. Each stream can contain one to multiple flows on multiple ports (see figure below).

SmartMetrics™ Measures QoS for each traffic flow. Up to millions of flows are supported.

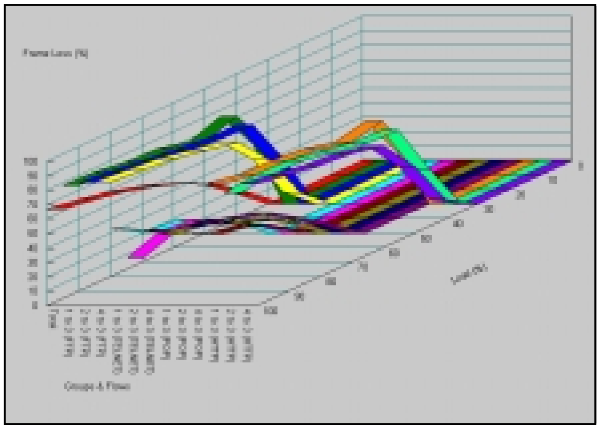

The key to the SmartMetrics approach revolve around it's unique receive analysis and tracking capabilities, allowing for extended testing. Each and every frame in the stream is stamped with a unique signature. Using that signature, SmartMetrics can track packets in millions of flows, simulating different classes of service (i.e., voice, multimedia, and data). The results of the QoS analysis are then presented per flow, per class of service, and as a total. Since the measurements are made in real-time (reported as histograms), the test can be run for extended periods of time at multiple loads. With streams and flows, and the ability to run the tests for hours, even days at a time, SmartMetrics is able to simulate real-life network operation (see below).

Transmit 64 million flows/port, track all flows in 15,000 streams, and analyze QoS over long periods of time.

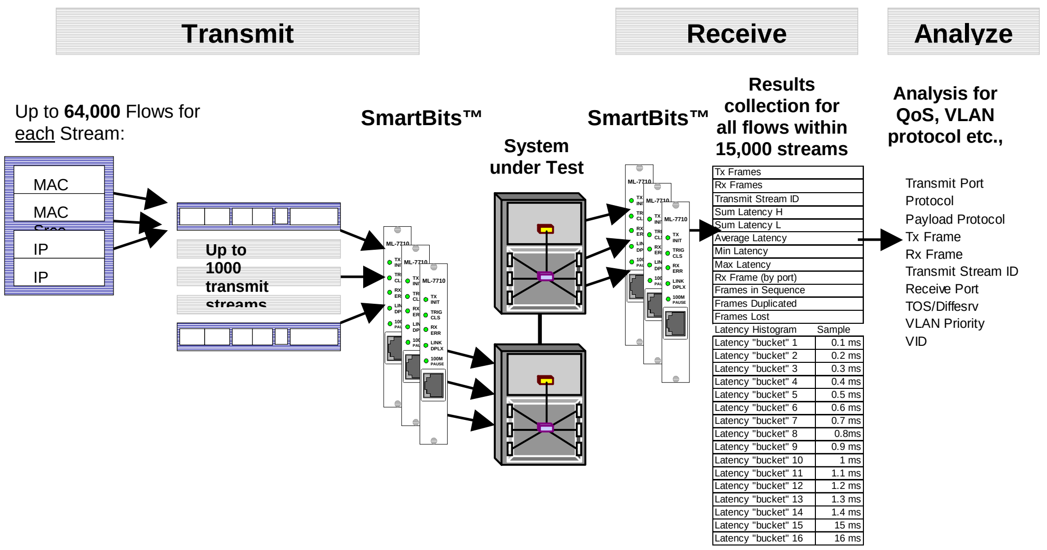

Tx Frames
Rx Frames
Transmit Stream ID
Sum Latency H
Sum Latency L
Average Latency
Min Latency
Max Latency
Rx Frame (by port)
Frames in Sequence
Frames Duplicated
Frames Lost

Latency Histogram	Sample
Latency "bucket" 1	0.1 ms
Latency "bucket" 2	0.2 ms
Latency "bucket" 3	0.3 ms
Latency "bucket" 4	0.4 ms
Latency "bucket" 5	0.5 ms
Latency "bucket" 6	0.6 ms
Latency "bucket" 7	0.7 ms
Latency "bucket" 8	0.8ms
Latency "bucket" 9	0.9 ms
Latency "bucket" 10	1 ms
Latency "bucket" 11	1.1 ms
Latency "bucket" 12	1.2 ms
Latency "bucket" 13	1.3 ms
Latency "bucket" 14	1.4 ms
Latency "bucket" 15	15 ms
Latency "bucket" 16	16 ms

Let's look at some important SmartMetrics per-flow measurements.

The network’s ability to prioritize traffic is key to providing QoS. It must be able to recognize which traffic needs priority and must be able to offer that priority, under all conditions. SmartMetrics offers the methodology to measure the network’s ability to provide the required priority for each flow of traffic.

Individual flows are generated and delivered at port densities and data rates that are high enough to stress the network. These flows represent various traffic priority treatments (e.g., voice and e-commerce transactions require high priority, while email suffices with low priority). Depending on the criteria, the high-priority traffic should get through within the required performance characteristics, no matter how congested the network becomes, while lower priority packets can be dropped. Per-flow priority can be assigned based on 802.1p, IP TOS, UDP source/destination ports, IP address, MAC address, or even application type.

SmartMetrics generates and analyzes traffic, tracking thousands of prioritized flows and accurately measuring performance for each priority level. This determines if the network or network devices are able to enforce QoS policies and meet the latency, latency variation, packet sequence, throughput, and packet loss performance levels that are required for each flow, to deliver new high-quality applications.

For example, when Voice over IP (VoIP) or video over IP packets are sent over the network, a manager or test engineer needs be able to determine the difference in latency (jitter) between the packets. This parameter is important in networks providing VoIP or multimedia multicast-over-IP because, while these applications are not very sensitive to delay across the network, variation in the delay can have a profound effect on quality (see figure below).

Latency variation affects voice and multimedia quality

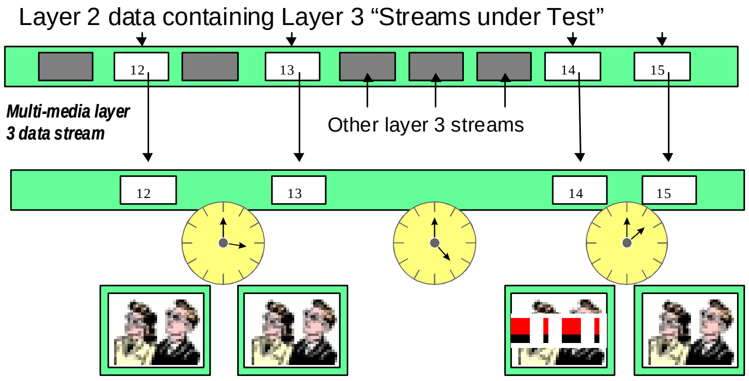
Layer 2 data containing Layer 3 “Streams under Test”
12
13
14
15
Multi-media layer 3 data stream
Other layer 3 streams
12
13
14
15

Consecutive packet loss is another key metric involved with networks sending VoIP. A loss of several packets in a row will also affect voice quality. (see figure below).

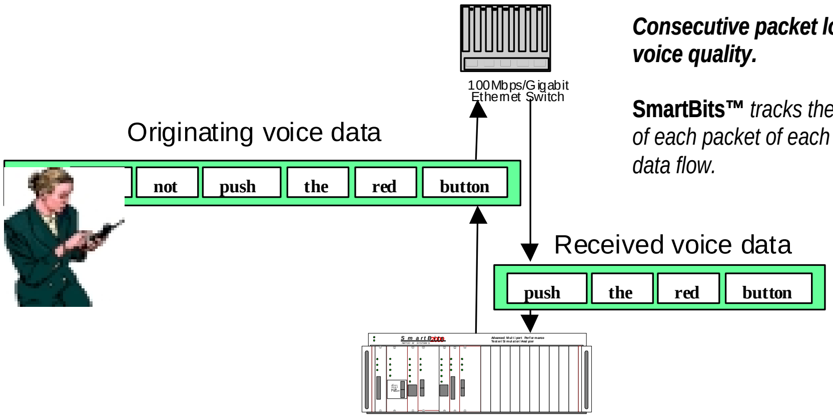

Consecutive packet loss affects voice quality.

SmartBits™ *tracks the sequence of each packet of each incoming data flow.*

With per-flow metrics, these parameters can be measured for each individual voice flow. This helps the manager see how well the network is delivering the QoS required for acceptable voice quality. It can also be determined how the voice traffic is doing in relation to other application traffic (e.g., it can be determined if less-sensitive e-mail or web traffic is taking up too much bandwidth, increasing network jitter, or causing packet loss). Once that is determined, perhaps a policy change giving the voice higher priority, can improve its quality. This requires an understanding of the different kinds of traffic and how they affect one another. The only way to make this measurement is to be able to track these metrics on a per-flow basis, for hours or even days. SmartMetrics tests make this possible because they are not limited by capture-buffer size, which could limit testing to only a few milliseconds when dealing with high-speed networks.

Another key measurement is per-stream-sequence tracking. Sequence tracking is the ability to make sure the packets sent from one end arrive in the right order at the receiving end. If packets start to arrive in the wrong order, they could be, depending on protocol, regarded as lost and will be retransmitted. This causes the network to retransmit many packets, which increases traffic and causes further degradation.

Even though a quick look at the simple layer 2 packet-level metrics might appear to show that all is well (because all it will show is the loss of one or two packets on a network that otherwise appears fine), this is misleading. What is hidden is the fact that out-of-sequence packets are causing retransmissions. While the layer 2 test will show data going across, much of that traffic could consist of retransmissions. The result is that overall data throughput plummets.

With per-flow metrics you can determine:

- The impact of prioritization on performance (priority-based traffic vs. non-priority-based traffic).
- The network's ability to deliver quality, high-priority traffic (can the network support high quality voice and multimedia?).
- What happens to low-priority vs. high-priority traffic under congestion or network over-subscription conditions (Does any low-priority traffic get through? What are the effects on the high-priority traffic? Does latency increase?).
- How many flows the network/device can handle.

- The network's (or device's) ability to forward very large numbers of flows and the rate of these flows.
- The network's ability to handle implemented policies correctly.
- How policy changes might effect the performance of the network before the change is rolled out permanently.

Connection Metrics

When performance is not what it should be and users are complaining, one is never sure if a problem is at the server, in the network, with the connections, or caused by the management of the network. Because it is so difficult to determine the root of a problem, it is critical to measure, in addition to the performance of individual flows, how many calls per second can be set up or taken down and how many calls can be sustained across the network. There have been numerous cases of catastrophic network failure because the system simply could not handle the call requests.

Connection metrics deal with such basic questions as "How many calls per second can be set up in a network?" or "How many connections per second can be set up through a device?" Among the key call connection metrics are the peak binding rate, peak session rate, session capacity testing, and session setup/tear down performance (see table below). In determining network QoS, it is vital to have a view of the sustained call rate performance of the network as well as its ability to set up and tear down sessions.

Call Connection Metrics

- *Connection Setup Rate - The rate of TCP connection establishment.*
- *Connection Teardown Rate - The rate of at which TCP connections are closed.*
- *Connection Setup Time - Connection setup time as the request rate increases. Can indicate any performance degradation which may occur as the number of open sessions increases.*
- *Session Capacity – The ability to sustain TCP sessions over a period of time and the maximum number of open sessions that can be supported.*
- *Session Rate - The rate of TCP session setup and teardown.*
- *Call Capacity - The number of concurrent connections that can be established and maintained.*
- *Peak Call Rate- This test measures the ability to sustain call setup and teardown loading over a period of time.*

SmartMetrics defines tests that generate the connection requests of many thousands of users or large servers, testing the ability of networks and devices to manage large numbers of sessions. It benchmarks both the rate and connection capacities of the network to establish, maintain, and tear down sessions (see figure below). Typically when connections are requested at a low rate, all sessions are established. However, as the request rate increases, the network may not be able to establish all the sessions, even if other network resources are available.

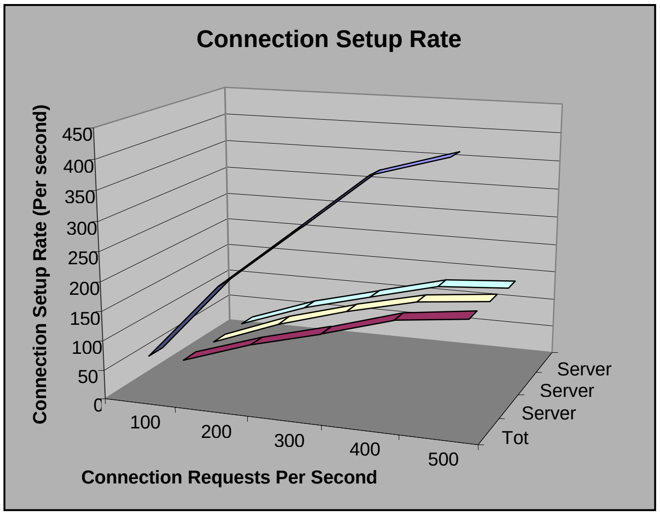

Connection Metrics determine the network's ability to create, maintain, and tear down sessions.

A failure to create sessions may occur when a request rate is too high for the network to keep up (it will reach a peak rate and then start to fall off); when the session capacity has been reached (the network cannot support any more sessions); or when the network cannot close sessions fast enough to open new ones. If a session cannot be established, such as when a phone call results in a busy signal, the user is not receiving either the connection or the QoS that is required.

There are two kinds of tests that apply in this situation: end-station and network performance tests. Both involve quality measurements. End-station checks include every-day questions such as "Can the user gain access to the network?," "Is the server still running?," or "Can the applications keep up with the demand?" Network performance tests are concerned with the ability to establish a session from the source all the way to the destination under all loads. Questions addressed by these tests are: "Are resources available to deliver the QoS required by the session request?" and "Can the network make enough connections in a short time, to handle a particular network event?"

Ideal examples for the use of call connection metrics include the testing of: a server load balancer's ability to create, maintain, and clear TCP sessions for users accessing a server farm; an ATM switch's capability to establish, maintain, and clear SVCs; the network's ability to manage PPP sessions over ATM virtual circuits; or the ability of the network to create voice over IP calls.

Testing the connection performance of a network helps determine:

- The maximum number of simultaneous users that the network can support.
- The network's ability to handle a large number of users accessing it in a short period of time.
- The network's ability to recover from a catastrophic event, such as the loss of a heavily used link or switch.
- The number of sessions that a network or device can sustain over an extended time period.
- How many ATM SVCs, TCP, PPP, or voice calls the network can manage.
- Whether or not the server can keep up with the demands of your E-commerce site.

Call connection metrics are vitally important when planning and expanding a network or preparing the network for an event where a large number of users will want to gain access in a short period of time.

Application Metrics

Through proactive analysis of the network's or device's ability to handle multicast traffic, the performance of firewalls/VPNs, the effect of router and DiffServ management processes, and how queue depth and processor loading will affect QoS, Application Metrics help managers evaluate how a new application, for example voice or video, will affect the network before bringing it on line.

As applications such as distance learning, video conferencing, and large-scale software distribution become more important, so does Multicast IP. Multicast IP conserves bandwidth by allowing transmission of a single copy of a data stream that is destined for multiple users. After initial transmission multiple copies of that stream are made, only as required, to deliver the data to the multicast group (individual recipients). The network must be tested in advance to determine if it can handle new multicast applications. The following questions need to be answered: "Can the network offer the QoS required by multicast applications?" and "How will multicast affect applications already running?" To answer these questions, measurements must be made on the network's ability to process and transmit multicast packets and the way the groups are managed through Internet Group Management Protocol (IGMP). Important test metrics for multicast IP include: throughput (used when sending a mix of multicast and unicast data [mixed-class throughput]); the latency of multicast packets being sent to multiple receivers; the time it takes to set up and tear down the multicast groups, and the number of simultaneous multicast groups that the network can support.

The SmartMetrics approach offers tests based on RFC 2432 (IP Multicast Benchmarking Terminology) to help you measure these metrics and determine the network's ability to deliver multicast traffic in a "real-world" situations

Firewalls and Virtual Private Networks (VPN) products protect your vital data. The VPN, through tunneling, encryption, authentication, and access control secures your data as it travels from remote users to corporate locations over a public network. The firewall protects your network from unauthorized access. How though does this protection affect the users who should have network access? The performance of the VPNs can have serious affect the QoS that each user/application receives. VPNs and firewalls can be implemented as add-ons to existing network devices (e.g., router/switch, server) or through standalone devices. If the VPN or firewall is implemented as an add-on, the most basic questions are "How do these features affect the performance of the device it is running on?" and "Can the device maintain wire rate performance while it has to manage VPN or firewall capabilities?" For standalone devices, performance, scalability, and reliability must be assessed. To determine how well these devices protect your data and affect overall network QoS, important metrics for VPNs need to be tested. These include; measurement of session creation and tear down; device throughput while managing encryption and tunneling services; VPN performance level changes that occur as data access moves to broadband speeds; data changes that occur as a result of encryption/tunneling; and the effect of the VPN on user response.

Important metrics for firewalls include: the maximum number of simultaneous sessions that can be supported; the session creation rate; the maximum and sustained throughput of the device; how the addition of new security rules or features such as network address translation affect firewall performance; and most importantly, whether or not any packets or connections that should be blocked, get through, under all traffic levels.

As convergence becomes a reality, voice will be a major network application. However, this will happen only when the data network can deliver voice with the quality and reliability offered by the switched telephone network. Can your network provide the QoS required by voice in terms of quality, scalability, and reliability? SmartMetrics can help answer this question. All SmartMetrics measurements come into play when testing your network for voice. Per-flow metrics are needed to determine if voice packets receive the proper priority or reserved bandwidth to provide the minimum delay, jitter, and packet loss levels required for quality voice, under all traffic conditions. Connection metrics are needed to ensure that the network can handle the signaling requirements under normal and extraordinary load conditions to enable a VoIP connection when requested. The network needs to be able to handle the voice management (signaling and bandwidth) without degrading overall performance. Also, voice quality needs to be measured under both lightly-loaded and congested network conditions. In addition, SmartMetrics tests can be used to isolate the cause of performance bottlenecks when QoS requirements are not met. Important metrics for voice include measurement of: one way latency and jitter under all traffic conditions; voice quality according to the ITU-T P.861 Perceptual Speech Quality Measurement (PSQM); and the Voice Signaling capability of the network.

In order to measure QoS, it is important to measure both the network's data delivery performance (data plane) as well as session control performance (control plane). These measurements include assessments of routing protocol performance (i.e., RIP and OSPF) and the network's ability to perform convergence in the event of an outage. It is also important to measure the effect that queuing management and additional policy requirements have on the processing power of the individual devices and ultimately on network performance (see figure below).

For application metrics, both control and data plane analyses are required (i.e., when using multicast IP, the group management is as important as moving data).

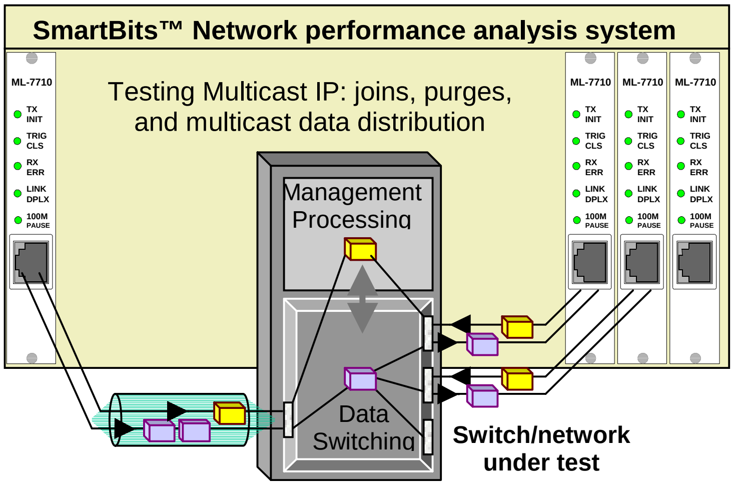
SmartBits™ Network performance analysis system
ML-7710
TX INIT
TRIG CLS
RX ERR
LINK DPLX
100M PAUSE
Testing Multicast IP: joins, purges, and multicast data distribution
Management Processing
Data Switching
Switch/network under test

Application metrics will be able to help you determine:

- If the individual devices in the network support new applications at the desired QoS.
- What will happen to the QoS of existing applications as new applications are added.
- If the network offers the consistent delay required by voice and if it can support the overhead of multimedia.
- The ability of the network to support and manage IP Multicast.
- The network's ability to maintain application performance as more multicast groups are added.
- How many simultaneous sessions a VPN support.
- If the network is a managed VPN offered through your service provider, and if it meets the agreed-upon performance level (Service Level Agreements or QoS commitments).
- How the added overhead of the VPN from encryption, tunneling, and authentication protocols affect overall network performance.
- If the firewall can really protect your confidential data without sacrificing performance and if the VPN can really keep your data secure.
- End-station application response time.
- What the "end-user experience" is like before the network goes live.

Edge Device and Cross Technology Metrics

Networks are not homogenous. On a typical network, traffic may travel over LANs, WAN links, the Internet, and a high speed backbone. Integration of different technologies (10/100/1000 Ethernet, ATM, Frame Relay, Packet over Sonet) and the addition of newer mediums (like cable modems and the various members of the xDSL family) present a complex and formidable network analysis challenge. Since QoS is determined end-to-end, every technology along the path must work together.

Typical technologies seen in today's networks that require testing include LAN-to-WAN edge (Ethernet-to-ATM, -POS, or -Frame Relay) equipment, high-performance backbone (ATM or POS) equipment, and finally, broadband access equipment (xDSL and cable modems).

Edge device and cross technology metrics determine how network QoS is affected as the traffic crosses technology thresholds. These metrics determine how per-flow latency, latency variation, sequence tracking, throughput, and packet loss are affected as the packet is converted to meet the requirement of one technology to the next. Testing is especially important at the LAN-to-WAN transition where traffic traveling on large local pipes is squeezed into smaller wide-area pipes. On the LAN, where bandwidth is less expensive, over-provisioning can aid in delivering the QoS, but on the WAN side where bandwidth is considerably more expensive, over-provisioning is not feasible. Finally, this testing cannot be accomplished on just a lightly-loaded network; it must be performed at all load levels, stressing each component to its maximum. SmartMetrics and the high density offered by the SmartBits systems offer the capability to stress any network.

New broadband access technologies such as Cable Modem, and xDSL, offer very high bandwidth to the user but present particular challenges to network design.

Cable modem bandwidth, from a head-end (central office equipment) to the drop at the customer site, is shared among all the users connected to that head-end. As more users are added and traffic grows, the potential for a loss of acceptable QoS increases. Additionally, with numerous cable modem products now available, compatibility between them is a problem. Although standards exist, not all products are in conformance.

xDSL offers cost-effective high bandwidth over existing twisted-pair telephone lines using ATM as the data link layer between the modem (user) and the DSLAM (central office equipment). With xDSL, several factors affect performance: the amount of traffic being put through the DSLAM; the distance from the user to the central office; and the transition from ATM to the protocols used over the complete path.

SmartMetrics edge device metrics help you determine and maintain cable modem and xDSL installation QoS capabilities by putting the entire installation to the test. For cable modem and xDSL environments, this implies stressing the central office equipment with hundreds or thousands of streams flowing downstream towards the users, while hundreds of other streams traverse the network in the opposite direction, moving from the users to the backbone network.

Cross-technology and edge device metrics will help you:

- Determine the effect of the technology transition as the frame traverses the network (e.g., how the transition from IP frames to ATM cells and back again affects latency, latency variation, and throughput, or how the ATM network handles IP traffic.
- Determine if priority schemes are mapped properly between technologies.
- How well edge devices can handle the transition from frames to cells and the transition affect this has on QoS.
- Determine if all the products you plan to deploy in the network are interoperable.
- Help you decide if ATM is the answer or if you should use Packet over Sonet as your backbone technology of choice.
- Test the effect that adding new users to the cable modem or xDSL installation will have on its QoS and the overall performance of the network.
- Simulate predicted endpoint demand so that you can plan the network to support growth and still offer the QoS that the applications need.
- Determine your cable modem's compliance with applicable standards (e.g., DOCSIS).
- Stress test cable modem and xDSL segments as you deploy them, before connecting live users.

A Word About Ping

Latency (packet delay created by the network) is a very important measurement. Of particular importance, is the latency created when devices are under load. Many network managers believe that the magical tool to evaluate the network latency performance is an extensive use of IP pings. Unfortunately, this is a major mistake. All that a ping really does is send a single packet across the network and measure how long it takes for it to get there and back. The ping test provides round trip latency only, while providing no indication of the delay at the remote end. It provides no measurement of streams of data (i.e., no voice or multimedia information), no view of the network components and it definitely does not load the devices under test (see figure below). Lastly, since pings are usually given a low priority they are virtually useless for measuring QoS.

Ping (pong) – it's all wrong

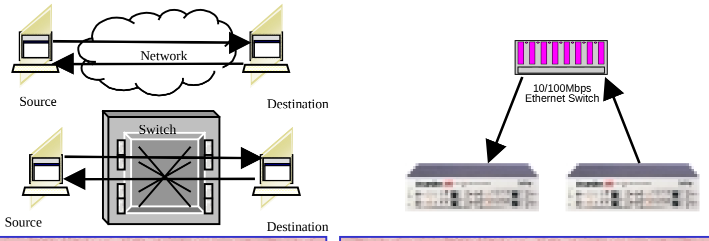

Ping test

- Round trip only
- No measurement of delay in remote end
- No measurement of streams of data (i.e., no voice or video information)
- No view of network components
- No loading of devices under test
- Low Priority

SmartMetrics™ testing

- One way latency, latency variation
- No end-station delay
- View of network components and individual traffic flows
- Latency testing of a device is only meaningful when device is under stress conditions

The only true metric to evaluate latency involves one-way latency and latency variation measurements, while the device is under stress conditions and when there is no end-station delay in the mix.

Who needs SmartMetrics

Aggressive performance testing is becoming a necessity for ensuring QoS. To be effective, QoS analysis must be performed at all stages of network development: during network design, as well as in network deployment and daily operation. SmartMetrics and SmartBits enable **NEMS** to accelerate the time to market while still delivering the best, most stable, and highest-performance products by:

- Allowing them to quickly and accurately evaluate key performance parameters under typical or extreme traffic load conditions.
- Shortening the development cycle, and helping produce reliable products. This is accomplished by providing ways to quickly isolate development problems, test new releases, test bug fixes, and stress the new product before having it undergo Quality Assurance testing.
- Testing sophisticated differentiating features such as QoS, latency minimization, IP services, bandwidth management, etc.
- Boosting product performance.

SmartBits and SmartMetrics help **NSPs**:

- Evaluate competitive products, making sure they purchase the right equipment to meet their needs.
- Predict exactly how a network component will work within different environments before deploying it in the network.
- Build test scenarios that can emulate the most demanding high-performance network environments.
- Test installations during network commissioning before actually offering service to live customers/users.

- Re-qualify equipment after firmware upgrades.
- Engage proactive testing during the network design and deployment stages, thereby saving millions of support dollars by avoiding service-related problems later on.
- Provision network services on predictable network capabilities by simulating and testing the most cost-effective configurations to match customer service-level agreements. This can be done prior to installing or revamping major network systems, then confirmed to ensure that the implementation fulfills the prediction.
- Offer consistent, reliable performance and stable service delivery.
- Optimize expensive bandwidth resources.
- Provide cost-effective solutions for voice data and multimedia.
- Give users proof that your network can handle their mission-critical data by verifying end-to-end QoS capabilities and providing deterministic SLAs for end-users.
- Provide intelligent insight on how effectively the network is handling real traffic, thereby enabling optimum network configuration so that the expected QoS is achieved.

SmartBits and SmartMetrics help **Enterprises:**

- Verify vendor quality and performance claims by testing new network devices at maximum load before equipment is deployed.
- Simulate the real-world requirements expected of the equipment.
- Predict exactly how a network component will work within different environments before deploying it in the network.
- Confirm interoperability with products from multiple equipment vendors and confirm compliance to prevailing standards.
- Evaluate higher performance technologies (ATM, 100Mbps, and Gigabit Ethernet).
- Test upgrades before they are deployed.
- Perform big-picture testing that looks at issues like how QoS works.
- Generate traffic flows of different types to recreate real-life situations to recreate and fix network problems.
- Determine how new applications will affect an existing network.
- Evaluate new technology to support new applications (e.g., Multicast IP and layer 4 products).

The Ideal Equipment for Network Performance Analysis

While tools that allow a network manager to monitor network status have been around for a while, the need for more in-depth analysis has established a new market for tools that perform network analysis before the network ever goes “live.” Proper network planning is 90% of winning the QoS battle. Proactive aggressive performance analysis of network components, both new and existing, is a key component of that planning. SmartMetrics QoS analysis represents the ability to create and generate the traffic of thousands of network-connected computers over a network or network device, and then capture and analyze the results to accurately measure its performance.

To properly access QoS, you need to understand what actually constitutes a decent response time across the network. There have been basic products on the market that provide simple information such as trouble reports. These information tools, however, do not provide information on scalability or give indications as to why the system failed. Likewise, these tools do not supply a method to pre-test the network, which if provided, would allow a manager to understand how systems scale if more end-stations were added, or how they would handle more incoming traffic.

The ideal performance analysis system offers:

- The capability of measuring a network's ability to provide differentiated services on a per- stream/flow basis. This capability is useful in measuring and configuring mixtures of multimedia, real-time, and background traffic (real-life scenarios) and to optimize the use of both local and wide-area devices and bandwidth.
- The ability to generate and analyze all classes of traffic and popular prioritization/ reservation methods including VLAN, IP TOS, DiffServ, MPLS voice, and multimedia traffic flows.
- Scalability to hundreds of ports and millions of flows so that it can generate and analyze traffic at full rate, no matter what technology is running in the network and no matter how many users it can support.
- The performance to generate session requests at a high enough rate to flood the highest performance device and to simulate the effects of new applications (e.g., Multicast IP, voice, and multimedia).
- Support important networking technologies and be able to test any combination of those technologies.
- No capture buffer limitations so that the test can run infinitely, not just for milliseconds. This is the only way to obtain a true view to a network’s QoS over time, under a variety of conditions.
- A design where the addition of a new card allows a network manager to upgrade or add new testing capabilities quickly and easily.
- Extremely accurate, deterministic, and repeatable tests that are remotely controllable.

And most importantly it needs to measure the new standard for QoS testing -- SmartMetrics.

You are very selective when purchasing network equipment and you need to be equally selective when you buy the equipment to test it. SmartBits systems are the industry’s only network performance analysis systems capable of generating and analyzing billions of flows of traffic, at full wire rate, for a wide range of technologies, including 10/100/1000Mbps (fiber and copper) Ethernet , Packet over Sonet, ATM, WAN T1/E1/V.35, Frame Relay, and Token Ring.

All SmartBits systems offer several Windows -based user control options from full-function graphical user interface software to fully-documented programming libraries. Netcom Systems applications for the SmartBits performance analysis systems are designed to address the specific testing needs of today’s most important network technologies.

With its SmartMetrics layer approach to network performance analysis, the SmartBits test scenarios can emulate the most demanding high-performance network to confirm that implementations fulfill QoS goals before you go live and while the network is up and running.

The Shift to SmartMetrics

Today's converged networks need more than traditional testing. If the network is going to offer high quality data, voice, and multimedia, it must be tested under a new class of metrics.

The SmartMetrics approach offers analysis of millions of data streams running in parallel. It allows testing of differentiated services for transporting traffic at various priority levels to help determine how different classes of service are actually treated. Such testing can help resolve questions regarding the handling of high priority traffic: “Should it always get 10% of the bandwidth or should it always get priority over everything else?” or “Can low-priority traffic be discarded?”

What constitutes good or bad QoS depends on the business needs of an individual network. However, proper testing is required to provide the data for purchasing and configuration decisions. Many other test areas also provide valuable insight into network performance under different scenarios. Routing switch testing, TCP testing, cross-technology evaluation, or checking out new technologies like packet over SONET, Multicast IP, and Voice/multimedia over IP are key areas for testing. Firewall and VPN testing are also important.

Once a network manager has a good grip on how the network is performing currently, it is possible to extrapolate and create more traffic following the same traffic pattern. A manager can then start to predict when a network might fail. Looking at individual streams of data will pinpoint which applications might fail. Stories abound about people going into a network operation center where half of the lights are flashing red. The network manager then calmly explains that this is a normal situation—and nothing is ever done until users complain. This kind of approach clearly states that the real issues of network performance are not known and that the configuration of networks that provide good QoS have not been possible.

Maintaining good network performance stems from: an analysis of the metrics related to traffic flow; connections and applications; the use of industry-accepted testing methodologies; network optimization of layers 2, 3 and 4; and an evaluation of the behavior of networks under stress.

Armed with this information, you can provide consistent, reliable performance and stable service delivery; optimize your expensive resources; provide cost-effective solutions; and give users proof that your network can handle their mission-critical data by verifying end-to-end QoS capabilities.