

---

# **Ensuring that the NewWAN Doesn't Lose its Luster**

## Application Note

---

A Hewlett-Packard  
Application Note by:  
Allan Campbell,  
Product Manager  
Dale Freelove,  
Regional Support Engineer  
Patrick Siu,  
Product Marketing Engineer

## **Introduction**

The current revolution of WAN<sup>1</sup>, LAN<sup>2</sup> and telecommunication technologies has led to leaner, more "fit" LAN/WAN hybrids spanning large geographical areas. This shift in focus from LAN to WAN has demanded more speed and reliability from the telecommunications sector. This trend further erodes the already blurred distinction between data communications and telecommunications. Because LAN, WAN, video and voice technologies are merging, the term data is no longer used only for computer communication payloads, but is now used for any digital payload. This is what is meant by the term NewWAN.

This application note provides an overview of the NewWAN technologies at play in this evolution, and shows how the design, implementation and deployment of NewWAN devices and services present fresh maintenance issues.

The virtual tester concept is one means to ensure that the NewWAN revolution does not lose its luster as the leading edge of WAN/LAN technology. The virtual tester represents a paradigm shift in testing philosophies by using stimulus-and-response testing as the foundation of Quality of Service measurements (QoS). Tester integration with network management and operational support systems is also discussed.

The real measure of the NewWAN will be how it performs and evolves towards the later half of this decade. Which technologies will still be in use and how will they be managed? Organizations can position themselves to benefit from the profit potential of the NewWAN by taking the right steps today.

## **An Overview of the NewWAN**

Early networks used noisy copper lines and analog-to-digital conversion equipment. The protocols used to transport data over this less-than-perfect media were designed to ensure data integrity and provide reliable delivery. For example, in a packet-switched network, the error detection and correction facilities inherent in each node consume execution time and message storage space. Each node on the path through the network must store the message and wait for acknowledgment before forwarding and deleting the message. This store-and-forward mechanism adds complexity to the intermediate devices and reduces the throughput of the WAN.


Fiber optic cable and digital lines are now used as the physical medium in transmission systems. Recent studies show that fiber and wireless connections are the two largest growth areas in communications mediums and installations because they tend to be much less error prone than their copper-wire predecessors.

When data is shuttled across the NewWAN, there is little need for the physical, data link, and network layers<sup>3</sup> to check and correct errors at each step along the way. Since this need has been removed from the lower three layers of the protocol stack, it has also been removed from the network itself. The retransmission of corrupted data is again handled by terminal devices, much like the Morse code operators of the past.

The foundation of the NewWAN revolution is the absence of need to establish the validity of the data in the transit path as the transmission medium and devices are deemed reliable. This brought about a subsequent refocusing on data throughput levels and greater bandwidth requirements for both user interfaces and backbone requirements. An example of this change is evident in the lightweight or "stripped down" protocol of Frame Relay, which does not perform some of the error checking included in older protocols..

Although existing networks will continue to flourish for many years, the movement towards stripped-down, faster WAN protocols has blurred the differences in speeds between LANs and WANs. Furthermore, the constant protocol data unit (PDU) rates of cell relay can carry more deterministic traffic such as voice and video. As a result, newWAN protocols will continue to merge separate voice and data networks into a single homogeneous network. This is truly the beginning of a leaner, more fit WAN.

Figure 1: A sample of the NewWAN



### **NewWAN Testing Challenges**

A whole new generation of testing challenges will evolve as the NewWAN grows. Testing has already shifted from the individual verification of the physical link, communications protocol and networked equipment to a methodology whereby one jointly tests the entire network as an operational whole for minimum throughput levels and guaranteed delivery rates. This requires a realm of new measurement techniques that can thoroughly exercise networks by placing calls, negotiating services, and generating traffic in the same fashion as an actual user. Clearly, traditional physical layer and protocol monitoring tests are not capable of doing this.

There are many Quality of Service<sup>4</sup> (QoS) standards by defined by standards bodies such as Bellcore, ANSI, the ITU-T<sup>5</sup> and others. ANSI X3.102, X3.141, and Bellcore T504 describe specific quality of service measurements which are suitable for existing networks and protocols. However, a new generation of QoS measurements requires definition because of changes in the mechanisms that deliver the services that network users consume.

### **Error Rate Testing**

The basic premise is that the medium will be error prone to some extent, and therefore we needed to ascertain the probability of erroneous data being received. A Bit Error Rate Test (BERT)<sup>6</sup> is used to determine how error prone the link is. Since signaling data is carried over the same channel as data, this one test quantifies the error rate for both data and signaling. In addition, a frequency response analysis of the link might be required to ascertain if carrier frequencies used by the modems are operational (frequency agile and rollback modems<sup>7</sup> were developed as a result of this type of service outage). Transmission Impairment Measurement Systems (TIMS) are oriented towards analog QoS measurements aimed at resolving problems below the bit error level. Note that a BERT requires a partially functioning layer one.

With protocols such as SS7<sup>8</sup> and ISDN<sup>9</sup>, data and signaling use separate channels or virtual connections. Both channels must be tested using bit error rate analysis. However, the call or connection must be up before bit patterns can be transferred. The BERT devices must now be capable of testing both data and signaling channels. This forces connection establishment with intact framing through call (switched services) or link setup (permanent services) being performed first, with the data channel then being assigned so we can use it for the BERT. In ISDN, the D channel call is placed, the B channel is assigned, and the terminal equipment then connects

## Ensuring that the NewWAN Doesn't Lose its Luster

to commence data communications. The BERT tester must have signaling capabilities and be able to switch to the assigned B channel before testing error rates.

We must also ascertain why any errors exist, not just if they are below an acceptable level. New issues can have negative effects on verifying the BERT sequence. Examples are:

- ◆ Exceeding the committed information rate
- ◆ Network congestion<sup>10</sup> in either direction
- ◆ Whether or not the stations' assigned priority is sufficient to get the next PDU available in order to transmit

These are all factors that can show up as errors in a BERT, unless the signaling aspects are taken into account. This is a prime example of a change in testing dictated by digital mediums, lightweight protocols and higher speeds.

### Error Detection and Correction

Protocols themselves must also be tested for implementation features and robustness. There are several type approval conformance test suites that test the CPE and terminal equipment, yet very few to test network behavior aspects. It may be argued that universal connectivity and interoperability issues are slowly forcing this to change. Thankfully, lightweight protocols have fewer parametric options and are therefore simpler to verify.

Protocols were historically oriented towards error detection and correction to ensure reliable delivery of the data in spite of errors being induced by the transmission medium. Many advances have been made reducing delivery and confirmation overhead. A customer will view several factors that represent the reliability of the call and subsequent success of the data transfer. They include, but are not limited to:

- ◆ Success rate of a call being placed
- ◆ Call termination for various reasons
- ◆ Rejection of frames and packets due to bit and framing errors (analogous to noise)
- ◆ Frequency of retransmissions (affects throughput)
- ◆ Call blockage due to congestion or a busy destination (can't use the system)
- ◆ Frame or packet rejection due to inappropriate protocol behavior (affects connection reliability and throughput performance)

There are two methods to determine the error detection and correction performance of the network as a whole. One is to monitor the link and make specific measurements pertaining to the ability of the device that uses the link. The second is to generate traffic and collect statistical information pertaining to the same set of measurements.

Reliability of the connection has been quantified as the number of frame/packet rejects incurred, the number of calls failed to connect.

The use of performance monitoring equipment and protocol analyzers to make these simple measurements will continue, but additional requirements, such as the ability to place a call, will be required when SVC facilities become widely available.

It is also desirable to measure access to a free cell from various priority levels and under variable load conditions, including full load and overload. The ability to get free cells for transmission can adversely affect throughput rate and end-to-end connection establishment. Test technicians need to be able to load the device or network with simulated traffic, and measure how long it takes to get a single cell and the duration of a transfer under various loads. This requires a multiport testing device with complete emulation capabilities, appropriate protocol behavior, timing synchronization between test ports, timestamp capability, and elapsed time measurements.

We must also be able to simulate congestion in a PVC environment to test the ability to establish a PVC from end-to-end, measure recovery times when frames are dropped, and measure how long it takes to transfer a known amount of data under various overload conditions such as bursts, sustained and periodic congestion levels.

## Ensuring that the NewWAN Doesn't Lose its Luster

Lightweight protocols will discard data if the maximum data rates is exceeded. If this occurs, we need to ascertain where the frame has been lost. A stimulus-and-response system can generate loads and simulate traffic patterns to benchmark the network. The resulting benchmarks can then be compared to *in situ* performance.

The following new QoS measurements should be also measured:

- ◆ Cell / frame loss measurements
- ◆ Address / routing table accuracy
- ◆ Congestion analysis
- ◆ Congestion location identification
- ◆ Free cell access rates
- ◆ Network to network interconnection

These tests require the generation of traffic with intelligent payload information imitating the reliability mechanisms so that retransmission can be forced and measured. The means of testing<sup>11</sup> cannot be simple traffic generators that do not adhere to protocol rules when they have been started<sup>12</sup>, or that cannot alter individual PDU payload contents. Information is injected into the network by the MOT at one end while statistics describing frame delivery are collected at the other end. If a PDU has not been delivered, a timer can be launched to measure the elapsed time between the notification of the

event and the redelivery of the dropped frame. This technique yields a reliability factor describing how often frames are dropped, and recovery timing parameters.

### Data Rates and Throughput Performance

Stated data rates for X.25 and Frame Relay, SMDS and ATM have different meanings. An X.25 service over a 9.6 Kb/s line will never be able to achieve a data delivery rate of 9.6 Kb/s since the associated overhead prevents this from occurring. However, Frame Relay has a guaranteed minimum data rate. Users wish to be able to determine if the leased service is indeed delivering the minimum rate that they require, and service providers wish to determine how much bandwidth a user requires.

There are many parameters that could be measured to depict the performance of a particular UNP<sup>9</sup>. Examples are end-to-end transit time and each virtual circuit's (VC) throughput as a percentage of the throughput class. The overall throughput of the link could also be measured by examining the sustained, maximum and minimum data rates versus the link speed. Every call should be examined to determine the data throughput rates (fps, pps, bytes per second), error rates (reject frames, frame/packet rejects), why the call was cleared, and who initiated it.

Stress testing the devices that will be deployed on the network ascertains performance aspects that will have a direct bearing on the network performance. Such tests includes:

- ◆ Device latency
- ◆ Frame / packet forwarding rate
- ◆ Frame / packet filtering rate
- ◆ Maximum number of simultaneous virtual connections
- ◆ Optimum routing table / address buffer size

In Frame Relay, random traffic and constant background load conditions (simulate LAN interconnect traffic) must be generated from end-to-end to measure the varying levels of throughput, versus the committed information rate (CIR). The CIR should be the minimum data rate, and is often the average data rate.

In SMDS or ATM, the ability to segment and re-assemble larger buffers using multiple short cells should be determined as well as the cell delivery timing. The size of the buffer should be varied starting from less than 48 octets in odd multiples to examine the effects of filling out partial cells<sup>14</sup>. Then, the optimum even multiple of 48 octets can be found so that maximum and minimum reference points can be located.

The MOT should have elapsed time measuring capabilities to measure the time between transmission of the frame or cell at one end and the validated receipt of the same frame or cell at the other end. This also means the MOT must be able to be able to

correctly simulate the signaling mechanism while transmitting load data, as well as interrogating the received frames or cells to determine if indeed the transmitted frames or cells were received. Since protocols are not designed to do this, the data fields in transmitters are used to carry a frame/cell identification mechanism.

The nature of traffic has changed from deterministic traffic patterns for voice, terminal traffic and file transfers to bursty non-deterministic traffic characteristic of LAN interconnect traffic. It is thus very important to establish a network performance baseline when commissioning services, and to monitor for deviations and remaining bandwidth. This permits a more proactive approach to handling impending overload conditions.

The effect of an increased load on the performance baselines (generate traffic and collect reliability/performance statistics) can be measured using these concepts. This could be used to rate the efficiency of the network services. Congestion levels can be foreseen, and steps taken to resegment and provide additional capacity if necessary.

### Frame Loss Measurements

With lightweight protocols, the onus has been placed on the CPE to perform the retransmission of lost frames. The user pays for these retransmissions under some new tariffing schemes<sup>16</sup>. The issue of who is dropping the frame then has economic implications. In such a situation, the user may have exceeded the bandwidth of the service, and would be better off to upgrade the service than pay for repeated retransmission.

This can be quantified through the use of load generators in conjunction with statistical analysis tools and a PDU recognition scheme. A background traffic level is set up with random traffic bursts to full load, and unique frames with a sequencing and identification mechanism are transmitted from end-to-end at each VC and line saturation levels. Each and every frame that is received is compared with that sent, and the previous PDU received, to discover if any frames have been lost. If so, further testing can isolate the problem. Accurate timestamping for each frame or cell can determine the travel time from transmitter to receiver and between test access points.

There are two methods of accomplishing this test. One could use several test access points, at strategic locations throughout the network path and an MOT with several distributed simultaneous test access ports operated from a

central location (distributed MOT). This would allow all access point to be monitored simultaneously, while test ports connected to each end inject the traffic and background load. A comparison for each frame in each leg of the network would immediately spot missing frames, and the MOT could trigger an alarm to signal the network operator at the central location that such an event has transpired, thereby isolating the cause. Test traffic could be injected at scheduled intervals to check for degradation in service, and if the MOT was coupled to a network management system or operational support system, alarms could be monitored for just such an event.

The alternative is to connect a dispatched portable MOT at isolated legs of the network. Each leg is tested separately by connecting the MOT, testing, disconnecting the MOT, and travelling to another access point. At best, this is time-consuming, inconvenient, and costly; and may simply be impractical for any but the smallest networks.

### **Configuration Errors**

Configuration errors can cause incompatibilities between services hired and the equipment on the customer premises. This is partly due to incompatibility in the messaging standards when new protocols are developed and deployed. Therefore, correct protocol behavior must be verified before products are deployed.

The compatibility of a device can only be proven by rigorously exercising its protocol implementation. This reduces service interruptions. Product evaluation testing is often part of the purchasing criteria for devices that do not conform by design.

The European NET standards are perfect examples of the way "permission-to-connect" tests have been used for this purpose (NET2 for X.25 PSN, NET3 for ISDN, etc.). These tests are used by PTTs to determine if a device may be connected to the network. This forces the manufacturer to pass a type approval test before they can sell and install devices in that particular country or network. Such tests assume that if the device passes the test, then it should not have a negative impact on the network.

However, passing the permission-to-connect test does not ensure that the device will successfully interoperate with other nodes and devices.

Manufacturers and service providers use regression testing to assure backwards compatibility with new features

and services, but this does not test the installation of new devices, nor does it verify configuration parameter errors that show up under heavy loading or growth conditions.

The ideal MOT would emulate several devices and equipment types, and connect to the network by placing a call and bringing up a link. All aspects of the protocol implementation and data transfers would be exercised to test compatibility, including traffic generation to simulate background traffic and various loading levels. However, developing an all-encompassing emulation device for every aspect of every device is an insurmountable task. The next best would be an MOT that emulates correct protocol behavior; has traffic/load generation capabilities as if it were another device such as a router or switch; has many ports such that it may appear as many user devices at the same time; and can emulate multiple protocols and load generators operating simultaneously.

Self-configuring equipment attempts to reduce the complexity of the installation process. This means novices may begin installing the equipment on the user side of the network, much the same as is current practice with LANs. Thorough testing will be required to ensure that default parameters will not cause network problems later when loads increase or new users are added. Traffic throughput, forwarding characteristics,

number of concurrent users, and protocol behavior must be thoroughly accessed.

Routing tables may exist at various points throughout the network, on customer premises and in the network itself. Routing protocols such as IGRP, RIP and OSPF may be used and the CPE equipment may connect through similar mechanisms to the network. This can result in several layers of tables and addressing mechanisms, increasing the odds of problems occurring.

One can test routing validity by injecting traffic destined for a particular location and monitoring the destination, as well as other access points, using a distributed MOT. If the traffic does not reach the destination, the distributed MOT can determine where it was erroneously routed. Periodic verification may be necessary after any network expansion and re-segmentation, and regression testing should be performed to verify the correct operation of hardware and software updates. The MOT must be able to specify various address schemes for several connections at the same time. The load generator must be able to use this mechanism and ensure that data is transmitted on all the circuits and connections being tested.

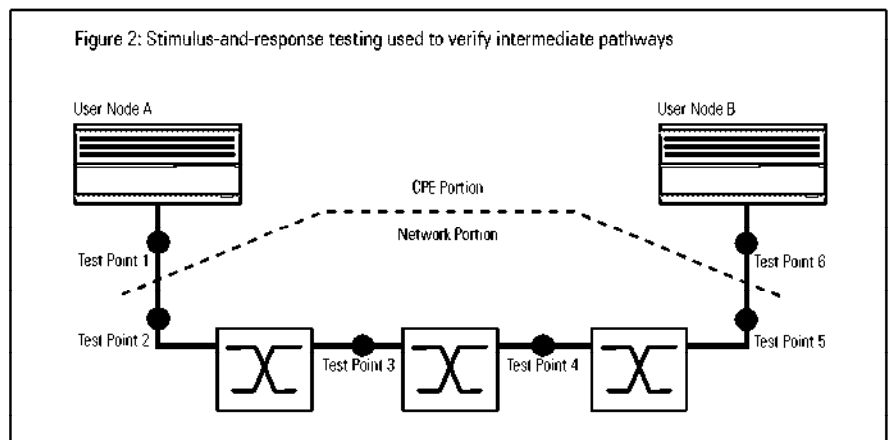
A similar function can be performed using a dispatched MOT, but the scope and scale of the test campaign is typically limited to a single device or a small geographical segment of the network. This approach is successful only with devices of very limited port capacity.

## Stimulus-and-Response Testing

Historically, data decodes have been used to troubleshoot new implementations and upgrades as well as to test working systems. Statistical analysis was used to indicate infrastructure health and helped to reduce a large amount of information to easy-to-understand figures or graphs. These tools were very effective for use in reacting to problems after they have occurred, which placed the emphasis on response and repair times. Reliability was measured in terms of mean-time-between-failures, and customer service in mean-time-to-repair. But are these tools enough in the face of new network architectures?

Stimulus-and-response testing has historically been used for type approval and device simulation in R&D labs. Yet this method of testing lends itself readily to improving the MOT so that pro-active maintenance techniques can be used rather than the traditional reactive ones.

One of the best analysis techniques is recreation — stimulating the device or network in an attempt to reproduce the error. If the captured data can be used to replicate the sequence of events leading up to the failure, the amount of time required to isolate intermittent errors can be reduced. Stimulus and response testing can be used to generate abnormal traffic on demand, which is essential in problem recreation.



While stimulus-and-response testing has been used very successfully to test protocols, it can also be used to verify networks. This is one of the most important differences in the NewWAN testing paradigm. It has immediate benefits in terms of reproducing errors and setting known baselines to which the network can be compared at some point in the future. It establishes a pro-active approach towards problem resolution through active participation, rather than patiently awaiting the next failure.

The pro-active scientific approach is to use the protocol capabilities of the MOT not only to decode and gather statistics, but to generate data of known quantity and quality. For example, by generating calls from one test access point to another throughout an X.25 based network, correct operating performance parameters can be ascertained at the same time. Communications between any two nodes can be readily verified through traffic simulation, from

the network control centre, rather than on a dispatched basis.

The MOT user could assign test capability to a number of points between node A and node B (see diagram). Then the test operator places a call (for example, SVC) between node A and test access point 1. If this works, the test operator progresses by placing a call between test access points 1 and 2. If this is operational and can sustain traffic, then he/she moves to points 2 and 3, and so on, until node B is reached. This approach alternates between placing calls and transferring bi-directional traffic with background load levels.

Each network segment and major device, including the routing tables, can be verified in this manner. This requires a distributed MOT, but response time is drastically reduced and there is no travel between sites. If this is done on a regular basis and contrasted to a previously established baseline, degradations in service could be predicted so that spare facilities are switched in before the failure occurs, eliminating downtime and crisis repairs.



## Ensuring that the NewWAN Doesn't Lose its Luster

In the NewWAN protocol environment, load generators and traffic simulators are used intensively rather than just simple call placing and data transfer. However, the SVC world is not far off for the new lightweight protocols, and so we may also one day use the same call placing principles coupled with the smart load generators and statistical analysis capabilities to verify frame arrival and search for missing or dropped frames.

### Device Testing

Switches and other devices used to build the NewWAN still require testing, but with a different viewpoint. Whereas switches were tested for routing and switching capacity, with specific measurements of store and forward and/or buffering capacity, now the switching fabric itself must be tested.

The reason for this is that cell relay switches are hardware-based, providing far superior performance, and yet do not provide the sort of buffering needed to ensure that a particular frame reached its destination. Indeed, switches may throw away data if forced to. An MOT with more than two test ports is required to see if the maximum traffic level can be sustained with correct routing, and determine which, if any, ports are dropping frames.

Traffic must be injected with various address patterns, and then be examined as it arrives at the destination ports. Statistical

analysis plays an important role as the reliability of delivery is measured.

The following is a partial list of performance measurements:

- ◆ Frame and packet forwarding rates; max, min, invalid and random frame sizes
- ◆ Cell forwarding rate (cps); valid and invalid cell sizes
- ◆ Cell / frame / packet arrival timing interval versus cell / frame / packet loss
- ◆ Probability of correct routing / addressing
- ◆ Routing table / address buffer optimum number of entries / size
- ◆ Routing table / address buffer flush and refill duration
- ◆ Routing table / address buffer error rate
- ◆ Number of concurrent connections

### Testing Protocols Versus Testing Services

So far, we have compared methodologies of testing the NewWAN versus older protocols. Generally speaking, we used to test the protocols to ensure that an unreliable medium can provide some degree of reliable data delivery, and now we test the service as a whole.

The fundamental shift is from testing and exercising the error detection and correction mechanisms to testing the performance minimum and measure future performance against a baseline. Out of sequence packets or frames; the myriad of parameters used for

connection establishment; the duration of call setups and call clears; busy conditions and link disconnects; inopportune frame and packet handling capabilities have all been augmented with a need to establish where the frames are dropped and active path assignment verification. This change has taken us from the passive monitoring and surveillance system days to active testing and traffic generation. Furthermore there is a real need to couple the MOT to a network management or operational support system for the centralized control that has so changed LAN management in the last few years.

### Network Management

A rapid development paralleling the NewWAN evolution is the Simple Network Management Protocol (SNMP). Lead by developments from the Internet Engineering Task Force (IETF) under the auspices of the Internet Advisory Board (IAB) and its members and contributors, SNMP has propelled the centralized management concept to the forefront. Several North American service providers have announced their intention to install, are installing, or have installed SNMP-based services

A desire to get the plug-and-play LAN environment under control has fuelled the development of network management. This is very beneficial to the NewWAN as it encompasses both old WAN and LAN technologies while embracing new advances. The line between datacommunications and telecommunications has blurred, as has the line between WANs and

## Ensuring that the NewWAN Doesn't Lose its Luster

LANs. Local ATM and fibre backbones are not limited to the local area any longer. We still await the lowering of the cost envelope and ready availability of service offerings, but their day will come soon.

Meanwhile, we are left with a quandary. As the SNMP world was based in a LAN with very different measurement parameters and characteristics, the NewWAN is a contrast to those metrics. Ring purges and ethernet collisions are of no importance to the WAN protocols and services. Congestion and lost frames, empty cells versus used cells and other required measures of service performance are not to be found in the many Management Information Bases (MIB) associated with SNMP or any other network management protocol, although there are ANSI standards for measuring performance in X.25 based networks! The RMON MIB is not useful for the all-encompassing NewWAN management tasks. The pressure to provide WAN management MIBs is becoming great, as the network management arena expands and encompasses all aspects of the NewWAN. Plans have been announced already to make an SNA/SDLC MIB and SNA/IBM Token Ring MIB and an X.25 QLLC MIB.

What is required is a comprehensive WAN MIB with various members such as X.25 and Frame Relay, Cell Relay and ATM so that the proliferation of proprietary WAN management

schemes can be replaced with a viable standard for the entire NewWAN, not just the LAN segments. This also holds true for the wireless explosion.

### The Virtual Test System

There are many examples of a centralized control architecture in use today. The process control industry has had to react to the same pressures that network suppliers and operators face: the need to right-size, reduce the time to resolve exception events, locate test equipment with the device under test instead of the expert, reduce operating expenses, and increase productivity.

The same holds true for NewWAN management. The virtual tester concept is an evolution of the distributed test architecture. The main difference is the incorporation of the next generation of testers to accomplish NewWAN management. Some of the required features include:

- ◆ Unique methods of testing reliability with BERT in the data channel and the use of frame dropping measurements.
- ◆ The concept of smart load generators used with statistical analysis
- ◆ The capability to transmit while receiving and monitoring
- ◆ The ability to saturate many lines with traffic destined for different addresses to verify routing tables

- ◆ The development of performance baselines during commissioning and the use of those baselines over time to weigh performance trends
- ◆ The use of call generation and traffic simulation to verify the end-to-end connection step by step
- ◆ The reduction of the time taken to respond to any given situation by remotely accessing multiple test points
- ◆ The centralized control of the network, including management tasks and coupling to network management or OSS systems
- ◆ The scalable system to allow the user to determine which test access points need be tested simultaneously, and thereby build the number of concurrent test ports by adding more test capacity and devices throughout the virtual tester, which in reality is a network of test components. It should be operated from a central location where the experts can be located. This has the added benefit of greater synergy between experts, as they learn from and help one another. The test access devices could be located in wiring closets, central offices and switching areas so that the power of the tester can be deployed where it is required.

The central site should be able to visualize each test access port as if it was a single port in a multipoint dispatched tester. There is no requirement for any test capability at the central site, so a computer workstation could be the pivotal device in the system. The friendly graphical user interface and power of these devices, plus the proliferation of network

## Ensuring that the NewWAN Doesn't Lose its Luster

management applications coupled with their multiple communications interface capability make these devices far superior to the terminals or other test devices that have been used in the older distributed testing systems.

There are barriers that may have to be overcome in order to implement the virtual tester.

One is the concept of transmitting maintenance traffic on the network, frowned upon in the past since it consumes bandwidth without generating revenue. This is why the determination of performance baselines at commissioning time is so important, and why the ability to distinguish maintenance traffic (payload decoding and recognition) from real traffic is important.

The ability to co-ordinate activity between multiple testers and have them behave as a virtual tester is paramount. Mere remote access to testers through a windowing mechanism does not suffice since users need more than remote keyboard and display capabilities. Timing measurements, triggering and messaging to the host computer, and broadcast command mechanisms that mold the entire network of test equipment into a co-ordinated unit is what the virtual tester is all about. Such features would allow the distributed tester to appear to the user to be a multipoint dispatched tester.

File transfers from the host to the remote testers becomes important for distribution of test scripts and operating system updates. A mechanism to remotely reboot and power cycle the testers is also necessary since if travel to the remote site to update testers, reboot them, or insert diskettes is to be avoided.

### Summary

The NewWAN has truly started another revolution in communications. The term data now no longer exclusively refers to computer information, but encompasses all forms of digital information from voice to video. The future holds the promise of faster pathways and omnipresent information highways.

The way we think about supporting networks and services must change to keep pace with this explosion. The futures network operators and service providers will survive today's competitive pressures by increasing service reliability, expanding points of presence, and offering innovative services with market-driven applications.

An integral part of this new, more reliable network will be the virtual tester. It will help the operations and maintenance staff realize far greater cost-effectiveness while assuring reliability and reducing downtime. Along with the growth and development of

the virtual tester will be a more proactive method of accurately measuring QoS parameters. Type approval and permission-to-connect testing will become common-place methods of regression testing, as well as a means of assuring the quality of new device, products and services.

Network operators and service providers can ensure that tomorrow's NewWAN does not lose its luster by deploying virtual tester-ready tools today.

## Ensuring that the NewWAN Doesn't Lose its Luster

---

### Footnotes:

- 1 Wide Area Networks.
- 2 Local Area Networks.
- 3 Layers 1, 2 & 3 respectively in the seven layer ISO OSI reference stack, also called levels 1, 2 & 3.
- 4 Quality of Service, or QoS, is a measure by which users of a given service may rate its performance attributes, therefore it is very customer oriented in terms of the services delivered.
- 5 The International Telecommunications Union, or ITU, is the CCITT's new name.
- 6 Bit Error Rate Tests are referred to as BERT, and one example is the ITU-T's G.781 standard.
- 7 The common POTS phone link has about 3000 Hz of bandwidth, so if the carrier frequency signal quality is insufficient to maintain a link, the modem changes carrier frequency by itself, or reduces the transmission speed in reverse steps from 19.2 Kb/s to 9.6 Kb/s to 4.8 Kb/s etc.
- 8 Common Channel Signaling System Number 7.
- 9 In ISDN, notice the network reference rather than the protocol name reference as with X.25. For example, Q.921 and Q.931 are not referred to as services.
- 10 Congestion is the term used to describe an overload condition from excessive traffic levels. It makes me think of severe head cold when one's mental capacity begins to shut down.
- 11 MOT is often used to refer to a formal type approval platform or collection of equipment and stringent testing steps. The term is used here to describe a test system in a much broader sense.
- 12 Dumb traffic generation ignores the signaling process and attempts to flood one or more channels or virtual circuits with predefined frames and predefined frame sizes.
- 13 UNI stands for the User Network Interface, the point of connection between user and network devices.
- 14 Each cell has 5 octets of overhead with a payload size of 48 octets; if our data size is only 31 octets, the remaining 17 octets are filled with a known bit pattern referred to as filler. Every cell must go empty (full of filler) or full (full of data, or a combination of filler and data).
- 15 In those service areas that have usage fees associated with the particular service, the user pays for retransmissions if the retransmissions are requested by the destination and sent by the source. In contrast, retransmission in X.25 networks often come from the network's store and forward buffers, costing the user nothing more than the initial data unit.