SONET/SDH VS. ETHERNET: MIGRATION AND TESTING ISSUES

152

APPLICATION NOTE

André Leroux, Systems Engineer
Bruno Giguère, Member of Technical Staff

In the past few years, the enthusiasm for Ethernet-based services has grown considerably. The advent of 10 Gigabit Ethernet (10GigE)—with its increased bandwidth and improved fault-detection schemes—has also fueled their popularity and further increased the trend towards the deployment of Ethernet technology on core transport networks.

Enterprise customers with an Ethernet infrastructure are upgrading their core transmission links to a 10GigE local-area network (LAN) to maximize the use of their existing infrastructure without losing functionality. Service providers are also migrating from a SONET/SDH-based infrastructure towards 10GigE wide-area network (WAN) technology, enabling them to offer real-time applications such as voice- and video-over-IP without having to upgrade their entire network.

However, 10GigE on its own does not guarantee the same level of performance as would be expected from a SONET/SDH network. Network operators must therefore test a number of network parameters to ensure that a deployed 10GigE system delivers an acceptable quality of service, and to show customers that the network complies with the performance metrics specified in their service-level agreements (SLAs).

When creating the 10GigE Standard, the IEEE was well aware that it was upgrading Ethernet to meet core-network capacity. Therefore, it needed to interoperate seamlessly with other core-network technology, the most prevalent being SONET/SDH. On the other hand, 10GigE also had to match the natural scalability of Ethernet, where 10GigE could transport 10 x 1GigE or 100 x Fast Ethernet.

Unfortunately, these two requirements cannot be easily supported with one technology since the effective bandwidth of an OC-192 SONET channel is only 9.58464 Gb/s. As a result, it is incompatible with a full 10 Gb/s bandwidth. This realization led to the development of two 10GigE varieties. One form—named 10GigE LAN—interoperates with previous Ethernet technology and offer a full 10 Gb/s bandwidth. The other form—10GigE WAN—interoperates with SONET networks by encapsulating Ethernet traffic inside an STS-192c channel and support an effective bandwidth of 9.58464 Gb/s. Although 10GigE WAN uses an STS-192c framing structure, to lower the costs associated with SONET, many constraints such as clock tolerance, protection switching, jitter and more, are not supported.

10GigE in the core

It is important to understand the key differences between the capabilities of SONET/SDH and Ethernet (see Table 1). These stem from the fact that SONET/SDH was originally created to transmit circuit-based voice traffic, while Ethernet was designed to carry packet-based data. Although SONET/SDH has since been modified to support data traffic more efficiently, it still lacks the flexibility inherent to the datacom world.

In contrast, the latest version of Ethernet offers enough bandwidth to meet the requirements of today's core networks but, in its most basic form, 10GigE lacks the high-level performance provided by SONET/SDH. Ethernet on its own provides no protection mechanisms, no end-to-end service monitoring, and no default quality-of-service (QoS), all of which are standard in SONET/SDH networks.

While extra features can be added to Ethernet networks to improve their performance, each one must be implemented and maintained by the network operator. As a result, the choice between SONET/SDH and 10GigE largely depends on the type of traffic to be carried by the network. For example, a pure 10GigE network might reduce the equipment cost, but could incur higher operational costs because of the need to support extra protocols that guarantee the required performance.

The necessity for additional protocols also makes it essential to perform network tests during installation and commissioning to ensure that a deployed Ethernet network provides the required level of service. In particular, these tests should ensure that the network is able to recuperate from network problems such as a protection switch, an inactive/malfunctioning link or inefficient traffic management.

Table 1: Comparison between the main capabilities of SONET/SDH and Ethernet

Criteria	SONET/SDH	Ethernet
Redundancy protection	– Automatic protection switching capability (50 ms) – Link capacity adjustment scheme (LCAS) for virtual concatenation	– Fast spanning tree (from 10 ms to 1 s, depending on network topology) – Link aggregation – Resilient packet ring (< 50 ms) – MPLS fast reroute (< 50 ms)
Operations, administration, maintenance (OAM)	– SONET/SDH OAM framework	– Point-to-point links covered by IEEE 802.3ah and ITU Y.17ethoam standards, while end-to-end services covered by IEEE 802.1ag – Metro Ethernet Forum service OAM
Fault detection	– Sectionalized error/defect monitoring and remote indications – Fault surveillance and threshold crossings in SONET, and performance monitoring in SDH	– Remote-link failure detection (10GigE only) – Remote monitoring with proprietary thresholds
Maintenance	– Loopback capability for out-of-service tests	– No loopback in Ethernet – Switch and router information can be obtained through remote monitoring or through IEEE 802.3ah services (point-to-point only)
Traffic engineering	– Virtual concatenation (VC)	– Nested VLAN (VLAN over VLAN) – MPLS label-switched paths – Pseudowire Emulation (PWE)
Scalability	– Transport rate up to 40 Gb/s – Granularity to VC level (1.5 Mb/s or 2 Mb/s)	– Transport rate up to 10 Gb/s – Granularity to any rate
QoS	– Deterministic	– Proprietary quality of service – Difficult inter-vendor support
Robustness	– 99.999% up time – Bit-error rate (BER) = 10^{-12}	– Up time based on redundancy/protection services implemented by the owner of the network – BER = 10^{-12}

Migration issues

Network operators must consider a number of migration issues before upgrading a fiber network to 10GigE. At the physical level, for example, it is crucial to understand the effects of chromatic dispersion and polarization mode dispersion, since these two parameters hold the key to maintaining optimum transmission performance. For more information on this particular topic, refer to our article entitled Dispersion and 10 Gigabit Ethernet Transmission.

At the interface

Besides these physical and optical issues, traffic engineering must also be carefully considered when migrating to 10GigE systems. Specific problems can arise at the interface between a LAN and WAN, and two important factors to consider are the difference in transmission frequency tolerance and the difference in bandwidth.

SONET/SDH networks are based on a master clock that is distributed to all network nodes, while Ethernet networks don't have a distributed clock and have a much lower clock tolerance. When these two worlds are connected, the clock difference can be as much as ±120 ppm in a LAN/WAN configuration or even up to ±200 ppm in LAN/LAN systems, causing Ethernet frames to be lost when traffic is sent at maximum throughput.

Additionally, a 10GigE WAN can only transmit an effective data rate of 9.58464 Gb/s. If the LAN network throughput is close to 10 Gb/s, Ethernet frames could be dropped at a rate of up to 415.36 Mb/s. Although a network operator could limit the ingress rate to 9.58464 Gb/s to address this problem, the difference in clock tolerance could induce an additional rate mismatch of 1.2 Mb/s. An RFC 2544 throughput test should therefore be used to determine the maximum throughput that a 10GigE WAN link can sustain.

Another important traffic engineering issue is the use of "pause packet", a flow-control protocol defined by IEEE 802.3x to backhaul the ingress traffic on point-to-point links when the internal buffers are close to saturation. When a threshold in the buffer set by the network operator is crossed, this scheme sends a pause packet to the downstream network element to stop the transmission of traffic until the local buffers are empty.

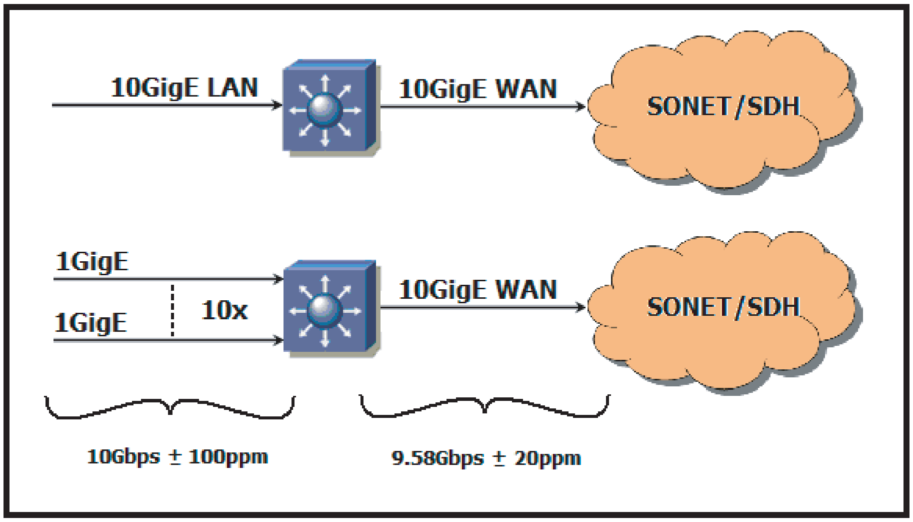

***Figure 1**. Connecting a LAN to a 10GigE WAN can generate significant clock differences (up to ±120 ppm), as well as a change in bandwidth. Both of these effects can lead to the loss of Ethernet frames.*

This pause threshold must be set at the right level to maximize throughput without losing any frames, and the effect of the pause frame is largely determined by the time taken for it to be transmitted and received at the downstream end. On fiber, one bit takes about 53 μs to travel a 10 km span. During this time, a maximum of 79 frames can be sent on a 1GigE link, while a 10GigE link is capable of transmitting about 788 frames. By the time a pause frame is sent and the traffic stopped, a 10GigE link will receive ten times more frames then a 1GigE link.

If the pause threshold is set too high, the reaction time will be too slow and the local buffers will overflow, causing frames to be lost. If too low, the pause frames will be sent too often, reducing the overall throughput and increasing packet jitter. Network operators should therefore exploit an RFC-2544 frame-loss test to determine the best threshold value.

The importance of testing

Testing networks during installation and commissioning provides network operators with a valuable record of network performance. These tests also form the basis of SLAs defined by telecom operators and their customers to stipulate the expected quality of service. The parameters used to define SLAs are typically network and application uptime/downtime and availability; mean time to repair; performance availability; transmission delay; link burstability; and service integrity.

Test equipment is essential for service providers to quantify most of these parameters. Nonetheless, certain tests are more important than others, depending on the type of network being tested.

In a core network, for example, the bandwidth is usually delivered on a transparent physical network, and the test methodology is quite different from the one required in a switched network where there is "processing" of Ethernet frames (see Figure 2).

For instance, next-generation xWDM networks carry 10GigE in a transparent way, but some service providers also use 10GigE to offer point-to-point wavelength services. Since these new applications are totally transparent, providers need a quick way to certify these services and to define parameters that can be used for SLAs.

This can be achieved with a BER-type test that measures a pseudo-random bit sequence inside an Ethernet frame. Such a test allows telecom professionals to compare results as if the tested circuit were SONET/SDH, not 10GigE. Most telecom engineers are more comfortable with terms like BER and Mb/s, as opposed to packet loss rate and packets/s, and this methodology provides the accuracy in bit-per-bit error count that is required for the acceptance testing of physical transport systems.

If the network to be tested is switched-based — and so includes overhead processing and error verification — the BER approach is not the best one. This is because a network processing element will discard frames or packets if an error is found, which means that most errors will never reach the test equipment. These lost frames are more difficult to translate into a BER value.

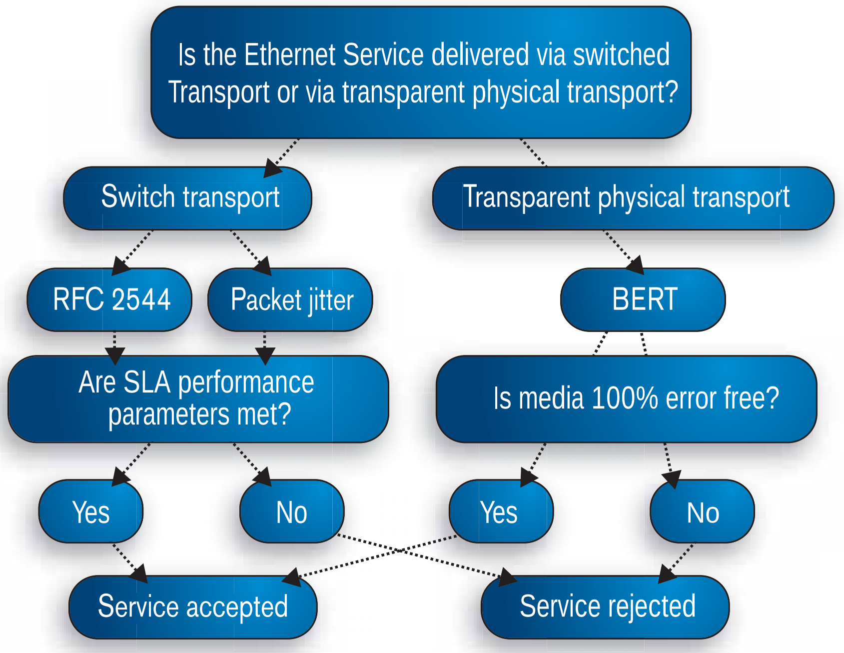

Figure 2. *Ethernet services require different test methodologies, depending on the transmission medium.*

The other way to test a network during installation or when commissioning a customer circuit is to use RFC 2544 (benchmarking methodology for network interconnect devices). RFC 2544 test procedures incorporate four major tests — covering throughput, back-to-back (also known as burstability), frame loss and latency — to evaluate the performance criteria defined in SLAs. The results are presented in a way that provides the user with directly comparable data on device and network performance, and also allows a service provider to demonstrate that the delivered network complies with the SLAs defined in the contract.

First, the throughput test determines the maximum rate at which the device or network can operate without dropping any of the transmitted frames, providing a clear indication of the capability of a network to forward traffic under heavy load. The test also puts the network elements under stress when processing the different checks required at different protocol levels. Changing the frame size also puts extra stress on the network processing capability.

Second, the back-to-back value is the number of frames in the longest burst that the network or device can handle, without losing any frames. This measurement is used to evaluate the network's capability to handle short to medium bursts of traffic from an idle state.

Third, the frame-loss measurement is the percentage of frames that is not forwarded by a network device under steady-state (constant) loads due to lack of resources. This measurement can be used to evaluate the performance of a network device in an overloaded state, which is a useful indication of how a device or network would perform under pathological network conditions such as broadcast storms.

Finally, the latency test measures the time taken by a frame to cross a network or device — either in one direction or round-trip, depending on the availability of clock synchronization. Any variability in latency can be a problem for protocols like voice-over-IP, where a variable or long latency can degrade the voice quality.

First and Foremost: Quality

As service providers introduce new real-time Ethernet-based services such as video and voice-over-IP (VoIP), they need quantitative measures of network behavior that can affect the service quality. Latency and frame loss can have a significant impact on these real-time applications, as can significant variations in the packet delay over time, also known as packet jitter (Table 2).

Table 2: Factors affecting video and voice quality on an Ethernet network.

Network problem	Impact on video services	Impact on voice services
Packet jitter (packet delay variation)	– Dropouts, loss of clarity, momentary, blanks, jumpiness	– Unclear speech, pops-clicks
Frame loss	– Graininess, dropouts, loss of clarity, momentary blanks, and in worse cases loss of session	– Garbled voice, pops-clicks, and in worst cases, loss of communication
Latency	– Momentary blanks, loss of session	– Delays and gaps in service

Packet jitter is usually caused by queuing and routing across a network, or buffering in switched transport networks. Although the variation in packet delay should be minimal in a core 10GigE network, jitter can be introduced as bandwidth demand increases and more buffering occurs.

Packet jitter should not occur in core networks at low utilization rates. Tests for packet jitter should therefore be performed under heavy load to provide enough information for the service provider to fully understand the behavior of the network and foresee any quality-of-service issues at maximum utilization. The maximum acceptable level of packet jitter is 30 ms for video applications and 10—20 ms for VoIP.

Since packet delay variation can result from queuing and buffering, it is important to measure jitter with test traffic that simulates a real application. Service providers must also ensure that the configuration of virtual LANs (VLANs) and other customer-defined services reflect typical end-user parameters. These measures will guarantee that the test traffic is processed in the same way as the customer traffic, and that the measurements can be used to determine the buffer level required to absorb packet jitter, and/or to define SLAs.

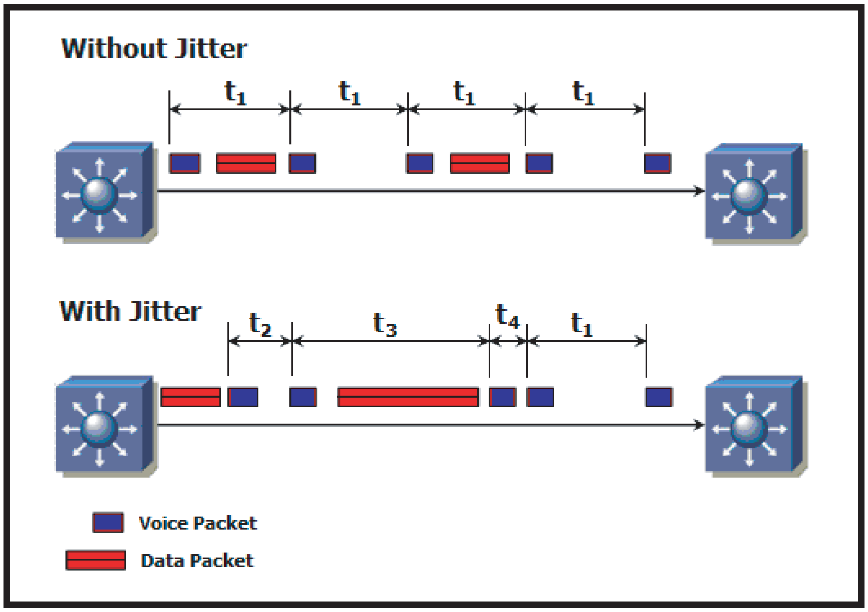

Figure 3. *Packet jitter causes variation in the time between packets, which can degrade voice and video.*

Peace of Mind for Service Providers

There is no doubt that 10GigE offers a more flexible alternative to SONET/SDH and can achieve the same reliability if extra protocols are added to the basic technology. However, a number of network parameters must be verified to ensure that a 10GigE connection provides an acceptable level of quality. Test equipment can provide all the data needed to quantify network behavior, ranging from dispersion measurement of the fiber plant to network performance testing with RFC 2544 and packet delay variation.

For service providers, this information is essential so that they can prove that they are fulfilling their SLAs and prevent potential loss in revenue. However, test equipment can also be used in a pro-active way to monitor the behavior and performance of a network. This can enable early detection of any degradation that could ultimately lead to a network outage, which in turn will improve network availability and help achieve the figures defined in the SLAs.

EXFO Corporate Headquarters > 400 Godin Avenue, Quebec City (Quebec) G1M 2K2 CANADA | Tel.: 1 418 683-0211 | Fax: 1 418 683-2170 | info@EXFO.com

Toll-free: 1 800 663-3936 (USA and Canada) | www.EXFO.com

EXFO Montreal	2650 Marie-Curie	St-Laurent (Quebec) H4S 2C3 CANADA	Tel.: 1 514 856-2222	Fax: 1 514 856-2232
EXFO Toronto	160 Drumlin Circle	Concord (Ontario) L4K 3E5 CANADA	Tel.: 1 905 738-3741	Fax: 1 905 738-3712
EXFO America	3701 Plano Parkway, Suite 160	Plano, TX 75075 USA	Tel.: 1 800 663-3936	Fax: 1 972 836-0164
EXFO Europe	Omega Enterprise Park, Electron Way	Chandlers Ford, Hampshire S053 4SE ENGLAND	Tel.: +44 2380 246810	Fax: +44 2380 246801
EXFO Asia	151 Chin Swee Road, #03-29 Manhattan House	SINGAPORE 169876	Tel.: +65 6333 8241	Fax: +65 6333 8242
EXFO China	No.88 Fuhua, First Road Central Tower, Room 801, Futian District	Shenzhen 518048, CHINA	Tel.: +86 (755) 8203 2300	Fax: +86 (755) 8203 2306
	Beijing New Century Hotel Office Tower, Room 1754-1755 No. 6 Southern Capital Gym Road	Beijing 100044 P. R. CHINA	Tel.: +86 (10) 6849 2738	Fax: +86 (10) 6849 2662

EXPERTISE REACHING OUT

Appnote152.1AN

Printed in Canada 06/07