125

NEXT-GENERATION SONET/SDH TECHNOLOGIES AND TESTING CONSIDERATIONS

APPLICATION NOTE

Claudio Mazzuca, Product Manager

SONET/SDH networks have stood the test of time and, until recently, they have been primarily used to transport aggregated voice signals and private-line frame relay and ATM services. The growth of these high-revenue-generating services over the years has resulted in large-scale deployments of metropolitan and long-haul SONET/SDH networks worldwide.

Today, however, service demand and consequently service provider revenue growth has shifted from these legacy-based services to Ethernet-based services. This shift in demand is largely due to two market factors. First, enterprise customers have expressed an increasing need for reliable Ethernet-based transport services and, second, Ethernet is becoming the Layer 2 transport technology of choice for both enterprise connectivity and access aggregation networks.

In addition, with the convergence of voice, video and data (e.g., triple-play access networks), Ethernet connections are now being used for delivery of mission-critical services, thus making 99.999% availability an absolute service requirement for these enterprise customers.

This application note provides an overview of next-generation SONET/SDH technologies designed for service providers to meet this demand shift. The paper also outlines some of the testing challenges that must be considered when qualifying next-generation systems and when activating and maintaining Ethernet-over-SONET/SDH circuits.

Next-Generation Technology Overview

The growing demand for Ethernet-based transport services has led to a 'rebirth' of SONET/SDH. Given the fundamental robustness of the technology and the massive capital investment that has been made over the years, service providers have been eager to find ways of using their existing SONET/SDH infrastructure to fulfill the growing market demand for Ethernet.

Consequently, in 1999, work was initiated within the ITU-T and ANSI standards bodies to define technologies that would help SONET/SDH technology evolve with the times by offering efficient means of transporting packet-based services over the widely deployed SONET/SDH networks.

Their answer was the definition and ratification of three key technologies that form the basis of next-generation SONET/SDH: generic framing procedure (GFP), virtual concatenation (VCAT), and link capacity adjustment scheme (LCAS).

GFP: A Common Data-Mapping Scheme for SONET/SDH

GFP, defined in ITU-T G.7041 specification, does exactly what its acronym implies; it provides a framing mechanism that supports the direct mapping of various traffic types into SONET/SDH frames, giving protocols like Ethernet and Fibre Channel the flexibility to be transmitted over long distances on existing SONET/SDH infrastructures.

Compared with other framing procedures such as Packet-over-SONET/SDH or Ethernet-over-Link Access Procedure (LAPS), also known as X.86, GFP has extremely low overhead requirements and robust frame-delineation qualities, making it a more efficient and robust data-services mapping scheme.

Figure 1 shows the mapping of GFP to SONET/SDH. IP data services that have previously been mapped to SONET/SDH through ATM and HDLC, both tightly coupled with SONET/SDH, can now be mapped through Ethernet and GFP.

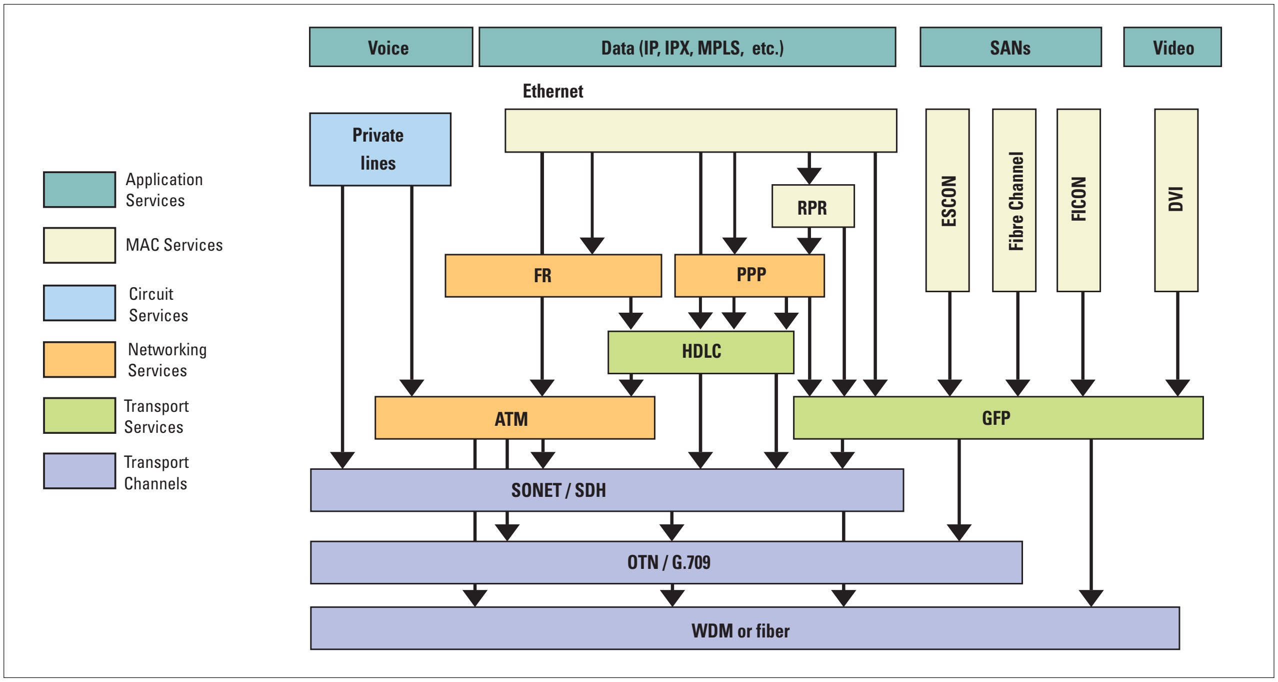

Figure 1: Data service mapping over GFP

There are two types of GFP frames that have been defined by the ITU standard: Framed GFP (GFP-F) and Transparent GFP (GFP-T). The decision on which mode to use depends on the underlying service to be transported.

Frame-Mapped GFP (GFP-F) is mapping mechanism in which one client signal frame is received and mapped in its entirety into one GFP frame. In this adaptation mode, the client/GFP adaptation function may operate at the data-link layer (or higher layer) of the client signal. Client protocol data unit (PDU) visibility is required. GFP-F mappings are currently defined for Ethernet MAC payloads (10/100/1000M/GbE) and IP/PPP payloads.

Transparent-Mapped GFP (GFP-T) frames are slightly different. GFP-T mapping provides a block-code-oriented adaptation mode in which the client/GFP adaptation function operates on the coded character stream rather than on the incoming client PDUs. Transparent GFP provides a way for a number of client data characters to be mapped into efficient block codes for transport within a GFP frame. With this type of mapping, block-coded client characters are decoded and then mapped into a fixed-length GFP frame and may be transmitted immediately without waiting for the reception of an entire client data frame (as in the case of GFP-F). Transparent mapping is intended to facilitate the transport of 8B/10B block-coded client signals for services that require very low transmission latency, making GFP-T ideally suited for Fibre Channel, ESCON, or FICON transport service.

A GFP frame consists of three main components: the core header, the payload header and the payload area. As shown in Figures 2a and 2b below, both mapping schemes share common aspects: core header, payload header, and optional payload FCS (pFCS). They differ solely in their payload adaptation mode as described above (GFP-T and GFP-F frames are identified via the user payload identifier (UPI) byte of the payload header).

The payload header carries information about the payload type (Ethernet, Fibre Channel, etc.) that it is carrying (see Figure 2c), while the core header carries information about the size of the GFP frame itself. Together, these components form the GFP frame, which can be mapped

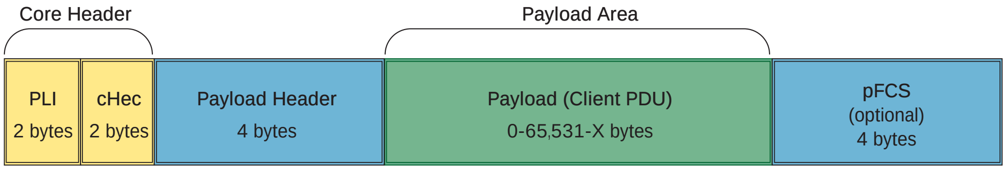

Figure 2a: GFP-F frame format

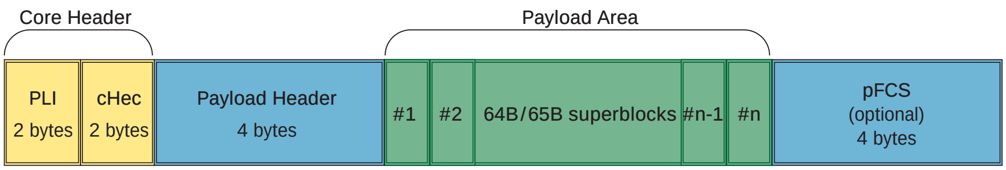

Figure 2b: GFP-T frame format

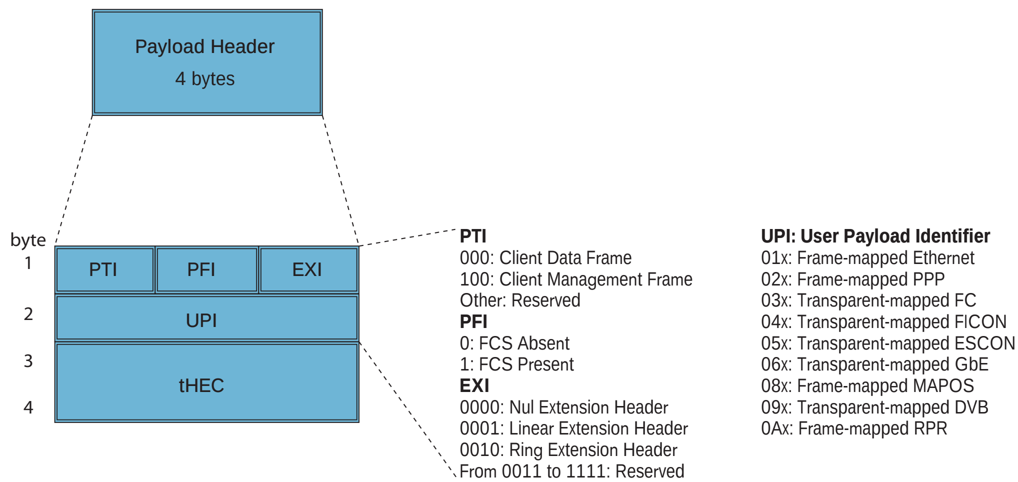

Figure 2c: Payload header format

into the T-Carrier/PDH channels that eventually go through the SONET/SDH signals. Each header contains a header error correction (HEC) calculation, allowing for the correction of single errors; that is, any errors that occur in the core header or in the payload header can potentially be corrected by the HEC. This creates a very robust mapping scheme, which ensures that GFP frames can get transported across a network without customer traffic loss.

In summary, GFP has been standardized to better optimize the transport of Ethernet and other data services over SONET. This technology has been embraced by network equipment and service providers as it provides an efficient way of providing interoperable data-services transport over the existing SONET/SDH install base. However, the answer to a truly efficient packet transport mechanism comes via the combination of both GFP and a bandwidth-optimizing technology such as VCAT and LCAS, as we will see in the next sections.

VCAT: Right-Sizing SONET/SDH Pipes

Contiguous concatenation has been part of SONET/SDH from its early days. It was conceived to accommodate high-speed data applications that use protocols such as ATM, Packet-over-SONET (PoS) and others that require high-bandwidth mappings over SONET/SDH. To enable this high-bandwidth transport, multiple SPEs are transported (and switched) across the SONET/SDH network as a single connection, with the first SONET container payload pointer set in normal mode and the subsequent payload pointers set to concatenation mode, thus linking all the units together.

Table 1 outlines supported contiguous concatenation for both SONET and SDH. For the SONET standard, these are denoted as STS-Xc, and for SDH as VC-4-Xc.

SONET	SDH	Payload Capacity (Mb/s)
STS-1	VC-3	48.38
STS-3c	VC-4	149.76
STS-12c	VC-4-4c	599.04
STS-48c	VC-4-16c	2,396.16
STS-192c	VC-4-64c	9,584.64

Table 1: Contiguous concatenation containers for SONET and SDH

Although contiguous concatenation has been successfully introduced and deployed for years, it presents some major deficiencies. Contiguous concatenation keeps the concatenated SONET payload through the whole SONET/SDH transport. Therefore, network elements must support contiguous concatenation from the source to the destination and at every intermediate node. In addition, data-service rates are not well-matched to these defined containers, so using GFP with current contiguous concatenation schemes results in sub-optimal use of the bandwidth because Ethernet and Fibre Channel data rates are not properly matched to these channels (e.g., 100M Ethernet service mapped over an STS-3c or VC4 results in approximately 33% of wasted bandwidth).

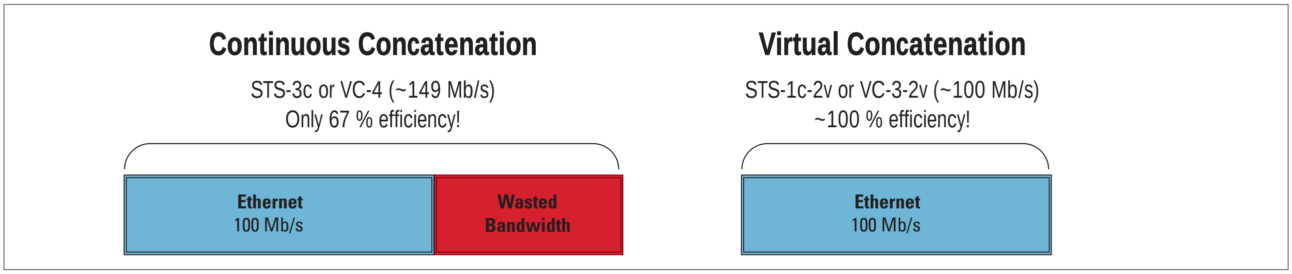

Figure 3: Contiguous concatenation vs. virtual concatenation efficiency when transporting 100M Ethernet service

To address these limitations, virtual concatenation (VCAT), defined in ITU standard G.707, was developed to allow individual SONET/SDH channels to be bonded to create a customized virtual concatenation group that can be any multiple of the basic rate (for example, for high-order VCAT, this basic rate would be an STS-1/STS-3 or VC-3/VC4, whereas for low-order VCAT this basic rate could be VC-11/12 or VT-1.5/2). For example, two STS-1s can be combined to form one virtual concatenation group that comes close to supporting a 100 Mb/s Ethernet interface (an STS-1 payload envelope of approximately 49 Mb/s creates an almost exact match to the 100 Mb/s circuit).

In essence, virtual concatenation is an inverse multiplexing procedure whereby the contiguous bandwidth is broken into individual SPEs at the source transmitter and logically represents them in a virtual concatenation group (VCG). Control packets, which contain the necessary information for reassembling the original data stream at its destination PTE, are inserted in some of the currently unused SONET/SDH overhead bytes (H4 byte for high-order, and Z7 (SONET) and K4 (SDH) for low-order). This information contains the sequence order of the channels and a frame number, which is used as a time stamp. The VCG members are transported as individual SPEs across the SONET/SDH network with all the intelligence required to handle virtual concatenation located at the end points of the connections (i.e., at the path termination equipment, or PTE). The receiving end-point (PTE) is responsible for reassembling the original byte stream. This allows SONET/SDH channels to be routed independently through the network without requiring any acknowledgement of the virtual concatenation. In this manner, virtually concatenated channels may be deployed on the existing SONET/SDH network with a simple end-point node upgrade.

Two forms of virtual concatenation have been defined:

- ***High-Order Virtual Concatenation (HO VCAT)***
  HO VCAT provides bandwidth for links that require greater than 51.84 Mb/s, but do not lend themselves to one of the standard contiguous concatenation bandwidth configurations. High-order virtual concatenation is realized under SONET and SDH by the PTEs, which combine either multiple STS-1/STS-3c SPEs or VC-3/VC-4, therefore making it ideally suited for 100M, Gigabit Ethernet, and Fibre Channel rates. HO VCAT rates are designated by STS-m-nv or VC-m-nv, where the nv indicates a multiple n of the STS-m/VC-m base rate.

- ***Low-Order Virtual Concatenation (LO VCAT)***
  LO VCAT provides capacity for links that require greater than 1.728 Mb/s (VT 1.5/VC-11), but less than 51.84 Mb/s. These concatenations are designated by VT-1.5/2-nv for SONET and VC-11/12-nv for SDH. LO VCAT is typically used for low-speed 10 Mb/s or sub-rate 100M Ethernet services.

As presented in Table 2 below, VCAT provides a much more efficient use of the transport bandwidth for data user interfaces. With VCAT, an OC-48 link can carry two full Gigabit Ethernet signals with 95% of the link used through seven virtual STS-3c/VC-4s each, instead of just one Gigabit Ethernet signal with 42% of the link used through an STS-48c/AU-4-16c.

Service	Bit Rate (Mbps)	Efficiency	
		Without VCAT	With VCAT
Ethernet	10	STS-1 (20%)/ VC-3 (20%)	VT1.5-7v (93%)/ VC12-5v (97%)
Fast Ethernet	100	STS-3c (47%)/VC-4 (67%)	STS-1-2v (99%)/ VC-3-2v (99%)
Gigabit Ethernet	1000	STS-48c (42%)/ VC-4-16c (42%)	STS-3c-7v (95%)/VC-4-7v (95%)
Fibre Channel (1x)	1062.5	STS-48c (43%)/ VC-4-16c (43%)	STS-3c-6v (98%)/ VC-4-7v (98%)

Table 2: Bandwidth efficiency using virtual concatenation

In summary, virtual concatenation enables SONET/SDH transport pipes to be filled more efficiently with data services by grouping individual SONET/SDH containers into a virtual high-bandwidth "link", matched to the required service bandwidth.

Link-Capacity Adjustment Scheme (LCAS)

LCAS, as defined by the ITU (per ITU-T recommendation G.7042), is a complementary technology to virtual concatenation. LCAS allows for the dynamic change of the bandwidth of a virtually concatenated group of channels. Signaling messages are exchanged within the SONET/SDH overhead (as with VCAT, H4 for HO VCAT and Z7/K4 for LO VCAT) in order to change the number of tributaries being used by a virtually concatenated group (VCG). The number of tributaries can be increased or decreased in response to an identified change in service bandwidth requirement, or in response to a fault condition of an existing VCG member.

LCAS works by ensuring synchronization between the sender (PTE) and receiver (PTE) during the increase/decrease of the size of a virtually concatenated circuit, in such a way that it doesn't interfere with the underlying data service. Should failures occur on individual member of a group, the size of the group can be reduced temporarily, instead of taking the entire group out of service (which would be the case if LCAS was not enabled – the entire VCG would be declared as "failed" in the event of a failure of one VCG member). With LCAS, once the defect is repaired, the group size can be restored to full bandwidth without affecting the underlying service.

In addition to providing a resiliency mechanism for VCAT, LCAS gives service providers the flexibility to tailor service bandwidth as required. For example, if a certain customer requires additional bandwidth in the late evenings for file transfers (i.e., banking institutions), the service provider can provide a value-added service by provisioning increased bandwidth for a predefined period.

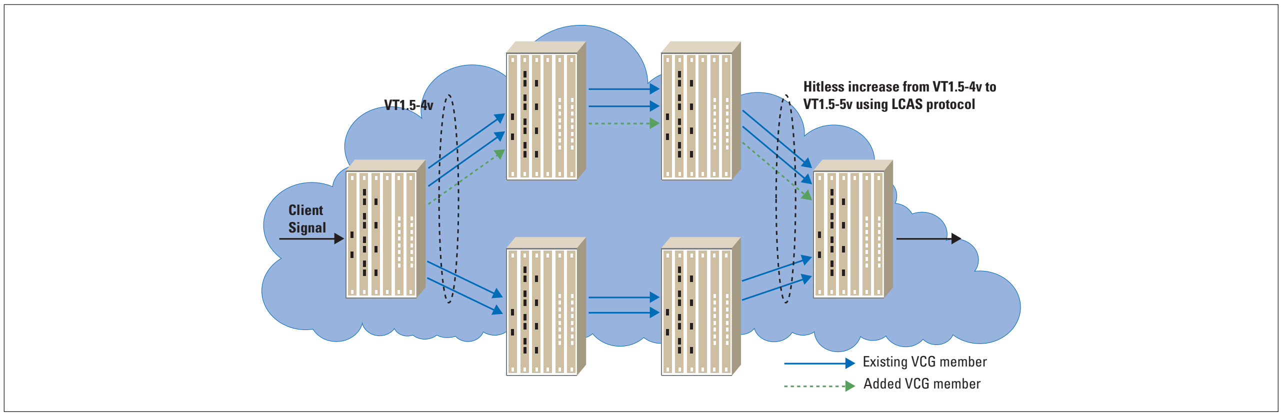

Figure 4: Increasing VCG size using LCAS

In summary, VCAT and LCAS combined can be extremely useful when provisioning data packet transport services over an existing SONET network. In this case, VCAT is used to provision point-to-point connections over the SONET network and LCAS provides resiliency to VCAT connections and allows for further "tuning" of the allocated VCAT bandwidth.

Multiservice Platforms: The Next-Generation ADM

Today, packet-aware SONET/SDH add-drop multiplexers (ADMs) supporting GFP, VCAT, and LCAS are commonly known as multiservice provisioning platforms (MSPPs) and multiservice transport platforms (MSTPs). These units are actively being deployed in service provider networks as they offer an efficient means of transporting packet-based client signals, such as Ethernet and Fibre Channel over existing SONET/SDH infrastructures. As stated above, the intelligence needed to support GFP, VCAT, and LCAS is located at the endpoints of the network; therefore, MSPPs/MSTPs need only be deployed at the edge of the transport network, where services are created. In many cases, MSPPs/MSTPs are simply upgrades to currently deployed ADMs, so with minimal capital investment, service providers can deliver high-growth services that offer significant value to their customers.

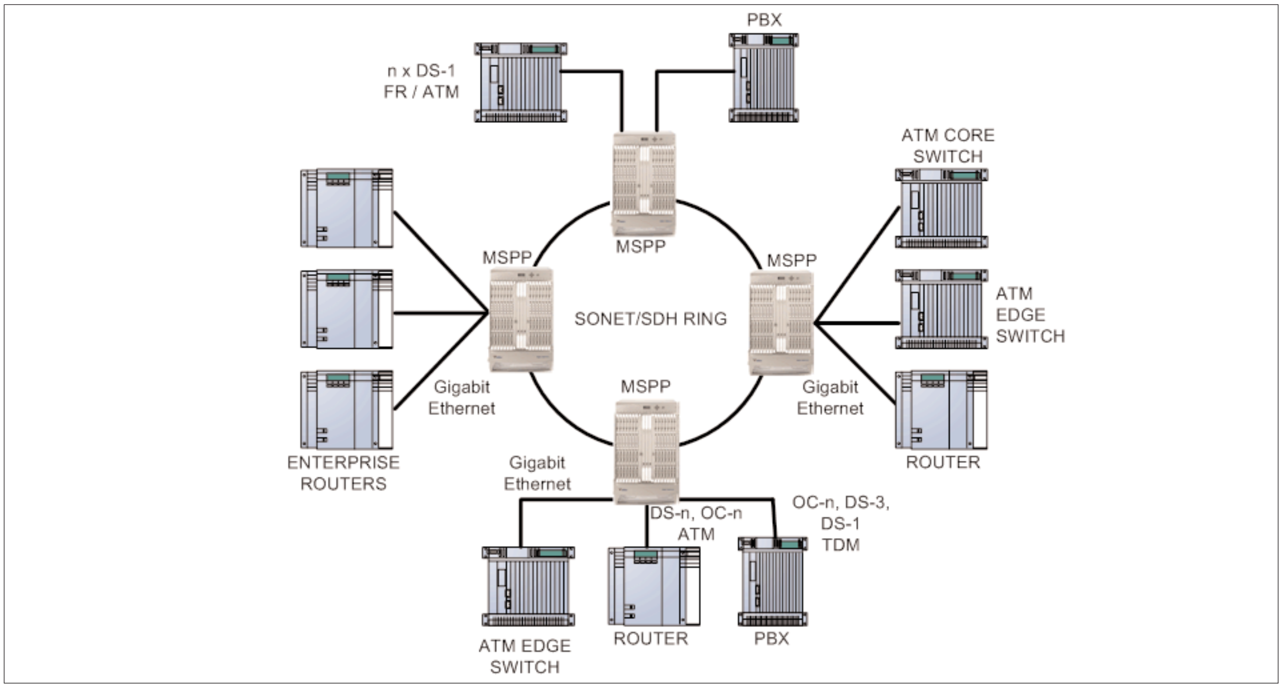

Figure 5: Multiservice provisioning platforms in a SONET/SDH ring

New Technologies = New Testing Challenge

Clearly, MSSPs/MSTPs offer obvious advantages to service providers as they help maximize the usage of the current install base, with minimal capital expenditure, while providing a new revenue stream to fuel growth for the coming years.

However, with these new technologies come new deployment challenges. GFP, VCAT and LCAS add another layer of complexity to accommodate the transport of different data services. In many cases, these next-generation functionalities are delivered as enhancements to existing SONET/SDH ADMs. Nevertheless, for a service provider, whether it is delivered by new network elements or enhancements to existing ADMs, this additional layer needs to be fully tested prior to wide-scale deployment in the network. GFP, VCAT, and LCAS conformance validation and performance characterization must be performed for all offered services (10M, 100M, 1000M, GigE, etc.). In addition, given the fact that multivendor networks are now commonplace and that these technologies are still in their early stages, it is important that proper vendor interoperability tests are performed so as to mitigate risks prior to deployment.

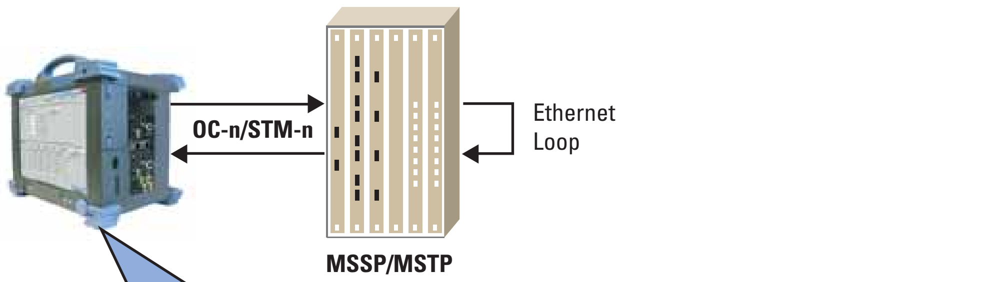

Tests:
- GFP mapping and demapping based on ITU-T G.7041, including support for extension headers (linear, ring)
- GFP mapping with variable rate Ethernet signal inputs (i.e. packet sweeps)
- GFP multiservice performance (mapper's ability to perform policing, scheduling, and classifying per client signal and service)
- GFP error injection and analysis (cHEC/tHEC/eHEC correctable/uncorrectable, pFCS)
- GFP alarm injection and analysis: (LOCS, LOCCS)
- GFP payload (i.e. Ethernet) error injection
- GFP overhead manipulation
- GFP optional payload FCS validation
- GFP mapping efficiency (per client signal frame size)

Purpose:
- Verify NE's conformance to GFP ITU-T specifications
- Verify NE's supported GFP functionaility – functional validation
- Verify NE's GFP mapper performance (i.e. characterization)

Figure 6: GFP conformance

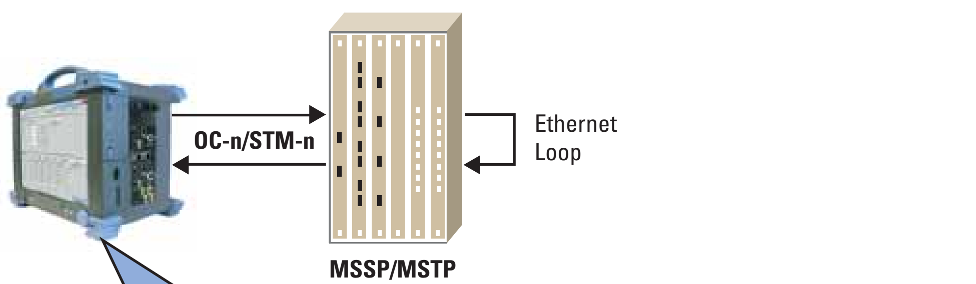

Tests:
- HO VCAT container support (all valid combinations)
- LO VCAT container support (all valid combinations)
- VCAT alarm injection and analysis (LOM, OOM1, OOM2, SQM, LOA)
- VCAT error insertion and analysis (all support SONET/SDH path errors – e.g. REI-P, B3, etc.)
- VCAT error and alarm injection: single and multiple member injections
- Sequence indicator manipulation and processing
- Differential delay injection and analysis – per member (up to 256 ms)

Purpose:
- Verify NE's conformance to LO and HO VCAT containers specified in ITU-T G.707
- Verify NE's performance with differential delay injection values specified in ITU-T G.707
- Verify NE's performance in the event of single and multiple alarm/errors events

Figure 7: VCAT conformance

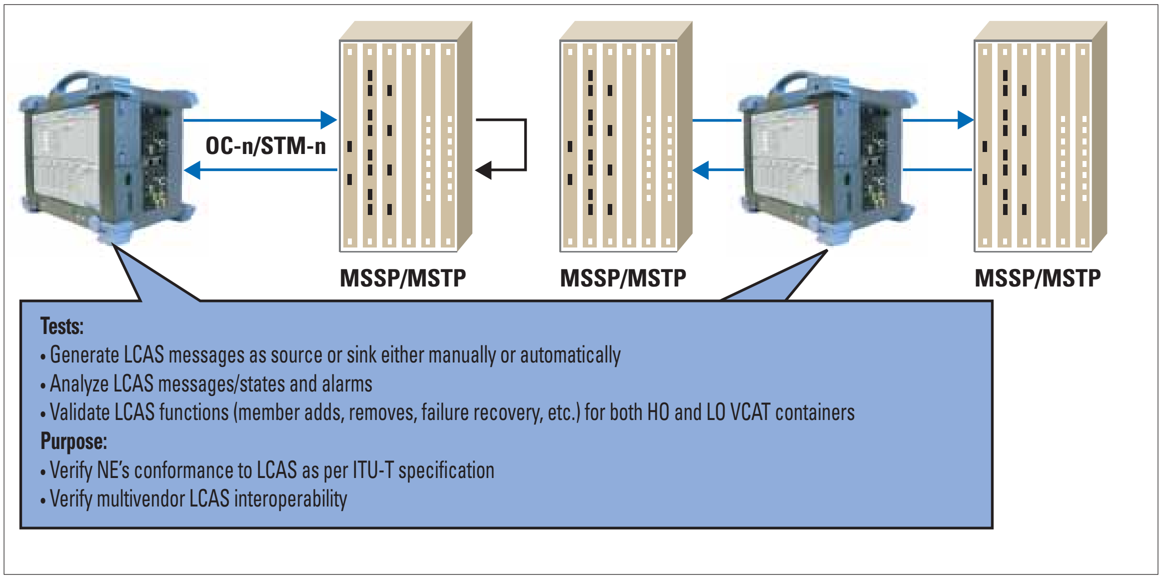

Figure 8: LCAS conformance

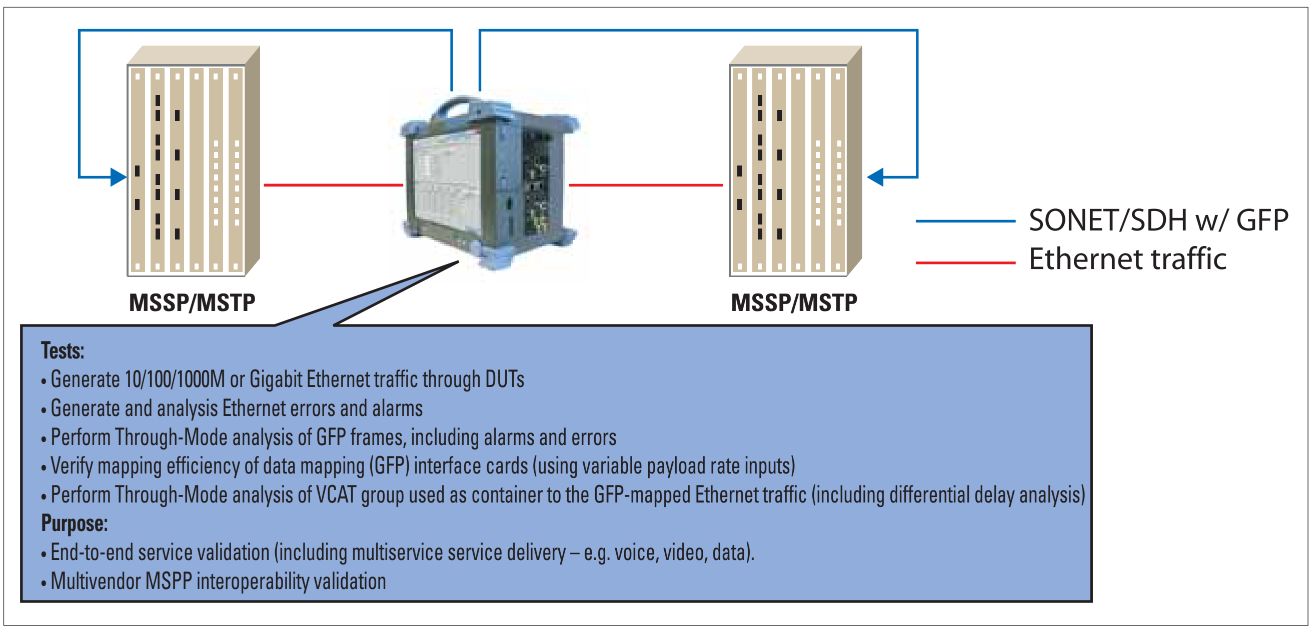

Figure 9: End-to-end and interoperability conformance

Once next-generation services are deployed, service providers are left with the challenge of ensuring that their end-to-end data services perform according signed service-level agreements. Post-deployment troubleshooting and maintenance activities must now not only include legacy SONET/SDH-layer visibility, but also visibility at the higher GFP, VCAT and LCAS layers. Service providers require a test solution that can identify and correlate any potential problems that can normally occur throughout the network but, more importantly, across multiple layers; namely, the physical, SONET/SDH, GFP, VCAT and data (Ethernet, Fibre Channel) layers.

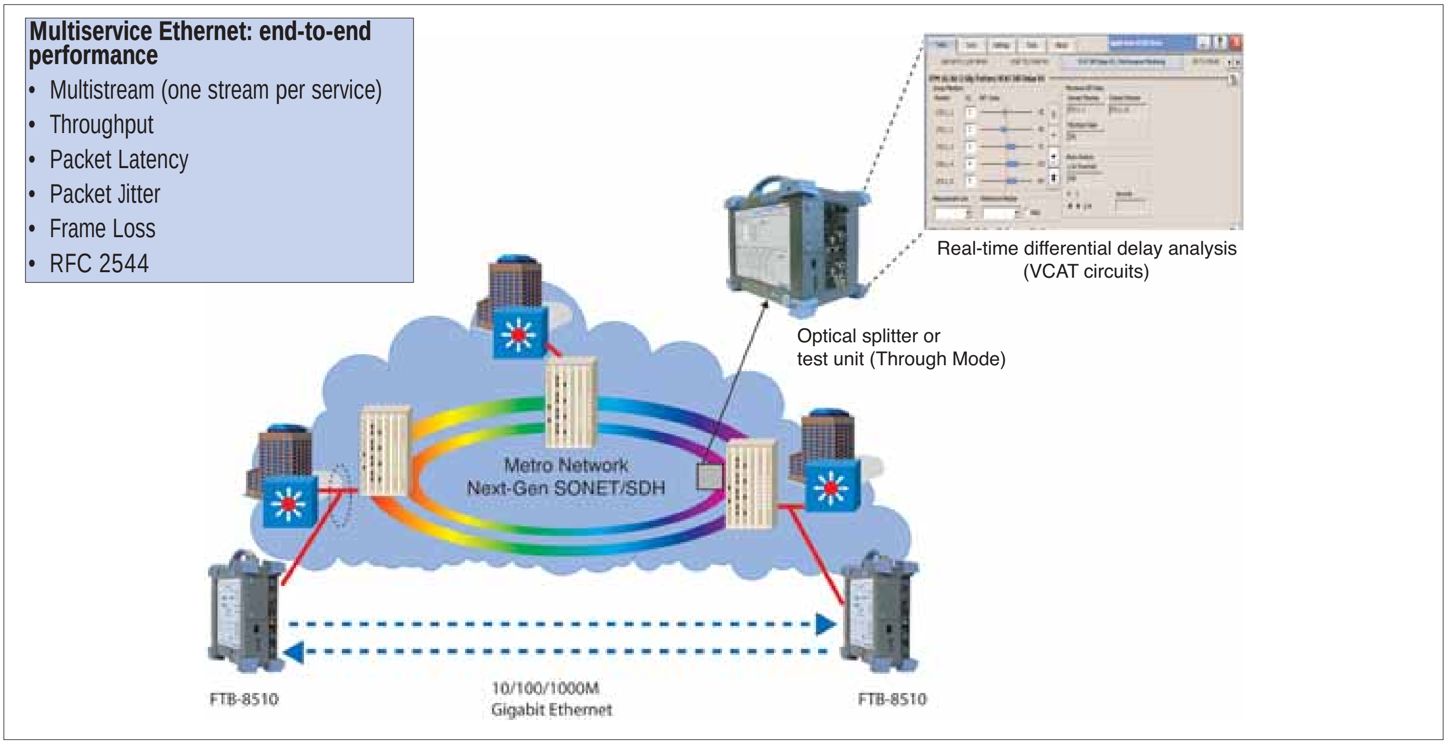

Figure 10: Packet services turn-up and troubleshooting

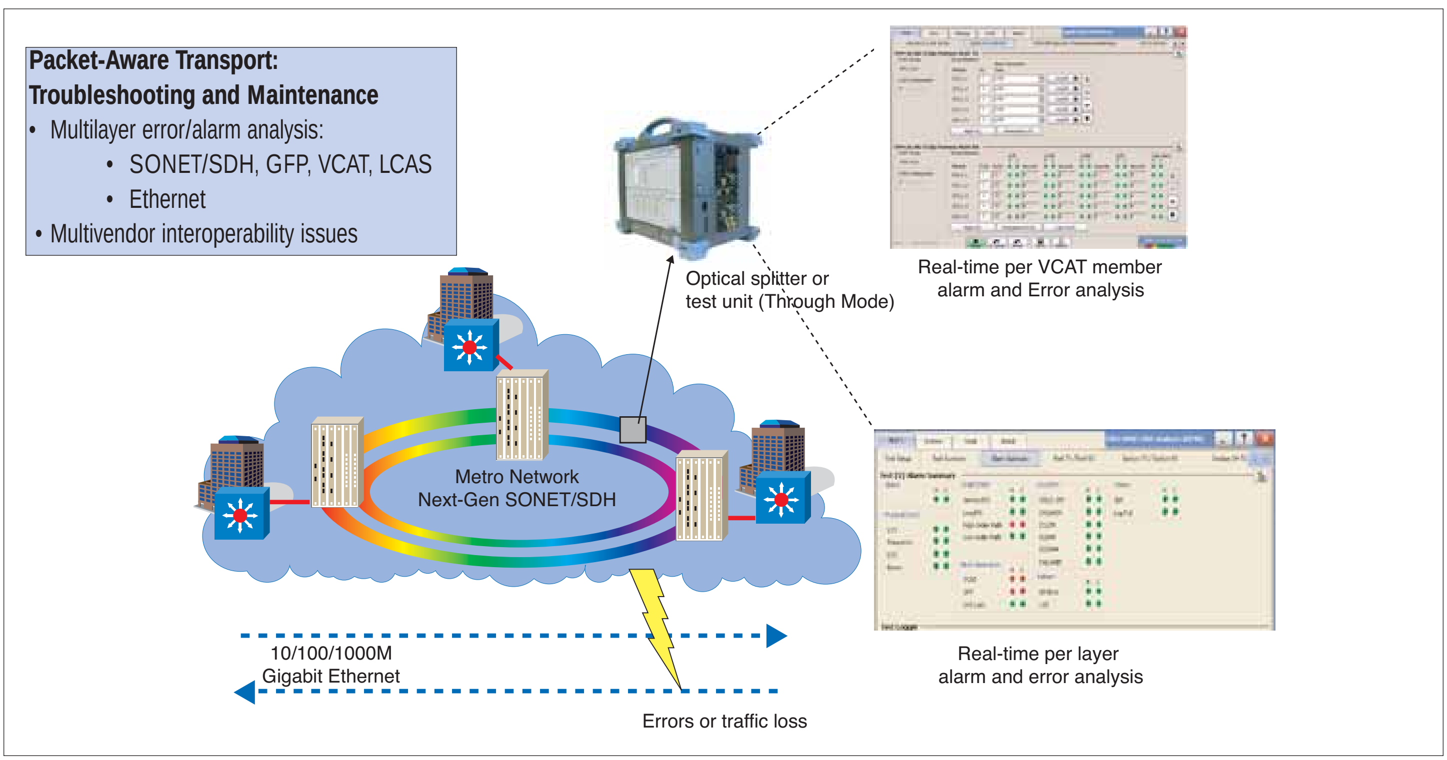

Figure 11: Ethernet-over-SONET/SDH troubleshooting and maintenance

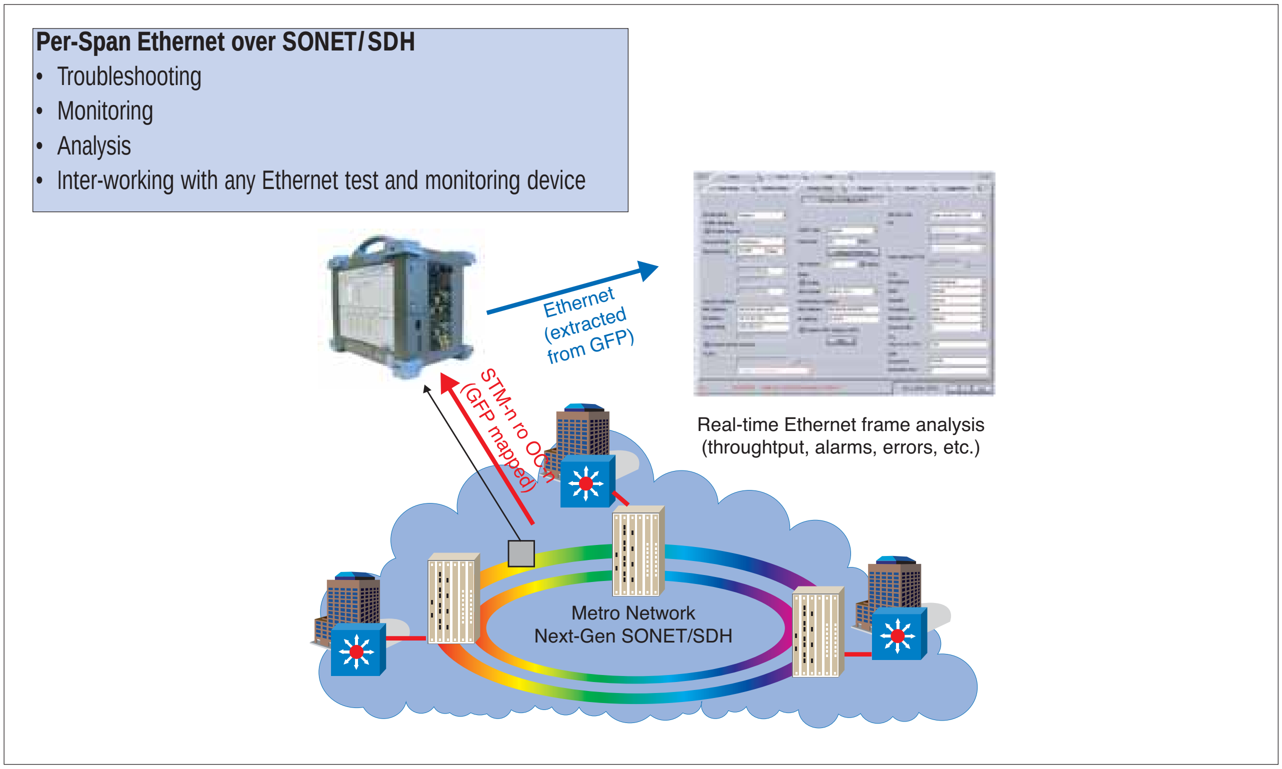

Figure 12: Ethernet signal analysis

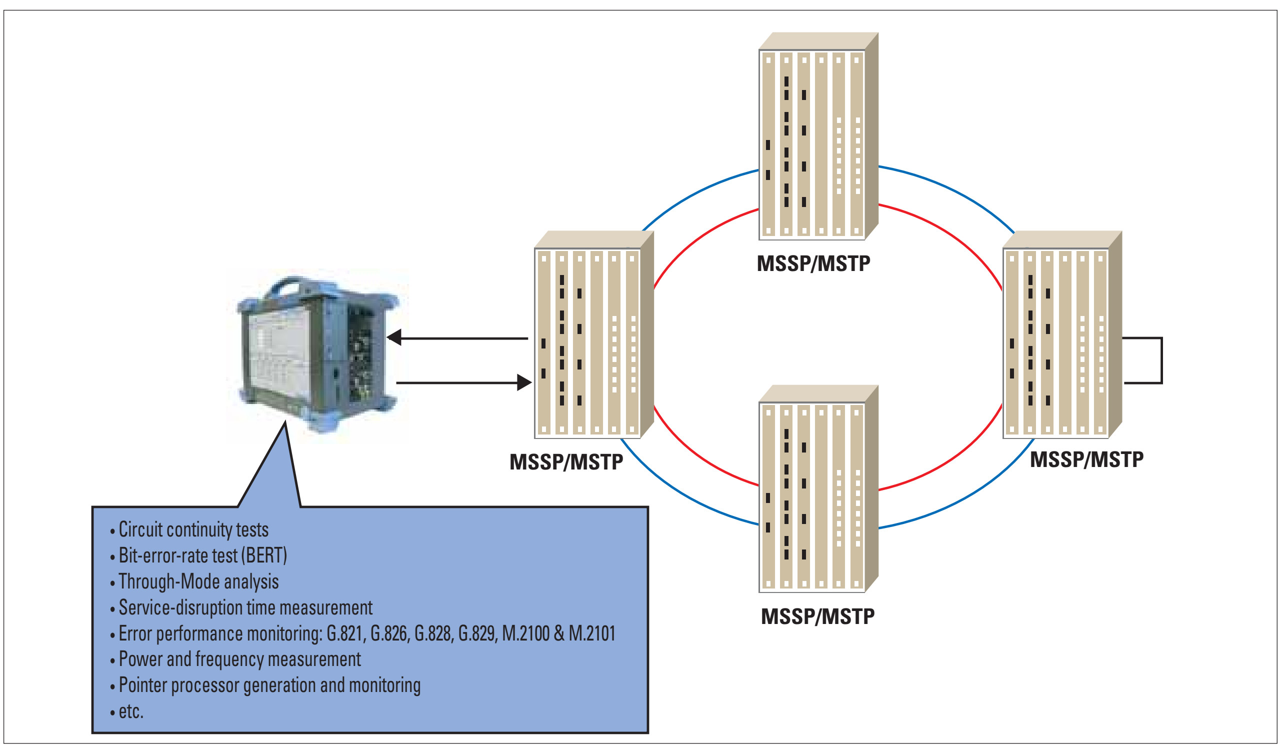

Figure 13: TDM circuit turn-up and troubleshooting

In addition, service providers cannot afford to supply each field technician with a multitude of disparate test equipment to perform their daily duties. They require versatile test equipment that is capable of supporting legacy TDM (electrical and optical) and next-generation packet-based circuit turn-ups, as well as common client signal test interfaces such as 10/100/1000M Ethernet and Fibre Channel.

With the introduction of the FTB-81x0 Next-Generation SONET/SDH Analyzer test modules, EXFO's widely deployed FTB-400 Universal Test System is now equipped with Ethernet, Fibre Channel and Next-Generation SONET/SDH features, providing the ideal solution for R&D lab test applications, as well as field installation and maintenance of next-generation SONET/SDH networks.

Corporate Headquarters > 400 Godin Avenue, Vanier (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 America	3701 Plano Parkway, Suite 160	Plano, TX 75075 USA	Tel.: 1 800 663-3936	Fax: 1 972 836-0164
EXFO Europe	Le Dynasteur, 10/12 rue Andras Beck	92366 Meudon la Forêt Cedex FRANCE	Tel.: +33.1.40.83.85.85	Fax: +33.1.40.83.04.42
EXFO Asia-Pacific	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
	Beiijing New Century Hotel Office Tower No.6 Southern Capital Gym Road, Room 1754-1755	Beijing 100044, P.R. CHINA	Tel.: 86.10.6849.2738	Fax.: 86.10.6849.2662

EXPERTISE REACHING OUT

Printed in Canada 06/02