



Engines for Global Connectivity

communicator

Summer 2002

From the President

The telecom industry is undergoing unprecedented change that will fundamentally affect customer needs along the entire supply chain.

Rapid deployment of new technologies and revolutionary new network architectures are now just thoughts of the past.

But one requirement which remains a "must" is that the telecommunication networks must be designed to be survivable. In doing so, the challenge for equipment suppliers is to reduce development costs while producing systems that foster network evolution. The survivability concept must extend to switching, transmission, and other elements of the total network, including the CPE equipment.

When we refer to the transmission networks,

Continued on page 2

Ethernet-over-SONET: Part 1

In order to support increasing data traffic levels, equipment developers must build systems that map Ethernet packets over SONET/SDH links. In this two-part series we'll lay out the encapsulation techniques required to make Ethernet-over-SONET come to life. The second part, we will continue the discussion of encapsulation techniques, looking at the generic framing protocol (GFP) and link access procedure for SDH (LAPS).

Today's incumbent service providers are seeking to accommodate increasing demands for bandwidth, security and service level management resulting from the proliferation of LAN-based applications. According to RHK, 2.4 million metro Ethernet ports distributed among 10/100/1000 Mbps rates are expected to be deployed by 2006. These Ethernet ports will primarily be used to support these emerging applications:

- Broadband Internet access for e-mail and web-based content;
- Packet-based video conferencing;
- Tunneled virtual private networks (VPNs);
- E-commerce hosting for outsourced applications services such as customer relationship management (CRM); and
- Storage area networking (SAN) for outsourced access/management of server content including disaster recovery services.

The big issue carriers face is how to handle Ethernet traffic. Depending on the cost, distance, bandwidth, and traffic management requirements of the WAN/MAN application, several metro approaches may be used to transport Ethernet data traffic. These include direct mapping of Ethernet over wavelengths (EoW), Ethernet over SONET/SDH (EoS),

Continued on page 2

In this issue

Ethernet-over-SONET: Part 1

President's Report.....Continued

survivability considerations span the fiber-based backbone, cross-connect systems, multiplexers, etc., including the operations, maintenance, and administration subsystems. The overall transmission network needs to be designed with redundancy and diversity in mind. These are important considerations in ensuring that the impact on a group of customers is kept minimal when a network element malfunctions.

TranSwitch has always believed that the overall communications network reliability ultimately depends on the reliability of the components used in the network, including the VLSI components. Our recently announced TXC-OMNI solution- the TXC-OMNI SE (Switch Element) and the TXC-OMNI TP (Transport Processor), incorporates all the necessary functions and features to enable equipment vendors to deliver a survivable system.

The TXC-OMNI is a silicon and software solution that enables telecom equipment vendors to rapidly and cost effectively develop next generation ADMs (add drop multiplexers), DCSs (digital cross connects), and hybrid TDM (time division multiplexing)/data switching systems. Some of the key features of the TXC-OMNI solution include: a.) Metro Evolution architecture; b.) supports TDM, data and hybrid applications; and c.) uniform Application Programming Interface (API).

Our customers have learned to rely on us because they know that we are as concerned as they are about the reliability and survivability of the communication networks which incorporate their equipment.

Santanu Dey

Ethernet-over-SONET.....Continued from page 1

optical Ethernet (i.e. native Ethernet over fiber for long haul, 1000BaseLX), or implementing an Ethernet in the first mile (EFM) solution over copper or fiber. Of all these solutions, EoS is the one gaining the most ground with today's developers.

In this two-part series, we'll lay out some of the key requirements for mapping Ethernet frames over SONET/SDH links. In part 1, we'll look at the advantages EoS brings to an equipment designer. We'll also begin a discussion on encapsulation techniques, looking at virtual concatenation and the link capacity adjustment scheme (LCAS). Part 2 will follow with a discussion on the generic framing protocol (GFP) and the link access procedure for SDH (LAPS).

EoS—Why it Should Win


EoS collectively represents a group of industry standard specifications that have been developed for optimal transport of Ethernet through battle-tested circuit-switched topologies. It embodies several related technologies and can be implemented using techniques such as virtual concatenation (VC), which is a recent extension

to the ITU-T G.707 SDH standards. Together these specifications account for the mapping, aligning, sequencing and delay compensation of the individual channels. The result is a very cost-effective way to provide flexible bandwidth in small to large increments with a rate adaptation benefit.

Historically, in order to compensate for transport inefficiencies that existed between Ethernet and SONET, service providers would over-provision their circuits to ensure reliable delivery of data. This traditionally made use of standard contiguous concatenation along with packet-over-SONET/SDH (POS) and was fixed in terms of bandwidth granularity.

POS, the transmission of IP data over SONET frames via PPP, has certain real advantages but in the past required that bandwidth be predetermined in a rigid and constrained manner. For example, a Gigabit Ethernet stream in POS would need a full OC-48 pipe with standard contiguous concatenation. Service providers have dealt with the resulting cost penalty by deploying extra router interfaces to sup-

Figure 1: Ethernet switches linked over a SONET/SDH network



port mesh and aggregation of links. On occasion they employ redundant routers to meet protection and load balancing requirements.

EoS takes a different approach to the problem. Unlike POS, which calls for rigid bandwidth requirements, EoS allows bandwidth to be shared among several Ethernet ports **Figure 1**. Using EoS in combination with VC, a Gigabit Ethernet channel can be built (STS-1-24C or 24 STS-1's concatenated) while the unused portion of the OC-48 bandwidth can be deployed for other Ethernet or TDM services.

Another problem with the POS approach is that a costly intelligent router function is required to strip Ethernet media access control (MAC) frame encapsulation before the payload is encapsulated in PPP for transport. The Ethernet MAC frame structure is lost, eliminating the layer 2 intelligence that has evolved over the last six years for use in security mechanisms (such as L2TP), multicasting, traffic prioritization (802.1p) and filtering (802.1q VLAN id), as well as ring protection schemes such as Ethernet protection switching (EPS) and resilient packet ring (RPR).

EoS, on the other hand, emulates a connectionless architecture with full-mesh connectivity that more synergistically supports the wave of emerging distributed applications. EoS capability in effect turns the SONET/SDH MAN/WAN infrastructure backbone into a transparent Ethernet segment for attached servers and clients. Bridging Ethernet traffic across the WAN at layer 2 removes the need for traffic to terminate at intermediate sites. Eliminating the need to maintain

route tables and exchange routing protocols between links improves end-to-end delay and greatly simplifies administration.

Ethernet layer 2 services (such as WLANs 802.11P priority) combined with efficient EoS mechanisms can work in partnership with underlying SONET/SDH networks to give service providers network edge flexibility. By using this capability in tandem with load-sharing switch/router topologies, allocation of protected vs. unprotected bandwidth according to the particular service level agreement (SLA) can increase and optimize a carrier's network revenue.

Encapsulation Techniques

OK, so we laid out why EoS is a better choice. Now let's talk about making EoS work. The real key to making EoS architectures come to life lies in the encapsulation schemes available. Currently, there are four encapsulation techniques available—virtual concatenation (VC) and the link capacity adjustment scheme (LCAS) techniques, which define the method for transport, and the generic framing procedure (GFP) and link access procedure for SDH (LAPS) techniques, which are layer 1 adaptation protocols for transport.

All four of these encapsulation techniques have been under development by various standards bodies. The main standards references for these encapsulation mechanisms have been developed by the T1X1.5 working group and formalized as recommended standards by the ITU-T Standards Committee Study Group 15. The relevant ITU-T documents are:

1. TU-T G.707/Y.1322, October 2000: Network Node Interface for the Synchronous Digital Hierarchy ([G707]);
2. TU-T G.783, October 2000: Characteristics of SDH Equipment Functional Blocks ([G783]);
3. ITU-T G.803, March 2000: Architecture of Transport Networks Based on SDH ([G803]);
4. G.805, March 2000: Generic Functional Architecture of Transport Networks ([G805]);
5. T G.7041/Y1303, January 2002: Generic Framing Procedure ([G7041]);
6. G.7042/Y1305, November 2001: LCAS for Virtually Concatenated Signals ([G7042]);
7. T Recommendation X.85/Y.1321, March 2001: IP over SDH Using LAPS; and
8. T Recommendation X.86, February 2001: Ethernet over LAPS.

Now, let's take a deeper look at how they work. In this part, we'll focus on VC and LCAS while in Part 2 we'll explore GFP and LAPS. Let's start with VC.

The Basics of VC

The SONET/SDH structures allow a flexible synchronous optical hierarchy to carry various client signals at different rates. The payload of the basic STS-1/STM-0 (51.84 MHz) signal may consist of independent lower order tributaries—virtual tributaries (VT) in SONET or tributary units (TU) in SDH—that provide for the transport of lower-rate services. Likewise, higher-rate OC-N SONET/SDH signals can be built using N

Continued on page 4

Ethernet-over-SONET.....Continued from page 3

units of the basic STS-1/STM-0. Payloads that exceed the standard payload container capacity (a VC-3) may be transported using a concatenation method.

There are two methods of concatenation for higher-order paths (STS-1/AU-3 and above): (1) pointer-based (always contiguous, with component STS-1s transported together and identical in phase; an example is the 155.52 Mbps STS-3c signal), and (2) VC, in which the payload is split up over multiple STS-1s/STS-3c's that may follow physically separate routes. For lower-order containers such as TU-3/TU-2/TU-12/TU-11 (SDH) or VT6/VT2/VT1.5 (SONET), pointer-based concatenations are not possible, so the mechanism used for aggregating containers for higher payload rates is VC.

VC is defined at two levels: high-order and low-order. High-order VC groups the payload of different

signals at 51.840 Mbps or 155.52 Mbps, while low-order VC groups the payload of different VTs/TUs at lower rates such as 1.544 Mbps or 2.048 Mbps. **Tables 1 and 2** show the capacity of virtually concatenated tributaries in SDH and SONET, respectively.

Bit rates for Ethernet LAN services may be 10 Mbps, 100 Mbps, 1 Gbps, or 10 Gbps. These can be accommodated in appropriately sized VC payloads, improving bandwidth usage. In addition, voice and data can be sent over the same transport structure using different containers.

It's important to note that VC requires concatenation functionality only at the path termination equipment,

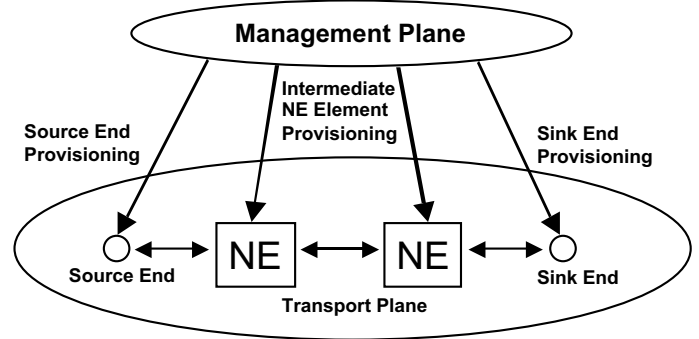


Figure 2: Provisioning a VCG (or LCAS add/remove) using the management plane

with older legacy equipment are left alone.

Tapping into the Management Plane

VC of multiple paths involves several steps. First, a link must be created with a source end (SE) and a sink end (SkE) for the virtually concatenated group (VCG) using the control/management plane. VC is a management plane-oriented protocol. **Figure 2** illustrates end-to-end transport connection provisioning using the management plane.

The management plane also identifies members of the VCG and assigns the physical path for each member, which includes tracking intermediate sub-network connection (SNC) switching.

After the link is established, the SE in all the constituent members of the VCG employs a path overhead (POH) byte to pass information to the SkE on the relative phase difference between arriving members of the VCG that indicates the physical delay incurred during transit of each member. The POH byte also delivers the sequence number of each arriving member that indicates to the SkE the original order in which the member appeared in the VCG at the SE.

Thus, incorporating VC functionality only involves upgrading the ends of the path, while the intermediate nodes

	Carried in	X	Capacity (kbit/s)	In steps of (kbit/s)
VC-11-Xv	VC-3	1 to 28	1600 to 44800	1600
VC-11-Xv	VC-4	1 to 64	1600 to 102400	1600
VC-11-Xv	Unspecified	1 to 64	1600 to 102400	1600
VC-12-Xv	VC-3	1 to 21	2176 to 45696	2176
VC-12-Xv	VC-4	1 to 63	2176 to 137088	2176
VC-12-Xv	Unspecified	1 to 64	2176 to 139264	2176
VC-2-Xv	VC-3	1 to 7	6784 to 47448	6784
VC-2-Xv	VC-4	1 to 21	6784 to 142464	6784
VC-2-Xv	Unspecified	1 to 64	6784 to 434176	6784

Table 1: SDH capacity of virtually concatenated VCs

	Carried in	X	Capacity (kbit/s)	In steps of (kbit/s)
VT1.5-Xv	STS-1	1 to 28	1600 to 44800	1600
VT2-Xv	STS-1	1 to 21	2176 to 45696	2176
VT3-Xv	STS-1	1 to 14	3328 to 46592	3328
VT6-Xv	STS-1	1 to 7	6784 to 47448	6784
VT1.5-Xv	STS-3c	1 to 64	1600 to 102400	1600
VT2-Xv	STS-3c	1 to 63	2176 to 137088	2176
VT3-Xv	STS-3c	1 to 42	3328 to 139776	3328
VT6-Xv	STS-3c	1 to 21	6784 to 142464	6784
VT1.5-Xv	Unspecified	1 to 64	1600 to 102400	1600
VT2-Xv	Unspecified	1 to 64	2176 to 139264	2176
VT3-Xv	Unspecified	1 to 64	3328 to 212992	3328
VT6-Xv	Unspecified	1 to 64	6784 to 434176	6784

Table 2: SONET capacity of virtually concatenated VTs

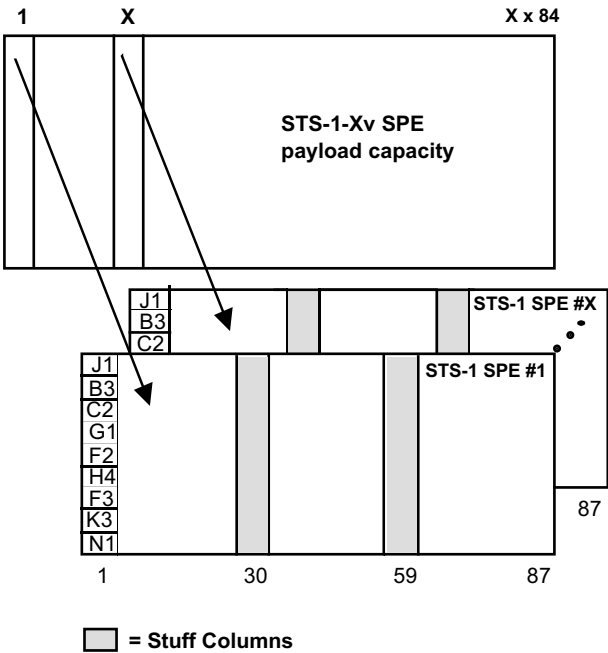


Figure 3: X virtually concatenated VC-n (n = 3) synchronous payload envelope (SPE) structure (higher-order VCG)

As an option, designers can also use LCAS to hitlessly adjust the payload allocated to VCGs. This enables the addition or removal of members to/from a VCG in response to requested increases or decreases in capacity requirements or link failure conditions. This option is laid out in the G.7042 specification.

VCGs are set up using higher order and lower order containers in SONET/SDH as shown in **Figures 3 and 4** (refer to ITU-T G.707, Section 11.2).

Now, the VC-enabled EoS system must determine if there is suitably-sized memory for buffering of all arriving members of a VCG at the SkE. By doing this, the system can achieve proper phase alignment (de-skewing) and sequence number matching (i.e., reordering members to match SE for the payload readout process).

The format of the VC protocol is described in the new Section 11 of the G.707 specification. For higher order VC there is a single H4 byte for each member, whether STS-1 or STS-3c. The H4 byte (refer to Sections 11.1 and 11.2 of G.707) in each higher order path carries a two-stage multiframe number.

microseconds, subject to a maximum of 256 milliseconds. Note that when a VCG is sourced, the combined multiframe number MF11 x MF12 is set equal to zero in every member.

The H4 byte also carries the 8-b long sequence number for the member spread over the MF11.

For low-order VC, the fundamental mechanism is the same. The only difference is that the MFI and sequence information are embedded in the K4 byte, bit 2 (refer to sections 11.3, 11.4 of the G.707 specification).

Performing Inverse Multiplexing

The higher-order VCG may be modeled as a distinct sub-layer in the SDH hierarchy that performs inverse multiplexing (refer to 5.3.3.1.2 of the G.805 spec). The inverse multiplexing sub-layer itself consists of a trail termination and an adaptation function for interleaving/de-interleaving the compound signal to/from the 'X' individual server layer VC-n (n = 3,4) trail termination functions (X individual containers in the VCG for transport) **[Figure 5]**.

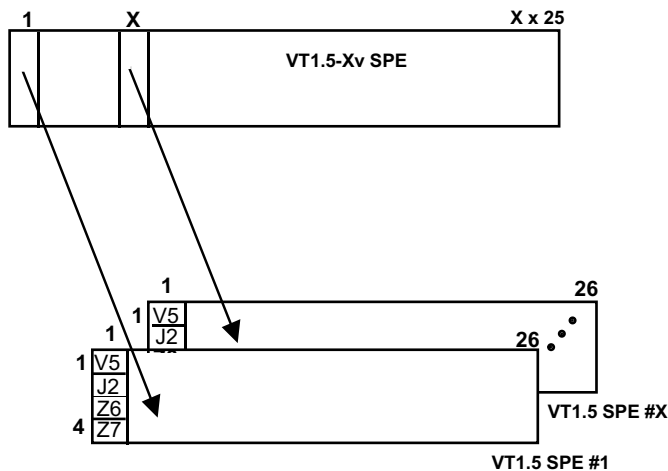


Figure 4: X virtually concatenated VT1.5 SPE structure (lower-order VCG)

The stage 1 multiframer (MF11) is based on a 16-frame count while the stage 2 multiframer (MF12) is based on a 256xMF11 count. This combination of MF11 and MF12 multiframer values then counts the phase delay of any one member with respect to another in steps of 125

The overall VC function shown on the left of Figure 5 combines 'X' VC-n's into a VC-n-X composite, which is then decomposed on the right side of Figure 5. The detailed VC function is comprised of the individual transport end VC-n (n = 3,4) trail termination functions, the VC-n/VC-n-X adaptation that handles byte interleaving/de-interleaving, the composite VCG trail termination on the client end that manages alignment, and reordering functions. VCG trail termination alarms may refer to Loss-of-Alignment or sequence

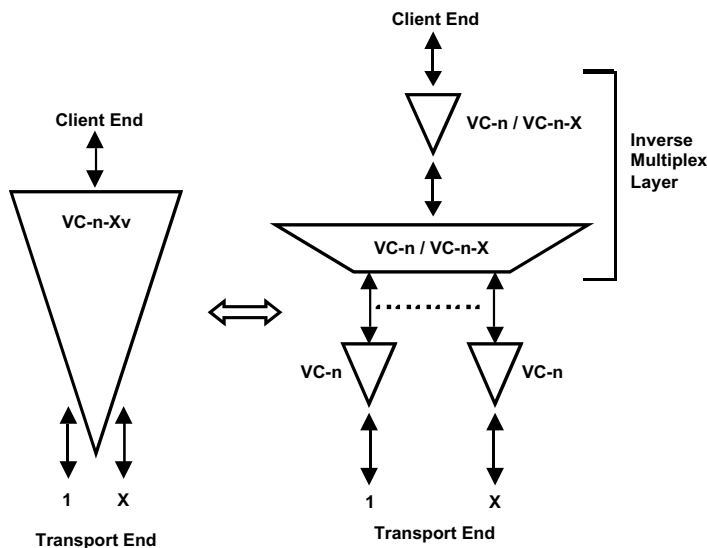


Figure 5: Basic architecture of the VC function

mismatch.

The low-order VC process is similar. Refer to section 5.3.2 of the G.803 spec for more on VC functional architecture.

Tackling LCAS

Defined by the G.7042 specification, LCAS is a means for the source and the sink to synchronize during addition or deletion of members to a VCG such that payload de-adaptation at the sink end may be done hitlessly under non-defect conditions. LCAS functionality can also restore temporarily unavailable members hitlessly. The synchronization mechanism is necessary because of the varying delays that each member of a VCG may incur.

Under the LCAS scheme, the management layer, as shown in Figure 2, issues add/remove commands for a given member separately to the source and sink ends of a VCG. The LCAS state machine at the source end then

signals to the sink end that it is ready to add a particular member to the VCG.

The LCAS sink end state machine then checks the “to be added” member for trail failures and signals to the source end that it is ready via an acknowledgement

signal. The source end then signals the start of the payload change and initiates the actual change. This entire “handshaking” process between the two ends takes place via the H4 byte for higher order VC.

Again, the LCAS state machine operation is identical for both higher and lower order VC. Note that for lower-order VC the LCAS handshaking process takes place via the K4 byte.

The importance of LCAS is pretty simple in a system architecture. By dynamically altering the bandwidth of SONET/SDH transport pipes, LCAS allows designers to adjust bandwidth based on quality of service (QoS) or other priority considerations.

On to GFP and LAPS

That wraps up Part 1 in our EoS tutorial. In the next issue of communicator, we will cover GFP and LAPS. Also, if you'd like more

information on the ITU specs, visit www.itu.ch.

Reference:

1. “Efficient Ethernet Data Transport over SONET/SDH Using Virtual Concatenation”, Mugica, D., Terradillos, E., and Areizaga, E., International Conference on Emerging Telecommunications Technologies and Applications (ICETA 2001), October 2001.

About the Authors:

Harpreet Chohan is manager of customer engineering at TranSwitch Corporation. Harpreet received his MSEE from Imperial College of Science and Technology, University of London, England.

Suvhasis Mukhopadhyay is a technical leader at TranSwitch. Asis holds a BSEE and MSc in Mathematics from Birla Institute of Technology and Science in Pilani, India. He also holds an MSEE from the University of Connecticut.

Rob Schwaber is product marketing manager for Ethernet/SONET products at TranSwitch. He has a BS Computer Science from Union College, Schenectady, NY.

This article was originally published by CommsDesign.com and reprinted with permission.

Copyright 2002 (c) CMP Media, LLC



Register to receive the
communicator online:
www.transwitch.com/subscribe