White Paper

Next-Gen SONET and SDH

Testing and Troubleshooting VCAT, LCAS, and GFP

By: Doyle Mills, Product Marketing Director, Digital Lightwave

Table of Contents

SONET and SDH have been the workhorses for international telecommunications carriers for over a decade. With network deployments well into the millions of dollars, they have been successful and have proven themselves to be more than capable of the promises that were made. The networks of today are reliable, robust, scalable and (mostly) flexible in their ability to adapt to different types of traffic in extremely high volume.

As SONET/SDH networks today carry a great deal of network traffic around the globe, it is appropriate to ask, "Is SONET/SDH a plan that came together, or not?"

SONET/SDH History Lesson

SONET and SDH were developed through the second half of the 1980s to be THE transport protocols used worldwide. Years of discussion, trials and experiments went into the development. Why then have we recently seen the emergence of the next-generation or colloquially "next-gen" SONET and SDH in the network? Quite simply, SONET/SDH was part of a plan that didn't quite come together.

A little-remembered fact is that SONET and SDH were envisioned as only the transport function of a future network. The bold future plans called for segregation of functions such that SONET/SDH would only have to be concerned with the transport of large volumes of data. More complex data routing, efficient use of bandwidth etc. would be handled by another technology developed at the same time, ATM (Asynchronous Transfer Mode). This was stated in the white papers and committee reports of the day. SONET, SDH and ATM were inseparable in the far-reaching plans of the day. Thus, some limitations were knowingly built into SONET and SDH to optimize their transport functions at the expense of other functions.

What are these limitations?

1. SONET traffic is carried in fixed bandwidth groups. The lowest data rate that can be carried on SONET is 1.544 Mbps. Slower rates can be carried but it requires tying up that much bandwidth even if less bandwidth were needed. The rates go up from there (approximate values): 2Mbps, 3Mbps, 6Mbps, 8Mbps, 34Mbps, 45Mbps, 139Mbps, 155Mbps, 622Mbps, 2.5Gbps, 10Gbps. Rates between those values could only be carried by using the next higher bandwidth, an inefficient arrangement. For example, the popular GigE circuits (at a bandwidth of 1Gbps) can only be carried on 2.5Gbps, thus wasting over half the available bandwidth.

2. SONET has no built-in capability of dynamically shifting bandwidth usage. Toward the goal of efficient use of bandwidth, it can be very advantageous to shift usage based on time of day or other factor. For example, a financial institution or merchant may need bandwidth only during business hours. If he were connected directly to a SONET/SDH circuit that bandwidth would be tied up all the time, even during nights and weekends.

3. Monitoring, error detection, etc. in SONET and SDH are extensive and are more than adequate for their intended transport function. Yet when smaller data streams are sent within the network, they are not directly monitored by the SONET and SDH equipment.

All the above limitations were handled very well, in the envisioned future network, by ATM. And, in fact, ATM was very capable of performing the functions. Regarding number one above, an ATM stream at any SONET/SDH compatible rate could carry dozens, hundreds or thousands of individual customer circuits, constantly shifting and prioritizing to squeeze out every last possible drop of bandwidth. ATM networks automatically compensated for number two above with the same function of constant shifting/prioritizing. If a daytime customer did not send any data during the night, he automatically did not use any bandwidth. Regarding number three above, there was an organized hierarchy of monitoring that would make it very easy to locate trouble in the imagined future network. Conditions were organized into "flows".

- SONET Section or SDH Regenerator Section = Flow 1
- SONET Line or SDH Multiplexer Section = Flow 2
- SONET or SDH Path = Flow 3

• ATM Virtual Path = Flow 4
• ATM Virtual Channel = Flow 5

These were precise technical terms, coughed up by the best engineers worldwide, yet never made it into common use. These terms may still be useful at this point in understanding the original plans and how the networks did evolve.

The Downfall of ATM

Why didn't ATM solve all the shortcomings and take its rightful place as the ubiquitous data format carried on SONET/SDH? There is no certain answer to this question. However, there are at least two significant factors that clearly had an effect. First is that ATM itself is not very efficient in its use of bandwidth. In ATM all payload data is separated into groups of 48 bytes called "cells". Each of the cells has a 5-byte header. The header is overhead. Right away, almost 10% of the bandwidth is lost. Certain types of data required even more ATM overhead. The second factor is just plain marketing. While not really being direct competitors, other protocols which performed some ATM-like functions crept in and grabbed market share. This includes IP, PPP and many others. Companies who bought into any particular protocol were unlikely to change.

ATM, intended as "the" format for data transfer, quickly became "one of" and interestingly enough was not in the favored category. IP, with its variable packet size and more efficient use of bandwidth, has emerged as one of the winners. Other protocols are common as well and must be carried by the networks of today.

ATM is not dead but is not implemented broadly enough to fulfill its intended role. Government had a hand in the development of ATM and remains one of its largest users. When we saw pictures of the Martian landscape, they had traversed millions of miles in the form of ATM cells. Battlefield pictures in real time are transported daily around the world using ATM. Yet, for most networks, ATM will not come to the rescue.

In Defense of SONET/SDH

Some textbooks and learned writings of today are blissfully unaware of the true intentions and the hope that existed for SONET and SDH in its formative stages. They define SONET and SDH with a bit of disdain, as if it is some kind of dated solution. One very learned and (otherwise) very excellent white paper published in the early 2000's contains the following line, "The protocol was developed primarily to transport voice traffic as fixed rate timeslots through the network, thereby leaving it inefficient for carrying bursts of traffic that are characteristic of data networks".

In truth, the fixed-rate-timeslot capability of SONET and SDH were intended only as an interim capability, to make it easier to integrate them into the networks of the day. That format was not intended to be their final usage. And, in defense of what some consider a dinosaur; SONET and SDH have been able to keep up with the times with some very clever new features, described below.

The Introduction of Next-Gen

As a pure matter of economic survival, the designers and builders of SONET and SDH equipment and networks, in liaison with the international standards bodies, have been forced to adapt. New features have been added to the old workhorses to enable them to effectively meet the needs of today's network.

These new features are often collectively referred to as "Next-Generation" or "Next-Gen", with or without hyphenation. This may or may not be a precise technical term but is in common use.

Next-Gen features do not increase the bandwidth of a SONET or SDH signal. Rather they render the signals more flexible, allowing a more efficient use of the available bandwidth.

Virtual Concatenation (VCAT)

Understanding how concatenation (without the "virtual") works in SONET may help to understand virtual concatenation.

The term concatenation has been a part of SONET and SDH since the very early days. The basic building block of SONET is the STS-1 (or OC-1 if on an optical interface) which operates at 51 Mbps. In SDH the basic building block is the STM-1 (at 155Mbps). Higher-rate SONET signals (OC-12, OC-48, etc.) and higher-rate SDH signals (STM-4, STM-16 and the like) are built up from combining the basic rates of STS-1 and STM-1. This places an artificial limit on the data that can be carried on the higher limit. An OC-192 could carry data at a rate of nearly 10 Gbps yet could only carry it in groups of 51 Mbps. Concatenation solved that by combining groups of the STS-1 timeslots into higher speed groups.

In SONET the standards defined how to combine 3, 12, 48 or 192 STS-1 timeslots into one "super-rate" path for carrying higher data rates. These were designated by placing a lower case "c" at the end. For example, STS-3c, STS-12c, STS-48c and STS-192c are the concatenated mappings in SONET. Sometimes the "c" was added to the name for the optical carrier, e.g. OC-12c and OC-192c though it was not technically correct to add the "c" to anything other than a mapping type. In SDH, general usage followed the correct terms much more frequently. The concatenated mappings in SDH were AU-4-4c, AU-4-16c and AU-4-192c. Only a limited list of rates for concatenation were defined, making it impossible to do arbitrary combinations such as an STS-2c, or STS-4c, etc.

Having concatenated mappings requires that all the equipment in the network be capable of handling them. The data had to be handled as just one path all the way through in order to function properly. This was challenging for two reasons. The first is that most equipment requires manual provision or set up to operate with concatenation. This entered in the factors of additional labor time and the possibility for human error. Also, some SONET and SDH equipment has been designed, built and sold without the added capability of concatenation. In today's world, nearly 100% of the network equipment does support concatenation. Yet there are still networks built from the early equipment or built without the option of concatenation.

Virtual concatenation gets around these problems and enables a better, more flexible use of bandwidth and, very importantly, can work on much of the existing network equipment that did not support concatenation.

Special equipment is needed at each end of the SONET or SDH path to enable virtual concatenation but the rest of equipment through the network does not require any special upgrade or provisioning.

In virtual concatenation, the user can elect to take any 2 or more STS-1s or STM-1s mapped through a network and use them to transport a higher rate signal. For example, two STS-1s can be used to make a signal of approximately 300 Mbps. This 300Mbps signal would enter the SONET network as a continuous data stream. The special equipment would split the data into two streams, sending each one down an STS-1 path through the network. At the receive end, the equipment would combine the two streams into the original single 300 Mbps data stream. Each path is called a "member".

A path created in that manner would be termed an STS-2v, with the "v" representing virtual concatenation. Any number of virtual paths are possible, with no theoretical restriction on which STS-1s are chosen for the purpose.

As the two (or more) STS-1s and STM-1s are carried completely separately there is often a matter of differential delay between the two streams. For example, different routing may cause data in one stream to arrive a few micro-seconds or milli-seconds later than the other. A special mechanism exists to compensate for this differential delay and to make sure the data is reassembled in the correct sequence. The circuitry is designed to compensate for 256ms of differential delay between the lowest and fastest streams of the Virtual Concatenation Group (VCG).

The usefulness of virtual concatenation can be readily seen by examining some of the common rates that must be carried on today's networks. The popular Ethernet signals, which are typically found at 10 Mbps, 100 Mbps, 1 Gbps or more recently 10

Gbps, is particularly a bad fit for the original SONET world. See below:

10 Mbps Ethernet

51 Mbps STS-1 or STM-0

100 Mbps Ethernet

155 Mbps STS-3 or STM-1

622 Mbps STS-12 or STM-4

1 Gbps Ethernet

2.5 Gbps STS-48 or STM-16

Finally, a rate that nearly fits: 10Gbps STS-192 or STM-64, Ethernet 10Gbps

As can easily be seen, it would not be efficient to carry Ethernet on SONET or SDH directly. A 1 Gbps Ethernet signal would have to be carried very wastefully on a 2.5 Gbps SONET or SDH signal.

However, with virtual concatenation, the 1 Gbps Ethernet signal could be carried on about 20 STS-1s efficiently. 51 Mbps x 20 = 1.02 Gbps. Any combination is possible, creating a very flexible scenario.

Virtual concatenation has proven to be very useful and is being widely implemented across the globe at the time of this writing.

Does VCAT have any shortcomings? There is at least one. If even one of the members in the group fails, the entire group is unusable. Provisioning multiple STS-1s or STM-1s through a network still involves some human interaction, which can be translated to mean that an occasional mistake will happen. Also equipment failures can result in a failure of one or more members and bring down the entire group. See below under the LCAS description to find out how this shortcoming may be addressed

VCAT Testing Considerations

Testing VCAT is generally quite simple, provided that one has test equipment available for whatever rate is in use.

In many cases it is very useful to test each and every one of the separate STS-1 or STM-1 members to confirm that each can traverse the network without error. They can be tested one at a time or, if direct access is possible, can be tested in a “daisy-chain”, with a single test signal running back and forth through the entire chain.

If the separate component signals have been confirmed as above, one should be able to virtually concatenate any number of them together. In theory any problem seen (error, data reassembled out of order, excessive delay) must be caused by the device that is performing the virtual concatenation function.

This is a workable process but not perfect. At least two conditions on the component signals can cause a failure of virtual concatenation even if they tested out fine individually. As stated above, virtual concatenation is designed to compensate for 256ms of differential delay between the slowest and fastest paths. A differential delay that is greater than 256ms will result in failure. Also, jitter and wander on the STS-1s/STM-1s may result in errors. If either is suspected it will be necessary to test latency and jitter on the individual paths.

In order to implement a system that will be robust and reliable into the future, it is imperative to stress the system’s ability to compensate for differential delay and jitter. Test scenarios should be set up where varying delay amounts, up to and including 256ms are intentionally placed into the different paths. Only doing so will make certain that the network will handle these challenges without failure.

Testing every possible VCAT combination is probably impossible or impractical most of the time. In developing the equipment itself, it is very important to verify every possible combination. In later manufacturing test, field implementation and ongoing maintenance, just doing several but not all combinations is usually sufficient.

Link Capacity Adjustment Scheme (LCAS)

LCAS is functional only as an additive to VCAT. LCAS contains two major functions, one of which addresses a shortcoming of SONET/SDH and one of which enhances the reliability and usefulness of VCAT.

One of the shortcomings listed for SONET/SDH is that they had no way of varying bandwidth distribution based on factors such as time of day or even a simple matter of whether a customer were actually sending data at any given time. LCAS makes this possible. The system can be set to provide a customer 10 STM-1s worth of bandwidth during the day and scale back to 1 STM-1 or even 0 STM-1s at night. This bandwidth allocation can also be done manually to route traffic around failures or to handle disasters.

The other feature of LCAS addresses the aforementioned shortcoming of VCAT. It is an ability to recover from failures of one or more VCAT members. If, for example, a VCG were made up of 9 members, a failure of any one member would result in a failure of the entire circuit. However, with LCAS enabled, the circuit can continue to operate with the bandwidth available in only 8 members. The data rate will be impaired but in most cases the data is still valuable and the circuit still usable.

LCAS Testing Considerations

LCAS testing would normally be integrated into a test plan for VCAT containing the conventions listed in the VCAT section. Additionally one should set up a series of automatic capacity adjustments to confirm that the addition/reduction in capacity does not result in errors or lost data.

A series of tests should be performed to assess what happens when individual VCAT members are lost. In this case, errors at the moment of the loss are acceptable but the equipment must quickly adapt to the new lower data rate and operate without errors. Part of the sequence would be restoring the members and confirming that the higher data rate is restored in the specified amount of time. The specifications for data rate adjustments can be found in the ITU-T publication G.7042.

Generic Framing Procedure (GFP)

GFP is defined in the ITU-T publication G.7041. The document defines a set of methods for framing different Layer 1 and Layer 2 data traffic types -- such as Ethernet, PPP, IP, Fibre Channel, ESCON, and DVB -- into a common structure for transport over SONET networks. The protocol is used in deploying metropolitan Ethernet services and bridging customer-facing Ethernet connections with the carrier-facing SONET network.

Other schemes have existed in the past to accomplish similar goals, notably POS (Packet over SONET/SDH). GFP has some key advantages, notably the ability to dynamically define the amount of bandwidth used for a given circuit.

GFP Testing Considerations

Testing GFP is mainly the same as testing any form of packet data. Various data rates, packet sizes, etc. are forwarded through the network to determine that the equipment or network under test meets its specifications.

Additionally it can be very useful to perform protection switching on the SONET or SDH system and determine the effect on the encapsulated data. The SONET/SDH is required to complete any protection switch in 50ms. It is important to know how the data carried on GFP will be affected by such a switch.

Summary

SONET and SDH have been commercial and technical successes. With these new enhancements and more to come, they will continue to endure and continue to provide efficient, robust and reliable service well into the future.

The visionaries, designers and standards bodies have done their jobs well. What remains now is the execution. Sound network planning, logical and thorough test plans fully executed and diligent ongoing maintenance efforts can truly make the difference in future successes and failures.

Digital Lightwave Inc. provides a full range of testing solutions for next-generation SONET and SDH. Visit www.lightwave.com or call +1 727.442.6677 for more information.

Digital Lightwave®
www.lightwave.com

Americas
Corporate Headquarters
15550 Lightwave Drive
Clearwater, FL 33760
Toll free: +1 877 442 DIGL
T: +1 727 442 6677
F: +1 727 442 5660

Europe/Middle East/Africa
Lakeside House
1 Furzeground Way
Stockley Park, UK
Uxbridge UB11 1BD
T: +44 (0) 02086 223023
F: +44 (0) 02086 223200

Asia/Pacific Rim
Braeside Grove
Unit 4, Sibthorpe Street
Braeside, Victoria,
Australia 3195
T: +61 3 9587 4900
F: +61 3 9587 4990

Asia/Pacific Rim
Room 308A, Tower B
He Qiao Bldg, Guanghua Rd.
Chao Yang District
Beijing, China 100026
T: 8610 65815317/5327
F: 8610 65815327

Digital Lightwave provides industry-leading products, technologies, and services for deploying and managing communications networks. With a presence in more than 80 countries, Digital Lightwave enables customers to successfully implement global communications networks worldwide. To find the nearest sales office, please visit www.lightwave.com.

© 2004 Digital Lightwave, Inc. All rights reserved. Digital Lightwave, and its logo are trademarks or registered trademarks of Digital Lightwave, Inc.

ISO 9001 CERTIFIED

NEXTGEN 7-2004A