Understanding ATM

By Jim Boston

Broadcasters have developed a bag of tricks over the years for moving audio and video from one location to another. Initially, the process was quite simple; the phone companies ran long runs of coax. Every so often, the video on those lines was amplified and re-equalized to restore it to a semblance of what was originally sent. The baseband nature of the video through coax meant that the signal quickly degraded over distance. Microwave systems were soon employed. The long distance carrier (it was only AT&T in the beginning) would frequency multiplex both voice and television over the same path. This path was the best transport option for broadcasters for more than 30 years. In the mid-1960s, satellite transmission was experimented with on transoceanic relays with the advent of the Telstar satellite. In the late 1970s, satellite use made the super-station and the cable channel possible. By the mid-1980s, the networks distributed their content via satellite. Local stations used satellites for back-hauling news. When microwave and coax were the only tools available for transport, most local and long-haul carriers didn't court the broadcasters' business; thus broadcasters learned how to use other transmission methods to accomplish long-distance program delivery.

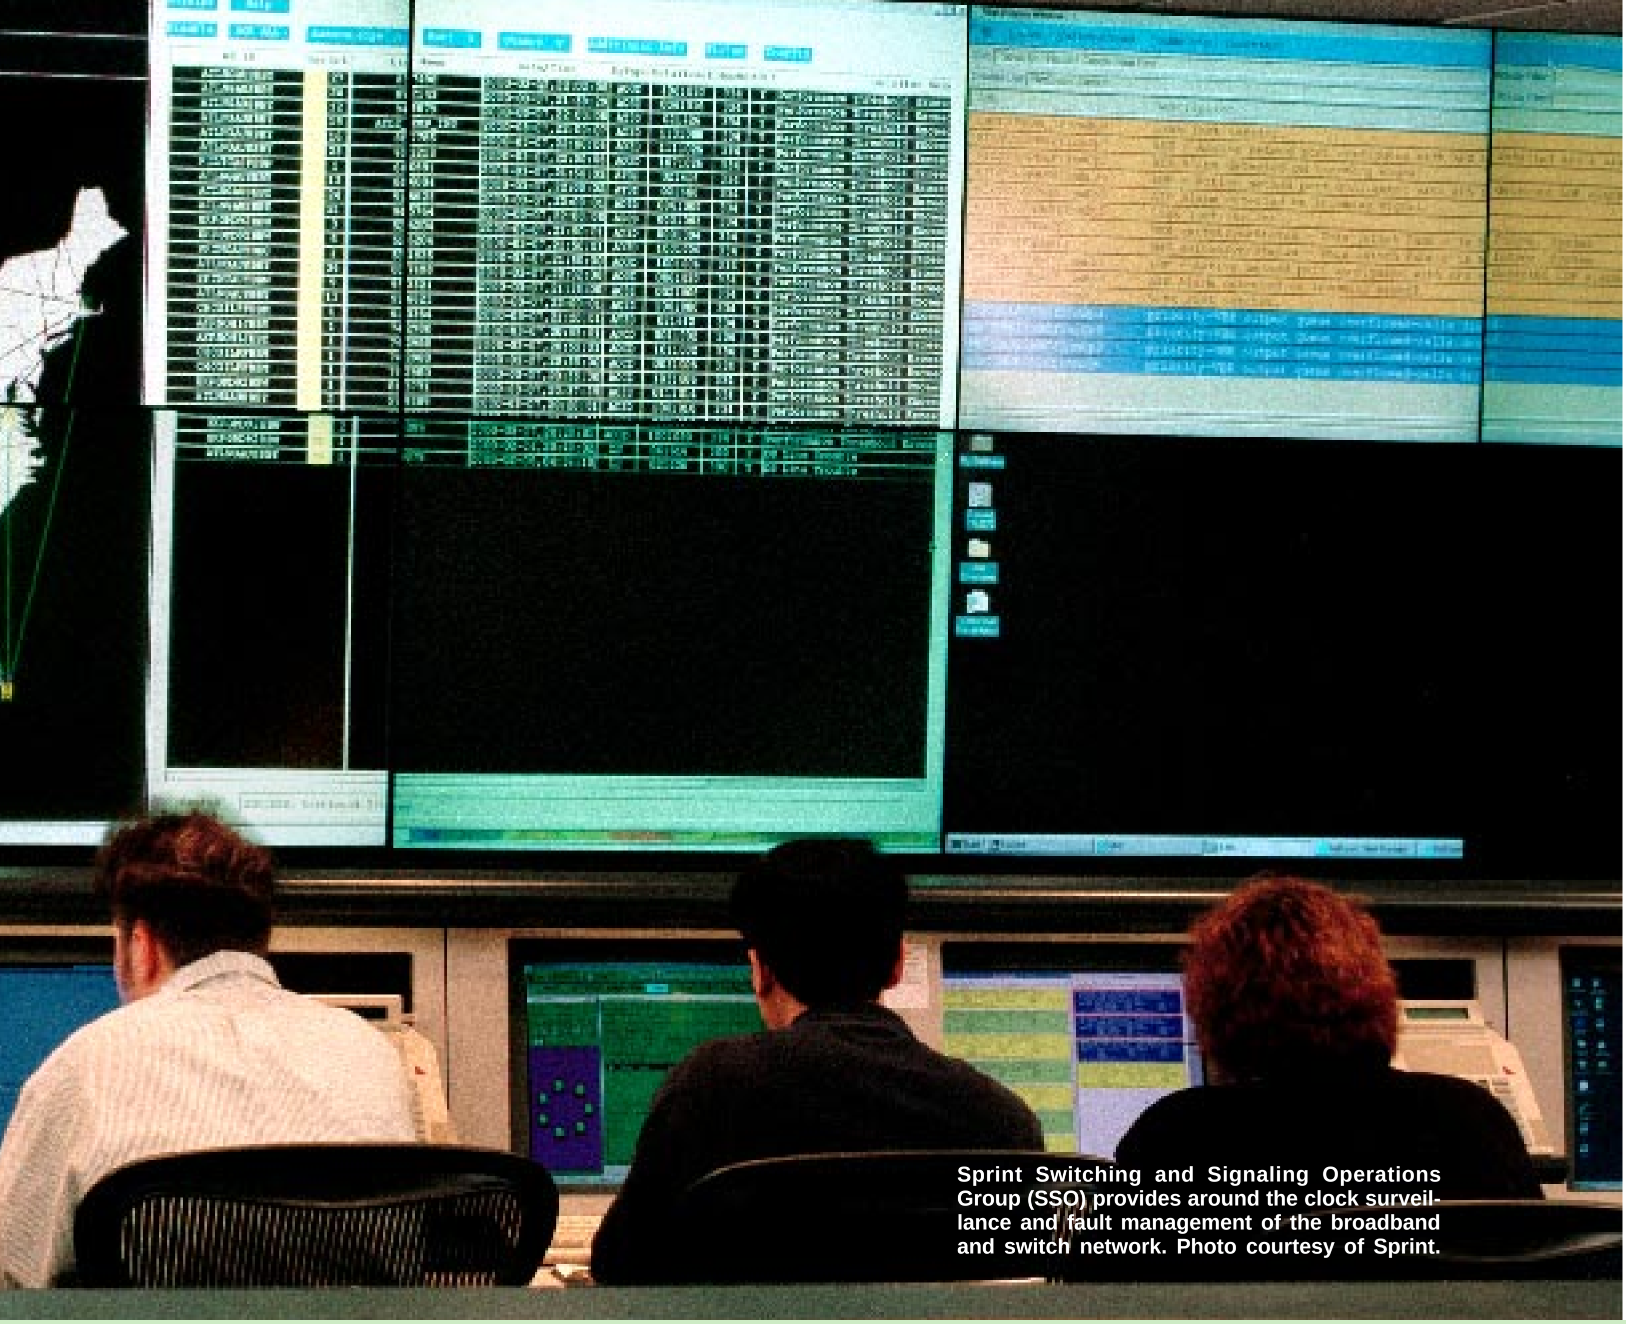

Sprint Switching and Signaling Operations Group (SSO) provides around the clock surveillance and fault management of the broadband and switch network. Photo courtesy of Sprint.

Telco tools

Although not interested in the video business through the 1980s, the telephony industry developed a large tool set to handle the emerging demand of moving data over long distances. With the large amounts of data carried by the Internet, the movement of data by the long-haul carriers is starting to rival the movement of voice. Because long-haul backbones are digital in nature, long distance voice traffic has been digital for years. This digital migration of voice traffic is making its way ever closer to the average home. Those homes and businesses with ISDN/xDSL already have a digital tentacle of that system. Now that the television industry is realizing that video and audio streams can be thought of as just another type of data, we can start to use the same data tools offered by the telephony people.

These tools come in three types: Frame Relay, IP, and ATM (asynchronous transfer mode). We will touch on Frame Relay and IP briefly here, but ATM is the main subject. These protocols use virtual paths, meaning a connectionless path. Instead of a physically switched and dedicated path between two locations, the data from one location is merged with other traffic headed the same direction via time multiplexing. Another common point of confusion is that IP (Internet Protocol) packets can be inserted as the data payload in ATM cells. Packets and cells are different names for specified chunks of data sent at a time. IP and Frame Relay call these chunks frames, whereas ATM calls them cells. These protocols emerged because long distance data traffic evolved from being mainly text-based to being graphics-based. Data traffic can come in bursts

Understanding ATM

and many data users require some guarantee that the data will arrive at its destination in a timely manner. Some earlier protocols existed, namely X.25, which was developed when paths were still largely analog. Because of this, X.25 had error correction capabilities not needed with digital paths. X.25 was an asynchronous system, whereas the newer protocols for the digital networks are synchronous. Additionally, an error-correction scheme known as TCP (Transport Control Protocol) was developed which wrapped the user's data, which could then be wrapped a second time in IP.

Frame Relay uses paths called Permanent Virtual Circuits (PVC) to connect end users. Many other data sessions share different links of the path, but the data carrier provisions sets up the virtual paths to be there all the time, whether the user has data to send or not. If no data was sent, the time slots (or more exactly packets) devoted to that data customer would be empty. This simplified the Switched Virtual Circuit (SVC) used with X.25. A SVC can be thought of as a phone call for data. The virtual path is set up only when data needs to be sent. One trade-off is that some time is spent setting up the call. The advantage of Frame Relay is that its frame lengths are long, generally 128 octets (1 octet = 8 bits) and up. This means that for lower bit rate paths there is less overhead. The tradeoff is that data requiring fast access to the network has to wait longer for its turn. ATM has shorter cell lengths in part to solve this.

IP is much cheaper to implement than ATM. But it is good for carrying time-sensitive material only in uncongested networks. IP traffic outside the carrier's SONET backbone is routed via routers, which are cheaper than switches. Routers are generally software-driven devices, making the propagation delay through them longer than the hardware-oriented switches used for ATM. IP doesn't generally have any Quality of Service (QoS) guarantees which are available with ATM (although some IP router manufacturers are implementing them, but they increase the overhead). QoS has various levels, as we will see. These levels provide assurances of bit rates, propagation delay through the ATM cloud, and the amount of allowable jitter. ATM is good for moving video through congested networks where QoS levels along the path need to be tunneled out. Most of us now are aware that TCP (transport layer in the OSI stack) usually rides on top of IP (network layer). But often User Datagram Protocol (UDP) is used in place of TCP, as it requires no re-send of cells that are lost (something that TCP requires). The re-sending of lost cells and the wait to assemble the re-sent cells in the proper order greatly hampers the high bit rate/real-time nature of television bitstreams.

Illuminet's Network Surveillance Control Center, Overland Park, KS. Control centers such as this monitor a network of public-switched telecommunications infrastructure, helping to ensure an even spread of traffic and the highest access speeds available. Photo courtesy Illuminet.

The telephony backbone

What actually comprises an ATM cloud? To use the ATM service of a backbone carrier, you need to gain access to the carrier. The access usually consists of using the Incumbent Local Exchange Carrier (ILEC), normally the local Baby Bell or a Competitive Local Exchange Carrier (CLEC) to provide connectivity from your facility to the ATM carriers Point of Presence (POP). Once your data is at the POP the ATM carrier charges a port charge. This is a subscription into the ATM cloud. One interesting thing about ATM tariffs is that you only pay for the amount of data inserted into the cloud, and not how far that data has to travel.

The access leg from your facility to the POP is often DS-3 (DS = Digital Service). DS-3 (often referred to as a M13 frame from a M13 multiplexer) when used for voice traffic, consists of seven DS-2 signals. Each DS-2 consists of four DS-1 (or T1 — which can carry 24 voice channels) signals. Hence a M13 multiplexer takes 28 DS-1s (the "1" part of M13) and muxes them into one DS-3 (the “3 “part of M13). When DS-3 is to be used to transport ATM, the ATM data is encapsulated into a Physical Layer Convergence Protocol (PLCP) frame. The PLCP frame has twelve 53-byte ATM cells plus framing, and parity information. PLCP allows for 40.7Mb/s of ATM cell data to be mapped into the 44.21Mb/s DS-3 payload rate. DS-3 is what usually arrives at the POP. At the POP the DS-3 is mapped onto a synchronous optical network (SONET) ring. A SONET ring’s capacity is measured in Optical Carrier (OC) rates. An OC-1 ring can carry a DS-3 data stream. Our PLCP frame is mapped into an STS-1 (Synchronous Transport Signal) frame. An STS-1 frame consists of data arranged in 90 columns with nine rows. This represents 810 octets of data. But the first three octets of each row, or 27 octets, are used for overhead. 8000 of these frames are sent per second, therefore each frame is 125 microseconds long. The 27 octets of overhead are used to point to where valid data starts in the frame, as the start of data does not always match the start of the STS frame (we will see why shortly). Also, control traffic and message traffic can be sent over the path using these overhead octets. The 783-byte payload can be used to carry ATM, or other types of traffic.

Originally DS-2 and above used Plesiochronous Digital Hierarchy (PDH). This means nearly synchronous, with multiple levels of multiplexing (DS-0 into DS-1, into DS-3). These are considered nearly synchronous because specific bits of information found in one frame may reside in a different location in the next frame. This is due to bit stuffing. As an example, PDH would be required if DS-1 channels from far-flung locations were being multiplexed into a DS-3 stream. The propagation time from the various DS-1 sources to the DS-3 multiplexer could vary as the weather changes, or other things slightly change along the path (sun rises and warms lines causing them to expand, lengthening the route). If the signal was taking a little longer than before to arrive, an extra bit would be added to make up for the delay. Therefore the size of the frame would vary. The converse would be true as the sun set, or went behind a cloud. DS-1 was always considered synchronous by telephone engineers because a particular bit for a particular channel is always found in the same place of a DS-1 frame. Telephone engineers have moved away from PDH and it’s bit stuffing approach towards the use of pointers. This pointer is placed in the header of a frame, and it tells the receiver how far into the frame the data starts. Because of this, Synchronous Digital Hierarchy (SDH) is now used at all levels in new installations. Almost all SDH signals travel over fiber, except when undergoing switching (but that is starting to change).

One DS-3 bitstream is inserted into an STS-1 frame, which in turn is inserted into an OC-1 (Optical Carrier) stream. An OC-1 stream is 51.840Mb/s with about 8Mb/s overhead on top of the 44.21 DS-3 stream. OC-3 (155Mb/s) carries 3 DS-3 streams, or 84 DS-1 streams, or 2016 DS-0 (voice) streams. A concatenated mapping where all the bandwidth is given to a single user (no digital hierarchy at all) is possible, such as to ATM users (a lower case c indicates a concatenated frame, ex. STS-1c).

OC-12 carries 12 DS-3 streams (622Mb/s), OC-48 carries an STM-16 bit stream, which handles 48 DS-3 streams. OC-192 is currently the highest bitrate. OC-192 can carry 192 DS-3 streams.

SONET rings

On a SONET ring, you often find different kinds of traffic. Some carriers have voice, IP, and ATM traffic all traveling over the same ring. Although the different traffic can be thought of as separate virtual networks, they all travel over the same physical fiber ring in STS frames. The carrier provisions the ring so that part of the bandwidth (time) is devoted to each segment on the ring. In addition, some of the bandwidth will be devoted to express paths on the ring. If two nodes on opposite sides of a ring had a lot of traffic between them, a nonstop path between them would be setup. SONET works on a drop and add system. At each node, time slots that are programmed to be dropped from the SONET ring for local distribution are de-multiplexed out of the ring, and new traffic destined for long distance delivery is added. But at most nodes, some of the traffic does not undergo the drop-and-add process as it is intended for nodes farther along the ring. This is sort of a local and express

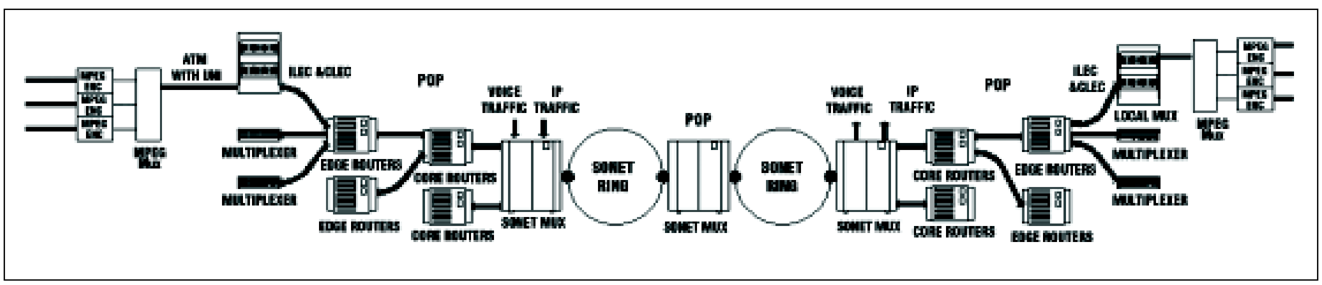

Figure 1. There are a variety of paths through the ATM cloud, but typically data will enter through an encoder, pass through some routers, through a POP and onto a SONET ring. It will then traverse one or more rings, be dropped into the POP nearest the destination and be routed on to its final destination.

train analogy. Some traffic only has to ride the ring for a couple of stops, while some have a long way to travel on the ring and it is inefficient to stop at each node.

Most traffic is destined for places not on the closest ring. The traffic must be handed off from ring to ring to complete its journey. Places where the rings are tangent to one another usually have a super POP. At these super POPs, traffic needing to transfer to a different ring is riding in a STS frame that is dropped off the ring. It is demultiplexed and routed if it is IP, or switched if it is ATM to a multiplexer that puts it back into a STS frame. It is then placed on the new ring. If the traffic is destined to ride additional rings, it probably will be placed into an express STS frame that will not undergo the add/drop process until it reaches the super POP tangent to the next ring to transfer to.

To reach its destination, most traffic will need to travel over multiple backbones owned by different entities. This means that the data must be handed off from one company's backbone to another. This is done at Network Access Points (NAP). NAPs are usually hosted by a telecom provider. Some handle all types of traffic, some just IP traffic. There are major NAPs in New York, Chicago, San Francisco, and Washington D.C. Although hosted by one backbone provider, many other companies can subscribe to a Service Interface (SI) at a particular NAP. The Chicago NAP, which is based on ATM and hosted by Ameritech, has 100 SI connections. These connections are to other telecom companies, Internet Service Providers (ISPs), universities, and government agencies. These entities hand off traffic to each other at these NAPs. Backbone providers and ISPs that exchange large amounts of traffic with one another often setup peering arrangements by running a path between one provider's backbone and another's backbone. NAPs and peering is what allows end-users using different access and backbone providers to communicate with one another.

When all the needed rings have been negotiated, the data is dropped at the closest POP to the destination. It then undergoes local switching/routing to arrive at the destination (See Figure 1).

AAL

ATM is intended to be used by different types of traffic, each with it's own set of requirements. Voice and video tend to produce constant bit rates, although video is usually at a much higher rate than audio. Data traffic is often bursty in nature. This means that although all bits are stuffed into ATM cells, these cells must be handled differently to satisfy the application. The OSI data-link layer is where the software code is specialized to handle the specific application. The data-link layer is what controls the hardware layer. The data (or video or voice) is put into ATM cells in the data-link layer, which is actually split in half. The lower half, known as the segmentation and re-assembly sub-layer, is where ATM cell generation takes place. The upper half, which is known as the convergence sub-layer breaks large blocks of data handed to it from the network layer above it into smaller (often 64kB) blocks and adds error detection and data recovery overhead. Some ATM is implemented without the upper intermediate step, instead relying on upper layer protocols (such as TCP) to handle error recovery. This process at the data-link layer is known as the ATM Adaptation Layer (AAL).

> **The movement of data by the long haul carriers is starting to rival the movement of voice.**

AAL is broken out into five types. AAL-1 specifies Constant Bit Rates (CBR) for real-time traffic. This is often used for voice traffic and WAN applications. AAL-2 specifies Variable Bit Rates (VBR) for real time traffic. AAL-3 is VBR but intended for non-real-time data traffic. AAL-1 through AAL-3 are connection orientated. This is data transmission with a pre-arranged connection. AAL-4 is the same as AAL-3 except it is a connectionless transmission. AAL-4 is used for LAN Emulation (LANE). LANE is when ATM is used to connect remote LANs so that they appear as one contiguous LAN. AAL-5 does not use the convergence data-link sub-layer. This is known as Simple and Efficient Adaptation Layer (SEAL). SEAL has the advantage of using all 48 bytes in the payload section of the ATM cell. Other ATM AALs use four of those bytes for convergence sub-layer overhead. Like the MPEG characteristic of having higher profiles and levels able to handle lower types, AAL-5 devices facilitate all the lower AALs. AAL-5 is used for LANE also. Different AALs are used to guarantee various levels of QoS.

The tariffs for the various levels of QoS can vary greatly. VBR-NT (variable bit rate — non-real time) can cost about two-thirds of what VBR-RT (VBR — real time) costs. The same ratio exists for VBR-RT verses CBR. Therefore, CBR pricing can be two and a half times what VBR-NT costs. PVCs and SVCs are priced differently also. PVC pricing is based on level of QoS and the amount of bandwidth requested, while SVC is based on QoS and the amount of data delivered to the destination. The two different services are usually priced so that PVC is cheaper if used more than 100 hours a month, up to that point SVC is more economical. PVCs have to be provisioned by the carrier. The user can not change QoS levels and bandwidths on their own. With an SVC, the user can change both attributes.

Although access and port charges can run a few thousand dollars a month for most users, the cost of sending standard-definition video (plus audio) isn't much more than a dollar and a half per minute in many cases. Besides the access and port charges, you need an ATM edge device. This can be an MPEG mux that talks ATM. Often the MPEG mux will go to another ATM mux downstream that is handling the merging of all ATM traffic within a facility, such as voice, and LANE activity. But if the MPEG mux is to directly generate ATM it must meet a number of standards:

ATM Cells

The simplest is to simply generate ATM cells. The 53 byte long packets have five octets (or 40 bits) of overhead in the form of a header. The header devotes the first four bits for flow control, although many ATM switches along a path combine these four bits with the next eight which are the Virtual Path Identifier (VPI). Each ATM switch encountered changes the VPI to manage handoff to the next switch. The VPI most likely will not be the same leaving the switch as it was when it arrived at the switch. A Virtual Path can carry more than one Virtual Channel. Separate services having the same source and destination, such as voice, LANE, video, audio, metadata, etc., would travel the same Virtual Path, but each would be separate Virtual Channels. Therefore the next 16 bits in the header are for the Virtual Channel Identifier (VCI). The VCI values also tend to change from switch to switch. The next three bits identify if the cell payload is user data or path setup and maintenance information. One bit is sent to indicate whether this cell has priority if congestion has mandated the dropping of cells. The final eight bits are for header error checking.

UNI

Next on the standards list is knowing how to negotiate an ATM cloud. The standard that encapsulates that is known as the User Network Interface (UNI). This standard was developed and is maintained by the ATM forum. A lot of it is based on various International Telecommunication Union (ITU) standards and recommendations. A common ITU recommendation often quoted in telecom circles is Q2931. A lot of Q2931 was incorporated into UNI version 3.1. The current UNI version is UNI 4.0. Q2931 specifies call setup and termination, VPI/VCI assignment, QoS requests, caller ID, error handling, and peak cell rate parameters. UNI communicates to the network the AAL level that is to be used. UNI also specifies that E164 addressing be used. E164 is essentially the addressing scheme we use when we manually dial a long distance phone number.

In summary, ATM technology has evolved to a state where it is another tool that broadcasters can add to the backhaul and program distribution tool set. Some SONET backbone providers have extremely wide fiber paths in terms of bandwidth. Some are literally terabytes wide. Even though the movement of program material is increasingly going to be non-real time in nature (such as FTP-type file transfers), real-time program movement can today be accomplished using ATM. ■

Jim Boston is the senior project engineer at Scripps-Howard, as well as author of a new McGraw-Hill book, The DTV Survival Guide. *The author would like to thank Qwest, especially Craig Brandetsas, Sylvia McKenna, and Richard Fiutem, and Rich Csontos of Marconi/Fore Systems for their guidance.*

The *DTV Survival Guide* provides readers with experienced guidance in the planning and building new digital facilities. The book also examines the many new opportunities and potential pitfalls that DTV technology brings to the industry. *The DTV Survial Guide* is available from the publisher by calling 800-262-4729 or through many booksellers.

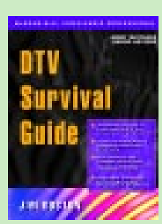

Common telco acronyms

AAL ATM	Adaptation Layer
ATM	Asynchronous Transfer Mode
CBR	Constant Bit Rate
CLEC	Competitive Local Exchange Carrier
DS-x	Digital Service
FTP	File Transfer Protocol
ILEC	Incumbent Local Exchange Carrier
IP	Internet Protocol
ISP	Internet Service Provider
LANE	LAN Emulation
NAP	Network Access Point
OC-x	Optical Carrier
PDH	Plesiochronous Digital Hierarchy
PLCP	Physical Layer Convergence Protocol
POP	Point Of Presence
PVC	Permanent Virtual Circuits
QoS	Quality of Service
SDH	Synchronous Digital Hierarchy
SEAL	Simple and Efficient Adaptation Layer
SI	Service Interface
SONET	Synchronous Optical Network
STS-x	Synchronous Transport Signal
SVC	Switched Virtual Circuit
TCP	Transport Control Protocol
UDP	User Datagram Protocol
UNI	User Network Interface
VBR	Variable Bit Rate (NT-non real time)(RT-real time)
VCI	Virtual Channel Identifier
VPI	Virtual Path Identifier

Conversion Chart

The number of voice channels and how they fit into the various services offered.

1 DS-0 = 1 voice channel

1 DS-1 = 1 T-1 = 24 voice channels

1 DS-2 = 4 DS-1s = 96 voice channels

1 DS-3 = 7 DS-2s = 672 voice channels = 44.21Mb/s (payload)

OC-1 = 51.840Mb/s 1 DS-3 fits into an OC-1, 192 DS-3s fit into an OC-192 (10Gb/s)