

**#1 Best Seller
Over 350,000 Sold**

NEWTON'S TELECOM DICTIONARY

**The Official Dictionary
of Telecommunications
Networking and
the Internet**

**16th
EXPANDED
& UPDATED
EDITION**

BY HARRY NEWTON

NEWTON'S TELECOM DICTIONARY

copyright © 2000 Harry Newton

Email: Harry Newton@TechnologyInvestor.com

Personal web site: www.HarryNewton.com

All rights reserved under International and Pan-American Copyright conventions including the right to reproduce this book or portions thereof in any form whatsoever.

Published by Telecom Books
An imprint of CMP Media Inc.
12 West 21 Street
New York, NY 10010

ISBN # 1-57820-053-9

Sixteenth Edition, Expanded and Updated, February 2000

For individual orders, and for information on special discounts for quantity orders, please contact:

Telecom Books
6600 Silacci Way
Gilroy, CA 95020
Tel: 800-LIBRARY or 408-848-3854
FAX: 408-848-5784
Email: telecom@rushorder.com

Distributed to the book trade in the U.S. and Canada by
Publishers Group West
1700 Fourth St., Berkeley, CA 94710

Manufactured in the United States of America

RingCentral Ex-1015, p. 2
RingCentral v. Estech
IPR2021-00574

NEWTON's TELECOM DICTIONARY

The Official Dictionary of
Telecommunications & the Internet

**16th Updated, Expanded and Much
Improved Edition**

Ionospheric Disturbance An increase in the ionization of the D region of the ionosphere, caused by solar activity, which results in greatly increased radio wave absorption. See ionosphere.

Ionospheric Focusing A variation in the curvature of the ionospheric layers can give rise to a focusing/defocusing effect at a receiving antenna. This may produce either an enhancement or attenuation in the received field strength due to signal phase variations.

Ionospheric Refraction The change in the propagation speed of a signal as it passes through the ionosphere.

IOP 1. Input/Output Processor.

2. Interoperability: The ability of equipment from different manufacturers (or different implementations) to operate together.

IOPS Internet Operators Group. On May 20, 1997 Nine of the nation's major Internet service providers announced the formation of IOPS.ORG, a group of Internet service providers (ISPs) dedicated to making the commercial Internet more robust and reliable. IOPS.ORG will focus primarily on resolving and preventing network integrity problems, addressing issues that require technical coordination and technical information-sharing across and among ISPs. These issues include joint problem resolution, technology assessment, and global Internet scaling and integrity. IOPS.ORG will provide a point-of-contact for these industry-wide technical issues.

The founding members of IOPS.ORG are ANS Communications, AT&T, BBN Corporation, EarthLink Network, GTE, MCI, NETCOM, PSINet, and UUNET, and it is expected that additional national and international Internet operators will join. IOPS.ORG will work with other Internet organizations, with Internet equipment vendors, and with businesses that rely on the Internet. IOPS.ORG members individually will continue to support other Internet organizations such as the Internet Engineering Task Force (IETF), the North American Network Operators Group (NANOG), and the Internet Society.

The Corporation for National Research Initiatives (CNRI), a Reston, VA-based not-for-profit organization which works with industry, academia, and government on national-level initiatives in information technology, will host the initial operations of IOPS.ORG. "IOPS.ORG will play a key role in the healthy technical and operational evolution of the Internet as an increasingly important component of the economy," said CNRI President Robert Kahn. www.iops.org

IOS 1. The International Organization for Standardization based in Geneva, Switzerland. The IOS develops and publishes hundreds of international ISO standards such as ISO-9000 for quality assurance or ISO 14000 for environmental performance. The IOS is comprised of more than 90 member standards bodies worldwide plus other international associations, government and non-government bodies. The U.S. member of the IOS is ANSI. See ANSI and ISO.

2. Internetwork Operating System from Cisco. This operating system runs the vast majority of routers now deployed in the core of the Internet. See also Junos Code.

IOTP See Internet Open Trading Protocol.

IP 1. The Internet Protocol. IP is the most important of the protocols on which the Internet is based. The IP Protocol is a standard describing software that keeps track of the Internet's addresses for different nodes, routes outgoing messages, and recognizes incoming messages. It allows a packet to traverse multiple networks on the way to its final destination. Originally developed by the Department of Defense to support

interworking of dissimilar computers across a network. While its roots are in the ARPAnet development, IP was first standardized in RFC 791, published in 1981, and updated in RFC 1349. This protocol works in conjunction with TCP and is usually identified as TCP/IP. It is a connectionless protocol that operates at the network layer (layer 3) of the OSI model. See IP Address, IPv4, IPv5, IPv6, the Internet, and TAPI 3.0.

2. Intelligent Peripheral. A device in an IN (Intelligent Network) or AIN (Advanced IN) that provides capabilities such as voice announcements, voice recognition, voice printing and help guidance. By way of example, MCI's 1-800-COLLECT makes use of IPs, which are specialized voice processing systems. The IP prompts the caller to enter the target telephone number and speak his or her name. The system then instructs the network to connect the call. Based on a spoken acceptance of the call by the called party, the system authorizes call completion.

3. Information Provider. A customer that offers recorded information on its listed numbers.

IP Address See Internet Address.

IP Address Mask Internet Protocol address mask. A range of IP addresses defined so that only machines with IP addresses within the range are allowed access to an Internet service. To mask a portion of the IP address, replace it with the asterisk wild card character (*). For example, 192.44.*.* represents every computer on the Internet with an IP address beginning with 192.44. See IP Addressing.

IP Addressing A networking term. IP (Internet Protocol) addressing is a system for assigning numbers to network subdivisions, domains, and nodes in TCP/IP networks. IP addresses are figured as 32-bit (four-byte) numbers. The high bytes constitute the "Class A" and "Class B" portions of the address, which denote network and subnetwork. The low bytes ("Class C" address segments) identify unique nodes — individual machines or (in the case of multi-addressing) individual node processes. The Class C address segment (two bytes) can represent 65,536 unique values — enough so that in most conventional TCP/IP LANs, sufficient values are available to afford each machine its own "fixed" IP address. In public internet-access, however, the number of fixed addresses available to a provider may not be sufficient to provide each dialup client with a permanent IP address. In such scenarios, available Class C addresses can be assigned dynamically, as machines log into network access ports — on the presumption that no more than N clients will attempt to log on, simultaneously (where N denotes the number of absolute addresses in the pool). Thus:

"Fixed" or "Static" IP address: a four-byte TCP/IP network address permanently assigned to an individual machine or account.

"Dynamically-assigned" IP address: a four-byte TCP/IP network address assigned to a machine or account for the duration of a single session.

IP Datagram The fundamental unit of information passed across the Internet. Contains source and destination addresses along with data and a number of fields which define such things as the length of the datagram, the header checksum, and flags to say whether the datagram can be (or has been) fragmented.

IP Device Control IPDC. See Simple Gateway Control Protocol.

IP Multicasting The transmission of an IP datagram to a host group, a set of zero or more hosts identified by a single IP destination address. A multicast datagram delivered to all

members of its destination host group with the same "best efforts" reliability as regular unicast IP datagrams, i.e., the datagram is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other datagrams.

IP Router A computer connected to a multiple physical TCP/IP networks that can route or deliver IP packets between networks. See also Gateway.

IP Security See IPsec.

IP Spoofing An attack whereby a system attempts to illicitly impersonate another system by using its IP network address. See IP, IP Address and IP Router.

IP Subnet All devices which share the same network address. Routers are boundaries between subnets so each connection to a router has a different network address.

IP Switching A term coined by Ipsilon Networks to describe a new class of switch it developed, combining intelligent IP routing with high-speed ATM switching hardware in a single, scalable platform. The IP switch implements the IP protocol stack on ATM hardware, allowing the device to dynamically shift between store-and-forward and cut-through switching based on the flow requirements of the traffic as defined in the packet header. Data flows of long duration, thereby, can be optimized by cut-through switching, with the balance of the traffic afforded the default ATM treatment, which is hop-by-hop, store-and-forward routing. Ipsilon suggests that first-generation IP Switches can achieve rates of up to 5.3 million PPS (packets per second) by avoiding ATM cell segmentation and reassembly, ATM overhead, and ATM switch processing of each cell header. Clearly, one of the advantages of IP Switching is the use of IP (Internet Protocol), which protocol is mature, well-understood, and widely deployed across a wide range of networks. Contrast with Tag Switching.

IP Telephony Here is Microsoft's definition, excerpted from their white paper on TAPI 3.0: IP Telephony is an emerging set of technologies that enables voice, data, and video collaboration over existing IP-based LANs, WANs, and the Internet. Specifically, IP Telephony uses open IETF and ITU standards to move multimedia traffic over any network that uses IP (the Internet Protocol). This offers users both flexibility in physical media (for example, POTS lines, ADSL, ISDN, leased lines, coaxial cable, satellite, and twisted pair) and flexibility of physical location. As a result, the same ubiquitous networks that carry Web, e-mail and data traffic can be used to connect to individuals, businesses, schools and governments worldwide.

What are the benefits of IP Telephony? IP Telephony allows organizations and individuals to lower the costs of existing services, such as voice and broadcast video, while at the same time broadening their means of communication to include modern video conferencing, application sharing, and whiteboarding tools. In the past, organizations have deployed separate networks to handle traditional voice, data, and video traffic. Each with different transport requirements, these networks were expensive to install, maintain, and reconfigure. Furthermore, since these networks were physically distinct, integration was difficult if not impossible, limiting their potential usefulness.

IP Telephony blends voice, video and data by specifying a common transport, IP, for each, effectively collapsing three networks into one. The result is increased manageability, lower support costs, a new breed of collaboration tools, and increased productivity. Possible applications for IP Telephony include telecommuting, real-time document collaboration,

distance learning, employee training, video conferencing, video mail, and video on demand. See the Internet, IP Telephony Algorithms, TAPI, TAPI 3.0 and TCP/IP.

IP Telephony Algorithms The major IP Telephony Algorithms in the market today (fall of 1997), according to a white paper, called "IP Telephony powered by Fusion" from Natural MicroSystems (www.nmss.com), include:

- **MS-GSM:** This algorithm, marketed by Microsoft, runs at 13kbps and is a derivative of the ITU (International Telecommunications Union) standard GSM work. GSM is used in 85 countries around the world as the standard for digital cellular communications. Microsoft's implementation varies from the standard in several ways including how the encoded data is represented and what aspects of the encoder are supported. Natural MicroSystems provides an MS-GSM encoder that is compatible with Microsoft's Win95/WinNT embedded product.

- **ITU G.723.1:** This algorithm runs at 6.3 or 5.4 kbps and uses linear predictive coding and dictionaries which help provide smoothing. The smoothing process is CPU-intensive, however (30Mips on an Intel Pentium), so don't expect a PC-based implementation to work well for lots of real-time activity.

- **VoxWare:** This is a proprietary encoder that has been bundled by Netscape with their Browser. It delivers 53:1 compression and very low jitter. VoxWare presents very low network bandwidth requirements; however, it also has lower speech quality.

Most speech encoder algorithms have a set of rules concerning packet delivery and disposition management. This is often called jitter buffer management. "Jitter" in this case refers to when the signal is put into frames. The decoding algorithm must decompress and sequence data and make "smoothing" decisions (when to discard packets versus waiting for an out-of-sequence packet to arrive). Given the real-time nature of a live connection, jitter buffer management policies have a large affect on voice quality. Actual sound losses range from a syllable to a word, depending on how much data is in a given packet. The first buffer size is often a quarter-second, large enough to be a piece of a word or a short word — similar to drop-outs on a cellular connection in a poor coverage area.

IPA Intellectual Property Attorney.

IPackage Installation Package.

IPARS The International Passenger Airline Reservation System. An IBM-originated term.

I-PASS An alliance of ISPs (Internet Service Providers) and IAPs (Internet Access Providers) to provide roaming capabilities for travelers. Based on proprietary standards, roamers are authenticated before being afforded Internet access. Usage is cross-billed through the I-PASS clearinghouse, with fees being set by each ISP for use of its facilities by roamers. I-PASS includes over 100 member ISPs in approximately 150 countries, and includes over 1,000 POPs (Points of Presence). I-PASS competes with GRIC (Global Reach Internet Connection). The IETF's Roamops working group is developing a standard for roaming, as well. See also GRIC and ROAMOPS.

IPC Interprocess Communications. A system that lets threads and processes transfer data and messages among themselves; used to offer services to and receive services from other programs. Supported IPC mechanisms under MS OS/2 are semaphores, signals, pipes, queues, shared memory, and dynamic data exchange.

IPCH Initial Paging Channel is the channel number used by your cellular provider to "page" the phones on the system.

The term "paging" refers to notifying a particular phone that it has an incoming call. All idle, turned-on phones on a system monitor the data stream on the IPCM. Non-wireline cellular carriers use channel 0333 as the IPCH, while wireline carriers (those operated by a telephone company use channel 0334).

IPCI See Integrated Personal Computer Interface.

IPCP IP Control Protocol; protocol for transporting IP traffic over a PPP connection.

IPDC Internet Protocol Device Control. See MGCP and Simple Gateway Control Protocol.

IPDS Intelligent Printer Data Stream. It's IBM's host-to-printer page description protocol for printing. You can now buy kits which let you use your present printer to emulate an IBM printer.

IPE Intelligent Peripheral Equipment. Northern Telecom's term for being able to extend all the features of its PBX over distances longer than a normal extension in a building. See Fiber Remote.

IPEI International Portable Equipment Identities. A wireless term.

IPEM If the Product Ever Materializes.

IPL Initial Program Load.

IPLC International Private Leased Circuit.

IPM Interruptions Per Minute or Impulses Per Minute.

IPNG IPng, IP Next Generation. Collective term used to describe the efforts of the Internet Engineering Task force to define the next generation of the Internet Protocol (IP) which includes security measures, as well as larger IP addresses to cope with the explosive growth of the Internet. The were three candidate protocols for IPng (CATNIP, TUBA and SIPP), were blended into IPv6, which is in trial stages at the time of this writing. See IPv6.

IPNS International Private Network Service. It actually international private line service and it's typically a circuit from 9.6 Kbit/sec up to T-1 or E-1. Domestically you would simply call it "Private line data service."

IPP IPP is the Internet Print Protocol, a collection of IETF standards developed through the Printer Work Group, www.pwg.org, that will make it as easy to print over the Internet as it is to print from your PC. IPP uses the HTTP protocols to "POST" a supported MIME Page Description Language file to a printer. Printers are given Internet addresses such as, www.mydomain.com/ipp/my_printer, so they can be located on the Internet. IPP has the support of all the major printer companies including, Xerox, HP, Lexmark, IBM as well as Novell and Microsoft. Since fax, at a sufficient level of abstraction, is "remote printing," work is under way to create a Fax Profile for IPP as well, so that IPP can duplicate the legal as well as common practices of fax transmissions. Richard Shockey, Rshockey@ix.netcom.com contributed this definition. Thank you.

IPO Initial Public Offering. Start a company. Some years later, take it public. Come out at \$12. A week later, your stock is at \$24. You're a success, and rich. IPOs are critical in saying "Thank You" to all your hardworking employees.

iPOD (Internet Protocol) Phone over Data. There tend to be two variations — emulation and driving. The emulation iPOD connects directly to digital station ports on a PBX and emulates a digital PBX feature phone. The emulation iPOD also enables the new PC IP PBX vendors to interoperate with enterprise PBXs. The driving iPOD drives digital PBX phones in the same fashion as if the phone were connected directly to a PBX station circuit card. The driving iPOD can enable the

new PC PBX vendors to use existing desk sets in the enterprise. Both versions of the iPOD provide a TCP/IP interface for the purpose of transporting the voice and call control signaling associated with a PBX digital station call over a packet network. Protocols, DSP algorithms, densities and different form factors all constitute possible platform variations.

IPR Intellectual Property Rights

IPRS Internet Protocol Routing Service. Defined by Bell Atlantic as "a low-cost access service for ISPs. This service supports basic dial, ISDN, and dedicated requirements for transparent connectivity from the end-user to the ISP."

IPS Internet Protocol Suite.

IPsec A collection of IP security measures that comprise an optional tunneling protocol for IPv6. IPsec supports authentication through an "authentication header" which is used to verify the validity of the originating address in the header of every packet of a packet stream. An "encapsulating security payload" header encrypts the entire datagram, based on the encryption algorithm chosen by the implementer. See also Authentication, Encryption, IPv6, and Tunneling.

IPT IP Telephony.

IPT Gateway IP Telephony Gateway. Imagine you and I work for a company which has a PBX — a telephone system. You dial 234 to reach Harry. You dial 9 and a long distance number to dial your biggest client in Los Angeles. Now imagine you want to call your company's branch office in London. You dial 22. You hear a dial tone. You then punch in 689. You hear another dial tone. Then you punch 123. Bingo, the boss of the London office answers. Here's what all those numbers mean. Dialing 22 dials you into a PC called the IP Telephony Gateway, which, on the one side, is connected to your PBX and on the other side is connected to a data line your company has between your office and your London office. Dialing 689 is you telling the IPT Gateway that you want to speak to the PBX in your London office. Dialing 123 tells the London PBX to dial extension 123.

That connection between your PBX and your London office's PBX might be anything from a dedicated private data line (e.g. part of your company's Intranet), to a virtual circuit on a Virtual Private Network (VPN) or it might be the public Internet. The IPT Gateway's major function is to convert the analog voice coming out of your PBX into VoIP (voice over Internet Protocol) and then send it on a packet switched data circuit which conforms to the IP. In short, an IPT Gateway allows users to use the Internet (or most likely an Intranet or Virtual Private Network) to talk with remote sites using (Voice over Internet Protocol).

IPTC On April 30, 1998, Ericsson Inc. released a press release which contained, inter alia, "Ericsson Inc. has developed a new IP telephony platform called Internet Telephony Solution for Carriers (IPTC) that raises the standard for IP telephony systems. IPTC offers phone-to-phone, fax-to-fax and PC-to-phone services over a TCP/IP network. It provides a superior operations and management (O&M) facility that moves IP telephony to a true carrier-class communications system. IPTC works by taking phone and fax calls that originate in the public switched telephony network (PSTN) and passing them to the IPTC platform, which carries them over the TCP/IP network to their destination where they are fed back to the PSTN network. PC-to-phone calls are taken directly from the TCP/IP network and carried to their destination in the same way...IPTC software runs on industry standard platforms that are based on Intel Pentium processors and Microsoft Windows NT...IPTC uses a Web-based management program to update and con-

trol multiple gateways. No longer is it necessary to change the parameters in individual gateways when IPTC can update all gateways within a network through one "netkeeper" applications program. The call and traffic control for individual gateways in a network is handled by sitekeepers. The sitekeepers connect to the netkeeper, which acts as a single point of control for the Q&M functions of the whole IPTC platform. The netkeeper is not involved in the processing of calls but stores the platform topology information, routing configuration and alarm information. Other features included in the IPTC platform are least-cost routing, dynamic route allocation, multiple IP networks support, and the ability to handle validated and un-validated traffic. Real-time billing with fraud prevention and call duration advice with integrated voice response software is also provided."

IPU Intelligent Processing Unit. Another way of saying CPU. See CPU. Also Intelligent Peripheral Unit, the hardware associated with an intelligent peripheral. Also Alcatel's parlance for an actual workstation that's associated mostly with one of Alcatel's applications called the local applications platform or LAP and a software applications package called the monitor reset controller-2 or MRC-2. In short, everyone is using IPU to mean whatever cool thing they want it to mean. Certainly sounds cool.

IPv4 Internet Protocol Version 4. The current version of the Internet Protocol, which is the fundamental protocol on which the Internet is based. Although its roots are in the initial development work for ARPAnet, IPv4 was first formalized as a standard in 1981. Since that time, it has been widely deployed in all variety of data networks, including LANs and LAN inter-networks. While IPv4 served its purpose for some 25 years, it has lately proved to be inadequate, largely in terms of security and limitations of the address field. The address field is limited to 32 bits; although 2 to the 32nd power is a very large number, we are running out of IP addresses just as we have run out of 800 numbers and traditional area codes. Hence, the development of IPv6. See IP.

IPv5 Internet Protocol Version 5. IPv5 is not exactly a missing link, although it might appear so. Rather, IPv5 was assigned to ST2, Internet Stream Protocol Version 2, which is documented in RFC 1819. ST2 is an experimental protocol developed as an adjunct to IP for support of real-time transport of multimedia data. See IP and IPv6.

IPv6 Internet Protocol Version 6. The new proposed Internet Protocol designed to replace and enhance the present protocol which is called TCP/IP, or officially IPv4. IPv6 has 128-bit addressing, auto configuration, new security features and supports real-time communications and multicasting. IPv6 is described in RFC 1752, The Recommendation for IP Next Generation Protocol, including the strengths and weaknesses of each of the proposed protocols which were blended to form the final proposed solution. At the time of this writing, IPv6 is standardized, but not widely deployed. It requires upgrades that are expensive. They will be fork-lift upgrades in many cases. Therefore, IPv6 is being deployed pretty much only in the NextGen carrier networks, which are being built from the ground up. IPv6 offers 128-bit addressing, auto configuration, new security features and supports real-time communications and multicasting. The 128-bit addressing scheme will relieve pressure on the current 32-bit scheme, which is nearly exhausted due to the widespread use of IP in the Internet and a wide variety of LAN, MAN and WAN networks. Clearly, 2 to the 128th power is a huge number, yielding a staggering number of IP addresses. According to Mark Miller of Diginet

Corporation, it equates to approximately 1,500 addresses per square angstrom, with an angstrom being one ten-millionth of a millimeter. Another way of looking at this is that IPv6 yields about 32 addresses per square inch of dry land on the earth's surface — in other words, we are not likely to run out of IPv6 addresses. (Don't be surprised to see your telephone assigned an IP address in the future.)

Autoconfiguration Protocol, an intrinsic part of IPv6, allows a device to assign itself a unique IP address without the intervention of a server. The self-assigned address is based in part on the unique LAN MAC (Media Access Control) address of the device, which might be in the form of laptop computer. This feature allows the user the same full IPv6 capability when on the road as he might enjoy in the office when the laptop is inserted into a LAN-attached docking station. IPv6 security is provided in several ways. Data integrity and user authentication are provided by any of a number of authentication schemes. Second, the Encapsulating Security Payload feature provides for confidentiality of data through encryption algorithms such as DES (Data Encryption Standard). Several different types of IPv6 addresses support various types of communications. Unicast supports point-to-point transmission, Anycast allows communications with the closest member of a device group, and Multicast supports communications with multiple members of a device group.

IPX Internet Packet eXchange. Novell NetWare's native LAN communications protocol, used to move data between server and/or workstation programs running on different network nodes. IPX packets are encapsulated and carried by the packets used in Ethernet and the similar frames used in Token-Ring networks. IPX supports packet sizes up to 64 bytes. Novell's NCP and SPX both use IPX. See also IPX.COM.

IPX Autodiscovery The ability of a network manager to discover the node address and functionality of network devices.

IPX.COM The Novell IPX/SPX (Internetwork Packet eXchange/Sequenced Packet eXchange) communication protocol that creates, maintains, and terminates connections between network devices (workstations, file servers, routers, etc.). IPX.COM uses a LAN driver routine to control the station's network board and address and to route outgoing data packets for delivery on the network. IPX/SPX reads the assigned addresses of returning data and directs the data to the proper area within a workstation's shell or the file server's operating system. See also Netware.

IPX/SPX Internetwork Packet Exchange/Sequenced Packet Exchange. Two network protocols. IPX is NetWare protocol for moving information across the network; SPX works on top of IPX and adds extra commands. In the OSI model, IPX conforms to the network layer and SPX is the transport layer.

IPXCP IPX Control Protocol; protocol for transporting IPX traffic over a PPP connection.

IPXWAN A Novell specification describing the protocol to be used for exchanging router-to-router information to enable the transmission of Novell IPX data traffic across WAN (Wide Area Network) links.

IR 1. Infrared. The band of electromagnetic wavelengths between the extreme of the visible part of the spectrum (about 0.75 μ m) and the shortest microwaves (about 100 μ m).

2. Internet Registry. See also Internet Assigned Numbers Authority.

IRAM Intelligent RAM. The idea is to put a microprocessor into a memory chip — a move that dramatically improve computer performance.

on their single line phones. This button gives the precise hookswitch signal for the precise length of time necessary — no more, no less. The Tap Button is also called a Flash button or a Tap Key.

Tap Key Also called Tap Button or Flash Key. A button on a phone that accomplishes the same function as a switch hook but is not a switch hook. See Tap Button.

TAPAC Terminal Attachment Program Advisory Committee. Body which recommends telecom standards to the Canadian Federal Government.

Tape Drive The physical unit that holds, reads and writes magnetic tape.

Tape Reader A device which reads information recorded on punched paper tape or magnetic tape.

Tape Relay A method of retransmitting TTY traffic from one channel to another, in which messages arriving on an incoming channel are recorded in the form of perforated tape, this tape then being either fed directly and automatically into an outgoing channel, or manually transferred to an automatic transmitter for transmission on an outgoing channel.

Tapered Fiber An optical fiber in which the cross section, i.e., cross-sectional diameter or area, varies, i.e., increases or decreases, monotonically with length.

TAPI Telephone Application Programming Interface. Also called Microsoft/Intel Telephony API. A term that refers to the Windows Telephony API. TAPI is a changing (i.e. improving) set of functions supported by Windows that allow Windows applications (Windows 3.xx, 95 and NT) to program telephone-line-based devices such as single and multi-line phones (both digital and analog), modems and fax machines in a device-independent manner. TAPI essentially does to telephony devices what Windows printer system did to printers — make them easy to install and allow many application programs to work with many telephony devices, irrespective of who made the devices. TAPI is one of numerous high-level device interfaces that Windows offers as part of the Windows Open Services Architecture (WOSA). TAPI simplifies the process of writing a telephony application that works with a wide variety of modems and other devices supported by TAPI drivers. See also Dial String, Microsoft Fax, TAPI 2.0, TAPI 3.0, Windows 95 and Windows Telephony for fuller explanations.

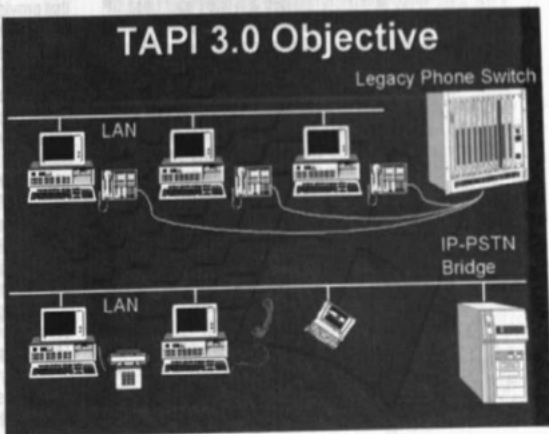
TAPI 2.0 The following is Microsoft's explanation: The Microsoft Windows Telephony API (TAPI) 2.0 ships as part of Windows NT Server 4.0 and Windows NT Workstation 4.0.


TAPI 2.0 is the latest release of the TAPI specification, introduced in 1993. TAPI helps bridge the gap between the telephone and computer. TAPI helps the PC to understand how telephone networks operate. With TAPI, programmers can exploit telephone network capabilities from within regular Windows-based applications. With TAPI, the jungle of proprietary telephony hardware is turned into standard programming interfaces. Application developers can wave goodbye to the complexity and variability of the underlying telephone network. They no longer have to hard-code their applications to a particular system's signaling or message set requirements. Instead, applications developers write code to the Telephony Applications Programming Interface. Developers can focus on their application without worrying about the nitty-gritty programming details of connecting to a specific telephone network. TAPI supports PBXs, key telephone systems, ISDN, the analog PSTN, cellu-

lar, CENTREX and other types of telephone networks. (So long as the various providers have written an SPI — Service Provider Interface, which is code to translate their commands into code which TAPI can understand. — Harry)

TAPI 2.0 includes these enhancements:

- 32-bit architecture. All core TAPI components are now based on the Win32 architecture. Non-Intel processors running Win NT Server 4.0 or Workstation 4.0 are supported.
 - 32-bit application portability. Existing Win32 apps currently running on Win 95 using TAPI 1.4 will run on NT Workstation 4.0 or Server 4.0 on Intel x86 microprocessors.
 - 16-bit application portability. Existing applications currently running on Win 95 and 3.1 using TAPI 1.3 will run on NT Workstation 4.0 or NT Server 4.0 on Intel x86 microprocessors.
 - Unicode support. Win32 apps can now call the existing ANSI TAPI functions or the new Unicode versions of functions. Unicode is a 16-bit, fixed-width character encoding standard. It encompasses virtually all of the characters commonly used on computers today.
 - Expanded feature support for call center applications. TAPI now supports an expanded set of features to better serve call center operations with Windows. New call center features supported include ACD queues, predictive dialing, and call routing.
 - Registry support. All telephony parameters are now stored in the Windows registry. All stored parameters can be updated across the LAN.
 - Quality of Service (QoS) support. Applications can request, negotiate, and re-negotiate QoS performance parameters with the network. Improved QoS support reduces or eliminates latency and other negative characteristics for applications, especially voice and data apps, over various networks.
 - Enhanced device sharing. Applications can restrict handling of inbound calls on a device to a single address. This supports features such as distinctive ringing.
 - Additions and changes to TAPI functions. Many new TAPI functions and messages are available with TAPI 2.0. In addition, several functions and messages already supported by TAPI 1.4 were changed in some measure to make them more consistent in their operation. See TAPI 3.0
- TAPI 3.0** Microsoft announced TAPI 3.0 in mid-September, 1997. At that time, Mitch Goldberg of Microsoft, said TAPI 3.0 will be available initially in NT 5 beta 1 in September, fully featured in NT 5 beta 2, and available commercially when the NT 5.0 operating system ships (around the middle of 1998).





multimedia traffic over any network that uses IP (the Internet Protocol). This offers users both flexibility in physical media (for example, POTS lines, ADSL, ISDN, leased lines, coaxial cable, satellite, and twisted pair) and flexibility of physical location. As a result, the same ubiquitous networks that carry Web, e-mail and data traffic can be used to connect to individuals, businesses, schools and governments worldwide. TAPI 3.0 is an evolutionary API that supports convergence of both traditional PSTN telephony and telephony over IP networks.


What are the Benefits of IP Telephony? IP Telephony allows organizations and individuals to lower the costs of existing services, such as voice and broadcast video, while at the same time broadening their means of communication to include modern video conferencing, application sharing, and whiteboarding tools. In the past, organizations have deployed separate networks to handle traditional voice, data, and video traffic. Each with different transport requirements, these networks were expensive to install, maintain, and reconfigure. Furthermore, since these networks were physically distinct, integration was difficult if not impossible, limiting their potential usefulness. IP Telephony blends voice, video and data by specifying a common transport, IP, for each, effectively collapsing three networks into one. The result is increased manageability, lower support costs, a new breed of collaboration tools, and increased productivity. Possible applications for IP Telephony include telecommuting, real-time document collaboration, distance learning, employee training, video conferencing, video mail, and video on demand.

What is TAPI 3.0? As telephony and call control become more common in the desktop computer, a general telephony interface is needed to enable applications to access all the telephony options available on any machine. Additionally,

it is imperative that the media or data on a call is available to applications in a standard manner. TAPI 3.0 is an architecture that provides simple and generic methods for making connections between two or more machines, and accessing any media streams involved in that connection. It abstracts call-control functionality to allow different, and seemingly incompatible, communication protocols to expose a common interface to applications. IP Telephony is a demand poised for explosive growth, as organizations begin an historic shift from expensive and inflexible circuit-switched public telephone networks to intelligent, flexible and inexpensive IP networks.

Inside TAPI 3.0: TAPI 3.0 integrates multimedia stream control with legacy telephony. Additionally, it is an evolution of the TAPI 2.1 API to the COM model, allowing TAPI applications to be written in any language, such as JavaT, C/C++ and the Microsoft Visual Basic programming system. Besides supporting classic telephony providers, TAPI 3.0 supports standard H.323 conferencing and IP multicast conferencing. TAPI 3.0 utilizes the Windows NT 5.0 Active Directory service to simplify deployment within an organization, and it supports quality of service (QoS) features to improve conference quality and network manageability.

Telephony Services Scenarios



TAPI 3.0 will also be available for Windows 98 ("Memphis") in the NT 5.0 time frame. microsoft.com/communications. At the same time, Microsoft released a White Paper titled "IP Telephony with TAPI 3.0." Here are key excerpts. For the full paper, www.ctexpo.com. TAPI 3.0 is an evolutionary API providing convergence of both traditional PSTN telephony and IP Telephony. IP Telephony is an emerging set of technologies which enables voice, data, and video collaboration over existing LANs, WANs and the Internet. TAPI 3.0 enables IP Telephony on the Microsoft Windows operating system platform by providing simple and generic methods for making connections between two or more machines, and accessing any media streams involved in the connection. TAPI 3.0 supports standards based H.323 conferencing and IP multicast conferencing. It utilizes the Windows NT 5.0 operating system's Active Directory service to simplify deployment within an organization, and includes quality of service (QoS) support to improve conference quality and network manageability. What is IP Telephony? IP Telephony is an emerging set of technologies that enables voice, data, and video collaboration over existing IP-based LANs, WANs, and the Internet. Specifically, IP Telephony uses open IETF and ITU standards to move multi-

There are four major components to TAPI 3.0:

1. TAPI 3.0 COM API
2. TAPI Server
3. Telephony Service Providers
4. Media Stream Providers

In contrast to TAPI 2.1, the TAPI 3.0 API is implemented as a suite of Component Object Model (COM) objects. Moving TAPI to the object-oriented COM model allows component upgrades of TAPI features. It also allows developers to write TAPI-enabled applications in any language, such as Java, Visual Basic, or C/C++. The TAPI Server process (TAPISRV.EXE) abstracts the TSPI (TAPI Service Provider Interface) from TAPI 3.0 and TAPI 2.1, allowing TAPI 2.1 Telephony Service Providers to be used with TAPI 3.0, maintaining the internal state of TAPI. Telephony Service Providers (TSPs) are responsible for resolving the protocol-independent call model of TAPI into protocol-specific call control mechanisms. TAPI 3.0 provides backward compatibility with TAPI 2.1 TSPs. Two IP Telephony service providers (and their associated MSPs) ship by default with TAPI 3.0: the H.323 TSP and the IP Multicast Conferencing TSP. TAPI 3.0 provides a uniform way to access the media streams in a call, supporting the DirectShow API as the primary media stream handler. TAPI Media Stream Providers (MSPs) implement DirectShow interfaces for a particular TSP and are required for any telephony service that makes use of DirectShow streaming. Generic streams are handled by the application. There are five objects in the TAPI 3.0 API:

1. TAPI
2. Address
3. Terminal
4. Call
5. CallHub

The TAPI object is the application's entry point to TAPI 3.0. This object represents all telephony resources to which the local computer has access, allowing an application to enumerate all local and remote addresses. An Address object represents the origination or destination point for a call. Address capabilities, such as media and terminal support, can be retrieved from this object. An application can wait for a call on an Address object,

or can create an outgoing call object from an Address object. A Terminal object represents the sink, or renderer, at the termination or origination point of a connection. The Terminal object can map to hardware used for

human interaction, such as a telephone or microphone, but can also be a file or any other device capable of receiving input or creating output. The Call object represents an address' connection between the local address and one or more other addresses (This connection can be made directly or through a CallHub). The Call object can be imagined as a first-party view of a telephone call. All call control is done through the Call object. There is a call object for each member of a CallHub. The CallHub object represents a set of related calls. A CallHub object cannot be created directly by an application. They are created indirectly when incoming calls are received through TAPI 3.0. Using a CallHub object, a user can enumerate the other participants in a call or conference, and possibly (because of the location independent nature of COM) perform call control on the remote Call objects associated with those users, subject to sufficient permissions:

Using TAPI Objects to Place a Call:

1. Create and initialize a TAPI object.
2. Use the TAPI object to enumerate all available Address

objects on a machine (for example, network cards, modems, and ISDN lines).

3. Enumerate the supported address types of each Address object (for example, a phone number, IP address, and so on).
4. Choose an Address object based on queries for support for appropriate media (audio, video, and so on) and address types.
5. Use the CreateCall() method of the Address object to create a Call object associated with a particular address.
6. Select appropriate Terminals on the Call object.
7. Call the Connect() method of the Call object to place the call. To Answer a Call:
 1. Create and initialize a TAPI object.
 2. Use the TAPI object to enumerate all available Address objects on a machine (for example, network cards, modems, and ISDN lines).
 3. Enumerate the supported address types of each Address object (e.g. a phone number, IP address, etc.).
 4. Choose an Address object based on queries for support of appropriate media (audio, video, and so on) and address types.
 5. Register an interest in specific media types with the appropriate Address object.
 6. Register a call event handler (i.e. implement an ITCallNotification interface) with the Address object.
 7. TAPI notifies the application of a new call through ITCallNotification and creates a Call object.
 8. Select appropriate Terminals on the Call object.
 9. Call the Connect() method of the Call object to place the call.
 10. Call the Answer() method of the Call object to answer the call.

Media Streaming Model: The Windows operating system provides an extensible framework for efficient control and manipulation of streaming media called the DirectShow API. DirectShow, through its exposed COM interfaces, provides TAPI 3.0 with unified stream control. At the heart of the DirectShow services is a modular system of pluggable components called filters, arranged in a configuration called a filter graph. A component called the filter graph manager oversees the connection of these filters and controls the stream's data flow. Each filter's capabilities are described by a number of special COM interfaces called pins. Each pin instance can consume or produce streaming data, such as digital audio. While COM objects are usually exposed in user mode programs, the DirectShow streaming architecture includes an extension to the Windows driver model that allows the connection of media streams directly at the device driver level.

These high-performance streaming extensions to the Windows driver model avoid user-to-kernel mode transitions, and allow efficient routing of data streams between different hardware components at the device driver level. Each kernel mode filter is mirrored by a corresponding user mode proxy that facilitates connection setup and can be used to control hardware-specific features. DirectShow network filters extend the streaming architecture to machines connected on an IP network. The Real-Time Transport protocol (RTP), designed to carry real-time data over connectionless networks, transports TAPI media streams and provides appropriate time stamp information. TAPI 3.0 includes a kernel mode RTP network filter. TAPI 3.0 utilizes this technology to present a unified access method for the media streams in multimedia calls. Applications can route these streams by manipulating corresponding filter graphs; they can also easily connect streams from multiple calls for bridging and conferencing capabilities. What is H.323? H.323 is a comprehensive International Telecommunications Union (ITU) standard for multimedia

communications (voice, video, and data) over connectionless networks that do not provide a guaranteed quality of service, such as IP-based networks and the Internet. It provides for call control, multimedia management, and bandwidth management for point-to-point and multipoint conferences. H.323 mandates support for standard audio and video codecs and supports data sharing via the T.120 standard. Furthermore, the H.323 standard is network, platform, and application independent, allowing any H.323 compliant terminal to inter-operate with any other.

H.323 allows multimedia streaming over current packet-switched networks. To counter the effects of LAN latency, H.323 uses as a transport the Real-time Transport Protocol (RTP), an IETF standard designed to handle the requirements of streaming real-time audio and video over the Internet. The H.323 standard specifies three command and control protocols:

1. H.245 for call control
2. Q.931 for call signaling
3. The RAS (Registration, Admissions, and Status) signaling function

The H.245 control channel is responsible for control messages governing operation of the H.323 terminal, including capability exchanges, commands, and indications. Q.931 is used to set up a connection between two terminals, while RAS governs registration, admission, and bandwidth functions between endpoints and gatekeepers (RAS is not used if a gatekeeper is not present).

H.323 defines four major components for an H.323-based communications system:

1. Terminals
2. Gateways
3. Gatekeepers
4. Multipoint Control Units (MCUs)

Terminals are the client endpoints on the network. All terminals must support voice communications; video and data support is optional.

A Gateway is an optional element in an H.323 conference. Gateways bridge H.323 conferences to other networks, communications protocols, and multimedia formats. Gateways are not required if connections to other networks or non-H.323 compliant terminals are not needed. Gatekeepers perform two important functions which help maintain the robustness of the network - address translation and bandwidth management. Gatekeepers map LAN aliases to IP addresses and provide address lookups when needed. Gatekeepers also exercise call control functions to limit the number of H.323 connections, and the total bandwidth used by these connections, in an H.323 "zone." A Gatekeeper is not required in an H.323 system-however, if a Gatekeeper is present, terminals must make use of its services.

Multipoint Control Units (MCU) support conferences between three or more endpoints. An MCU consists of a required Multipoint Controller (MC) and zero or more Multipoint Processors (MPs). The MC performs H.245 negotiations between all terminals to determine common audio and video processing capabilities, while the Multipoint Processor (MP) routes audio, video, and data streams between terminal endpoints.

Any H.323 client is guaranteed to support the following standards: H.261 and G.711. H.261 is an ITU-standard video codec designed to transmit compressed video at a rate of 64 Kbps and at a resolution of 176x44 pixels (QCIF). G.711 is an ITU-standard audio codec designed to transmit A-law and E-law PCM audio at bit rates of 48, 56, and 64 Kbps. Optionally, an H.323 client may support additional codecs: H.263 and

G.723. H.263 is an ITU-standard video codec based on and compatible with H.261. It offers improved compression over H.261 and transmits video at a resolution of 176 x 44 pixels (QCIF). G.723 is an ITU-standard audio codec designed to operate at very low bit rates.

The H.323 Telephony Service Provider (along with its associated Media Stream Provider) allows TAPI-enabled applications to engage in multimedia sessions with any H.323-compliant terminal on the local area network. Specifically, the H.323 Telephony Service Provider (TSP) implements the H.323 signaling stack. The TSP accepts a number of different address formats, including name, machine name, and e-mail address. The H.323 MSP is responsible for constructing the DirectShow filter graph for an H.323 connection (including the RTP, RTP payload handler, codec, sink, and renderer filters).

H.323 telephony is complicated by the reality that a user's network address (in this case, a user's IP address) is highly volatile and cannot be counted on to remain unchanged between H.323 sessions. The TAPI H.323 TSP uses the services of the Windows NT Active Directory to perform user-to-IP address resolution. Specifically, user-to-IP mapping information is stored and continually refreshed using the Internet Locator Service (ILS) Dynamic Directory, a real-time server component of the Active Directory.

The following user scenario illustrates IP address resolution in the H.323 TSP:

1. John wishes to initiate an H.323 conference with Alice, another user on the LAN. Once Alice's video conferencing application creates an Address object and puts it in listen mode, Alice's IP address is added to the Windows NT Active Directory by the H.323 TSP. This information has a finite time to live (TTL) and is refreshed at regular intervals via the Lightweight Directory Access Protocol (LDAP).
2. John's H.323 TSP then queries the ILS Dynamic Directory server for Alice's IP address. Specifically, John queries for any and all RTPPerson objects in the Directory associated with Alice.
3. Armed with Alice's up-to-date IP address, John initiates an H.323 call to Alice's machine, and H.323-standard negotiations and media selection occurs between the peer TSPs on both machines. Once capability negotiations have been completed, both H.323 Media Stream Providers (MSPs) construct appropriate DirectShow filter graphs, and all media streams are passed off to DirectShow to handle. The conference then begins. What is IP Multicast Conferencing? IP Multicast is an extension to IP that allows for efficient group communication. IP Multicast arose out of the need for a lightweight, scalable conferencing solution that solved the problems associated with real-time traffic over a datagram, "best-effort" network. There are many advantages to using IP Multicast: scalability, fault tolerance, robustness, and ease of setup. The IP Multicast conferencing model incorporates the following key features:

1. No global coordination is needed to add and remove members from a conference.
2. To reach a multicast group, a user sends data to a single multicast IP address. No knowledge of the other users in a group is necessary.
3. To receive data, users register their interest in a particular multicast IP address with a multicast aware router. No knowledge of the other users in a group is necessary. Routers hide the multicast implementation details from the user.

Traditional connection-oriented conferencing suffers from a number of problems:

1. User complexity: Users must know the location of every user they wish to converse with, limiting scalability and fault-

tolerance and rendering it difficult for users to add and remove themselves from a conference.

2. **Wasted bandwidth:** A user wishing to broadcast data to N users must send data through N connections.

The total bandwidth required for multiparty conferences in which all users are sending data goes up as N squared the number of parties involved, leading to huge scalability problems. IP Multicast takes advantage of the actual network topology to eliminate the transmission of redundant data down the same communications links.

IP Multicast implements a lightweight, session-based communications model, which places relatively little burden on conference users. Using IP Multicast, users send only one copy of their information to a group IP address that reaches all recipients. IP Multicast is designed to scale well as the number of participants expands. Adding one more user does not add a corresponding amount of bandwidth. Multicasting also results in a greatly reduced load on the sending server. IP Multicast routes these one-to-many data streams efficiently by constructing a spanning tree, in which there is only one path from one router to any other. Copies of the stream are made only when paths diverge.

Without multicasting, the same information must either be carried over the network multiple times, one time for each recipient, or broadcast to everyone on the network, consuming unnecessary bandwidth and processing. IP Multicast uses Class D Internet Protocol addresses to specify multicast host groups, ranging from 224.0.0.0 to 239.255.255.255. Both permanent and temporary group addresses are supported. Permanent addresses are assigned by the Internet Assigned Numbers Authority (IANA) and include 224.0.0.1, the "all-hosts group" used to address all multicast hosts on the local network, and 224.0.0.2, which addresses all routers on a LAN. The range of addresses between 224.0.0.0 and 224.0.0.255 is reserved for routing and other low-level network protocols. Other addresses and ranges have been reserved for applications, such as 224.0.13.000 to 224.0.13.255 for Net News (for more information, see RFC 1700, "Assigned Numbers" at <http://ftp.internic.net/rfc/rfc1700.txt>).

The transport protocol for IP Multicast is RTP (Real-time Transport Protocol), which provides a standard multimedia header giving timestamp, sequence numbering, and payload format information. Applications for IP Multicast include video and audio conferencing, telecommuting, database and Web-site replication, distance learning, dissemination of stock quotes, and collaborative computing. At present, the largest demonstration of the capabilities of IP Multicast is the Internet MBONE (Multicast Backbone). The MBONE is an experimental, global multicast network layered on top of the physical Internet. It has been in existence for about five years, and presently carries IETF meetings, NASA space shuttle launches, music, concerts, and many other live meetings and performances. www.mbone.com.

The IP Multicast Conferencing TSP is chiefly responsible for resolving conference names to IP multicast addresses, using the Session Description Protocol (SDP) conference descriptors stored in the ILS Dynamic Directory Conference Server. It is complemented by the Rendezvous conference MSP, described later. The IP Multicast Conferencing MSP is responsible for constructing an appropriate DirectShow filter graph for an IP multicast connection (including RTP, RTP payload handler, codec, sink, and renderer filters).

TAPI 3.0 uses the IETF standard Session Description Protocol to advertise IP multicast conferences across the enterprise.

SDP descriptors are stored in the Windows NT Active Directory — specifically, in the ILS Dynamic Directory Conference Server. In contrast to the Dynamic Directory servers used by the H.323 TSP, there is only one ILS Conference Server per enterprise, since conference announcements are not continually refreshed, therefore consuming little bandwidth.

TAPI 3.0's IP multicast conference mechanism is illustrated in the following scenario, in which John wishes to initiate a multicast conference:

1. John's TAPI 3.0-enabled application uses the Rendezvous Controls to create an SDP session descriptor on the ILS Conference Server. The SDP descriptor contains, among other things, the conference name, start and end time information, the IP multicast address of the conference, and the media types used for the conference.

2. Jim queries the ILS Conference Server for SDP descriptors of conferences matching his criteria.

3. Mary and Alice perform similar queries and use the SDP information they receive to decide to participate in John's conference. Armed with the multicast IP address of the conference, they join the multicast host group.

The Rendezvous Controls are a set of COM components that abstract the concept of a conference directory, providing a mechanism to advertise new multicast conferences and to discover existing ones. They provide a common schema (SDP) for conference announcement, as well as scriptable interfaces, authentication, encryption, and access control features.

The user may add, delete, and enumerate multicast conferences stored on an ILS Conference Server via the Rendezvous Controls. These controls manipulate conference data via the Lightweight Directory Access Protocol (LDAP). The conferencing application uses the Rendezvous Controls to obtain session descriptors for the conferences that match the user's criteria. Access control lists (ACLs) protect each of the stored conference announcements, and whether or not an announcement is visible and accessible depends upon the user's credentials. Once the user has chosen a conference, the user application searches for all Address objects that support the address type "Multicast Conference Name." The application then uses the conference name from the SDP descriptor as a parameter to the CreateCall() method of the appropriate Address object, passes the appropriate Terminal objects to the returned Call object, and calls Call->Connect(). The Rendezvous Controls store the conference information on an ILS Conference Server in a format defined by the Session Description Protocol (SDP), an IETF standard for announcing multimedia conferences. The purpose of SDP is to publicize sufficient information about a conference (time, media, and location information) to allow prospective users to participate if they so choose. Originally designed to operate over the Internet MBONE (IP Multicast Backbone), SDP has been integrated by TAPI 3.0 with the Windows NT Active Directory, thereby extending its functionality to local area networks. An SDP descriptor advertises the following information about a conference.

A session description is broken into three main parts: a single Session Description, zero or more Time Descriptions, and zero or more Media Descriptions. The Session Description contains global attributes that apply to the whole conference or all media streams. Time Descriptions contain conference start, stop, and repeat time information, while Media Descriptions contain details that are specific to a particular media stream.

While traditional IP multicast conferences operating over the

MBONE have advertised conferences using a push model based on the Session Announcement Protocol (SAP). TAPI 3.0 utilizes a pull-based approach using Windows NT Active Directory services. This approach offers numerous advantages, among them bandwidth conservation and ease of administration.

TAPI 3.0's conference security system addresses the following needs:

- Controlling who can create, delete, and view conference announcements.
- Preventing conference eavesdropping.

TAPI 3.0 utilizes the security features of the Windows NT Active Directory and LDAP to provide for secure conferencing over insecure networks such as the Internet. Each object in the Active Directory can be associated with an Access Control List (ACL) specifying object access rights on a user or group basis. By associating ACLs with SDP conference descriptors, conference creators can specify who can enumerate and view conference announcements. User authentication is provided by the Windows NT security subsystem.

Session Descriptors are transmitted from the ILS Conference Server to the user over LDAP in encrypted form, via a Secure Sockets Layer (SSL) connection, ensuring that the SDP is safe from eavesdroppers. IP Multicast makes no provision for authenticating users. Any user may anonymously join a multicast host group. To keep conferences private, TAPI 3.0 allows an IP multicast conference to be encrypted, with the encryption key distributed from within the conference descriptor. Only users with sufficient permissions have access to a conference's SDP descriptor, and therefore the Multicast Encryption Key. Once an authenticated user fetches the encryption key, he or she can participate in the conference.

In contrast to traditional data traffic, multimedia streams, such as those used in IP Telephony or videoconferencing, may be extremely bandwidth and delay sensitive, imposing unique quality of service (QoS) demands on the underlying networks that carry them. Unfortunately, IP, with a connectionless, "best-effort" delivery model, does not guarantee delivery of packets in order, in a timely manner, or at all. In order to deploy real-time applications over IP networks with an acceptable level of quality, certain bandwidth, latency, and jitter requirements must be guaranteed, and must be met in a fashion that allows multimedia traffic to coexist with traditional data traffic on the same network.

Bandwidth: Multimedia data, and in particular video, may require orders of more bandwidth than traditional networks have been provisioned to handle. An uncompressed NTSC video stream, for example, can require upwards of 220 megabits per second to transmit. Even compressed, a handful of multimedia streams can completely overwhelm any other traffic on the network.

Latency: The amount of time a multimedia packet takes to get from the source to the destination (latency) has a major impact on the perceived quality of the call. There are many contributors towards latency, including transmission delays, queuing delays in network equipment, and delays in host protocol stacks. Latency must be minimized in order to maintain a certain level of interactivity and to avoid unnatural pauses in conversation.

Jitter: In contrast to data traffic, real-time multimedia packets must arrive in order and on time to be of any use to the receiver. Variations in packet arrival time (jitter) must be below a certain threshold to avoid dropped packets (and therefore irritating shrieks and gaps in the call). Jitter, by determining receive buffer sizes, also affects latency.

Coexistence: In comparison with multimedia traffic, data traffic is relatively bursty, and arrives in unpredictable chunks (for instance, when someone opens a Web page, or downloads a file from an FTP site). Aggregations of such bursts can clog routers and cause gaps in multimedia conferences, leaving calls at the mercy of everyone on the network (including other IP Telephony users). Multimedia bandwidth must be protected from data traffic, and vice versa.

Public-switched telephone networks guarantee a minimum quality of service by allocating static circuits for every telephone call. Such an approach is simple to implement, but wastes bandwidth, lacks robustness, and makes voice, video, and data integration difficult. Furthermore, circuit-switched data paths are impossible to create using a connectionless network such as IP.

QoS support on IP networks offers the following benefits:

1. Support for real-time multimedia applications.
2. Assurance of timely transfers of large amounts of data.
3. The ability to share the network in a manner that avoids starving applications of bandwidth.

Quality of service in TAPI 3.0 is handled through the DirectShow RTP filter, which negotiates bandwidth capabilities with the network based on the requirements of the DirectShow codecs associated with a particular media stream. These requirements are indicated to the RTP filter by the codecs via its own QoS interface. The RTP filter then uses the COM Winsock2 GQoS interfaces to indicate, in an abstract form, its QoS requirements to the Winsock2 QoS service provider (QoS SP). The QoS SP, in turn, invokes a number of varying QoS mechanisms appropriate for the application, the underlying media, and the network, in order to guarantee appropriate end-to-end QoS. These mechanisms include:

- a. The Resource Reservation Protocol (RSVP)
- b. Local Traffic Control: Packet Scheduling; 802.1p; Appropriate layer 2 signaling mechanisms
- c. IP Type of Service and DTR header settings

The Resource Reservation Protocol (RSVP) is an IETF standard designed to support resource (for example, bandwidth) reservations through networks of varying topologies and media. Through RSVP, a user's quality of service requests are propagated to all routers along the data path, allowing the network to reconfigure itself (at all network levels) to meet the desired level of service. The RSVP protocol engages network resources by establishing flows throughout the network. A flow is a network path associated with one or more senders, one or more receivers, and a certain quality of service. A sending host wishing to send data that requires a certain QoS will broadcast, via an RSVP-enabled Winsock Service Provider, "path" messages toward the intended recipients. These path messages, which describe the bandwidth requirements and relevant parameters of the data to be sent, are propagated to all intermediate routers along the path. A receiving host, interested in this particular data, will confirm the flow (and the network path) by sending "reserve" messages through the network, describing the bandwidth characteristics of data it wishes to receive from the sender. As these reserve messages propagate back toward the sender, intermediate routers, based on bandwidth capacity, decide whether or not to accept the proposed reservation and commit resources. If an affirmative decision is made, the resources are committed and reserve messages are propagated to the next hop on the path from source to destination.

Packet Scheduling: This mechanism can be used in conjunction with RSVP (if the underlying network is RSVP-enabled)

or without RSVP. Traffic is identified as belonging to one flow or another, and packets from each flow are scheduled in accordance with the traffic control parameters for the flow. These parameters generally include a scheduled rate (token bucket parameter) and some indication of priority. The former is used to pace the transmission of packets to the network. The latter is used to determine the order in which packets should be submitted to the network when congestion occurs.

801.2p: Traffic control can also be used to determine the 802.1 User Priority value (a MAC header field used to indicate relative packet priority) to be associated with each transmitted packet. 802.1p-enabled switches can then give preferential treatment to certain packets over others, providing additional quality of service support at the data link layer level.

Layer 2 Signaling Mechanisms: In response to Winsock 2 QoS APIs, the QoS service provider may invoke additional traffic control mechanisms depending on the specific underlying data link layer. It may signal an underlying ATM network, for instance, to set up an appropriate virtual circuit for each flow. When the underlying media is a traditional 802 shared media network, the QoS service provider may extend the standard RSVP mechanism to signal a Subnet Bandwidth Manager (SBM). The SBM provides centralized bandwidth management on shared networks.

Each IP packet contains a three-bit Precedence field, which indicates the priority of the packet. An additional field can be used to indicate a delay, throughput, or reliability preference to the network. Local traffic control can be used to set these bits in the IP headers of packets on particular flows. As a result, packets belonging to a flow will be treated appropriately later by three devices on the network. These fields are analogous to 802.1p priority settings but are interpreted by higher layer network devices. www.microsoft.com/communications. See H.323.

TAR TAR is the UNIX standard program for combining and compressing files. It's like a UNIX Winzip program. I believe it stands for "Tape Archiver". The UNIX System V manual lists the command description as "tape file archiver". One no longer uses it just for tape archiving, but that was its initial usage.

TARGA Truevision Advanced Raster Graphics Adapter.

Target A SCSI device that performs an operation requested by an initiator.

Target Host Number The number that identifies the destination software program during the user logon process.

Target Token Rotation Time. TTRT. An FDDI (Fiber Distributed Data Interface) token travels along the network ring from node to node. If a node does not need to transmit data, it picks up the token and sends it to the next node. If the node possessing the token does need to transmit, it can send as many frames as desired for a fixed amount of time.

Tariff Documents filed by a regulated telephone company with a state public utility commission or the Federal Communications Commission. The tariff, a public document, details services, equipment and pricing offered by the telephone company (a common carrier) to all potential customers. Being a "common carrier" means it (the phone company) must offer its services to everybody at the prices and at the conditions outlined in its public tariffs. Tariffs do not carry the weight of law behind them. If you or the telephone company violate them, no one will go to jail. The worst that can happen to you, as a subscriber, is that your service will be cut off, or threatened to be cut off. Regulatory authorities do not normally approve tariffs. They accept them — until they are successfully challenged before a hearing of the regulatory

body or in court (usually Federal Court). Many tariffs were accepted by regulatory commissions only to be struck down in court as unlawful, discriminatory, not cost-justified, etc. Monies collected under the tariff have been refunded and unnecessary equipment removed. In these new, competitive days, many telephone companies are violating their own tariffs by charging less money than their tariffs say they should, or bundling services together at a discount. They are also providing service and equipment on terms less onerous than outlined in their tariffs. Many users now regard tariffs as starting bargaining points, rather than ending bargaining points.

Tariff 12 A user-specific long distance tariff of AT&T. Tariff 12 gives AT&T the ability to price its long distance services for one company practically any which way it feels — giving them a mix of services at stable prices over the long term with significant volume discounts. As this dictionary was going to the printer, a federal appeals court overturned the Federal Communications Commission's April 1989 decision allowing AT&T to offer custom networks and ordered the FCC to reopen its investigation into the legality of the Tariff 12 deals. There are still some users. They are "grandfathered" until we get a final say on the tariff. And, as we go to press, AT&T can offer Tariff 12 customized services to any company — but cannot include 800 services in its Tariff 12 pricing.

Tariff 15 A user-specific long distance tariff of AT&T. Tariff 15 gives AT&T the ability to price its long distance services for one company practically any way it feels. Tariff 15 is single-customer discounting. Some of AT&T competitors claim the tariff is "illegal."

Tariff Rebalancing Largely an initiative of the FCC in the US, national, regional (e.g., EU) and international (ITU) regulatory authorities and policy-making bodies are considering the rebalancing of tariffs in order that they be more closely related to the costs of providing the various telecommunications products and services. While most attention is focused on the accountings rate for long distance, both domestic and international, tariff rebalancing encompasses all tariffed products (including equipment rentals) and services (e.g., local service), and across both the business and consumer domains. Historically, tariffs for individual products and services were designed in the context of the overall tariff structure, which addressed the full range of such products and services. Within this overall structure existed a complex set of cross-subsidies which generally resulted in low tariffs for basic consumer services, such as residential local service. Relatively high tariffs existed both for optional services, such as long distance, international and custom calling features. The primary justification for this arrangement was that of the desire to gain "universal service" — a phone in everybody's home. Mixed in with this was the concept of "ability to pay." Further, individual consumers vote, while companies do not. As consumer rights advocates became more vocal during the past twenty years or so, the pressure to retain these cross-subsidies increased. It generally is recognized, however, that the traditional tariff structure places incumbent carriers and service providers at a decided disadvantage in a competitive environment, as the pricing policies of the newer competitors are not constrained by tariffs, nor strange societal concepts of universal service and ability to pay. Further — and this is perhaps a major impetus — there exists a clear imbalance in the accounting rates for international long distance calls. For instance, a call from Argentina to the U.S. is much more expensive than is a call in the reverse direction, even though the costs for call origination are roughly equal to the costs of