MadgeOne

Architecture for Multiservice Switched Networks

A TECHNOLOGY WHITE PAPER

AUTHOR: Martin Taylor, Madge Networks, July 1996

Executive Summary

This white paper describes Madge's approach to the provision of sophisticated network services to the desktop, over a switched LAN infrastructure with switched WAN access. The need for these new services arises from requirements to guarantee the performance of mission-critical data applications, the desire to provide support for interactive desktop multimedia conferencing over the LAN, and the emerging requirement to integrate telephony services over the LAN. Networks that support all existing applications and, in addition, offer these new services are described as "multiservice networks". The paper explains how Madge will deliver multiservice networking capabilities, focusing on the exploitation of ATM technology. Topics covered include support for existing data applications, transport of real-time voice and video streams, call control for person-to-person communications applications and WAN connectivity for data, voice and video.

About the Author

Martin Taylor is Vice President, Network Architecture for Madge Networks. He is responsible for the strategic planning of the company's products, including LAN switching hubs, ATM switches and LAN-ATM access switches. Martin joined Madge Networks in 1991, and formerly held the position of Director of Product Marketing. Prior to joining Madge, Martin held business development positions in Local Area Networks and in fiber optic cabling systems with GPT Ltd., the UK's leading supplier of telecommunications equipment. Martin has a total of 10 years experience in the communications industry, and a further 6 years experience in MIS management. He can be reached on email at mtaylor@madge.com.

Today's Single-Service Networks

Today, the majority of desktops in large enterprises are equipped with two kinds of communications terminals – the PC and the telephone. These devices are supported by two entirely separate kinds of network connection – the PC being connected to a LAN, while the telephone is connected to a PBX. Each of these networks provides a single "service". The telephone network delivers bi-directional streams of data at 64 kbps (typically digitally-encoded voice) while the LAN delivers packets of data on a "best effort" basis.

The technologies that lie behind the telephone network and the LAN are very different. Each has evolved over the years according to the specific needs of the service they provide. LANs have become faster, cheaper and more reliable. Telephone networks connect calls faster than they used to, and have acquired many sophisticated call control functions. Both technologies continue to provide essentially the same, single service as they have always done.

It is becoming clear, however, that the single services delivered to our desktops respectively by LANs and telephone networks do not offer the best solution to current business needs and will fall far short of meeting emerging needs.

- Today's LANs (and LAN internetworks) give the same priority to all users regardless of their application. As a result, time-critical business applications suffer from highly variable performance when they share a common infrastructure with bandwidth-hungry file transfers and Intranet Web browsers. Network planners have learned the hard way that throwing bandwidth at the problem is costly and doesn't solve the problem – fatter pipes soon fill up with more traffic, leading to another round of network upgrades.

- More and more businesses are discovering the productivity gains and travel savings that can come from effective use of video-conferencing, but today's room-based systems are often inconvenient and incompatible with spontaneity. Desktop video-conferencing is increasingly affordable, but today's networks don't support video – the LAN can't provide real-time communications while the telephone network can't provide the bandwidth. Most network planners balk at the idea of a third, ISDN-based network to every desktop that needs video.

- The telephone and the PC are both vital tools of business, but the separateness of the two networks is often a frustrating obstacle to productivity. How often do we look up a phone number on a PC-based address book, and then have to punch the number on a phone keypad? See the calling number of an incoming call, but wonder who it is? Wish we didn't have to deal with two message queues, voicemail and email? Businesses benefited on a massive scale when PCs put processing power on every desktop and the mainframe was no longer the only tool for handling data – and we can expect another big round of benefits when that mainframe for telephony, the PBX, is joined by PCs supporting telephony services at the desktop. Today's LANs are not up to handling telephony, and we aren't about to run our data applications over the phone network.

The multiservice network is emerging as the answer to these challenges. A multiservice network is a single network infrastructure that is capable of delivering a range of different services, each designed to meet the needs of a particular application. For a typical corporate desktop, a multiservice network would deliver the familiar best effort datagram delivery service for data, and the bi-directional real-time 64 kbps service for voice. In addition, the multiservice network would provide a range of guaranteed data delivery services for high priority time-critical data applications, as well as real-time delivery of video, voice and data for interactive conferencing applications.

Real business applications have a range of different needs for services from the network, and the diversity of needs is increasing with new and emerging applications. Today's single service networks are ill-matched to meet this diversity of needs. The time has come for multiservice networks.

Applications for Multiservice Networking

The multiservice network of the future will provide a high performance solution for all existing data networking needs, as well as three new kinds of capability which have great potential value for many kinds of enterprises:

- Ability to reserve network resources for critical data applications
- Support for new kinds of interactive real-time multimedia applications
- Opportunity to integrate telephony with the data network

Addressing Existing Data Networking Needs

As its base level of functionality, a multiservice network will provide transparent support to all existing data applications over a high performance transport for all commonly used networking protocols, including IP, IPX, NetBIOS, SNA etc. The infrastructure for multiservice networking may include frame switching and cell switching capabilities to provide scaleable bandwidth, and routing or Layer 3 switching capabilities to move packets between subnets. Additional kinds of network intelligence, such as broadcast control, may also be provided.

Reserving Resources for Data Applications

Networked business applications are not all created equal, but they are all treated equally by today's network infrastructures. In the early days of the industry, applications requiring only terminal access to a mainframe used a tiny proportion of the bandwidth of a 10 Mbps LAN. Today, networks support a wide variety of different application types including client/server, file and print services, Intranet Web access, database replication and server backup, all of which demand far more bandwidth than mainframe access. Furthermore, the processing power available at the desktop is capable of generating bursts of network traffic at rates of 100 Mbps or more, if the bandwidth is available. Users of interactive applications can suffer unacceptable variations in response time as a result of other network activity.

The traditional approach to dealing with network performance problems in the LAN has been to over-provision bandwidth. The drawback to this approach is that wider highways tend to generate more traffic, and network planners have learned that the performance benefits of each upgrade are short-lived.

With multiservice networking, the possibility exists for network administrators to identify those applications that are time-critical and to allocate network resources to these applications. The response time for these applications is now insulated from the effects of other network activity. The remaining bandwidth that has not been reserved is available for all other traffic to share on an equal basis.

This does not remove the necessity to upgrade the total capacity of the network as overall demand for bandwidth grows. It avoids the need for massive over-provisioning, while ensuring much more consistent response times to users of mission-critical interactive applications.

Supporting Interactive Multimedia Applications

The value of videoconferencing, both for improving intra-enterprise communications and for reaching out to customers and prospective customers in new ways, has been well proven. Real-time video requires a modest amount of network bandwidth – in the range of 128 kbps to 1.5 Mbps typically – but requires guaranteed end-to-end network delay of a few tens of milliseconds for acceptable interactivity. Existing LANs and internetworks cannot guarantee such low delays, and existing phone networks cannot generally support bandwidths above 64 kbps at the desktop. Today, desktop video is best supported with a separate network infrastructure providing ISDN services at the desktop, but it would clearly be preferable to deliver video over a common LAN infrastructure.

A new standard has just emerged for desktop videoconferencing over the LAN. The standard is known as H.323 and is published by the International Telecommunications Union (ITU). This standard has been developed to support videoconferencing over LANs that specifically do <u>not</u> support Quality of Service – i.e. existing LAN technologies. But the performance of H.323 applications will only be as good as the LAN infrastructure allows. When the infrastructure is under stress from many competing demands for bandwidth, videoconferencing performance is likely to suffer from the break-up of both image and voice, or from excessive end-to-end delays.

Multiservice networking will allow bandwidth to be reserved for the voice and video streams generated by H.323 applications and will guarantee low end-to-end delays. This will ensure that the full potential of any investment in H.323 videoconferencing applications is realized by maximizing the perceived quality of real-time communications.

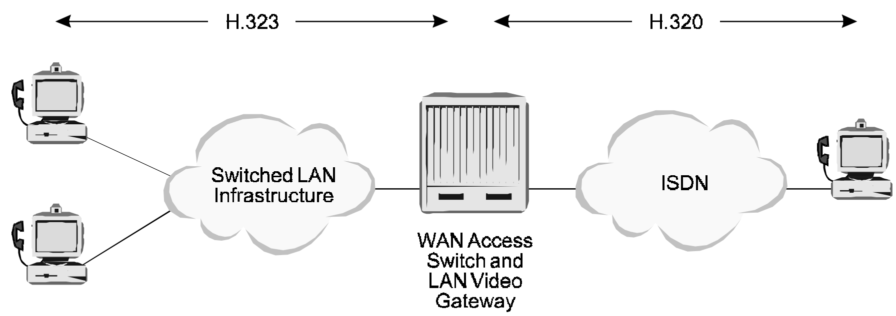

Figure 1: Videoconferencing protocols in a mixed LAN/WAN environment

Opportunity for Telephony Integration

Telephony is the most pervasive tool of business activity, and the use of the telephone to support the business of the enterprise is often closely associated with the use of networked data applications. But telephony and data exist separately everywhere in the LAN except perhaps in call centers, where typically Computer Telephony Integration (CTI) links exchange information between data servers and voice switches.

In call centers, the CTI link may be used to direct incoming calls to the right agents based on the phone number of the calling party, and present information about the call to the agent – sometimes known as a “screen pop”. But enterprises for whom call centers are an important focus of business can benefit from a much tighter integration of telephony and data. For example, this could make it easier to build a database providing records and analysis of incoming calls, to support business decisions. It could also make it easier to automate the “user interface” presented to the caller with Interactive Voice Response (IVR).

A multiservice network that supports voice alongside data on the same infrastructure would allow telephony to be treated as simply another application on the network, and allow phone calls to be treated as data objects. This would allow the businesses to apply, for the first time, all the power of client/server techniques to telephony. Further, the integration of voice and data onto a single infrastructure can result in substantial capital cost savings.

In a multiservice network, general purpose PCs, employed by a company's knowledge workers and administrative staff, can be used for telephony applications. The added processing power PCs feature, compared to desktop telephones, promises to make these workers more productive by improving the way they perform existing telephony applications and spurring the creation of new applications. Previously, the only workers that could benefit from the application of PC processing power to telephony applications were call center staff.

A LAN's ability to provide advanced voice services to users is increasingly being recognized by large organizations. According to a March 1996 survey of Fortune 1000 companies by the Yankee Group, voice traffic on the LAN, including real-time applications such as desktop videoconferencing and standard telephone calls, non-real-time applications such as voice mail retrieval and distance learning, is today present on about 17 percent of corporate networks. In five years, voice applications on the LAN will be implemented by 42 percent of Fortune 1000 companies. In fact, according to the Yankee Group survey, voice traffic for the next five years will grow faster than any other type of traffic on LANs except for video. Other types of traffic quantified by the survey include traditional LAN and host applications. In addition, Tom Nolle of CIMI, Corp. states, "Corporate key workers, dependent on their personal computers for productivity support, are now looking to use their computers to mediate in their relationships with other workers. Voice data integration on a LAN opens a new set of collaborative opportunities for these workers."

Bryan Van Dussen, director of telecommunications services for Yankee Group adds, "In a reversal of the data-rides-for-free phenomenon that was characteristic of enterprise networks in the 80s, the Yankee group believes that voice communications will move back onto private enterprise networks from VPNs within the next three years. Supported by advances in broadband- and LAN-telephony technologies, and the economic efficiencies they create, voice will ride for free on ATM- and cell-based enterprise networks."

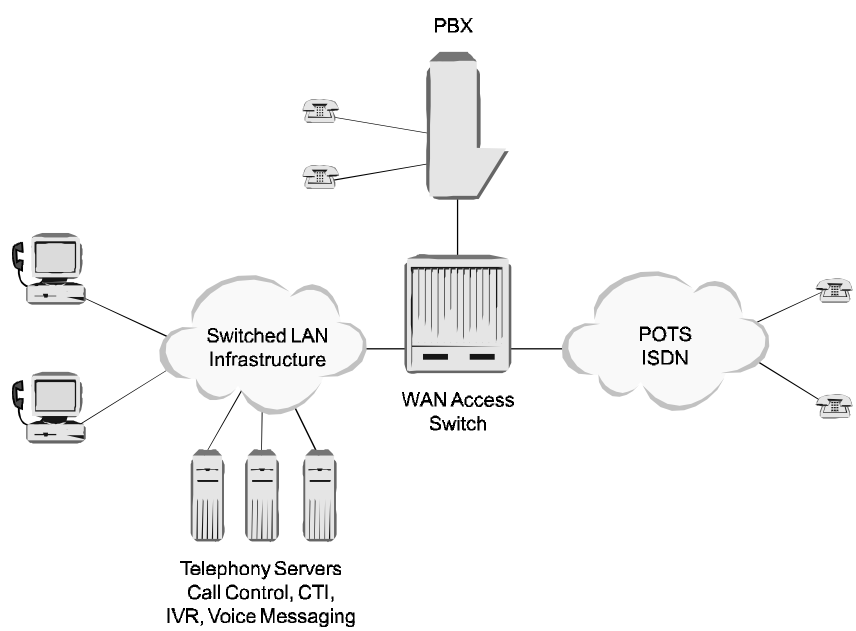

Figure 2: Integrated Telephony on the LAN

Technologies for Multiservice Networking

Today's LANs and internetworks based on Ethernet and Token Ring with LAN switches and routers do not inherently provide multiservice capability. These networks provide a single service whereby all packets are treated with the same priority and transferred as quickly as possible across the network. But the time taken for a packet to reach its destination will vary according to the instantaneous loading on the network – when there is a lot of other activity, then delivery may take quite a lot longer. There are two important implications of this:

- All data applications are treated the same. There is no way to distinguish between mission-critical applications that need a rapid response from the network and other tasks that are less important. The performance of all applications, no matter how important, is therefore subject to variability.

- There is no way reliably to run applications that need real-time communications like telephony or videoconferencing over the network, because these applications rely on a steady stream of data being delivered with low end-to-end delay.

In seeking a technology solution for multiservice networking, there appears to be a choice between enhancing existing frame-based networking technologies, or looking to a new technology that already has the required capabilities: ATM. Before we discuss these two approaches, we need first to establish the requirements for acceptable performance of multimedia applications, in terms of maximum end-to-end delay across the network.

End-to-End Delay Requirements for Multimedia Applications

Today, voice and video services are generally delivered over circuit-switched networks such as the public phone network or ISDN. We take it for granted that digitized voice and video information will be delivered to the far end in a steady stream with low delay. When we put digitized voice or video into packets or cells and send it across a frame-switched or cell-switched network, we introduce some additional delays that affect the perceived quality of the communication.

End-to-end delay for voice and video streams in packet-based or cell-based networks has three main components:

- Packetization delay – the time taken to collect enough voice or video information to fill a reasonably-sized packet for transmission.

- Network delay – the minimum time that a packet can take to traverse the network from end-to-end, including propagation delay plus store-and-forward and queuing delays in the switching or routing elements in the network.

- Receive buffer delay – the delay introduced by the play-out buffer at the receiving station which compensates for delay variations introduced by the network.

Packetization delay is easy to quantify. For example, if we put a digitized telephony-quality voice stream into ATM cells, we are taking a 64 kilobit/second stream and using it to generate 48-byte cell payloads. The time taken to fill each cell is (48 x 8)/64,000 seconds, or 6 milliseconds.

Network delay includes propagation time plus best-case store-and-forward delays. Propagation time amounts to about 1 millisecond per 200 kilometers. Store-and-forward delays in switches and routers depend very much on the design and architecture of the particular devices used, and can vary from fractions of a millisecond to hundreds of milliseconds in some cases.

The final major component of delay is that due to the receive buffer, which is necessary to compensate for variations in the arrival time of packets or cells across the network. Due to fluctuations in network loading across the various links and switching points that make up real networks, there is always some variability of delay introduced. This means that packets or cells transmitted at exactly regular intervals will arrive at the receiving end at somewhat irregular intervals – a phenomenon known as "jitter".

In order to ensure that information is presented in a steady stream to the audio or video output device at the receiving end, the receiving station must buffer the information in a "First-In, First-Out" (FIFO) buffer. Information is extracted from arriving packets and pushed into the buffer, and at the far end the information is played out to the appropriate output device. The buffer has to be sized to allow for the worst case delay variation or jitter that the network can introduce so as to ensure that the station never runs out of information to send to the output devices.

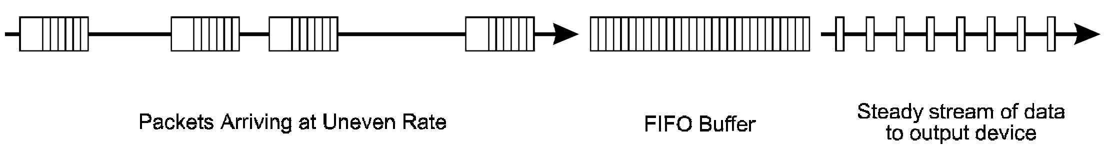

Figure 3: First In, First Out Buffer for Smoothing Delay Variation

The issue with delay is that it introduces degradation to the perceived quality of person-to-person communication. Delay does not impact on the fidelity of reproduction of voice or video, but it does make interaction between people noticeably more difficult. The issue is familiar to anyone who has made an international phone call that has been routed over a satellite link. The delay introduced by the satellite link means that both parties hear an awkward silence when one party stops speaking, and both then start speaking at the same time. Then both back off and there is an awkward silence again.

Delays of this magnitude are quite unacceptable for regular business voice communications in a local environment, where the existing phone network sets the benchmark for quality and delays of less than 20 milliseconds are the norm. Interactive video communications are a little less sensitive to network delay, since the process of digitizing and compressing the video signal adds a delay of perhaps 150-200 milliseconds. For video, a network delay of less than 75 milliseconds yields acceptable performance. One-way voice and video communications, such as audio or video broadcast feeds, involve no person-to-person interaction and so are much less sensitive to network delay. For these kinds of applications, up to 2,000 milliseconds of network delay is tolerable.

Enhancing Frame-based Technologies

The need for multiservice networking in the enterprise is now widely recognized by the networking industry, and much effort is being expended on developing new techniques and new standards to provide some level of multiservice capability based on existing technologies. However, the problem is extremely complex and requires new solutions in the following areas:

- How to identify priority traffic at both Network Layer (for example, IP) and Data Link Layer (for example, Ethernet).
- How to reserve bandwidth and network resources for traffic that requires guaranteed end-to-end delay.
- How to police and control the reservations of bandwidth over the LANs and on the inter-router links.
- How to take account of bandwidth reservations in determining the optimum routing for any particular session through the router network.
- How to apply all of these techniques in a multicast environment.
- How to handle and manage multiple queues containing traffic of different priority in both routers and LAN switches.

All of these questions (except the last, which is a vendor implementation issue) are now under consideration by the various standards bodies. The Internet Engineering Task Force (IETF) is taking care of the enhancements needed to the IP family of protocols, while the IEEE 802 project is working on the implications for the 802.x family of Ethernet and Token Ring protocol standards.

Considerable progress has already been made. For example, the IETF have defined the Resource ReSerVation Protocol, which allows an end station to request from routers that a specified Quality of Service should be applied to a particular stream of data that it wishes to receive. However, there is still a great deal to be accomplished. Some areas of the work that is required may involve fundamental changes in existing technologies. For example, it will probably be necessary to change the basic Ethernet frame format to accommodate the necessary priority level indications, and to alter the way that the Ethernet CSMA/CD media access algorithm works. This is likely to require new silicon in Ethernet switches and perhaps even in Network Interface Cards (NICs).

Furthermore, adding multiservice capabilities to routers and LAN switches will not be easy. Determining the priority of any given packet or frame and managing multiple output queues per router or switch port will require a great deal of additional per-packet processing. In order to maintain acceptable performance, new silicon will be required to assist in this process. This means that routers and LAN switches are likely to need hardware upgrades to support multiservice networking.

Building multiservice capability on the basis of existing frame-based LAN and internetworking technologies is a huge task. It will not be easy to reach agreement on new standards among the many vendors that have an interest, and it will be a long and complex job to implement the necessary new capabilities in routers and LAN switches. In summary, working multiservice networking solutions based on existing frame-based technologies are several years away from reality.

Cell-based Technologies: ATM

Unlike the frame-based technologies of Ethernet, Token Ring and IP, ATM was conceived from its very inception as a technology for delivering multiple services, including voice, video and high-speed data communications, over one single integrated network.

All of the complex needs identified above were taken into account from the very beginning of the development of ATM.

Working with small, fixed-size units of data called “cells” allows ATM to deliver very low end-to-end delay for time-sensitive traffic like voice and video. Delay times for ATM cells passing through ATM switches are measured in microseconds, not milliseconds, and packetization delays are small because each cell only carries 48 bytes of traffic payload.

Applying the appropriate Quality of Service to each traffic flow is made easier by the fact that ATM is connection-oriented. Each ATM cell is associated with a particular connection, and the ATM switches that make up an ATM network store details of the required Quality of Service of each connection. Bandwidth and network resources are reserved as necessary when an end-station signals the ATM network that it needs a connection to some particular destination. The signaling protocols needed to do this (User-Network Interface, or UNI) are fully standardized, and requests to reserve bandwidth are policed by an established technique known as Call Admission Control. The standardized Private Network-to-Network Interface (PNNI) protocol deals with connection routing, taking bandwidth reservations into account. Finally, multicasting is fully supported through the capabilities of Point-to-Multipoint Switched Virtual Circuits, which in turn are supported by all the various ATM standards.

Naturally, the development of all the standards necessary for ATM to solve these complex problems has taken a considerable time. The process is now nearly complete. Since early 1996 ATM has been able to provide practical, standardized solutions for data networking, including transparent interoperability with Ethernet and Token Ring LANs and support for existing LAN-based applications. Standards for video applications over ATM already exist, and standards which will allow voice telephony applications to run over ATM are expected to be complete by the end of 1996.

MadgeOne: Integrated Multiservice Networking

The MadgeOne architecture sets out to meet the full set of data, voice and video networking requirements across the end-to-end network span, desktop to desktop. It embraces switched connectivity to existing premises networks for data, voice and video: Ethernet, Token Ring, PBXs and video room systems. And it also embraces switched or routed connectivity with external services such as ISDN and the public switched telephone network, Frame Relay, and the Internet.

Initially, MadgeOne uses the capabilities of ATM to deliver multiservice networking, because today ATM is the only technology that is truly capable of handling all data, voice and video networking needs in one single integrated approach to networking. Wherever possible, MadgeOne extends the reach of integrated data, voice and video networking over other technologies such as switched Ethernet and Token Ring. And as frame-based technologies evolve to support Quality of Service, MadgeOne will embrace the new enhancements to Ethernet, Token Ring and IP routing to provide a greater choice of migration paths to multiservice networking.

ATM: The Foundations of MadgeOne

ATM comprises a set of technologies by which data, voice and video may be transported in the form of fixed-size cells over a network comprising switching devices linked together at a range of different speeds. This network is supported and controlled by signaling protocols which enable end stations to request from the network a connection with characteristics that are appropriate to the type of traffic – data, voice or video – that needs to be carried. Thus ATM provides the core functionality needed for the realization of MadgeOne.

However, while ATM in principle provides a complete solution to all aspects of data, voice and video networking, in practice ATM will be deployed in situations where existing and disparate networks are already in place, supporting various kinds of networked applications and voice and video communications. MadgeOne utilizes a range of standardized techniques to provide for interoperability and interconnection between these existing networks and ATM-based solutions.

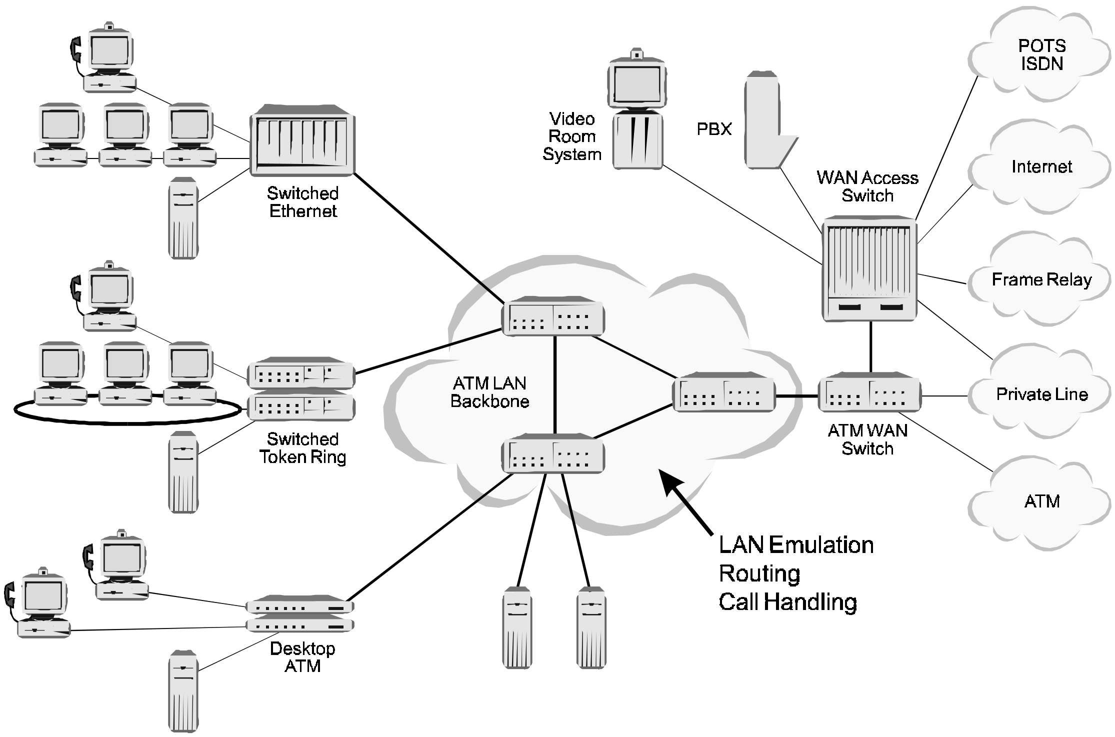

Figure 4: Elements of the MadgeOne Architecture

LAN Emulation over ATM (LANE)

Most existing LAN-based data applications are designed to run over Ethernet or Token Ring networks. LAN Emulation over ATM was conceived as a means to support such existing applications to run over ATM, and to make it easy to interconnect ATM networks with Ethernet and Token Ring networks.

The essence of LAN Emulation is that the ATM network carries Ethernet or Token Ring data frames, segmented into ATM cells. Switching devices that link physical Ethernet or Token Ring LAN segments to ATM can use LAN Emulation to move Ethernet or Token Ring data frames across the ATM network. End stations connected directly to ATM can also talk in the familiar language of Ethernet or Token Ring data frames. With LAN Emulation, interoperability between end stations attached to physical Ethernet or Token Ring LAN segments and end stations attached directly to ATM is easily achieved.

LAN Emulation over ATM requires some central services to be provided in the ATM network, usually in the form of software running on ATM switches. This is because Ethernet and Token Ring are connectionless protocols, whereas ATM is connection-oriented, and the ATM LAN Emulation services are needed to take care of setting up the required connections as and when LAN stations initiate communication between each other. Also, ATM does not support the any-to-any modes of broadcasting that most LAN protocols use, for example, to enable end stations to find servers. To cover this need, ATM LAN Emulation services take care of forwarding broadcasts to all stations that belong to the emulated LAN.

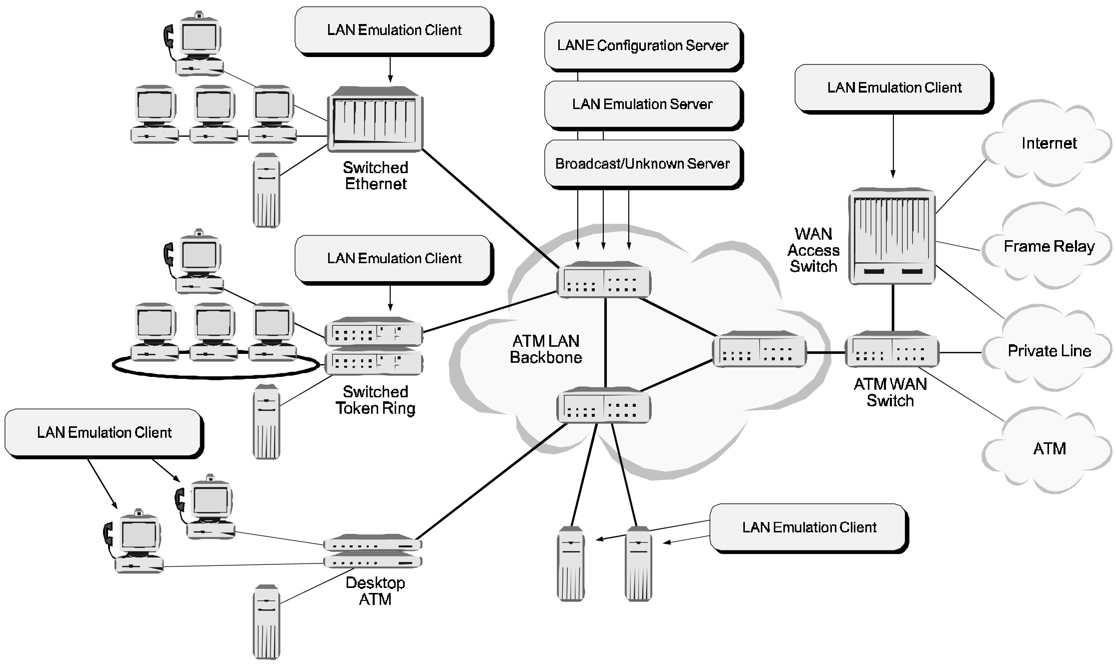

Figure 5: Software Components of LAN Emulation over ATM

Multi-Protocol over ATM (MPOA)

With LAN Emulation, many of the concepts of switched LANs are carried over into the ATM domain. While this has the virtue of simplicity, it also has certain drawbacks. For example, there is a practical limit to the number of stations that can be allocated to a single emulated LAN, because of the amount of broadcast traffic that will be generated and seen by all the stations. Therefore large network installations must be logically segmented into multiple emulated LANs, and routers are needed to perform packet transfer between these emulated LANs.

The technique of Multi-Protocol over ATM is intended to deliver a more sophisticated mapping of LAN and internetworking protocols such as IP over ATM networks. Instead of converting Ethernet and Token Ring frame formats into cells at edge devices or end stations, MPOA places IP datagrams directly into ATM cells. MPOA also takes care of setting up Switched Virtual Circuits between end-points of an ATM network to enable end stations or edge devices to intercommunicate directly without having to pass through any routers – even if the IP addresses of these end-points belong to different IP network numbers or subnets.

Like LAN Emulation, MPOA is based on central software services in the ATM network to resolve addresses, forward broadcasts where necessary and set up Switched Virtual Circuits. However, because of the greater scope of the MPOA solution, these central services are necessarily much more complex that those needed for LANE. LANE version 1.0 was standardized by the ATM Forum in May 1995, while MPOA is not expected to be emerge as a standard until 1997.

Data Networking over ATM with MadgeOne

Until MPOA is standardized, not just for IP but for other important LAN protocols such as Novell's IPX, LANE remains the only practical solution for interconnecting Ethernet and Token Ring LAN segments with ATM. Since LANE enables the ATM network to transport Ethernet and Token Ring frames directly, regardless of the internetworking protocol that may be contained within the frames, LANE makes it possible to connect ATM easily to Ethernet and Token Ring LAN switches. The ATM network is simply treated by the switch as a very fast Ethernet or Token Ring port, and switching decisions are made in the normal way based on MAC address or Source Routing information as appropriate.

The MadgeOne architecture makes use of LAN Emulation both in Ethernet and Token Ring to ATM access switches, and in the drivers on ATM Network Interface Cards, to support data applications over ATM. The MadgeOne implementation of LAN Emulation will include Madge's Active Broadcast Control technology, which employs intelligent self-configuring broadcast filters to achieve major reductions in broadcast traffic flow. This in turn allows for larger emulated LANs to be constructed, to provide flatter and faster switched networks with a greatly reduced requirement for router capacity. Routing between emulated LANs can be carried out with third party routers, or in the future with router modules that will be options for Madge ATM switches. To support data applications which require a guaranteed response time from the network, MadgeOne will allow Quality of Service parameters to be applied to emulated LANs.

The deployment of MPOA changes somewhat the nature of the connection between ATM and Ethernet or Token Ring LANs. The data transfer that takes place in an MPOA edge device involves extracting the IP or IPX datagram from Ethernet or Token Ring frames, and then segmenting it into ATM cells. In this sense an MPOA edge device behaves like a router, except that it does not need to run any complex routing protocols – the central MPOA services take care of locating the destination end station or edge device and resolving the target IP address to an ATM address, and the ATM switches determine the routing of the Switched Virtual Circuit.

MPOA represents a further advance of switching technology over routers, since with MPOA the need for routers to link emulated LANs effectively disappears. MPOA figures strongly in the evolution of the MadgeOne architecture: In the future, Madge ATM access switches for Ethernet and Token Ring will be upgradable to support MPOA, and Madge ATM switches will support MPOA central services as part of their embedded software load. MPOA will also support data applications that require guaranteed response from the network by allowing Quality of Service parameters to be specified on a session-by-session basis.

Both LANE and MPOA will support data applications over ATM in the WAN. With standards in place for ATM to operate over T1/E1 and even lower speeds, this provides a complete solution for the interconnection of large and medium sites. The Active Broadcast Control capability in MadgeOne’s LAN Emulation is particularly valuable here, since it greatly reduces the amount of WAN bandwidth that is consumed needlessly by broadcast traffic.

For smaller sites and telecommuters where a T1/E1 connection is not appropriate, MadgeOne provides for router-based connectivity over Frame Relay, X.25, private lines, dial-up connections or the Internet. Router-based connectivity will be supported by means of an optional router module for the Madge WAN AccessSwitch.

MadgeOne and Virtual LANs

By implementing standard techniques for LAN Emulation over ATM, MadgeOne supports the definition of multiple Ethernet and Token Ring emulated LANs within the ATM infrastructure. The concept of an emulated LAN in the ATM domain is an extension of the concept of a Virtual LAN (VLAN) in the Ethernet or Token Ring domain. The mapping of ATM end stations and ports on Ethernet and Token Ring access switched to these emulated LANs is configured at the management console, and made easier by VLAN management tools that learn about stations on the network – which subnets they belong to, and which server resources they connect to.

VLAN and ELAN management offers the means to divide the network logically into broadcast domains, and provides far greater flexibility than the physical broadcast segmentation that is most commonly implemented today, with routers separating switched or shared LAN segments. Unlike other vendors’ approaches to logical segmentation, which demand the definition of many separate VLANs linked by routers to prevent broadcast traffic getting out of control, MadgeOne’s Active Broadcast Control technology provides for highly effective limitation of broadcast traffic within the VLANs and ELANs. This means that much larger VLAN/ELAN broadcast domains can be defined, which reduces both the administrative burden of managing VLAN membership lists and reduces the total amount of routing capacity needed in the network to support packet transfer between VLANs.

MadgeOne for Voice and Video

Networked data applications are supported in PCs today by low-level programming interfaces which enable Ethernet or Token Ring frames to be sent to, or to be received from, the LAN. The technique of LAN Emulation over ATM enables end stations directly connected to ATM to support these applications without any change because the ATM adapter card driver software emulates an Ethernet or Token Ring programming interface from the point of view of the application.

The relatively small number of networked voice or video applications in use on LANs today are mostly based on adaptations of existing data-oriented protocol stacks. For example, if we use IP-based Internet phone software on a LAN-attached PC, the voice is being sent as the payload of an IP datagram over an Ethernet or Token Ring network.

Unfortunately, the performance of voice and video applications running over IP is compromised by the inability of today's networks to guarantee end-to-end delay. We have already discussed above the requirements that need to be met for end-to-end delay if we are going to provide business-quality voice and video communications – and these requirements cannot be met by today's IP-based Ethernet and Token Ring networks.

Until frame-based network technologies are enhanced to provide guaranteed end-to-end delay, the only viable solution is to use ATM-based techniques. This is the approach taken by MadgeOne – real-time voice and video streams generated at ATM end stations are placed directly into ATM cell payloads, bypassing the transport and network layers of the protocol stacks such as TCP/IP that support data applications. This implies the need for "real-time stream" driver software for ATM adapter cards, that runs alongside the LANE or MPOA drivers used by data applications. Madge will deliver the driver software needed to support real-time streams over ATM, and will engineer this software to support third-party applications over standard APIs such as Winsock 2.0.

Despite the superiority of native ATM modes of transport for real-time streams, there will still be many voice and video applications designed to run over IP. Since MadgeOne provides a highly effective solution for the transport of standard IP traffic, these applications are fully supported. And as the techniques evolve for providing Quality of Service over frame-based networks, the scope of MadgeOne will broaden to include support for these also.

The Voice and Video Connection to Ethernet and Token Ring

We have discussed the use of ATM networks to handle voice and video by taking advantage of the inherent capabilities of ATM to set up connections with reserved bandwidth, which guarantee the delivery of real-time streams with very low delay. This clearly means that we can bring toll-quality telephony to desktops that are equipped with direct ATM connections.

Currently, the costs of bringing ATM directly to the desktop are somewhat higher than the costs of other technologies such as switched Ethernet and Token Ring. Furthermore, there is a huge installed base of Ethernet and Token Ring at the desktop. It would be desirable, therefore, to support telephony at the desktop over not only direct ATM connections, but also Ethernet and Token Ring connections. This can be achieved with a combination of LAN switching in the workgroup and ATM in the backbone.

As the costs of LAN switching fall, and the demand for bandwidth in the LAN continues to grow, we can expect to see more and more networks bringing the ATM backbone to the wiring closet and providing connections to this backbone via switched Ethernet or Token Ring. The challenge then is to support telephony over Ethernet and Token Ring only between the desktop and the wiring closet.

Neither Ethernet nor Token Ring support the reservation of bandwidth to guarantee low-delay delivery of real-time traffic, but there are ways to overcome this difficulty. With Ethernet, we can provide dedicated 10 Mbps switched connections to each desktop economically. The 64 kbps bandwidth of a voice channel to the desktop is a tiny fraction of the available 10 Mbps: even so, there is some chance that a burst of data packets, for example caused by a file transfer, may delay and disrupt packets containing voice. This problem can largely be overcome by tuning the Ethernet access hardware in LAN switches and by providing an adequate play-out buffer to even out any variable delays.

Token Ring supports a prioritization mechanism, which makes things somewhat easier. A Token Ring station with a real-time packet to send can designate this as high priority, in which case it can grab the token ahead of other stations on the ring that have lower priority data packets to send. Taking advantage of the priority mechanism should provide good voice and video performance on shared rings with perhaps 10 to 20 users, connected to a Token Ring to ATM access switch.

Extending the ATM Connection to Legacy Desktops

Dedicated switched Ethernet, and Token Ring's priority mechanism, make it theoretically possible to transport real-time voice and video from non-ATM desktops to the ATM backbone. However, there are a couple of issues for which we must find a solution:

- We need an appropriate protocol and frame format for encoding voice over Ethernet or Token Ring, that makes it easy to transcode the voice or video traffic in the wiring closet onto voice or video over ATM.

- We need some way of extending ATM's signaling protocols all the way to the desktop, so that the desktop telephony or video application can request a connection over the ATM network that has the appropriate Quality of Service for the type of communication required.

A highly effective solution is provided by a technique for carrying ATM cells as the payload of Ethernet or Token Ring frames, known as "Cell-in-Frame". This technique allows an Ethernet or Token Ring connection to an ATM network to be treated as an extension of the ATM network itself, with Ethernet or Token Ring acting as another kind of physical transport layer for ATM cells.

With the Cell-in-Frame technique, a voice stream generated at an Ethernet or Token Ring-connected PC would be encoded into a stream of ATM cells, just as it would if the PC were connected directly to ATM. One or more ATM cells is placed into an Ethernet or Token Ring frame and sent to the access switch in the wiring closet, where a special function (known as the CIF Attachment Device) unwraps the cell or cells from the frame and sends them out over the ATM network.

Cell-in-Frame brings ATM protocols directly to PCs connected by Ethernet or Token Ring. This includes ATM signaling protocols as well as ATM transport. Therefore stations connected by the Cell-in-Frame technique are capable of signaling the set-up of Switched Virtual Circuits across the ATM network.

By bringing ATM protocols to Ethernet or Token Ring desktops, CIF neatly solves both transport and signaling problems for real-time voice and video traffic. By contrast, the use of CIF for data applications is somewhat clumsy. The technique for supporting data applications over ATM is LAN Emulation, which effectively involves taking Ethernet or Token Ring frames generated by the application and segmenting them into ATM cells. With CIF, these cells would then be placed into Ethernet or Token Ring frames. The frame-to-cell-to-frame conversion process involved in this adds a good deal of additional overhead and is likely to compromise performance.

The MadgeOne architecture provides the optimum solution by employing a combination of CIF for real-time voice and video traffic, with conventional Ethernet and Token Ring protocol support for data traffic. This implies dual protocol stacks in the desktop PC, and it also implies that the access switch which links Ethernet or Token Ring connections to ATM must support both the LAN Emulation and CIF Attachment Device processes. Ethernet or Token Ring frames arriving from the desktop at the access switch are directed to the LAN Emulation process if they are conventional data frames containing, for example, IP datagrams. Frames containing ATM cells that carry real-time voice or video traffic are directed to the CIF Attachment Device, for forwarding over the ATM network.

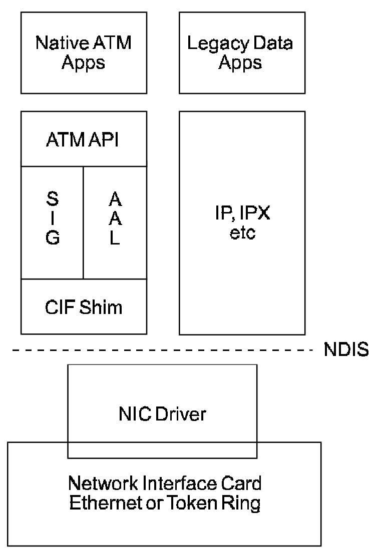

Figure 6: End Station Protocol Stack Supporting CIF and Legacy Data

Madge will provide CIF protocol support software for both Ethernet and Token Ring end stations. This software will be compatible with standard adapter card software interfaces such as NDIS and will therefore run over any standard adapters. Madge will also support the CIF Attachment Device process in future Ethernet and Token Ring to ATM access switches.

CIF-based communications over Ethernet or Token Ring links to the desktop will never deliver the quality of service possible with physical ATM connections at 25 Mbps or 155 Mbps, but are likely to offer an economical solution for bringing ATM's real-time traffic handling to existing desktops.

The WAN Connection

So far we have shown how MadgeOne combines the handling of data with real-time voice and video over an ATM-based local and wide area network, and how we can extend the ATM-based transport for real-time traffic to Ethernet or Token Ring-connected desktops.

Unlike other networking technologies, ATM operates with equal facility in both the LAN and the WAN. This means that data, voice and video integration can be extended seamlessly from the LAN into the WAN. ATM WAN connections can be made over private leased lines, or over public cell-switching services that are starting to become widely available. When deploying the MadgeOne architecture across the ATM WAN, there is one additional consideration that needs to be taken into account: the cost of WAN bandwidth. In the future, MadgeOne will incorporate new capabilities to help make the most efficient use of WAN bandwidth, including enhanced broadcast control capabilities (such as local broadcast spoofing), voice compression and silence suppression on voice connections.

As ATM becomes more widely deployed in the wide area, and as lower speed ATM services become more widely available, it will be possible to extend the span of MadgeOne-based communications over the WAN to all but the smallest sites in an enterprise network. In the meantime, however, it will be necessary to provide connectivity with packet-based data services such as Frame Relay, with the Internet, and with low-speed leased line and dial-up connections that support data-only communications. MadgeOne employs conventional router-based approaches to address these needs.

For real-time voice and video, MadgeOne supports switched connectivity with ISDN, telephony access lines and private lines. This will be provided by Madge's WAN AccessSwitch which will be enhanced by the addition of an ATM interface module, with the ability to convert voice and video streams over ISDN or private line WAN connections to streams of ATM cells for transport over the ATM network. The interface to the ATM network is treated by the WAN AccessSwitch as a trunk connection that supports a large number of 64 kbps or N x 64 kbps channels. The control and routing of calls through the circuit switch and across the ATM network is accomplished by cooperation between the WAN AccessSwitch control software and the call control server application (described in the next section). This cooperation involves the exchange of ISDN signaling protocols between these two devices.

The WAN AccessSwitch will also be enhanced by the addition of a LAN Video Gateway module which will provide translation between LAN-based H.323 video and H.320 video over ISDN and private line connections in the WAN.

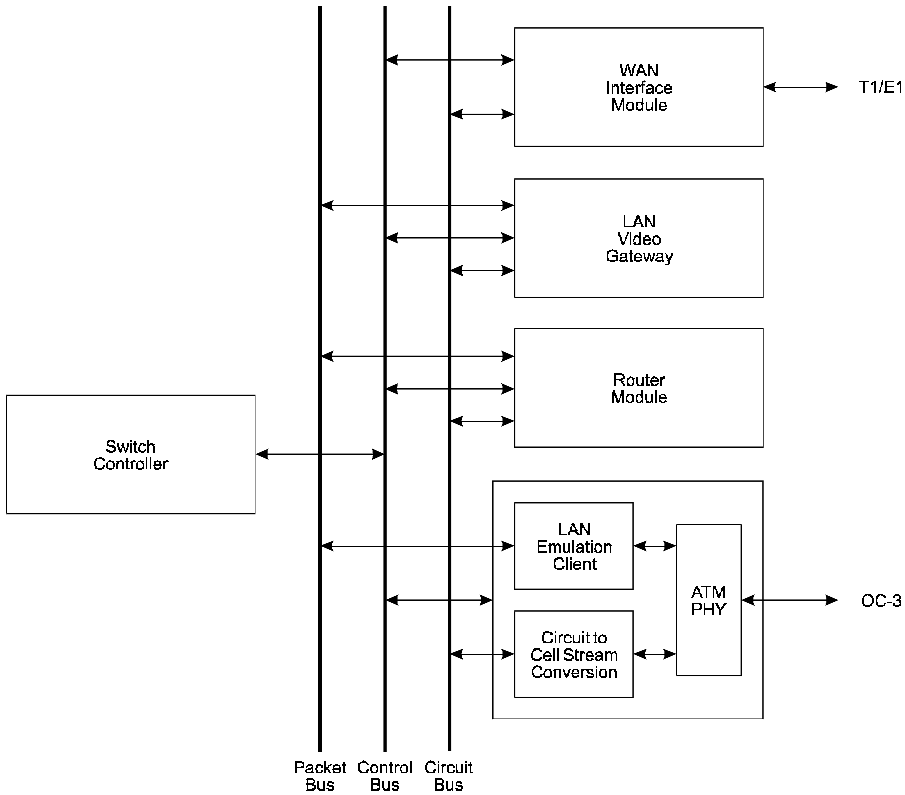

Figure 7: WAN AccessSwitch architecture, showing ATM interface module

MadgeOne Call Control

MadgeOne uses real-time modes of ATM transport physically to carry voice and video streams to and from the desktop, but the handling of voice and video involves much more than just transport. Real-time voice and video involve person-to-person or multi-party calls, and therefore we need a control capability within the network that supports the set-up and control of such calls.

The need for call control barely exists in the data-only environment. Users log on to servers, sometimes with the aid of directory services such as Domain Name Server, and log off when their work is done. But with voice and video calls, we need to deal with events such as incoming calls which are not answered, support features like intrusion when a call arrives for a user who is busy on another call, allow calls to be transferred and diverted, and provide a means for parties to be added to conference calls. These, and many other call control features, are provided by PBX systems today for voice calls. Call control for video calls is still an emerging technology.

ATM is a connection-oriented technology which requires end stations to make use of signaling protocols to set up Virtual Circuits before they can send information. It is possible for a telephony application to make direct use of the ATM signaling protocols to set up a voice call to another user on the network, and to tear down the call when finished. However, telephony users expect a rich set of call control features which are not supported by ATM signaling protocols. For example, there is no way that the current ATM signaling protocols could be used to divert a call to another user or to voice mail if the called party did not answer.

MadgeOne sets out to provide telephony users on the LAN with a full set of call control capabilities which match, or improve upon, those available today from PBXs. Since these capabilities cannot be provided directly by means of ATM signaling protocols, MadgeOne makes use of a variety of telephony control protocols running transparently over the ATM network between end stations and a call control server application. These telephony control protocols are not used as a substitute for the ATM signaling protocols, but as a means of supplementing them.

Client/Server Call Control

With a conventional PBX-based voice system, both call control and switching are carried out in one central location. All requests to connect or control calls are sent to the PBX, which determines what action to take and then controls its own internal switching matrix to build or tear down connections between its ports.

With dumb telephone handsets replaced by PC-based telephones, it is possible to imagine a fully distributed approach to call control, where each PC makes its own local control decisions as to how to set up outgoing calls and deal with incoming calls. However this approach suffers from some serious drawbacks: for example, it provides no way of dealing with an incoming call which is directed at a PC phone that is not currently up and running.

For this reason, MadgeOne implements a combination of centralized and distributed call control. This is provided by means of a call control server application which cooperates with call control intelligence in each PC phone. All call control requests are sent from the PC to the call control server, which then directs the agent software in the PC to set up and tear down ATM Switched Virtual Circuits as appropriate. Voice and video traffic does not, in general, pass through the call control server – it is passed directly between the end stations by Switched Virtual Circuits. In other words, the switching of the calls is done by the ATM network in a distributed manner. The call control server directs the setting up of the calls, not by controlling the ATM switches in the network, but by telling the end stations what connections they are to set up. The complete call control framework is described as the MadgeOne Stream Control Architecture™.

Stream Control Architecture

The MadgeOne Stream Control Architecture is based on two distinct software components at the end station. First is the call control application, typically provided by third party software written to one of the standard telephony application programming interfaces such as Microsoft's TAPI or Novell's TSAPI. This application submits requests for call setup and other control features to the call control server over the network, using standard call control protocols. The second end station software component is the MadgeOne Stream Manager, which responds to commands issued by the call control server, sets up or tears down Switched Virtual Circuits to other end stations as required, and directs voice and video streams to and from the appropriate input and output devices in the end station.

Each end station requires two logical links to the call control server, one to exchange call control requests between the PC telephony application and the call control server application, the other being used to send commands to the Stream Manager. These links are set up automatically when the telephony application in the PC is initialized, usually when the PC is started. The PC telephony or video application "registers" with the call control server, identifying itself by a unique name or extension number. The call control server then knows that this particular user is ready to receive calls, and knows where to send calls that are intended for this user.

The WAN AccessSwitch is treated as a special type of end station by the Stream Control Architecture. The Call Control Virtual Circuit Connection (VCC), which carries the telephony signaling protocols between end stations and the call control server, is connected directly to the switch controller function in the WAN AccessSwitch, while the Stream Client is connected to the Circuit Bus in the switch where it picks up 64 kbps or N x 64 kbps channels as directed by the switch controller and the call control server application. (See Figure 8). Similarly, voice and video application servers can be supported as end stations within the Stream Control Architecture, providing support for applications such as voice and video messaging, interactive voice response and text-to-speech conversion.

The call control server application is supported by software components which provide a standardized application interface, and which include message translation, messaging protocol support and a state replication function which supports the provision of multiple redundant call control servers. These software components, comprising the MadgeOne Stream Server, are designed to support third-party call control applications and multiple different call control protocols, including TSAPI, distributed TAPI, CSTA and Q.931.
To provide further flexibility, the MadgeOne Stream Server supports connections to third-party CTI applications. These interact with the call control server application in much the same way as end station telephony applications, but CTI links can be designated as "master" control applications to which all call routing decisions are referred. A variety of different standard CTI interfaces will be supported.

MadgeOne will extend the sophisticated call control concepts which are widely used today to support voice telephony, and apply these concepts also to multimedia calls, comprising any mix of video, voice and data sharing.

Madge will provide “middleware” components including the MadgeOne Stream Manager and MadgeOne Stream Server to enable third party call control applications in both client and server to be operated over an ATM infrastructure. The first such call control application is likely to come from Mitel, a leading PBX vendor, with whom Madge is working to create a fully-featured general purpose telephony solution based on MadgeOne.

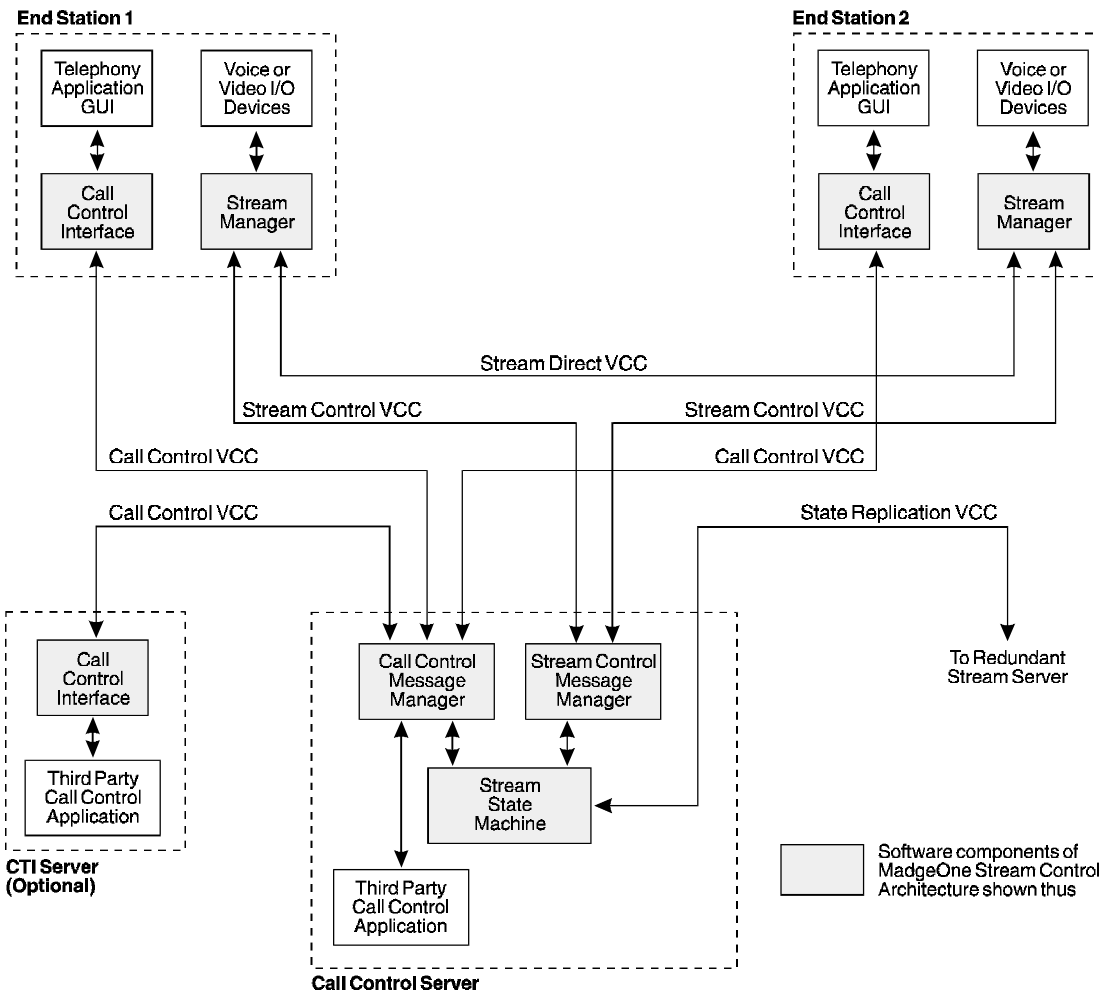

Figure 8: MadgeOne Call Control Architecture

Reliability and Fault Tolerance

Today, voice communications in the enterprise are generally handled by PBX-based networks, and enjoy a level of reliability and availability which is typically somewhat better than that seen with a conventional LAN. It is therefore necessary to pay careful attention to reliability when considering an integrated LAN solution which handles mission-critical voice traffic alongside data.

The reliability of LANs has improved vastly since the early days. The introduction of intelligent hubs did much to ensure isolation of badly behaved network nodes, and further improvements are being achieved as shared LANs give way to switched LANs, where users benefit from dedicated bandwidth on a point-to-point link into the network, and switches isolate faulty nodes.

As a fully switched technology with sophisticated connection routing that supports load-sharing across multiple paths through the network, ATM offers scope for highly redundant network designs that can provide very high reliability. With a backbone based on ATM switches interconnected as a mesh, using dual-homed access switches with two separate connections to the backbone, the failure of any one device in the network will affect few users, if any.

The MadgeOne architecture takes full account of the requirement for very high levels of reliability. Madge ATM switches offer a number of key features to improve availability including redundant power supplies, hot swap line card modules and internal temperature and fan speed monitoring. Madge ATM access switches support optional dual-homed uplinks. Furthermore, the message-passing protocol on which the MadgeOne Call Control architecture is based supports multiple redundant call control servers.

The target for MadgeOne is clear: to match the availability characteristics of today's best PBX networks. This need has been taken into account at every stage of both architecture and individual product design.

Network Management

The evolution of networking through shared LANs to switched LANs to ATM poses a number of new challenges for network management. As shared LANs give way to switched LANs, the ability to see all network traffic with an analyzer or RMON probe is lost. The move from physical to logical network segmentation with Virtual LAN and Emulated LAN techniques places new demands on network management as a tool for network configuration. And the use of ATM to handle multiple classes of traffic on both LAN and WAN links requires new approaches to bandwidth management.

In the past, network management has often been regarded as important, but it use has been optional. Shared media hubs require little in the way of configuration, and LANs have become so reliable that monitoring of network health is often given low priority. Management features have typically been added to networking equipment as an afterthought, taking advantage of the fact that the relatively low data rates of current LANs allow monitoring to be carried out by software running on powerful processors.

Switching and ATM demand a new level of commitment to network management in the core design of networking equipment. Because traffic travels on point-to-point links, placing probes on the network provides only a very limited view of what is happening. Only the switches see all the traffic – and so the switches themselves must provide the required visibility of network activity. Furthermore, the sheer rate at which switches are capable of handling traffic demands specialized hardware to perform the monitoring. Switch throughputs of 1 Gbps or more defeat the software capabilities of even the most powerful processors available.

As a result, effective network management solutions can only be provided when the core elements of the network – the switches themselves – have been designed from the outset to provide management visibility and control. This has been the approach taken by Madge in the design of all the switching products that form the basis of MadgeOne.

The MultiMan™ family of network management applications forms the basis for comprehensive end-to-end management in the MadgeOne architecture. MultiMan transcends the limits of traditional device management systems by providing visibility across the network for performance, topology and configuration management.

For performance management, traffic visibility is provided through MultiMan SMON (Switched Network Monitoring) and AMON (ATM Network Monitoring) agents. These agents report traffic activity in frame-switched and cell-switched portions of the network to MultiMan applications which consolidate the information to provide a “whole network” view.

Topology management is provided at the level of both physical connectivity and virtual circuit connectivity. MultiMan’s topology views also show the bandwidth reserved on links between switches for services that require guaranteed response time.

And finally, MultiMan’s network-wide configuration management deals with Virtual LAN and Emulated LAN management, and provides powerful tools for management of VLAN and ELAN membership, including automatic tracking of moves and changes.

Conclusion

MadgeOne is a network solution that provides scaleable bandwidth with multiple service levels to support all current and emerging data applications, with the additional capability to support mission-critical real-time voice and video communications, all in one single integrated infrastructure.

MadgeOne is based on appropriate use of ATM technology, because only ATM provides guaranteed Quality of Service and low end-to-end delay that is an essential pre-requisite for effective real-time communications.

By supporting all existing LAN protocols with the aid of LAN Emulation over ATM, and by providing a migration path to future standards for supporting current internetworking protocols directly over ATM – such as MPOA – MadgeOne provides a highly scaleable solution for data networking needs and provides fully transparent connectivity with existing LAN infrastructures based on Ethernet or Token Ring, interconnected with switches and/or routers. Furthermore, MadgeOne supports the extension of the LAN environment seamlessly across the WAN, to provide faster and more efficient inter-site networking.

Unlike other approaches to network evolution which are based exclusively on router-based internetworking, MadgeOne takes advantage of native ATM support for real-time voice and video. This approach allows voice and video communications to travel efficiently on end-to-end switched connections, bypassing routers and making it unnecessary to install costly and complex router upgrades. And by making intelligent use of new Cell-in-Frame technology, MadgeOne also supports the extension of ATM-based real-time services to Ethernet- and Token Ring-attached desktops.

MadgeOne goes beyond solving the problem of transporting real-time voice and video, and addresses also the requirement for sophisticated call control. This is achieved with middleware components that support the use of third-party telephony and video call control applications in a client/server implementation over the ATM-based network. By providing for the connection of existing PBXs through Madge's WAN AccessSwitch, and by supporting a range of PBX control protocols, MadgeOne will offer transparent operation of call control features between the PC-based telephone terminals and conventional PBX-attached phones.

MadgeOne represents the first serious attempt by any networking vendor to bring together the disparate worlds of voice, video and data communications over a single network infrastructure all the way to the desktop. This paper has described some of the technical approaches that Madge is taking to make MadgeOne a reality. MadgeOne is not just concerned with solving complex technical problems: the solution to these problems will bring a wide range of cost benefits and opportunities for service innovation to those enterprises for whom the network is, and will continue to be, an instrument of competitive advantage.

