

Lucent Technologies
Bell Labs Innovations



Product information begins on page 2.

Lucent and Ascend have merged.

With the Lucent-Ascend merger, customers gain a broader and more powerful portfolio of next-generation data, voice, fax, and video services and products. To access up-to-the-minute information about our products, see page 2.

We also invite you to contact us with your questions directly at: info@ascend.com

Cascade 500

High-Scalability ATM Switch Applications & Technology

Introduction

Much has been written about the promise of Asynchronous Transfer Mode (ATM) technology and how it is the answer to network growth problems in the LAN and in the WAN. However, when the Enterprise/Service network planner examines existing ATM switching technology, there is a yawning chasm between currently available products and what is *really* needed to anchor the network. This need is true whether the network be an Enterprise network or a Data Service such as Frame Relay. Existing ATM technology is appropriate for test beds and trials, but it does not have the scalability required to deploy a production ATM network. The Enterprise/Service planner knows that planning for network growth is in truth a multi-dimensional problem. It is a problem that requires much greater complexity than the simple dimension of bandwidth that has popularly been used to proclaim the merits of existing ATM switches.

This article reviews the imperatives behind deploying scalable technology and introduces the Cascade 500, the first production-ready ATM switch. The important deployment scenarios for ATM switches are outlined. These scenarios demonstrate how Cascade's family of products allow data service and Enterprise networks to keep pace with the mushrooming demand and to build the fusion of ATM-based Enterprise and WAN data service networks.

Imperatives: Data Services Growth

While many data networking technologies offer large projected growths, Frame Relay and IP data services enjoy the important distinction of demonstrating real growth now. For this reason these data services offer an excellent example for understanding the challenges associated with network growth. Enterprise customers readily understand the economic advantages of Frame Relay services and are migrating leased-line, multiprotocol and SNA networks to Frame Relay services in record numbers. Numerous customer-survey studies point to continued Frame Relay growth. Figures 1 and 2 illustrate the dramatic growth of Frame Relay and IP data services in terms of service revenues. These figures also show the trend lines for the continued growth of the service.



CASCADE




Figure 1: IP Data Services Market Revenues (Source: Forrester Research Inc. 1995)




Figure 2: Worldwide Frame Relay Service Market Revenues (Source: Vertical Systems Group 1995)

The data service demand growth mandates the deployment of an ATM backbone because ATM is the only technology that can carry aggregated customer traffic at speeds of one hundred megabits and above. It is also the only technology that can provide the migration to new data service applications such as video conferencing and usage-based billing. Figure 3 illustrates how the data service network uses

ATM technology in its core. ATM is the backbone for carrying aggregated data services traffic. Because data services are a “real” application generating billions of dollars of annual revenue, the backbone ATM switching equipment cannot be “trial deployment” equipment. It must be service-ready with the scalability to handle the growing demand.

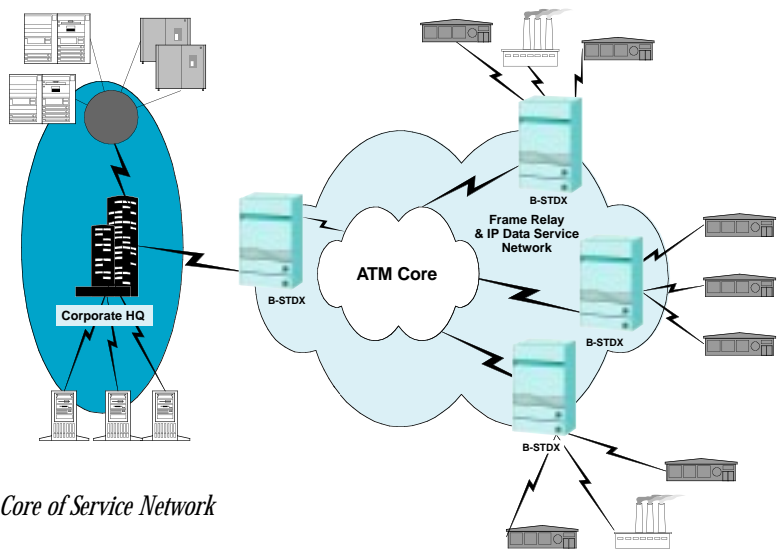


Figure 3: ATM in Core of Service Network

The timely development of Frame Relay to ATM Interworking standards (Frame Relay Forum FREL5 Network Interworking and FREL8 Service Interworking) ensures the ready availability of interworking technology to convert Frame Relay user traffic to ATM network traffic. IP data services can likewise migrate to an ATM future via Frame Relay to ATM Interworking or via RFC 1577 "Classical IP and ARP over ATM" or via Multi-Protocol over ATM (MPOA).

Multiservice data switches, such as Cascade's B-STDX family, implement Frame Relay interworking features and feed data service traffic in ATM-converted format to

the core ATM network. Interworking functionality is key to ensuring that ATM migration is seamless to the customer and also ensures the coexistence of ATM technology with existing packet-based Customer Premise Equipment (CPE). Interworking functionality is best implemented at the ingress port from the CPE because it requires the least packet processing. Implementing interworking functions within the core ATM network itself would require needless additional packet reconstitution operations. Figure 4 illustrates ATM interworking functionality (IWF) that the service switches provide.

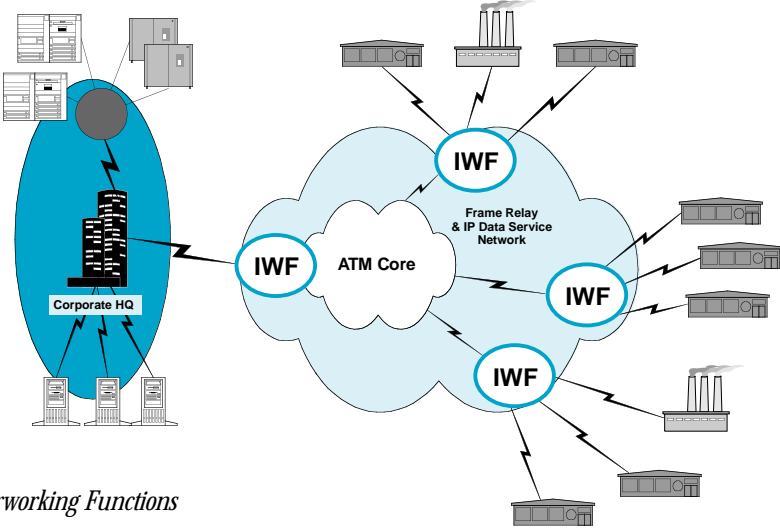


Figure 4: ATM Interworking Functions

Imperatives: Enterprise Integration with ATM

Coincident with the deployment of ATM technology in the data service network, Enterprise network managers are implementing ATM technology in their own networks. The Enterprise network uses ATM for data-volume growth reasons as well as for the integration of voice, video and data.

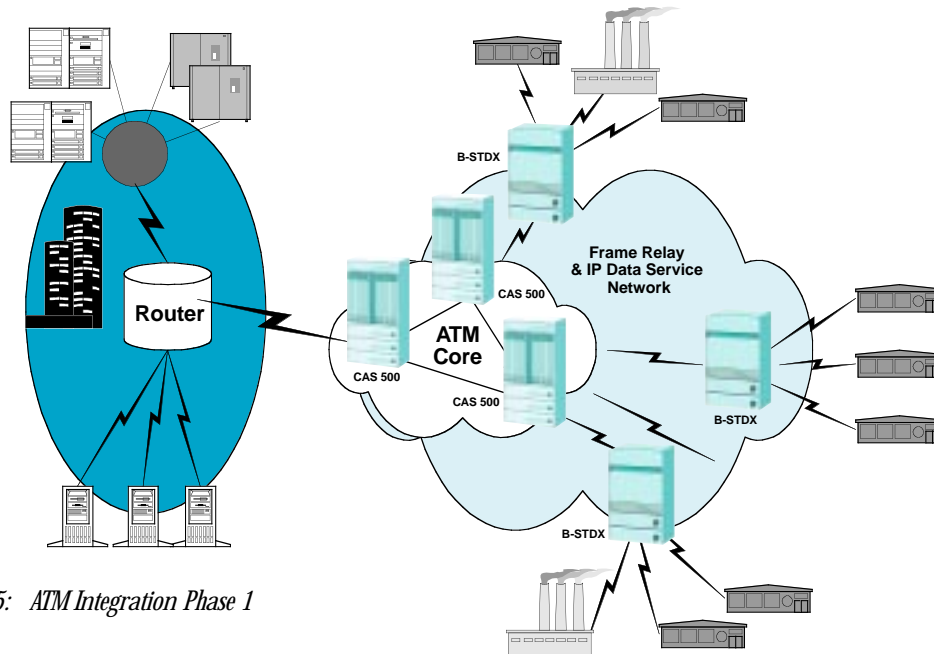


Figure 5: ATM Integration Phase 1

Enterprise and service provider ATM network interconnection will happen first for data transport, and then for voice/video integration in the following three stages as the technology matures.

◆ Phase 1: 1996

Service Providers deliver access to the ATM core of their network to enable high-speed network access and to differentiate their services in the highly competitive data services marketplace. Customers maintain the Collapsed Backbone Router model for their Enterprise data networks. These backbone routers connect to the ATM WAN network via a native ATM connection at DS3 or OC3c speeds. Figure 5 illustrates Phase 1 ATM Integration.

◆ Phase 2: 1996-1997

Enterprise managers deploy Campus ATM backbones to increase network performance. In this phase the collapsed backbone router is now serving more of a “route server” role than bandwidth provider role. The campus ATM switches serve as traffic aggregators and administrative and traffic-domain administrators. Figure 6 illustrates Phase 2 ATM Integration.

◆ Phase 3: 1997-1998

Enterprise managers integrate voice and video networks and applications with the data infrastructure to realize significant line savings and new application integration, such as video conferencing/collaboration and Computer Integrated Telephony (CIT). Figure 7 illustrates Phase 3 ATM Integration.

In the 2nd and 3rd phases of ATM integration, the role of the Campus ATM switch becomes very significant. The Enterprise backbone is a "real" application, supporting billions of dollars of annual revenue. For this reason, the Enterprise ATM switching equipment cannot be characterized as "trial deployment" equipment. It must be service-ready with the scalability to handle the growing demand. It is now the production vehicle for the Enterprise network. The backbone must have the reliability and scalability to ensure the successful deployment and continued growth of the Enterprise ATM network.

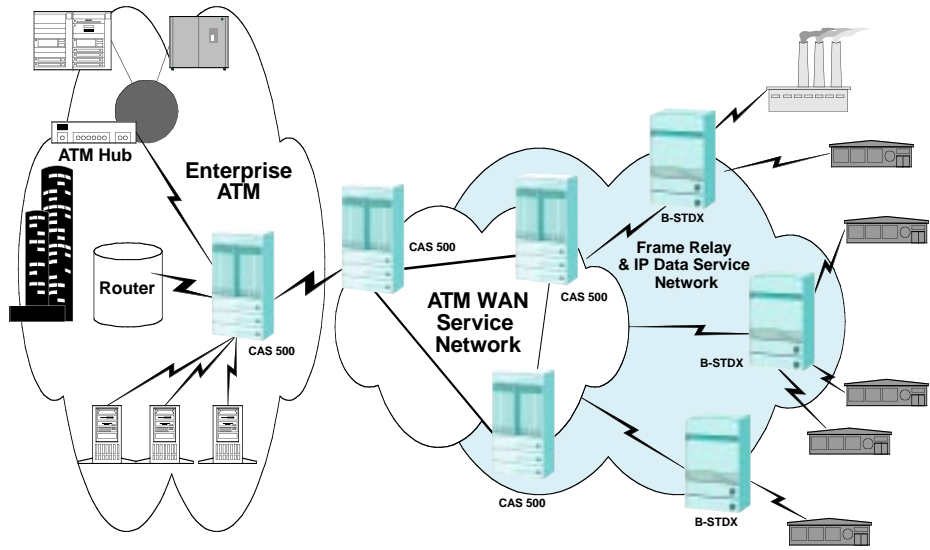


Figure 6: ATM Integration Phase 2

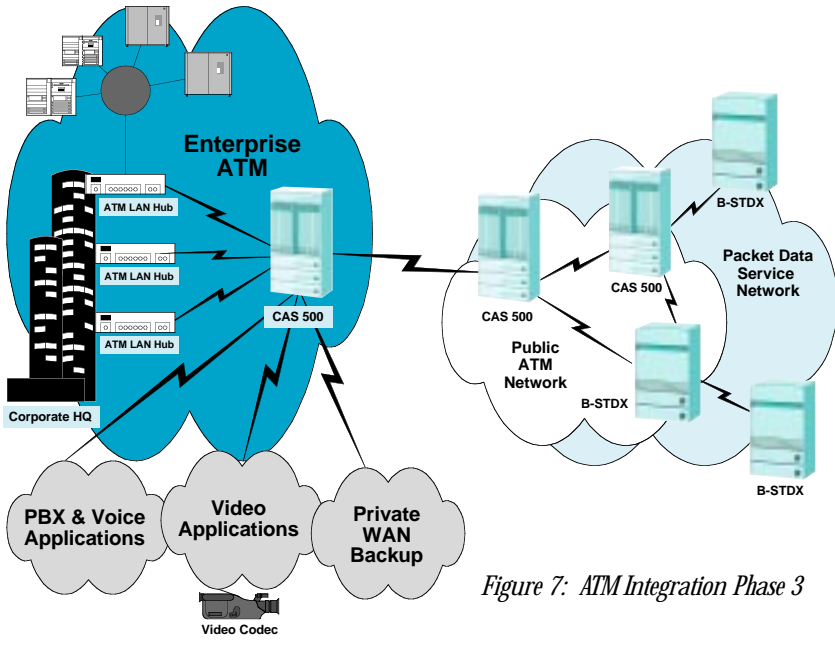


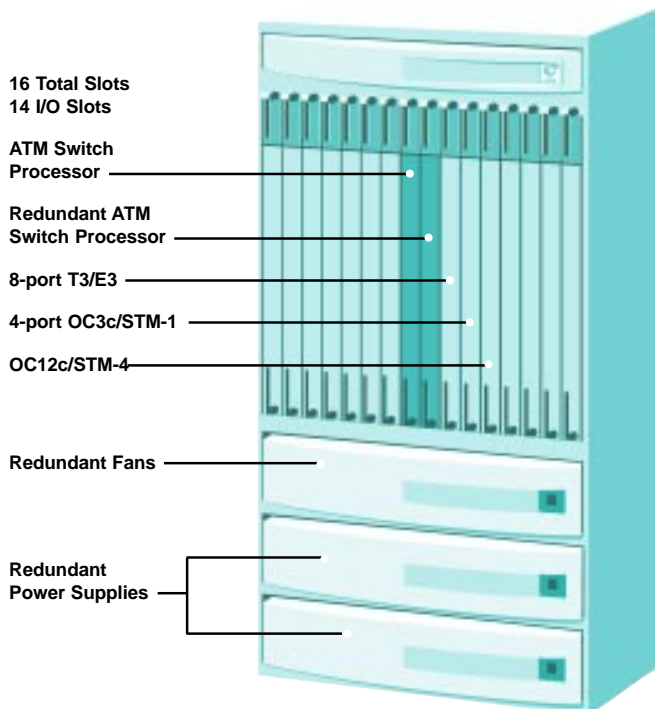
Figure 7: ATM Integration Phase 3

Production ATM: Cascade 500 High Scalability ATM Switch

Seeing Through the Hype

ATM has received wide industry exposure, but the hype has unfortunately exceeded reality. ATM has been largely over promoted along a single dimension of bandwidth capacity. Hundreds of gigabits of capacity are touted to be the answer to all problems. However, the Enterprise/Service provider knows the multi-dimensional challenge for ATM technology must include:

- ◆ Sophisticated topology management services for large networks with 100,000 ports and auto re-route around failed links and nodes
- ◆ Call Admission Control dynamically coupled with topology management to guarantee Quality of Service end-to-end in the network
- ◆ Ability to support tens of thousands of virtual circuits per port
- ◆ Signalling performance to scale with virtual circuit capacity must equal thousands of SVCs/sec/switch
- ◆ Hardware Multicast & Quality of Service guaranteed via fabric-integrated buffers
- ◆ Expansive cell buffers per line card for bursty data traffic
- ◆ 622 Mbps (OC12c) port speeds directly in the switch fabric
- ◆ Large fan-out densities for OC12c/STM-4, OC3c/STM-1 and T3/E3 ports
- ◆ Extremely high system availability and serviceability
- ◆ Usage-based billing capability



The current generation of ATM Enterprise and WAN switches do not meet the above requirements because they were largely designed for the research and pilot phase of ATM networking. The Cascade 500 was designed from the outset for the production phase of ATM networking. The Cascade 500 is a 16-slot platform with 2.5/5 gigabits of switching capacity, expandable in the future to 10 gigabits. The Cascade 500 differs from available ATM switches in that it emphasizes high performance on the many dimensions that real service performance must be measured. Figure 8 outlines some of the important characteristics of the Cascade 500.

- 56 OC3c Ports
- 112 DS3 Ports
- 14 OC12c Ports

Figure 8: Cascade 500

Production ATM: Virtual Network Navigator Topology Navigation Services

ATM technology will be deployed as the core backbone technology for larger feeder networks, which in turn will migrate to native-ATM operation over time. This vision holds for Enterprise networks and for WAN service networks. As such, ATM switches must support sophisticated topology navigation services that will allow all of the switches, feeder and backbone, to build dynamic network maps for routing traffic. These topology services must encompass not only quasi-static physical characteristics, such as link connectivity and raw bandwidth, but also must be able to find the route that meets the Quality of Service (QoS) contract for requested virtual connection end-to-end through the network. Thus, Topology Navigation Services must not only find the path, but they must also meet ATM cell delay and cell loss metrics.

Cascade's Virtual Network Navigator™ (VNN) technology starts with the industry standard Open Shortest Path First (OSPF) algorithm, invented by Cascade's own John Moy, and builds upon it to bring new levels of sophistication to ATM network topology management. VNN incorporates the following additional link state characteristics:

- ◆ End-to-End Delay — End-to-end delay values will be measured at link initialization and circuit setup requests may specify the minimum allowed delay.
- ◆ Cell Delay Variation (CDV) and Cell Loss Ratio (CLR) — Circuit setup requests may specify their CDV and CLR.
- ◆ Administrative Path Control — Users can steer traffic along preferred paths based on bandwidth or delay metrics.
- ◆ Virtual Private Networks — Inter-switch trunk circuits may be dedicated to specific end-user ports, thereby creating multiple Virtual Private Networks out of a single network of Cascade switches.

VNN includes sophisticated support for point-to-multipoint connections in an ATM network. These connections are characterized by a “branching tree” topology. VNN ensures maximal use of network resources by finding the optimal last branch points in the distribution tree.

These enhancements to OSPF make Cascade's VNN Topology Navigation Services without peer in the ATM industry. The ATM Forum is in fact utilizing the same approach of augmenting OSPF for ATM use in the PNNI Working Group. Cascade will also support the ATM Forum's PNNI IISP, PNNI Phase I and B-ICI requirements. VNN is a Cascade value-added feature for the interim period while the standards and solutions mature.

Production ATM: Connection Admission Control Services

ATM networks operate by establishing virtual circuits between nodes. Virtual circuits will be deployed on an end-user subscriber basis and, as such, the success of an ATM network depends crucially on the efficient operation of the Connection Admission Control (CAC) Services. Service provider/Enterprise equipment must efficiently employ the available network bandwidth, while simultaneously meeting the Quality of Service metrics established for each customer connection.

Cascade's Connection Admission Control Services offer unprecedented efficiency in matching requested connections incrementally against available network resources. These services provide the most efficient utilization of service provider capacity, while maintaining the required Quality of Service contracts for both new and existing connections. CAC translates traffic load and burst characteristics into a link-by-link additive metric that is used by Cascade's VNN services to ensure end-to-end connection viability. CAC simultaneously performs Connection Admission Control for all ATM QoS types. CAC and VNN together find all available network bandwidth and ensure the connection route validity across all ATM metrics.

Production ATM: Virtual Circuit Capacity

While gigabit bandwidth capacity is one important metric in ATM switches, perhaps a more appropriate measure of switch capacity in the Enterprise/Service context is the number of Virtual Connections that a switch can handle. More Virtual Connections ensure that the switch can support more simultaneous end-user connections. Stated conversely, a nominal high-bandwidth capacity switch that is limited in the number of Virtual Connections that it supports will not be able to use the available bandwidth once the Virtual Connection capacity is exhausted. Virtual Connection capacity is the truest measure of the revenue-generating capacity of an ATM switch.

The Cascade 500 is capable of supporting 16,000 Virtual Connections (Virtual Path or Virtual Channel) per I/O module and 16,000 point-to-multipoint virtual circuits per Cascade 500. With fourteen I/O modules in a single chassis, that translates into 224,000 Virtual Connections per ATM Switch! This capacity is implemented on a per-I/O-card basis to ensure the scalability of the platform as additional capacity is added.

Production ATM: Scalable Signalling Performance

While it is generally accepted that ATM-based networks require Switched Virtual Circuit (SVC) capabilities, most ATM switch implementations off-load SVC processing to a centralized processor that all ports on the switch share. The ability of existing switches to process dynamic end-user connection requests is limited by the signalling capacity of the shared resource. Native-ATM data services, such as LAN Emulation and Multi Protocol Over ATM (MPOA), will initiate frequent SVC requests. Because incremental SVC requests are processed too slowly, processing may be unable to effectively support native ATM services in a network of any appreciable size ATM switches with shared SVC.

The Cascade 500 employs a distributed multi-processor implementation with a dedicated i960 SVC processor for each I/O card. This signalling implementation offers a dramatic improvement over conventional implementations. Every I/O card has a signalling processor that can independently broker connection requests across the switch fabric and across the larger network. While signalling performance will be a strong function of network conditions, the Cascade 500 will deliver 3000 SVCs/sec/switch. Cascade's customers will not find their ATM service capacity inhibited by inferior SVC capacity that must be shared regardless of the port count growth.

Production ATM: Quality of Service, Buffering & Available Bit Rate Data

ATM networks differ from existing technology most dramatically by offering different Qualities of Service (QoS) for individual connections. Individual Virtual Connections are characterized by numerous attributes including: their cell rate or bandwidth, the delay experienced, the variation in mean delay, the cell burst sizes and spacing, and the cell-loss maximum guarantees that are provided. Clearly ATM links are complex entities! The ATM Forum has defined four distinct classes of service to categorize different combinations of connection attributes. These classes are: Constant Bit Rate (CBR), Variable Bit Rate Real Time (VBR-RT), Variable Bit Rate Non Real Time (VBR-NRT) and Available Bit Rate (ABR)/Unspecified Bit Rate (UBR).

The Cascade 500 switching fabric is actually four parallel switch fabrics, one for each class of service. Cells are switched into one of four parallel dedicated buffer structures that constitute the switch fabric. The buffers supporting Variable Bit Rate traffic (B and C) are further subdivided with four individual threshold markers creating a total of ten distinct QoS classes: one CBR class, four VBR-RT classes, four VBR-NRT classes and one ABR/UBR class.

CAC and VNN ensure that these classes are available and guaranteed end-to-end in the ATM network. The ATM switch fabric has a total of 128,000 cell buffers partitioned into four planes to support the QoS classes. Figure 9 illustrates the Cascade 500 ATM switching matrix.

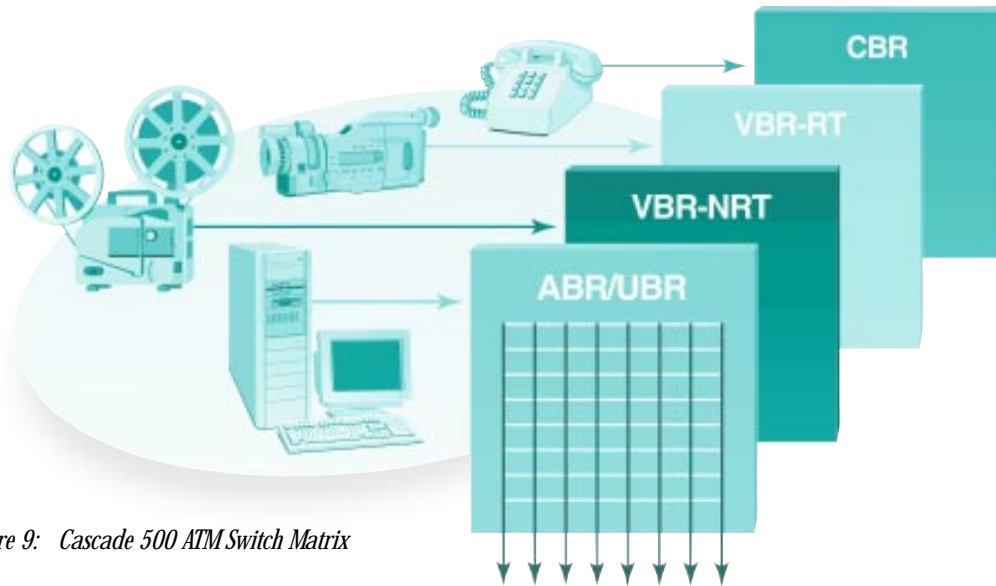


Figure 9: Cascade 500 ATM Switch Matrix

Cascade's I/O modules provide additional levels of cell buffering for each I/O module, reaching a maximum of 96,000 cell buffers per OC3c/STM-1 module. The cell buffer capacity is offered as an incremental option on the OC3c module. Thus, the deployment of buffer capacity can be tailored according to the necessary requirements. This hardware support for QoS classes in the switch fabric ensures that video traffic is not disrupted by bursty data traffic.

Buffer capacity is most important for supporting the ABR/UBR class of service for bursty computer traffic. ABR means that a given VC connection can increase its cell transfer rate to utilize all of the available bandwidth on a particular link. Available bandwidth is a dynamic attribute that changes as new connections are initiated. ABR requires a complex rate-throttling protocol to operate hop-by-hop across the switches in a connection path. ABR also requires ATM end stations (be they workstations in ATM LANs or customer-premise routers in WAN Internetworks) to actively participate in the end-to-end rate setting algorithms. The ATM Forum has been diligently working on this topic and it is clear that the final requirements will not be stable until 1996

at the earliest. It is also clear that ABR Flow Control will require substantial protocol processing to dynamically negotiate the allowed burst rate on each hop and to schedule the traffic flow out of the buffers according to the allowed rates. The processing required by ABR Flow Control processing cannot, however, subtract from the available processing power for other ATM operations such as Signalling and Topology Navigation.

UBR is a more recent development from the ATM Forum that suggests that bursty computer data connections should have an unspecified bit rate and should burst at their maximum line speed to maximally utilize available link bandwidth and switch buffers. When UBR is coupled with two algorithms called Early Packet Discard (EPD) and Fair Queuing, it can provide a very effective data transport service. EPD allows a switch to discard an entire packet (composed of multiple cells) when congestion conditions start. Fair Queuing ensures that UBR connections get throttled proportionately to their transfer rate so that high cell rate users do not monopolize the available bandwidth. UBR with Early Packet Discard and Fair Queuing makes ATM UBR services equivalent to existing datagram networks.

In addition, Fair Queuing relies on the operation of higher-layer protocols (e.g., TCP/IP) to recover from network congestion conditions. UBR is considered a more likely near-term solution to ATM congestion control that will be used until the installed base of workstations and Customer Premise routers are upgraded to participate in the ABR Flow Control algorithms.

The Cascade 500 supports UBR and ABR Flow Control with a daughter card hosted on the I/O modules. The I/O cards are thus designed to support UBR/ABR Flow Control at the outset via a segmented design that separates these developing functions and supports them with their own microprocessor-based subsystem card. Network planners can thus deploy these features with assurance that the performance will scale with the growth in the network and that the implementation can be incrementally modified without obsoleting existing investments.

Production ATM: OC12c (622 Mbps) Port Direct into Switch Fabric

ATM technology has been in the limelight for over three years now. In technology life cycles three years is a full generation. Many currently available ATM switches were designed at the start of the ATM epoch when 155 Mbps OC3c connections seemed very fast. Such switches are architecturally precluded from supporting OC12c speed links. The Cascade 500 ATM switch fabric supports OC12c links at the outset. Enterprise/Service providers know that their backbone networks will soon migrate to OC12c links, and any technology they invest in must provide 622 Mbps connections.

Production ATM: Highest Port Densities

WAN bandwidth has traditionally been very expensive, limiting the commercial potential of DS3 and OC3c transmission services. But this expensive bandwidth paradigm is likely to change in the face of increased competition. Soon customers will receive inexpensive bandwidth and improved value-added services will deliver the revenue premium. The service provider must be ready to meet this challenge in a cost-effective manner.

The Cascade 500 offers the highest density T3/E3, OC3c/STM-1 and OC12c/STM-4 port concentrations in a single platform in the industry. We offer a single-slot 8-port T3/E3 module, a single-slot 4-port OC3c/STM-1 module and a 1-port OC12c/STM-4 module. With these options, the Cascade 500 provides an industry leading fan-out of 112 T3/E3 ports, 56 OC3c/STM-1 ports and 14 OC12c/STM-4 ports-per-chassis. Note that these densities come in a single 34-inch tall chassis for a standard 19-inch rack-mountable system.

Cascade's unique midplane I/O architecture, which separates the physical layer components from the cell processing components, ensures that customer investments in I/O technology will be able to migrate with the evolving standards for copper and fiber-based OC3c/OC12c link technology.

Production ATM: Availability by Design

Enterprise/Service planners understand that availability is a fundamental prerequisite for service-ready technology. The Cascade 500 provides unprecedented levels of availability in its very design. The Cascade 500 provides redundant switches in the same system chassis with failover benchmarked at only 200 microseconds. Connections operating on the primary switch are transferred to the backup switch without requiring connection re-establishment. Power supplies, I/O modules and system clocking support full redundancy as well. Cascade also extends this availability paradigm to the network connection level providing redundant Virtual Connections and auto reroute on service failure in our VNN technology.

Conclusions

We have reviewed how both data service and Enterprise networks will incorporate ATM technology in the next three years. These networks will require production capable ATM products as the networks already support the “killer application” of the existing business environment. We have also reviewed the major features of the Cascade 500 that show it to be a true production generation ATM switch, ready for the multiservice integration challenge in the WAN data service and Enterprise Environment.



WHITE PAPER

Cascade Communications Corp.

5 Carlisle Road
Westford, MA 01886 U.S.A.
(800) 647-6926 (within U.S.A.)
Main number: (508) 692-2600
Fax: (508) 692-5052
Internet: mktg@casc.com
Internet Home Page:
<http://www.casc.com>

Latin America

Cascade Communications Corp.
Pembroke Pines Professional Center
9050 Pines Boulevard, Ste. 360
Pembroke Pines, FL 33024 U.S.A.
Phone: 954-450-9922
Fax: 954-450-9733

Asia

Cascade Communications Corp.
555 Twin Dolphin Drive
Suite 340
Redwood City, CA 94065 U.S.A.
Phone: 415-595-2520
Fax: 415-595-2404

Singapore

Cascade Communications Corp.
MRCENTRE PTE LTD
9 Raffles Place #27-00
Republic Plaza
Singapore 048619
Phone: +65-435-0414
Fax: +65-339-1676

Japan

Cascade Communications Asia Corp.
MS Bldg. 5F
5-11 Shiba 4-chome
Minato-ku, Tokyo 108 Japan
Phone: +81-3-3798-5181
Fax: +81-3-3798-5180

Europe

Cascade Communications Ltd.
5B Foundation Park
Roxborough Way
Maidenhead
Berkshire
SL63UD U.K.
Phone: +441-628-828500
Fax: +441-628-825002

Germany

Cascade Communications Corp.
Hessenring 121
D-61348 Bad Homburg
Germany
Phone: +49-6172-925858
Fax: +49-6172-925825

Spain

Cascade Communications Corp.
Serrano 240 Planta 5,
28016 Madrid
Spain
Phone: +34-1-345-7276
Fax: +34-1-457-2079



The Cascade logo, CascadeView/UX, Cascade 500, and B-STDx are trademarks of Cascade Communications Corp. All other product names are trademarks of their respective companies. ©1996 Cascade Communications Corp. Printed in the U.S.A. All Rights Reserved. Some products, features, and options described may still be in development.

696.10000
500WPA4