

FORE
SYSTEMS.

white paper

Building an IP Service Network

December 1998

BUILDING IP SERVICE NETWORKS WITH THE FORE ADVANTAGE	1
The FORE Advantage	2
SERVICE PROVIDER MARKET TRENDS	2
NETWORK REQUIREMENTS FOR A COMPETITIVE IP SERVICE PROVIDER	3
IMPLEMENTING THE IP SERVICE NETWORK: TECHNOLOGY CHOICES	4
Today: IP Routing over ATM Switching	4
Tomorrow: Unified Routing and Switching with MPLS	5
MEETING THE ENGINEERING CHALLENGE WITH FORE SYSTEMS	5
The Functionality and Flexibility of a Layered Architecture	6
Maximizing Network Utilization Through Traffic Engineering	7
Configuring a PVC-Based ATM Core	8
Configuring an ATM Core with Capacity Aware Routing	9
Integrating Traffic Engineering with IP Routing	
Using MPLS	10
Network Resiliency	11
PNNI and Failure Recovery	11
Ensuring Router Scalability with MPLS	12
Developing a High-Speed Transmission Layer	14
Providing Service Availability with APS Protection	14
Maximizing Link Efficiency with Frame-ATM	14
Reducing Deployment Costs with DWDM Transmission Technology	15
CONCLUSION	16
Acronyms	16

BUILDING IP SERVICE NETWORKS WITH THE FORE ADVANTAGE

As the world's first public data network, the Internet has fundamentally changed business practices. Business transactions that once occurred over the public voice network or Public Switched Telephone Network (PSTN) now occur over the Internet. Suppliers replenish stock, retailers offer consumers expanded inventories with advanced inventory searching tools, and businesses market their latest products and services. Service providers and equipment vendors are responding quickly to the Internet's opportunity.

Today, private data networks are becoming as important to businesses as the PSTN and the Internet. Network managers have recognized the cost savings of converged voice and data Wide Area Networks (WAN) and are now looking to improve costs by moving their private networks to the Internet using Virtual Private Network (VPN) services.

As a pioneer in next generation telecommunications networks, FORE Systems, Inc. developed the ATM switching technology used to transition the Internet from a research to a commercial network. FORE is a leader in developing the hardware and software required to build commercial IP service networks. The FORE solution combines ATM and IP technologies to build fast, efficient networks capable of supporting a wide range of voice, video, and data services.

Here we describe the FORE advantage and how you can harness it to achieve the sustained growth of your network service business. We look at trends, network requirements, technology choices, and how the FORE advantage meets the engineering challenges that must be overcome for your business to succeed.

Unprecedented levels of traffic, constantly changing technology, deregulation, well-financed competition...no one ever said building an IP service network would be without its challenges.

The FORE Advantage

- Superior ATM-based core switching technology and products.
- Functional network implementation based on open standards and a flexible layered architecture, for low-cost, component level upgrades.
- Advanced traffic-engineering techniques to ensure maximum utilization of switch and link capacity, for best return on infrastructure investment.
- *Capacity Aware Routing*TM, a FORE-developed implementation of ATM's PNNI protocol, providing dynamic routing guided by the state of the network connections themselves.
- *Dynamic Protection Switching*TM, providing network resiliency and failure recovery without the high cost of redundant links.
- Router scalability and optimization through innovative forwarding adjacency techniques.
- A clear, cost-effective, software-only migration path to full MPLS implementation.

Add it all up and you have the FORE advantage — a robust, highly reliable IP service network that uses proven technologies to deliver voice, data, and video services... today and tomorrow.

SERVICE PROVIDER MARKET TRENDS

If a single word can describe today's Network Service Provider environment, it is "unsettled." Change is a daily occurrence — networks have more users, and they are mixing their legacy applications with newer and more bandwidth-hungry, network-aware applications. Technology constantly changes to meet this demand. And if you aren't there to meet your customers' requirements with carrier-class quality of service, one of your competitors surely will be.

Then there's the Internet. Access to the Internet and its applications have driven traffic volumes to unprecedented levels. Along with, or because of, the growth of the Internet, its IP protocol has become the dominant protocol for data traffic and is making inroads into voice as well. The trend is clearly running towards high-capacity, IP-based networks.

Deregulation has changed the SP landscape forever. The incumbent carriers (RBOCs, PTTs) are pouring their PSTN revenues into offering new services, such

as Internet access and Virtual Private Networks (VPNs). The alternate service providers, who have formed more recently, have constructed networks with enormous switching capacity. They intend to leverage their network capacity to offer competitive private voice and data networking services in addition to Internet connectivity. Both emerging and incumbent suppliers recognize the cost savings and performance advantages of converging voice, data, and video services onto a common network infrastructure. Converged networks allow more competitive service offerings at a lower cost. However, it isn't clear what the converged network will look like.

Thus, it is imperative for you to develop a strategy for converging voice and data traffic on your network as part of becoming a competitive IP service provider with staying power. A sound first step is to consider just what that entails.

NETWORK REQUIREMENTS FOR A COMPETITIVE IP SERVICE PROVIDER

If your goal is to manage sustained growth in an IP service business, your service infrastructure must support the expansion from Internet/IP connectivity into the broader realm of private business communications. To accomplish this, your IP network must:

- Be developed on a network core infrastructure based on open, standard components that can be upgraded or interchanged without disrupting network services.
- Provide capacity sufficient to meet changing customer demands over time. It is essential that the network architecture be able to scale to meet future demand in a cost-effective way.
- Maximize the available capacity of the network infrastructure using current technology by optimizing link and switch utilization through traffic engineering and intelligent routing.
- Provide the highest level of network resilience and reliability with robust service availability and consistent end-to-end performance.
- Provide and support the service-driven bandwidth guarantees required to meet Service Level Agreements (SLAs), and be able to monitor and track those Agreements.
- Employ scalable transmission technology that exploits the highest available bandwidth and supports advances in repackaging IP packets to provide faster transmission.

What an SP should demand of its technology partners: high-capacity core switching, optimized resources, bandwidth guarantees, and a scalable infrastructure.

To meet these criteria — and deliver competitive commercial services — SPs are evaluating and deploying new technologies that don't trace their lineage back to copper wire. These new technologies combine the stability and reliability of the traditional carrier networks with the speed and capacity of the digital age. Among them are:

- High-capacity ATM switches running the Private Network to Network Interface (PNNI) protocol in a connection-oriented network.
- Terabit routers, running the IP protocol in a non-connection based configuration.
- Packet Over SONET (POS) high-speed interface, which bundles IP packets into frames at the network interface, running at speeds of up to OC-48 (2.5 Gbps) today.
- The emerging Multi Protocol Label Switching (MPLS) standard for implementing IP routing over ATM switching.
- Dense Wave Division Multiplexing (DWDM), for increasing the traffic delivery capacity of fiber.

IMPLEMENTING THE IP SERVICE NETWORK: TECHNOLOGY CHOICES

While future technology choices always appear promising, the fact is you need to run a business today. And there is technology available today that is highly efficient and cost-effective, offering a sensible business choice — one that is stable today and scalable tomorrow. FORE advocates a business strategy based on IP routing overlaid on ATM switching to meet your service needs today, with a migration to MPLS to meet your needs tomorrow.

Today: IP Routing over ATM Switching

The major IP network vendors have embraced the idea that the backbone of the converged network will be powered by ATM core switching. Delivering new services to your subscribers will require more networking capacity than the existing PSTN or an IP router backbone can offer. ATM, with PNNI, provides the QoS capabilities required to provide guaranteed bandwidth and premium services, and ATM switches have consistently proven their reliability under heavy traffic loads.

With ATM running on an optical physical layer such as SONET, you can provide voice, high-speed data, and even video services like CATV over a single infrastructure. And you can do so with distinct operating advantages — predictable

Your best choice today is implementing IP routing over ATM switching. And for tomorrow? The emerging MPLS standard promises efficient IP routing united with ATM connections, and network provisioning in a single step.

latency (delay), high-capacity non-blocking throughput, fast network recovery, and low cost. This is the core switching choice that has produced some of the biggest successes among the emerging Service Providers.

Tomorrow: Unified Routing and Switching with MPLS

The emerging MPLS standard is viewed as a cornerstone of delivering high-bandwidth IP service. This framework architecture provides the connectionless packet switching of IP with a controlled connection-switched infrastructure. In the networking industry, MPLS is seen as an important way to meet bandwidth provisioning and QoS needs for IP-based services. It bridges the gap between ATM connections and IP services.

Keep in mind, however, that MPLS is a standard, not a product, and it is not something you can order from a supplier. ATM, however, is here today, and FORE has built into its ATM switching two of the fundamental advances promised by MPLS: capacity optimization and traffic engineering. The benefit for SPs is that an investment in ATM today is the shortest path to full MPLS tomorrow. As more of MPLS becomes available, FORE's migration to it can be achieved through software upgrades, and migration occurs at a significantly lower cost than altering the infrastructure.

MEETING THE ENGINEERING CHALLENGE WITH FORE SYSTEMS

You want to concentrate on building your network services business, and to do so you want to feel confident in the decisions you've made regarding your technology base. An investment in FORE's hardware and software should help you attain that level of confidence. There are significant engineering challenges involved in deploying a state-of-the-art WAN, and if you want to have confidence in a vendor's technology, it's in your interest to know how that vendor meets those challenges. FORE's objective is to provide:

- Functionality and flexibility through a layered software architecture.
- Maximum utilization of network resources through traffic engineering.
- Maximum efficiency in transporting traffic through Capacity Aware Routing.
- Availability and resiliency through Dynamic Protection Switching.
- Router scalability for future growth through MPLS.
- Transmission efficiency through maximizing the use of available capacity.

FORE's layered architecture is based on open systems concepts and presents a solid approach to delivering the different functions of service, switching, and transmission. Such layering accommodates in-service upgrades and ensures high performance.

The Functionality and Flexibility of a Layered Architecture

Implementing IP over ATM requires a flexible networking strategy that will accommodate prudent change without disrupting or replacing the substantial infrastructure already in place. FORE advocates the use of a layered architecture, based on vendor-independent open standards as a way to localize the impact of such change. A successful IP-over-ATM implementation uses the three-layer architecture shown in Figure 1:




Figure 1

- **Service (IP routing).** The service layer should be extendable to offer timely services (either first to market with a new service, or quick to respond to a competitor's successful offering). Among the critical responsibilities of the service layer are provisioning, billing, and accounting.
- **Switching (ATM infrastructure).** The switching layer handles service transport and traffic engineering. It must be capable of delivering multiple services, optimizing bandwidth utilization, and recovering quickly from network outages. The switching layer employed by FORE places much of this functionality within the network connections themselves, using the power of ATM switching and network control.
- **Transmission (SONET or SDH).** The transmission layer lights up the fiber, delivering the capacity you need to support the bandwidth demands of your subscribers. Best service delivery depends on maximizing available capacity through dense transmission technologies such as WDM, a solid point-to-point framing technology like SONET/SDH, and failure recovery capabilities, such as those provided by Automatic Protection Switching (APS).

There are two SP benefits to such a functional approach. First, use of an independent service layer allows the flexible introduction of new service gateways without network disruption. Second, layering allows you to capitalize on new technology. For example, it allows ATM switching to replace switching based exclusively on IP routers. When it is cost-effective to do so, layering allows you to

replace SONET or SDH in the transmission layer with even greater bandwidth, such as Wave Division Multiplexing (WDM). In a business sense, the increase in network capacity afforded by WDM translates directly into the ability to offer your subscriber base a greater number of revenue-generating services.

The key point to remember about layering is that regardless of the technology employed, your network *must* have the functionality each layer provides. The delivery technology employed can and does change... but ten years from now, your network will still need transmission, switching, and service functionality. Investing in a system like FORE's, based on industry standards and using the open systems approach to flexibly interchanging components as necessary, is your best defense against the high cost of replacement technology. You can combine the three layers within a single product as long as they can be separated as the network grows. Without separation, any upgrade entails complete software replacement (which is more expensive than replacing single components), and might disrupt network operations to the point of affecting availability (which can provoke a customer backlash). The use of standard interfaces between functional layers, as in the FORE system, is the key to holding down costs and ensuring better-than-satisfactory performance.

Maximizing Network Utilization Through Traffic Engineering

In traffic engineering, network designers map end-to-end traffic flows, match network capacity to those flows, and assign policies to different types and classes of service. Traffic engineering is an acknowledged strength of ATM. Another is the use of PNNI, a switch-to-switch protocol that supports the dynamic routing of Switched Virtual Circuit (SVC) requests. PNNI is a major breakthrough in routing: It is scalable, plus it is the only routing protocol that implicitly understands QoS.

Working together with our customers, FORE has built upon these strengths to develop traffic engineering methods that make optimal use of your infrastructure, and routing methods guided by the state of the network itself. We designed and implemented this networking software based on years of monitoring and analyzing the operations of customer-deployed WANs across the world.

A brief example illustrates why traffic engineering is a problem that must be solved if we are to sustain any kind of traffic volume over time. Unmanaged traffic on the network core inevitably leads to an imbalance in link utilization, with some underutilized and others overutilized to the point of congestion. And this costs the network operator money, because the overutilized links must be upgraded with additional bandwidth to relieve the congestion. Given the current state of routing technology, the existing underutilized link capacity cannot be tapped, which essentially forces upgrades to meet traffic demands.

The ATM PNNI protocol addresses the imbalance in link utilization that results when traffic is left unmanaged on the network core.

The simple topology shown in Figure 2a (below) illustrates the traffic engineering problem. There are two paths between B and C: B-C, and B-A-C. The routing protocol is programmed to choose the least-cost link as its primary path, and in this case the least cost link is B-C. There is a high volume of traffic between 1-6 and 2-5, which is routed across the primary link. The result is serious contention for resources on the B-C path, while the B-A-C path is underutilized. The result is low throughput, high delay, and a low rate of return on the SP's bandwidth investment.

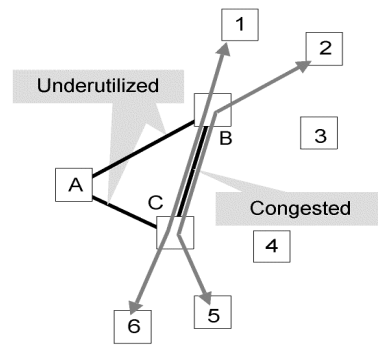


Figure 2a

Configuring a PVC-Based ATM Core

Because traffic management is so critical, many ISPs have already deployed switched core networks using ATM switches. Using offline path calculations, a traffic engineer configures explicit end-to-end routes across the ATM network using Permanent Virtual Connections (PVCs) and maps IP flows to the PVCs on the routers. Connections direct the traffic across links to optimize network capacity. In our traffic example, replacing core routers with ATM switches allows the network to use the available capacity on the B-A-C path. Figure 2b (below) demonstrates that provisioning a connection for traffic between 1-6 across the B-A-C path relieves congestion without adding links or switches to the network.




Figure 2b

In a PVC-based ATM core, if the end points are defined, the connections themselves direct traffic across links to optimize network capacity.

In a typical configuration, multiple PVCs are provisioned for each flow to provide alternate paths in the event of a link failure, and IP routing protocols reroute traffic onto the alternate PVCs. Although a PVC core brings all the advantages of ATM traffic engineering, it has service provisioning disadvantages. IP routers and ATM switches must be provisioned separately. All PVCs must be provisioned end-to-end (a significant task in a large network), and alternate PVCs must be provisioned between routers to work around failure. This doubles the number of connections required in the network. An additional disadvantage is that the routers are responsible for failure recovery, which can result in long recovery times.

Configuring an ATM Core with Capacity Aware Routing

These are daunting disadvantages from the SP point of view. However, FORE overcomes these problems with two different techniques:

- Directed Soft PVCs (D-SPVCs) allow the network manager to retain control of flows, plus it adds fast recovery by switching over to specified alternate paths.
- *Capacity-Aware Routing* introduces a true dynamic, “undirected” SPVC that makes routing decisions based on the state of the network itself. It is the more powerful of the two methods.

With D-SPVCs, the network engineer identifies explicit primary and up to three secondary end-to-end paths for each IP flow. Like a PVC configuration, these paths are assigned PVC identifiers on the ATM switch ports connected to the routers, and IP flows are mapped to the PVCs on the routers. Unlike a PVC configuration, the ATM switches use PNNI signaling to complete the end-to-end connection along the primary path.

The D-SPVC approach does lessen the complexity of configuration, but the real FORE advantage is in *Capacity Aware Routing*. Capacity Aware Routing is FORE’s implementation of PNNI, and it breaks significant new ground with its undirected, automatic routing. You provision the endpoints and assign a capacity to each connection, and the switches choose the best path between them, all in real time.

Use of PNNI clearly reveals the differences between unmanaged traffic and managed traffic on the network core. IP routing protocols have some built-in intelligence; they know how to find the open shortest path in response to a network outage. However, they do not know anything about the conditions on that path. PNNI does. The IP protocols will queue traffic on an already congested link. PNNI will not, because it understands the capacity available on each link and switch, and will choose the best open path to the endpoint.

FORE’s Capacity Aware Routing introduces dynamic connections that make routing decisions based on both link optimization and traffic QoS requirements.

With Capacity Aware Routing, FORE takes PNNI a step further: We've built QoS understanding into the network. Not only do you benefit through the simplified configuration — you identify the endpoints, PNNI chooses the route — but with Capacity Aware Routing, you also identify the service requirements for the end-to-end flow (bandwidth needed, service level, priority, tolerance for delay). PNNI finds not only the best path based on the topology, but also the best path based on topology and service requirements. And only FORE offers this routing innovation today.

The SP benefits derived from Capacity Aware Routing are significant. Simplified configuration and maintenance lessen network complexity, even greatly diminishing the “N-squared” problem associated with ATM networks by eliminating the need for redundant connections between routers. Intelligent routing improves service reliability, thus delivering on SLA bandwidth guarantees. And it lessens the impact of outages on network availability. Capacity Aware Routing is an essential piece of the FORE advantage, and it's your network that benefits from it.

Integrating Traffic Engineering with IP Routing Using MPLS

One goal of MPLS is to provide direct communication between the service and switching layers — thus producing a vast improvement in how network services are provisioned.

If Capacity Aware Routing and PNNI can offer such benefits today, what more can fully implemented MPLS add? To answer that question, consider a nagging inconsistency that Capacity Aware Routing alone cannot remedy: the IP and ATM protocols cannot communicate with each other directly. It is the reason why switches and routers are provisioned separately, and why the advantage of Capacity Aware Routing is isolated in the switching core. MPLS will provide, finally, the means for direct service-to-switching communication. MPLS will eliminate two-stage provisioning and greatly simplify configuration because the network manager, at the router, can provision the service and identify the endpoints, and MPLS provides the interface between the service and switching layers. Provision services at the router, and they're provisioned for the whole network. This is a significant step forward. Every service is mapped to a connection at the edge of the network, closest to your interface to the network.

The conclusion of this examination of traffic and how FORE handles it can be summarized as “cost-effective migration.” You can start today with the best high capacity networking available: IP over ATM, driven by FORE's Capacity Aware Routing-enhanced PNNI. Then migrate to full MPLS as it becomes available in an orderly way, by software upgrade, at a cost lower than modifying the infrastructure.

Network Resiliency

Network resiliency is a measure of network availability and protection against failure. It is a key building block in the IP service network. Maximum uptime is essential to building and maintaining a high level of customer satisfaction and reducing maintenance costs.

In a communications network like the PSTN, the resiliency is built into the switch and the link, in a pattern of one-for-one redundancy. Install two links between the same endpoints with fully redundant switches at an endpoint, and if one goes down, immediately switch to the other. Although this can provide reliable 100 connections, it is costly: Adding infrastructure in this way adds significant expense and may not be necessary for all services. FORE has taken a software approach that provides resiliency and lowers cost, using a technique called *Dynamic Protection Switching*.

To enable Dynamic Protection Switching, a network manager configures an end-to-end mesh, and defines policies that prioritize routing for different types and classes of service. But FORE switches have immediate access to real-time knowledge about the network state, and will route around link failure automatically. The software can also prioritize the order in which connections are re-established after failure. This offers enormous benefits to the SP:

- Automatic rerouting ensures resiliency and helps you meet QoS goals.
- Dynamic Protection Switching is fast: it can route around a failure in as little as seven ms. This reinforces the competitive choice of ATM for core switching.
- By building resiliency into the network layer, you save the money you would otherwise spend on building redundant links into the infrastructure. You can invest that hardware money elsewhere: in expanding into new markets, or in providing new services to your existing customer base.

PNNI and Failure Recovery

During an outage, PNNI quickly identifies the failed link or switch and establishes the alternate path for all connections affected by the outage. PNNI can identify the failure very quickly because of the nature of ATM connections. PNNI maintains state for each connection in the network. When a failure occurs, the PNNI software on switches adjacent to the failure immediately recognizes the failure and sends a notification for each connection using the failed link. When notified, switches at the endpoints of the D-SPVCs establish the alternate path configured by the network manager.

FORE's Dynamic Protection Switching lowers the cost of network resiliency by incorporating failure recovery in software, rather than requiring hardware redundancy. Working with PNNI, failures are resolved in as little as 7 ms.

The entire failure recovery process occurs at the ATM layer and is transparent to IP routing. In the FORE switch, failures are resolved by PNNI working together with DPS in as little as 7 ms. (Refer to Figure 3 below to see how the FORE switch compared to other ATM switches in recent tests conducted by an independent laboratory.) This low number is important from the service provider point of view, because fast network recovery results in high service availability.

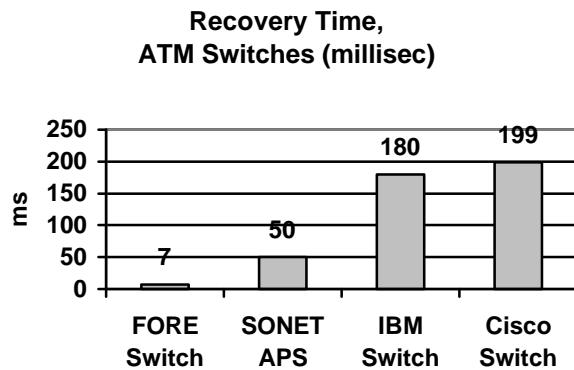


Figure 3: Source: Data Communications¹

Ensuring Router Scalability with MPLS

FORE's MPLS implementation sidesteps adjacency requirements and improves router efficiency by opening a one-hop tunnel between the source and destination routers — both increasing throughput and reducing overhead.


Cost-effective IP routing scalability is a major concern to the growth-oriented SP. FORE's implementation of MPLS functionality within PNNI can assist the SP in streamlining the amount of network maintenance required of IP routers. And freed resources can be applied to other, more profitable areas, such as network growth or service expansion.

path (forwarding adjacency), and it sends and receives network state information to and from its adjacent routers to maintain an accurate routing table (flooding

1. Test results reported in Figure 3 were reported in the article, "ATM: Brains and Brawn," by Robert

Data Communications Test Lab, where switches from FORE, IBM®, and Cisco® were tested.

Despite the fact that it knows about addresses at a distant router, a source router has no forwarding adjacency to that distant router, so it follows the flooding adjacency path, hop by hop. Figure 4a (below) shows the complexity of this arrangement. In a large network, handling the forwarding adjacency traffic can consume substantial processing power as traffic makes its way across the full mesh.



able to open a “label switched path” or tunnel to the destination router, despite the fact there is no forwarding adjacency between the two. Traffic traverses the , greatly reducing the complexity in favor of Figure 4b’s simplified routing (shown below). The result is a dramatic decrease in network-




Figure 4b

is necessary today because the IP and ATM signaling and routing protocols aren’t the same — a central issue that MPLS is designed to resolve.

FORE's ForeRunner ASX-4000 switch ensures the high level of service availability provided by APS through its use of redundant load-sharing power supplies, a redundant switch control processor, and port-level redundancy.

Developing a High-Speed Transmission Layer

The transmission layer performs numerous networking functions: maximize capacity, provide reliable point-to-point delivery, ensure correct timing for network-wide synchronous services, and ensure fault detection and recovery. Many of the advances necessary to keep pace with growing traffic volumes are being made at the transmission layer. These advances boost both the transmission capacity of each fiber, and the transmission efficiency. Technologies include Automatic Protection Switching (APS), Frame-ATM for optimizing the delivery of data traffic, and DWDM for moving more traffic over the same fiber.

Providing Service Availability with APS Protection

Automatic Protection Switching (APS), part of SONET, provides two paths for optical signals: a working path and its backup, the protection facility. The protection facility mirrors everything that is transmitted over the working path, and if failure occurs at either end of the working path, the protection facility comes into immediate use. The network benefits through a higher level of service availability.

FORE's ForeRunner® ASX-4000 switch ensures this level of availability through its use of redundant load-sharing power supplies, an optional redundant switch control processor, and port-level redundancy using APS. This permits in-service software upgrades or maintenance with no associated down time. In its "30/10" configuration, the ASX-4000 can be configured as a 30 Gb switch, with 10 Gb fully protected using a traditional redundant physical link and switch port, and another 20 Gb configured using Dynamic Protection Switching. This approach delivers best of both worlds to SPs: traditional redundancy with network layer resilience on the same switch, at the same time.

By combining the ASX-4000's less costly level of product and link redundancy with ForeThought's sophisticated network resiliency, service providers can achieve high availability without the cost and limited capacity of fully redundant links and switches.

Maximizing Link Efficiency with Frame-ATM

One of the criticisms of ATM is that it is less efficient than IP routers with POS links for a data-oriented IP service network because it uses a 53-byte cell, of which 48 bytes are user data. The remainder is sometimes called the ATM cell tax — overhead tacked on to the cell that enables it to traverse the ATM network. The overhead can consume up to 15 percent of the available bandwidth. The small cell,

however, processes quickly and is less subject to delay, which makes it ideal for multi-service networks transmitting voice and video services. In most service networks, the multi-service and bandwidth management capabilities of ATM greatly outweigh the ATM cell tax. However, where the line cost is significant, such as with transcontinental or transoceanic links, the cost of the cell tax is amplified.

There is, however, a solution for data transmission: Frame-ATM, which uses a variable-sized cell and follows the frame relay model. Frame-ATM reduces the cell tax by sending a bigger payload with the same number of overhead bytes used in the 53-byte cell. Using the larger frame can move data faster in resource-constrained situations, for example when only one port serves a particular destination.

Reducing Deployment Costs with DWDM Transmission Technology

DWDM technology allows very high capacity transmission over a single optical fiber, thus enabling a service provider to meet the increasing capacity demands posed by Internet traffic using the existing infrastructure. FORE has incorporated DWDM transmission technology into the ASX-4000 switch to enable very high transmission capacity without the associated cost or complexity of separate transmission equipment. This is particularly useful when deploying a new service network on a fixed capital budget.

FORE's WDM interfaces use ITU standard wavelengths to ensure interoperability with other standard DWDM equipment. Using FORE's approach, up to four 2.5 Gbps links can be optically combined over a single fiber to deliver a total of 10 Gbps, equivalent to OC-192 speeds. Using standard ITU wavelengths also makes possible clean upgrades to higher density WDM in the future. The output from the FORE switch can be optically inserted into a DWDM link, without the need for costly electronic conversion hardware.

Incorporating DWDM dramatically boosts the transmission capacity of the existing infrastructure. The ASX-4000 switch can deliver a total of 10 Gbps over a single fiber — a speed equal to OC-192.

CONCLUSION

FORE is focused on providing service provider solutions, and FORE's ATM with PNNI architecture presents SPs with the best core switching layer available today. In every facet of building an IP service network — implementation of MPLS, traffic engineering, Capacity Aware Routing, network resiliency and failure recovery, router scalability, and frame-based ATM — the FORE advantage can lead to cost reduction, revenue enhancement, and maximum return on your networking investment.

The future is filled with opportunity, and FORE will continue to develop its industry-leading switching hardware and ForeThought networking software product lines to seize these opportunities and expand our presence in the emerging SP market. Anticipated advances include integrating the transmission and switching layers with WDM technology, and completing the addition of MPLS to ATM switches to serve as the interface to the IP Service Layer. FORE is committed to providing an IP-over-ATM architecture that will meet the subscriber demand for bandwidth and services well into the next century.

ACRONYMS

APS	Automatic Protection Switching
ATM	Asynchronous Transfer Mode
DWDM	Dense Wave Division Multiplexing
IP	Internet Protocol
IS-IS	Intermediate System–Intermediate System
ISP	Internet Service Provider
MPLS	Multi Protocol Label Switching
OSPF	Open Shortest Path First
PNNI	Private Network to Network Interface
POS	Packet Over SONET
PSTN	Public Switched Telephone Network
PVC	Permanent Virtual Connection
PVP	Permanent Virtual Path
QoS	Quality of Service
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement
SONET	Synchronous Optical Network
SP	Service Provider
SPVC	Soft Permanent Virtual Connection
VPN	Virtual Private Network
WDM	Wave Division Multiplexing

© Copyright 1998 FORE Systems, Inc. All rights reserved. FORE Systems, ForeRunner, and ForeThought are registered trademarks and Capacity Aware Routing and Dynamic Protection Switching are trademarks of FORE Systems, Inc. IBM is a registered trademark of International Business Machines Corp. Cisco is a registered trademark of Cisco Systems, Inc. Please be advised that technical specifications set forth in this product literature are correct as of December 1998 and are subject to change by FORE Systems, Inc. at any time without notice.