September 1997

Ascend

Competitive Analysis

Comparing GX 550 to Competing Core ATM WAN Switch Products:

Cisco BPX-ST
Newbridge 36190
“Other Players”

ASCEND

Table of Contents

ATM WAN Switch Market Evolution

In the latter half of 1996, the industry witnessed substantial growth in ATM WAN switch deployments. The pilot networks and test labs installed in 1994 and 1995 evolved into large-scale service networks. In 1997, service providers have recognized that in order to provide a successful and profitable data service offering, they need to provide higher density and higher speed access and trunk interfaces, while also improving overall system availability and resiliency at the core of their networks. Additionally, service providers are now looking to cost-effectively integrate lower-speed services such as Frame Relay and new services such as XDSL into these highly available, high-speed core WAN switching solutions.

Service providers realize that only the core ATM WAN switching platforms offering the lowest cost of ownership, while also providing the most feature-rich technology will satisfy their critical business objectives. As a result, they require core WAN switches that provide high-density connectivity coupled with high levels of system availability, seamless connection establishment, extensive diagnostic features and robust end-to-end traffic management.

For service providers to meet these evolving market requirements, the underlying core WAN switching solution must support the technologies that provide a high level of switch intelligence and feature richness. A successful WAN switch offering will not only provide sufficient bandwidth, but will also provide detailed service measuring and customer specific Quality of Service (QoS).

Competing Products

Most competing core WAN switch products are built on architectures that promise to scale to large and bandwidth densities, but offer little in regards to system availability, connection establishment, diagnostic features and end-to-end traffic management. Competitors will promise to include functionality that their first-generation edge WAN switches do not provide. As a result, most of these new competing core WAN switch platforms will be replacements to their edge switches rather than a logical migration from them.

Primary Competitors

There are a number of traditional WAN switch vendors and start-up companies attempting to develop a solution for service provider needs at the core of the network. This document will provide detailed analysis of Ascend's two primary competitors in the emerging core ATM WAN segment: Cisco BPX-ST and Newbridge 36190.

Other Competitors

Other competitors attempting to compete in the ATM core switch space include Nortel, Newbridge, GDC, fore and Lucent. These players will be covered in less detail for various reasons. Fore has plans to deliver 20 and 40 Gb/s ATM switches, but recent financial problems have caused many canceled developments. Nortel's Concorde switch has been available for quite sometime and will reportedly scale to 80 Gb/s, but its limitations have prevented it from gaining acceptance in the marketplace. The Newbridge 36170 is an edge switch that competes with the Ascend CBX 500, not the GX 550. The 25.6 Gb/s GDC Strobos is undelivered and the company is facing bankruptcy. Lucent remains a mystery as to what they may or may not deliver. In mid-1998, Lucent plans to deliver an ATM switch that will scale to 160Gb/s, but there are few known details.

The Ascend Advantage

The Ascend GX 550 is a scalable, high-capacity core ATM WAN switching system providing from 25 Gb/s to 100 Gb/s of bandwidth. The switch will support high density OC-3 and OC-12 access, with OC-48 trunking capabilities at its introduction. Trunk speeds will soon scale to include OC-192 interfaces.

The GX 550 is designed to cost-effectively scale today's integrated data networks and provide large-scale native ATM service deployment. The GX 550 is highly distributed and autonomous in design. Not only does the GX 550 provide the highest levels of port density and circuit scaling, more importantly the GX 550 offers the most future-proof core WAN switch technology in the industry – offering the highest levels of system availability, robust end-to-end QoS, seamless point-and-click provisioning, and extensive diagnostic features.

The GX 550 protects the Service Providers' current investment in WAN switch equipment, and provides cost-effective scaling to grow their current networks. With OC-48 or OC-192 trunk interfaces providing high-speed ATM backbones for CBX 500 and B-STDX 9000 switches, the GX 550 provides Internet Service Providers (ISPs) with the ability to continue scaling their frame-based networks. Public Carriers can use the GX 550 to deploy widespread OC-3/OC-12 native ATM service, while offering a multiservice QoS capable backbone. Additionally, the GX 550 provides a high-speed cell-based backbone for the aggregation of ADSL service muxes for delivery to appropriate content providers.

Connection routing and failure recovery decisions are hosted on each GX 550 and not dependent on a centralized Network Management System (NMS) server. In this manner, the Ascend network operations capabilities are quite unlike traditional central-weighted management systems. Because a service's quality is only as good as the measurement system behind it, the GX 550 is designed with rigorous attention to gathering the port, circuit and service level statistics. Ascend products deliver the service statistics and billing systems required to migrate customers to a service offering where service level guarantees are universal across voice and data offerings.

Core WAN Switch Competitive Matrix

Features	Ascend GX 550	Cisco BPX-ST	Newbridge 36190
Architecture			
Bandwidth	25 Gb/s shelf; scalable to 100 Gb/s	20 Gb/s (9.6 Gb/s “useable”)	20 Gb/s shelf; plans to scale to 640+ Gb/s
Number of Physical Interface Slots	40	12 on broadband shelf 16 on narrowband shelf	16 per shelf
Flexible “Mixed” Interfaces per line card	Yes	No	No
NEBS Compliant	Yes	Yes	Yes
Maximum Port Density: “Single Shelf”			
ATM DS1/ E1	Future	80: narrowband shelf (1280: total per BPX)	128 per shelf
ATM DS3/E3	Future	36 today 144: future promise	64 per shelf
OC3/STM-1	160	24 today 96: future promise	32 per shelf
OC-12/STM-4	40	24: future promise	Future: 16 per shelf
OC-48/STM-16	10	Not capable	Future promise
OC-192/STM-64	Future	Not capable	Plans unknown
Logical Interface Support			
UNI 3.0 and 3.1	Yes	3.1 only	No
Signaling 4.0 (UNI 4.0)	R2.0: phase 1	Future	Yes
IISP	Yes	Claimed	Yes
PNNI 1.0	Yes	Future	Future
BICI: V1.1 (PVC), V2.0 (SVC)	Yes	Promised	Yes

Core WAN Switch Competitive Matrix (Continued)

Features	Ascend GX 550	Cisco BPX-ST	Newbridge 36190
Circuit Scaling			
VC Support	2.5 million/switch	24,000/switch today 384,000/switch future	Unknown
SVC Calls/Second	5,000/switch	50-100/switch today 1,000 future	Unknown
Multicast	192,000/switch	None today 4000/switch future	Unknown
Cell Buffering	10.2 million/switch	~ 144,000/switch today 2.4 million/switch future	Unknown
System Redundancy			
Switch Fabric	Yes	Yes	Yes
Lossless Switchover	Yes	Unknown	Unknown
Node Processor	Yes	Yes	Yes
Interface Cards (Intercard)	Yes	Claimed: DS3/E3; OC-3 OC-12 future	Yes
Port Level (Intracard)	Yes	No	Yes
System Availability Features			
Mid-plane Design	Yes	Yes	Unknown
Redundant PCMCIA Hard Disks	Yes	Yes	No
Redundant FLASH Drives	Yes	No	No
Redundant Timing Module	Yes	No	No
Distributed Processing	Yes	No (needs Fulcrum)	Claimed
Distributed Signaling	Yes	None today, future	Yes

Core WAN Switch Competitive Matrix (Continued)

Features	Ascend GX 550	Cisco BPX-ST	Newbridge 36190
Service/Operation Features			
Minimal Disruption (Redundant Hard Disks)	Yes	Yes	No
Point-Click End-to-End Provisioning	Yes	Claimed	Unknown
Self-Learning Reroute Capability	Yes	Claimed	Claimed
Statistics Gathering	Yes	Level unknown	Level unknown
Performance Monitoring	Yes	Future promise	Future promise
Fault Management	Yes	Future promise	Future promise
Integrated Billing Servicing	Yes	Future promise	Future promise
Traffic Control Features			
ATM Service Classes	CBR, rt-VBR, nrt-VBR, UBR, UBR+, ABR	CBR, rt-VBR, nrt-VBR, UBR, ABR	CBR and rt-VBR only at FCS
Quad-Plane Buffer Architecture	Yes	No	No
Per VC Queuing	Yes	Future promise	Yes
Per VC Policing	Yes	Future promise	Yes
Dynamic Back-Pressure Cell Buffering	Yes	No	No
Dual Leaky Buckets	Yes	Unknown	Yes
NTM/NDC	Yes	Unknown	Unknown
TM 4.0	Yes	Future promise	Unknown
Packet Level Discard	Selective cell discard, EPD, PPD, RED	Unknown	EPD and PPD only
Traffic Shaping	Yes	Unknown	Unknown
Policy-Based Routing	Yes	Unknown	Unknown
VPN	Yes	Unknown	Unknown

Core WAN Switch Competitive Matrix (Continued)

Features	Ascend GX 550	Cisco BPX-ST	Newbridge 36190
ATM SVC Service Features			
SVC Closed User Groups	Yes, up to 1024 per switch	No	No
SVC Security Screens	Yes	No	No
SVC Addressing Formats	Native E.164 E.164 AESA ICD AESA DCC AESA custom AESA	Unknown	Unknown

Competing Platform Overviews

Cisco BPX-ST

How to Beat BPX-ST	
Cisco BPX-ST	• Superior bandwidth and density scaling • Superior system availability features • Robust statistics-gathering and diagnostics
GX 550	• Most complete traffic control features available • Fully distributed processing provides superior performance • Flexible “mixed” interface options per line card • Functionality and feature promises continuously delayed • No port-level redundancy today • Expensive and external route/topology processing
BPX-ST	• Inferior density and Virtual Circuit scaling • Minimal statistics gathering/diagnostic capabilities • Lacks bandwidth required for a “Core” WAN switch

Cisco Core WAN ATM Solution Description: BPX-ST

The majority of BPX-ST switches have been deployed in carrier-based Frame Relay services: AT&T, the largest single BPX customer, uses the platform almost exclusively for Frame Relay. Cisco claims that new frame developments will continue on the platform. However, Cisco has positioned its 75xx platform as a multiservice router/switch that is now their lead Frame Relay platform. The BPX-ST, with its newly “re-announced” ATM line cards, has been repositioned as their core ATM WAN switch, with much less emphasis on Frame Relay.

Cisco’s repositioning of the BPX-ST was driven in part by its need to protect the tremendous installed base of Cisco routers. Today, Cisco will generally position the existing 75xx to ISPs for low-speed Frame Relay switch/routing. The GSR will be positioned to ISPs for POS (Packet over SONET) trunking and alleviating the congestion problems associated with the 75xx. The BPX-ST is positioned to carriers for ATM switch requirements.

Cisco previously justified this repositioning with its September 1996 re-announcement of a new switch fabric and high-density, high-bandwidth ATM cards for the BPX. The BPX-ST platform provides a 20 Gb/s backplane switch fabric (“ST” signifies the 20 Gb/s fabric), and the ***future*** ATM line cards—previously code-named Monarch—are referred to as Broadband Switching Modules (BXM).

BPX-ST ATM Line Card Offering

Today	**Access Cards** • 4-port T1/E1 UNI (Narrowband shelf via external AXIS shelves) • 2-port T3/E3 UNI (Broadband shelf) • 2-port OC-3 UNI (Broadband shelf)
	Trunk Cards • 3-Port T3/E3 NNI (Broadband shelf) • 2-port OC-3 NNI (Broadband shelf)
~ 3 Months	**Trunk Cards** • 8-port T1/E1 IMA NNI (Narrowband shelf)
~ 6 Months	**BXM Cards*** • 6-port T3/E3 (Broadband shelf) • 12-port T3/E3 (Claimed via Broadband shelf) • 4-port OC-3 (Broadband shelf) • 8-port OC-3 (Broadband shelf) • 1-port OC-12 (Broadband shelf) • 2-port OC-12 (Broadband shelf)

** Cisco claims that BXM cards provide either access or trunk connectivity.*

BPX-ST ATM Product Pricing

BPX-ST Chassis*	Availability	List Price
Non-Redundant BPX-ST	Now	$27,450
Redundant BPX-ST	Now	$49,500

AXIS Shelf*	Availability	List Price
Non-Redundant AXIS	Now	$23,500
Redundant AXIS	Now	$37,500

** Several bundled BPX/AXIS packages available*

Narrowband Shelf Line Cards** (via AXIS)	Availability	List Price
4-port T1/E1 UNI	Now	$15,500
8-port T1/E1 IMA NNI (Narrowband shelf)	2H97	$18,000

*** Includes cost of back card.*

Broadband Shelf Line Cards**	Availability	List Price
2-port T3/E3 UNI	Now	$14,000
2-port MM OC-3 UNI Access	Now	$18,000
2-port OC-3 UNI (IR SM fiber)	Now	$24,000
2-port OC-3 UNI (LR SM fiber)	Now	$28,000
3-port T3/E3 NNI	Now	$27,000
2-port OC-3 NNI	Now	$22,000
2-port OC-3 NNI (IR SM fiber)	Now	$28,000
2-port OC-3 NNI (LR SM fiber)	Now	$32,000

*** Includes cost of back card.*

Future Broadband Shelf BXM Cards	Availability	List Price**
6-port T3/E3	1H98	$24,000
12-port DS3	1H98	$48,000
4-port OC-3*	1H98	$28,000
8-port OC-3*	1H98	$56,000
1-port OC-12*	1H98	$22,000
2-port OC-12*	1H98	$44,000

** Assumes MM fiber (pricing for SM fiber will be more)*

***** All future BXM card pricing is estimated: Cisco is reluctant to quote actual cost.***

Required Extras		
Fulcrum INS	Now	$40,000-$100,000
Foresight License for BPX	Now	$10,000 for each T3/E3 Trunk card
AXIS ABR Foresight	Now	$4,000 for each 4-port ASM card
IMA Software	3 months	$8,000 for each 8-port T1/E1 card
Configuration Save and Restore License	Now	$2,000 per BPX-ST
Virtual Trunking License	Now	$10,000 BPX-ST
Multiuser Config. Sessions	Now	$4,000 BPX-ST
Priority Bumping Software License	Now	$5,000 BPX-ST
AXIS Service Redundancy License	Now	$8,000 per shelf

Cisco BPX-ST Overview

In September 1996, Cisco announced new BXM ATM interfaces and a 20 Gb/s switch fabric to help better position the BPX-ST as an ATM-core WAN switch for the carrier market. The 20 Gb/s switch fabric is designed to optimize operation on the new BXM cards, which consist of ATM DS3, OC-3, and OC-12 options. ***The BXM interfaces are not shipping,*** and there are still many unknowns revolving around these future ATM offerings—pricing is extremely vague, and time frames for feature enhancements have not been committed to. One fact is clear: ***The ATM products available on the BPX-ST today do not meet carrier requirements.*** Today, the BPX-ST platform provides limited density, scaling, and performance, and lacks sufficient system availability, end-to-end QoS, and diagnostic features.

BPX-ST Chassis Description

The BPX-ST provides connectivity via a Broadband shelf and a Narrowband shelf. The Broadband shelf can be considered the engine of BPX—it houses the 20 Gb/s switching fabric. The BPX switch employs a mid-plane design, and each front card has a corresponding back card that provides the physical interface to the transmission media.

The Broadband Shelf provides 15 slots for line cards. Three slots are reserved for common control modules, and 12 slots are available for high-speed (34 Mbps and above) interface cards. The Common Core Group modules consist of the Alarm Status Monitor (ASM) card and the Broadband Control Card (BCC). The ASM card provides alarm and status-monitoring capabilities. The BCC can be configured redundantly (second BCC card required) and supports the following functions for the BPX-ST: ATM cell switching; internal and remote node communication; node synchronization; network management communication; shelf management communication.

The Narrowband Shelf provides 16 slots for external AXIS shelves. The AXIS shelves provide low-speed access (below 34 Mbps) and are connected to the Narrowband Shelf of the BPX-ST via a proprietary DS3 connection. ***All ATM traffic going into the AXIS shelves must be passed to the Switch Fabric on the Broadband shelf, even if the destination is another port within that same AXIS shelf. As a result, the DS3 connection is a bottleneck.***

Each AXIS shelf provides 16 slots for line cards. Six slots are reserved for control and the uplink modules to the Narrowband Shelf, and the 10 remaining slots are available for the low-speed interface cards. The control cards consist of an Access Shelf Controller (ASC) and a Service Redundancy Module (SRM). The DS3 uplink card is referred to as a Broadband Network Module (BNM). One of each of these three cards is mandatory, while a second card for each can be installed for redundancy.

BPX-ST Limitations

Availability of BXM Product

Cisco "re-announced" the new BXM products and features back in September 1996. For nearly a year before that time, Stratacom had been alluding to BXM to its customers under nondisclosure agreements developments (then referred to as Monarch technology). ***Cisco, and previously Stratacom, has been positioning new ATM products for the BPX-ST for over two years—and the product is still not available!*** It is believed that Cisco is still six months away from delivering "production quality" BXM product and features.

Without the BXM cards, the BPX-ST is severely limited and cannot meet today's ATM WAN switch market demands: DS3 and OC-3 densities are limited to 36 and 24 ports respectively, and OC-12 is not available. Additionally, virtual circuit scaling, SVC performance, cell buffering, traffic management, and diagnostics are insufficient.

Hidden Costs

Today, Cisco's BPX-ST technology requires several add-ons to provide the density and features they promise: AXIS Shelves, extra line cards, separate line cards for trunk Vs access, Fulcrum INS server, Foresight congestion control software and bundled software options. Even with the future BXM cards, Cisco will still require many of these expensive extras.

The Fulcrum INS is just one example of an extra cost that Cisco avoids mentioning. The Fulcrum is essentially a rack-mounted SPARC Station with a 155 Mbps connection to the BPX-ST. It is required to support SVC signaling as well as routing/topology functions for the ATM cards available today. ***Although Cisco is claiming that the future BXM cards will provide onboard distributed call processing, the Fulcrum will still be required to provide routing/topology processing.***

The cost of Cisco's signaling and routing/topology implementation is substantial: The price of a Fulcrum INS ranges from $40,000 to over $100,000, depending on software. Additionally, the Fulcrum requires an expensive OC-3 connection to the Broadband Shelf of the BPX-ST. ***The excessive costs of Cisco's add-on functionality significantly increase the overall cost of the company's ATM switch offering.***

Usable Bandwidth

Cisco positions the BPX-ST as providing a 20 Gb/s switch fabric. However, each of the 12 Broadband shelf line cards only has a single 800 Mbps "inbound" rail to the switch fabric (two 800 Mbps "outbound" rails attach from the switch fabric to the line card). ***This design restricts usable bandwidth of the BPX-ST to only 9.6 Gb/s (12 x 800 Mbps).***

Blocking Architecture

The BPX-ST limitation of 9.6 Gb/s of usable bandwidth will create a ***blocking architecture for Cisco's future BXM OC-3 and OC-12 density claims.*** Cisco positions that the BPX-ST will support 96 OC-3 ports and 24 OC-12 ports. However, for non-blocking configurations, these port densities are reduced to 60 and 12 respectively. ***Furthermore, until the BXM cards are available, the high-speed ATM port densities on the BPX-ST are limited: 36 T3/E3 ports, 24 OC-3 ports, and no OC-12 support.***

The external AXIS Shelf technology creates another bottleneck. The AXIS Shelves are connected to the Narrowband Shelf of the BPX-ST via a proprietary DS3 connection. Cisco claims that each AXIS Shelf supports up to 40 ATM T1/E1 ports (80 with future 8-port cards). However, all ATM traffic going into the AXIS shelves must be passed to the Switch Fabric on the Broadband Shelf, even if the destination is another port within that same AXIS shelf. ***The result is that the DS3 connection from the AXIS Shelf to the BPX-ST is a bottleneck: Only 29 ATM T1 connections or 21 ATM E1 connections can be supported for non-blocking configurations.***

Conclusion

Stratacom announced the BPX in early 1993 but did not begin shipments until December 1993. The limitations on the BPX-ST platform have prevented it from making significant inroads into ATM WAN switch environments. Only 500-600 BPX/BPX-ST switches have been installed since the product's inception – over four years ago.

We know of ***less than 10 accounts in which the BPX/BPX-ST has been installed*** (AT&T, PacBell, WorldCom, CompuServe, Ameritech, Telecom Italia, Sita/Equant and Halifax Building Society via BT). Additionally, most of these installations cannot be considered significant *ATM* WAN switching networks. As an example, AT&T uses the BPX/BPX-ST platform exclusively for Frame Relay services. Also, Ameritech and WorldCom use the BPX-ST only for low-speed ATM. Some may argue that Telecom Italia is a significant ATM WAN switch account for the BPX-ST (potential of up to 250 switches is estimated), but this network is still in its infancy stage.

Until the BXM cards (which promise increased density and functionality improvements) are available, Cisco's BPX-ST simply does not provide the solutions to meet the needs of today's core networks. The GX 550 offers the scalability and robust feature set that today's core WAN switching environments require.

Newbridge 36190

How to Beat 36190	
GX 550	• Superior bandwidth and density scaling per shelf • Superior system availability features • Robust statistics — gathering and diagnostics • Most complete traffic control features available • Fully distributed processing provides superior performance • Flexible “mixed” Interface options per line card
36190	• Poor redundancy and availability implementation • Does not support All ATM service classes • QoS is not adequately addressed • Maximum density requires up to 32 peripheral shelves • Minimal statistics gathering/diagnostic capabilities • NMS responsible for centralized routing

Newbridge Core WAN ATM Solution Description: 36190

On January 20, 1997 Newbridge and Siemens announced that their networking alliance had resulted in the 36190 core ATM WAN switch. Newbridge/Siemens is positioning the 36190 as the largest ATM core infrastructure switch in the industry – with future scaling capabilities exceeding one terabit per second (Tbps). In January, Newbridge claimed that FCS of the 36190 would occur in mid-1997. However, the 36190 is still not ready for production networks, and it is believed that only a couple of sites are currently evaluating the switch.

36190 ATM Line Card Offering	
At FCS	• 8-port DS1/E1 UNI • 2-port OC-3 UNI • 4-port DS3/E3 UN
3-6 Months after FCS	• 3-port E3/DS3 CE with DS granularity • 1-port OC-12 NNI
6+ Months after FCS	• 3-port E3/DS3 CE with DS granularity • 1-port OC-48 NNI

36190 ATM Line Product Pricing

This is unknown at this time.

Newbridge 36190 Overview

The 36190 is nothing more than a redesign of the Siemens EWSXpress 3100 switch, and therefore will have many of the same limitations. The Siemens EWSX switch was intended to provide bandwidth options that scale from 2.5 Gb/s to 320 Gb/s. However, the EWSX has never scaled to those high-end claims, and severe feature limitations have prevented it from gaining acceptance as a credible ATM WAN switch. Lack of OC-12 support, poor buffering, limited traffic management, CBR-only services, minimal diagnostic features, no SVC or multicast support are just some of the limitations of the EWSX switch.

The "new" jointly developed 36190 will attempt to rectify some of the limitations of the EWSX switch: The design of the switching fabric has changed to scale from 20 Gb/s to 640 Gb/s; higher speed interfaces will be offered; hardware redundancy will be improved and additional service class support are all promised by Newbridge.

Siemens is hoping that the inclusion of the Newbridge product labeling (36190), as well as new engineering resources, will provide a new life into their EWSX switch. Newbridge is looking to Siemens to provide the "big company" credibility they are looking for.

The 36190 is positioned as a jointly developed switch. Newbridge is quick to include in all their marketing documents that the 36190 is "the result of 1500 collaborative Newbridge/Siemens ATM development engineers." However, in reality ***Newbridge is providing little in regards to the 36190 development.*** Newbridge has provided some "consulting engineers" to the project, however, it is believed that they have little impact and responsibility in regards to the actual development process.

36190 Release Schedule

Newbridge is claiming that the initial release of the 36190 will occur in Q497. Upon this release, the 36190 will be limited to 20 Gb/s bandwidth; will not provide interface speeds higher than OC-3 and QoS will be limited to CBR and rt-VBR.

Newbridge promises that the next "major" release will occur within ninety days of the initial release. When this occurs, the 36190 will supposedly scale to 160 Gb/s; include ATM DS3 circuit emulation with DS1 granularity; add OC-12 interfaces and provide PNNI support.

The third "major" release will allow the 36190 to scale beyond 160 Gb/s (by just how much has not been identified); include ATM DS3 circuit emulation with DSO granularity; and provide OC-48 interfaces. Newbridge is claiming that this release is in the planning stage, so no time frame as to expected availability is provided.

Newbridge 36190 Limitations

Density

In addition to tremendous bandwidth promises, Newbridge is positioning the 36190 as providing the industry's highest densities. However, to obtain these numbers the 36190 will require as many as 32 separate peripheral shelves (this does not include the core "switch fabric" shelves, shelves for control units or shelves for fan boxes and breaker units). Because the *per-line card* port densities are so poor (see above), Newbridge requires an excessive amount of expensive floor space to obtain their density claims.

Advanced Services and QoS

The 36190 does not effectively address ATM advanced services or QoS enhancements. Instead, Newbridge/Siemens focuses on the fact that the 36190 will offer the "world's largest core ATM carrier-class switch." It is unknown if the 36190 will ever attain the 640 Gb/s bandwidth levels that Newbridge/Siemens promise. Regardless of these bandwidth claims, today's service providers need more than the promise of high bandwidth. To meet their critical business objectives, Service Providers require the ability to provide advanced services, a complete QoS offering, high system availability, robust diagnostics and superior circuit scaling.

NMS Responsible for Centralized Routing and Control

Separate control unit shelves house line cards that store configuration information. ***A centralized 46020 Network Manager is actually responsible for configuring the route topology of the 36190 network, not just managing the network.***

The 46020 Network Manager workstations can attach to the network via an Ethernet connection. Newbridge recommends that each 46020 be connected directly to their switch via an OC-3 connection. Newbridge also requires that a second 46020 be installed for redundancy. In this scenario, the backup 46020 would become active if there was a failure to the primary 46020 itself or its connection to the network.

Newbridge claims that a single "pair" of 46020 Network Managers (primary and backup) can provide route topology setup and signaling for the *entire* Newbridge network. However, multiple 46020 Network Managers are required to provide adequate scaling for large networks consisting of several Newbridge switches. In these configurations, several "delegate" 46020s must be installed to effectively support the entire network. Additionally, multiple "En Point" software licenses are needed for each switch to provide the "local intelligence" to the various line cards installed within that switch. If this is not confusing enough, separate network management applications are required to support other services and Newbridge switch platforms.

The end result is ***that Newbridge requires the customer to invest in an extremely expensive management and control solution.*** Each 46020 Network Manager costs $10,000 ($20,000 for the primary and backup pair). Each "delegate" 46020 costs $12,000 and each 36170 can require several thousands of dollars of En Point licenses. Every 46020 requires the additional cost of $22,000 for SUN or HP OV support and every 46020 requires the cost of a workstation.

Expensive Switching System Configuration

Although pricing is unknown at this time, it is believed that the higher bandwidth 36190 configurations will be very expensive. The limited per-line card densities will require that many separate peripheral shelves be added to scale the configuration. These configurations will require expensive external cabling to connect the shelves, additional control unit shelves for management and control, as well as a tremendous amount of floor space.

A fully configured, redundant 36190 Switching System will exceed 35 total shelves (for the planned 640 Gb/s configuration). Not only are there excessive costs associated with purchasing multiple switches, but the shelf-connecting cable costs alone will be significant.

Conclusion

The joint development of the 36190 is a defensive move for both Siemens and Newbridge. Siemens hopes that adding the "Newbridge label" to their product will provide them with the ATM WAN switch that the EWSX could not. Newbridge is looking to Siemens to provide them with credibility and starting point for the high-bandwidth core switch that their own internal R&D resources cannot provide.

However, relying on Siemens-driven development efforts will postpone the delivery of the 36190, and any future enhancements to the platform. Siemens has a poor reputation in the data communications industry and virtually no marketshare presence. Allowing Siemens to absorb the costs of adding new functionality to an existing product will save Newbridge millions in R&D expenses, and provide them with a core ATM switch much sooner than if they developed it themselves. However, probable shipment delays and lack of robust features on the 36190 will result in a lack of widespread customer acceptance.

The Other Players

Newbridge: 36170

The 36170 is a multishelf switch configuration that currently provides up to 12.8 Gb/s bandwidth. For a detailed description of the 36170 and its limitations, please refer to the Core Switching Unit Competitive Web page located on the Ascend Intranet.

Today, the 36170 is positioned against the Ascend CBX 500. However, for several months, Newbridge has been promoting that the 36170 will scale to 51.2 Gb/s. Because of this promised increase in bandwidth, Newbridge might attempt to position the 36170 against the Ascend GX 550. However, even though Newbridge marketing documents currently refer to the 36170 as providing 51.2 Gb/s, as recently as May, 1997, Newbridge engineers had not decided how they would engineer this bandwidth increase. Therefore, do not expect the 36170 to actually provide an increase in bandwidth until late 1998, or even 1999.

GDC: Strobos

Few details are known about the Strobos switch. GDC introduced Strobos at N + I in September 1996, claiming that the switch would ship in Q497. However, it is believed that GDC has already slipped FCS into 1H98.

Strobos will provide a 25.6 Gb/s ATM non-redundant switch fabric (via two separate 12.8 Gb/s fabrics). The Strobos switch fabric can be configured redundantly, but this will reduce the available bandwidth to 12.8 Gb/s. There will be two processors (switch controllers) for switching/topology processing, and two additional processors for management and signaling. Although these processors can be configured redundantly, they are still restricted to providing centralized processing – this underlying limitation will effectively limit the Strobos scaling and performance capabilities.

There will be 16 slots available for I/O cards. GDC plans to provide 8-port OC-3 and 2-port OC-12 cards, however, it is believed that only the OC-3 cards will ship at FCS. Each I/O slot has two 800 Mbps threads to the switching fabric - one to each 12.8 switch fabric. This will provide enough bandwidth for the OC-3 and OC-12 cards, but it is doubtful that this architecture will effectively support OC-48 capabilities.

Nortel: Concorde

Although the Concorde switch has been shipping for about two years, it has gained minimal acceptance. It is believed that the Concorde has been installed in less than six accounts worldwide.

Product Description

Nortel claims that the Concorde will provide bandwidth from 10 to 80 Gb/s. To obtain this scaling, the architecture of the Concorde requires a multishelf configuration: a centralized core switch fabric shelf is designed to connect separate 10 Gb/s peripheral shelves. However, it is believed that the current release of the Concorde can only scale to 40 Gb/s (four peripheral shelves). Connections between the core switch fabric and the peripheral shelves are believed to be via proprietary OC-12-type links.

Each 10 Gb/s peripheral shelf provides 16 line card slots. It is unclear if all of these slots can be populated with I/Os or if some of them are dedicated to processor and/or trunking cards. Available I/O cards include 2-port DS3, 4-port OC-3 and 1-port OC-12. Nortel is promising that OC-48 will be supported in the future.

Little is known in regards to performance and feature support of the Concorde. Although Nortel claims that the Concorde supports all ATM service classes, it is believed that it is still restricted to providing only CBR support. Nortel promised that SVC support would be available in early 1997, however, it is believed that availability of SVCs has slipped into 1998. The future SVC performance is unknown but believed to be limited to less than 500 calls/sec per Concorde.

Lack of Integration

The Concorde lacks key integration with Nortel's other WAN switches: specifically the Vector and Passport switches. Development and marketing of these three switches fall under either the "Magellan" or "Concorde" teams.

The Concorde switch is developed and marketed under the Concorde family name. The Vector switch (5 Gb/s ATM switch OEMed from Fore) is marketed under the Magellan family name, but the product's development is actually driven by the Concorde marketing and development team. The Passport is developed and marketed under the Magellan family name.

The attempts to integrate the Vector with Passport have been halted by the marketing requirements to integrate the product with the Concorde. Both Concorde and Vector are SNMP managed devices, whereas the Passport is still struggling to become SNMP capable. Passport is managed under the Magellan NMS which was originally developed to manage X.25 packet networks (read: DPN heritage). A customer who has an existing passport network will always require the addition of a new network management system to manage the Vector or Concorde due to this incompatibility.

Lack of End-to-End QoS

Nortel has also harmonized the traffic prioritization and signaling systems between Vector and Concorde, but not between the Concorde/Vector and the Passport. The Vector implements the Fore *Spans SVC code*, and this has been ported onto the Concorde. However, the Passport uses the *Trillium SVC code*. Furthermore, the traffic prioritization system that operates on the Passport has a different number of prioritization levels (queues for QoS buffering) from the Vector and Concorde , and the necessary intelligent routing system/CAC procedures do not interoperate between the Concorde/Vector and the Passport.

Fore Systems

Fore has been rumored to be working on at least three different "next-generation" ATM WAN switches. However, due to their recent financial troubles, some development efforts have either been canceled or delayed. Therefore the availability of any one of their new switches will not be available until well into 1998.

"SuperGarfield": 20 Gb/s switch

Fore has been working on delivering a 20 Gb/s version of the existing ASX1000. Code named "SuperGarfield" (and also called ASX2000), development of this switch has experienced many problems, and may actually be canceled in favor of delivering a 40 Gb/s switch (see "LeapFrog" description below). It is rumored that SuperGarfield will be announced in Q497, but FCS will not occur until 1998.

It is believed that this 20 Gb/s switch will look very similar to the ASX1000, and may actually be delivered in the same form factor, with the same dimensions and interface slots. However, the larger switch fabric will support higher speed and higher density interfaces.

“LeapFrog”: 40 Gb/s switch

Fore is also working on a 40 Gb/s switch code named "Leapfrog." It is to be announced shortly after SuperGarfield is announced (assuming that SuperGarfield is not cancelled). Expected delivery of this 40 Gb/s switch is late 1998.

Although the 20 Gb/s SuperGarfield is expected to look like the existing ASX1000, the 40 Gb/s Leapfrog will have a different look. Fore will most likely develop the LeapFrog so that it will support existing I/O cards designed for the ASX1000. However, LeapFrog is expected to be significantly larger than the ASX1000, allowing it to support more interfaces as well as higher speed and higher density interfaces.

“T2”: 80 Gb/s switch

After reportedly working over 18 months on an 80 Gb/s ATM switch (code named “T2”), it is believed Fore has canceled development on this machine. Fore was simply unable to bring this development to completion. It is believed that Fore convinced the defense agency that funded the project with $6 million that they do not need a “T2” type switch, and instead sold them more ASX1000 switches.

Lucent

Lucent is planning to release their “second generation” GlobeView 2000 ATM switch in mid-1998. This switch is believed to provide bandwidth scaling from 20 Gb/s to 160 Gb/s. The initial focus will be on high-density OC-3 access, with OC-12 trunking added at a later date. No information on pricing, specific hardware components or software features are available at this time.

Comments

Additional competitive information about the subject matter in this document will be disseminated as it is obtained. For comments, questions or additional input, please contact:

Mark Gibson
Ascend Communications, Inc.
1 Robbins Road
Westford, MA 01886 USA

Phone: (508) 952-1210
Fax: (508) 392-1484
E-mail: mgibson@ascend.com-

Worldwide and North American Headquarters
One Ascend Plaza
1701 Harbor Bay Parkway
Alameda, CA 94502, United States
Tel: 510.769.6001
Fax: 510.747.2300
North American E-mail:
info@ascend.com
International E-mail:
air-info@ascend.com
Toll Free: 800.621.9578
Fax Server: 415.688.4343
Web Site: http://www.ascend.com

European Headquarters
Rosemount House
Rosemount Avenue
West Byfleet
Surrey KT14 6NP, United Kingdom
Tel: +44 (0) 1932.350.115
Fax: +44 (0) 1932.350.199

Japan Headquarters
Level 19 Shinjuku Daiichi-Seimei Bldg.
2-7-1 Nishi-Shinjuku
Shinjuku-ku, Tokyo 163-07, Japan
Tel: +81.3.5325.7397
Fax: +81.3.5325.7399
Web Site: http://www.ascend.co.jp

Asia-Pacific Headquarters
Suite 1908, Bank of America Tower
12 Harcourt Road
Hong Kong
Tel: +852.2844.7600
Fax: +852.2810.0298

Latin, South America and the Caribbean Headquarters
One Ascend Plaza
1701 Harbor Bay Parkway
Alameda, CA 94502, United States
Tel: 510.769.6001
Fax: 510.747.2669

15-13-CA
09/97