Ascend

COMPETITIVE ANALYSIS

Ascend GRF 400 vs. IP Switch ATM 1600

Summary

Key selling points against Ipsilon

- The Ipsilon product is not designed for Internet-level IP traffic environments
- The ATM 1600 uses Ipsilon proprietary software
- The switch fabric used is ATM only, so any workstation or host must have an ATM interface and run proprietary Ipsilon software to be connected
- Ipsilon requires end-to-end Ipsilon hardware and/or proprietary software running on workstations, hosts and gateways
- Ipsilon's GSMP and IFMP software is nonstandard and incompatible with other standard ATM Forum specifications
- Connecting standards-based media like Ethernet or FDDI requires Ipsilon's PC-based IP gateways
- 5.3 million packet per second throughput is actually 5.3 million ATM cells/sec. (cell is not a packet); the estimated throughput is a theoretical number based on the aggregate bandwidth of the ATM switch
- 80% of data traffic must be flows to achieve high level of performance
- $3,000 price per port only includes a direct connection to the ATM 1600 from an ATM workstation; the IP gateways required to connect the the ATM 1600 to standard media networks is not included in price per port
- Conventional router problems are not solved by Ipsilon and the problems actually increase by relying on a PC-based platform where route processing and cache-based route look-up are centralized and shared

The following table is a brief point-counterpoint of the Ipsilon architecture versus the GRF 400 architecture. A description of most of the major comparisons follow the table.

ASCEND

Ipsilon ATM 1600 versus Ascend GRF

	Ipsilon ATM 1600 Architecture	Ipsilon Limitations	Ascend GRF 400 Architecture	Ascend GRF Benefits
Bus/Switch	Nonblocking switch	• ATM only switch • Inefficient with packets not aligned with cells • IP gateways needed for non-ATM links	**Crosspoint Switch** • Nonblocking, 1 Gb/s dedicated path for each IP Forwarding Media Card	**Crosspoint Switch** • Contention eliminated • Full aggregate bandwidth achievable
Aggregate Bandwidth	• 2.5 Gb/s aggregate bandwidth	• Bandwidth is ATM cell based • Cell tax is 15% and is wasted overhead	• 4 Gb/s aggregate bandwidth • Each media card gets dedicated 1 Gb/s	• Full 4 Gb/s aggregate bandwidth achievable • Performance scales linearly with additional media cards • Full slot capacity usable • Best price/performance ratio
Packet Throughput Rate	• 5.3M pps, is actually 5.3M ATM cells	• Performance linked to high percent of flows vs. non-flows	• 2.8M pps • Linear scaling	• 2.8M pps achievable
Route Processing and Packet Forwarding Engines Media Cards	• Primarily media-specific receiver/transmitters • Rely on shared CPU(s) and memory for route processing	**Media Cards** • ATM only • Media cards do not route, but rely on centralized PC processor	**Media Cards** • Independent, full-functioning router on each card • All Layer-3 decisions for packets made independently on each card • Support full range of high-speed media • Do not rely on any shared resources	**Media Cards** • Eliminate performance limitation of dependence on shared resources • On-board CPU(s), memory and software allow each media card to make Layer-3 decisions • Customer free to choose media type • Adding new cards to chassis linearly scales performance
Route Table Design	**Route Table Design** • Cache designed for optimal “fast path” plus shared full route table as slow path • Cache compensates for slow path poor performance • Route look-up process: Cache lookup first; if no cache hit, use software-based “slow path” • Single, full route table shared by all media cards used; main CPU manages lookup	**Route Table Design** • Achieved performance gain of cache is minimal, if any, in high-traffic IP environments • Cache misses delay packet forwarding until destination address is converted to next hop by “slow path” route table lookup	**Route Table Design** • Cache and its problems eliminated • Full, 150K entry route table on each media card • Hardware-assisted full route table lookup for each packet • 1 microsecond lookup time for next hop; under 2 microseconds for tables as large as 150K routes	**Route Table Design** • One stop lookup for next hop adds no delays • Layer-3 decisions at switching speed • Route table capacity (150K entries) handles future address growth

Ipsilon ATM 1600 versus Ascend GRF (cont'd)

	Ipsilon ATM 1600 Architecture	Ipsilon Limitations	Ascend GRF 400 Architecture	Ascend GRF Benefits
Flow Schemes	**Flow Schemes** • Flows used to improve IP forwarding performance over that of a PC-based router	**Flow Schemes** • 80% of data traffic must be flows to achieve high performance • Flow determination done by same software and hardware used for routing store/forward decisions • VCs to handle flows set up for 60 seconds, regardless of flow duration (flow length not determinable from packet) • If flows are less than 80% of all traffic, then the slow PC-based routing functions dominate and overall performance degrades dramatically	**No Flow Schemes Needed** Unnecessary with hardware lookup and low switch latency	**No Flow Schemes Needed** No unnecessary overhead for little performance gain
Media Types; Port Density Scalability	• Single port ATM OC-3c only • 15 ports available	• ATM OC-3c only supported, no OC-12c or IP over SONET • No direct connect for other media; Ipsilon IP gateways required • $3K cost per port does not include IP gateways	All media can be directly connected • 2-port ATM OC-3c • 4-port FDDI • 2-port HSSI • 8-port 10/100Base-T	• Two- to four-times port count advantage • Low cost per port • ATM OC-12c and IP over SONET being added for future performance and scalability

IP switching

The need for IP switching came out of the inadequacies of conventional shared-bus architecture routers. An IP switch will be defined as a hybrid networking device with the intelligence and the ability to handle multiple IP routing and management protocols using a switching matrix to provide transport between media interfaces. IP switching takes the intelligence of an IP router and performance of a switch in a single networking device. IP switching uses the IP address of the destination to determine the best path a packet should take through the switch to get to its next hop. What separates the two main implementations of IP switching today are compliance with existing standards. The GRF, unlike the Ipsilon IP switch, is based on open standards. All high-speed media types are supported, so the customer has freedom of choice. With Ipsilon only ATM is supported.

Ipsilon's IP switching architecture and issues

The Ipsilon software architecture is designed to simplify the complexities of the ATM Forum protocols for running IP over ATM. It also was designed to reduce the number of lines of software code necessary for ATM workstations and route servers to run. Simplifying ATM and improving the performance of IP over ATM infrastructure are the main Ipsilon marketing messages.

Taking on the ATM forum

The ATM Forum has produced, to date, two competing protocols for IP over ATM: LAN Emulation (LANE) and Multi-protocol Over ATM (MPOA).

LANE is a bridging protocol that works at the MAC Layer-3 only. Essentially, LANE uses software and workstation-based servers to emulate Ethernet or token ring protocols over ATM. The emulation is done by assigning every ATM-based workstation in a network a unique 48-bit MAC address identifier.

MPOA works on the Layer-3 protocol and adapts Layer-3 functionality to ATM. For example, MPOA handles ARP calls using a route server to convert the IP address to an ATM E.164 address for SVC call setup or VPI/VCI in a PVC environment. MPOA, like LANE, is designed to handle multiple protocol environments (e.g., IPX, DecNet, Appletalk, and IP) and is not optimized for IP or IP based traffic.

Both LANE and MPOA are complex, and not easily managed. They require what Ipsilon describes as an excessive number of lines of software code to be run on ATM workstations and route servers.

Ipsilon realized the weaknesses of the ATM Forum solutions and created its own competing software solutions to take advantage of the market's indecision on implementing MPOA and LANE. Ipsilon has been successful in this strategy, with several major vendors implementing Ipsilon software on their switches. Two new protocols were created and defined in informational RFCs by Ipsilon: the General Switch Management Protocol (GSMP) and the Flow Management Protocol Specification (FMPS).

The GSMP software allows an ATM switch to handle the Ipsilon proprietary call setup, tear down and status of call state functions. GSMP is the interface between the FMPS protocol and the ATM switch hardware. The GSMP currently only runs on some ATM switches.

The second protocol is FMPS. According to Ipsilon, a flow is an extended IP conversation. More specifically, a flow is a sequence of IP packets sent from the same source to the same destination sharing the same protocol type (such as UDP or TCP), type of service and other characteristics as determined by information in the packet header.

The idea behind FMPS is to examine traffic patterns coming through the Ipsilon switch. When a pre-defined set of criteria is met, FMPS sets up a Virtual Channel (VC) inside the ATM switch to handle the remainder of data flow that matches the criteria. Criteria for setting up a flow is very subjective and adds significant overhead and processing time to the shared central processor in Ipsilon's PC-based route manager.

The FMPS has no way of knowing exactly how long a flow will be, so it holds the VC open for 60 seconds, regardless of flow length. In most cases, VCs may be setup for short flows whose data traffic concludes well short of the 60 second timeout. If you believe Ipsilon's packet throughput claims of 5.3 million pps, all VCs by default are setup to handle 318 million packets — a very significant flow.

The question becomes whether the performance benefits of FMPS can be achieved when they are offset by the overhead necessary to detect a flow, to setup and teardown a VC, and to reserve switch bandwidth for longer than may be necessary. The flow setup, when performed on short messages, actually increases the inefficiency of the total transmission. By Ipsilon's own performance claims, useful performance is not approached unless 80 percent of the data traffic is in the form flows. This is not realistic in an Internet environment.

Ipsilon proprietary software required end-to-end

Ipsilon requires its software to be installed on workstations and hosts connected directly to the ATM 1600 IP Switch. Note that the Ipsilon ATM 1600 has ATM interfaces only.

If the Ipsilon switch needs to be interfaced to any other MAC or Data link protocol (e.g., Ethernet and FDDI), an Ipsilon IP Gateway running Ipsilon software is required. The Ipsilon switch also does not allow ATM Forum software to coexist with its propriety software on the ATM 1600. This prevents the ATM 1600 from interoperating with other ATM-based hardware as well as ATM Cell Relay Service providers.

Ipsilon price/port misleading

Ipsilon emphasizes its $3,000 cost per port for ATM interfaces to the ATM 1600. Consider that the $3,000 does not include IP gateway costs necessary to link non-ATM networks to the ATM 1600. The $3,000 per port refers to the cost per workstation.

Performance is in cells *not* packets

According to its literature, the ATM 1600 can pass 5.3 million IP packets per second through the switch. However, Ipsilon is referring to ATM cells not IP packets. A standard IP packet with its IP, TCP and AAL-5 headers takes up a little over 48 Octets . This would require the entire ATM Payload, allowing no room for user's data in any form. The aggregate throughput of 5.3 million IP Packets would have to be about 10.6 million cells with an efficiency of about 50 percent . The GRF switches on the IP level, so native IP packets transverse the switching matrix, allowing performance of up to 2.8 million IP packets per second.

Problems of shared resources haunt Ipsilon's architecture

Ipsilon uses a centralized PC platform to run switch controller software to identify flows, to setup the switching process, and to update and maintain route cache and full route tables. These are the problem areas IP networks face today with conventional routers. There is no solution here for the major cause of Internet brownouts, packet loss and downtime. Just as the shared resource routers work well in multiprotocol LAN environments, so will the Ipsilon ATM 1600. Neither is designed for the Internet, however.

The only way to alleviate these shared resource problems is to use a distributed architecture like that of the GRF 400. A distributed architecture allows processors to be used more efficiently, freeing each to perform their unique duties and allowing fewer bottlenecks and greater throughput on the switch.

Selling the GRF versus the Ipsilon ATM 1600

Ipsilon is trying to position itself to sell into the large campus market. The following is a list of some concerns network managers should consider:

- **Not standards-based** — Ipsilon is not based on standards. An end user will have trouble connecting the ATM 1600 to outside ATM providers or external ATM switches because Ipsilon does not support the ATM Forum's UNI 3.1 software.
- **Centralized processing and cache-based routing** — Ipsilon will experience the same problems as conventional routers like the Cisco 7000 series. As IP traffic increases and becomes more diverse, the shared CPU and route processing can become overloaded, dropping or delaying packets beyond allowable network and application tolerances.
- **Flows** — Predicting a flow in today's fast changing and high-performance IP networks is difficult. Ipsilon's flow prediction software overhead may actually decrease overall performance rather than increase it. When flows are not available, Ipsilon's PC-based routing functions become a major bottleneck.
- **Additional Ipsilon hardware needed** — In a campus environment where non-ATM media is dominant, Ipsilon IP gateways will be necessary to build an end-to-end solution. A single GRF will replace not only the Ipsilon ATM 1600 but also the Ipsilon IP gateways.
- **Paying ATM cell tax** — An ATM-based architecture automatically forfeits up to 15 percent of the available data traffic bandwidth to ATM overhead (ATM and AAL overhead). New technologies like IP over SONET without ATM, supported by the GRF, recover that lost portion and return it to usable data bandwidth.

Pricing comparison

	Ipsilon			GRF		
Base Chassis Price	• ATM 1600 - Fully configured 15-slot chassis: $45,000			• GRF 400: $15,650 (4 media card slots, plus software)		
Base Software Price	No Charge			No Charge		
Media Support	Price	Ports	Price/Port*	Price	Ports	Price/Port*
Ethernet - 10 mbps	Not available	Not available	Not available	See Ethernet 10/100Base-T entry		
Ethernet - 100 mbps	Not available	Not available	Not available	See Ethernet 10/100Base-T entry		
Ethernet 10/100Base-T	Not available	Not available	Not available	$20,000	8 ports	$2,500
				14,000	4 ports	3,500
HSSI	Not available	Not available	Not available	17,500	2 ports	8,750
FDDI	Not available	Not available	Not available	19,000	4 ports	4,750
ATM OC-3c	$3,000	1 port	$3,000	20,000	2 ports	10,000
ATM OC-12c	Not available	Not available	Not available	25,000	1 port	25,000

[1] This is assuming a IP header of 20 Octets, a TCP header of 20 Octets and AAL-5 overhead of 8 Octets.
[2] The low efficiency of IP over ATM when using small packets from the padding that is required to fill the remainder of the second packet.

Feature comparison

Comparison	GRF 400	Ipsilon ATM 1600	Comparison	GRF 400	Ipsilon ATM 1600
Chassis			**Routing Protocols**		
Dimensions	19x5.25x19	17.50x3.5x18	OSPF	✔	✔
Backplane	Switched	Switch	MOSPF	✔	✔
Aggregate Bandwidth of Backplane	4 Gb/s	2.5 Gb/s	IS-IS	✔	✔
Dedicated Bandwidth per Card	1 Gb/s	155 mbps	RIP - v1/v2	✔	✔
Number of Slots	4 Slots	16 Slots	Hello	✔	✔
Processing Architecture	Distributed	Centralized	Router Discovery	✔	✔
OS Kernel	Ascend Embedded OS	GSMP	EGP	✔	✔
System Processor	Pentium 166 MHz	Intel Pentium on PC	BGP3/4	✔	✔
System RAM	64 MB	INA[1]	BGP4 - Communities	✔	No
Performance			MTU Configurable	✔	✔
Latency Through Device	20 Microseconds	<100 Microseconds	ICMP	✔	✔
Full Route Table Updates per Second	50	INA[1]	Multicast	✔	✔
Route Table Size	150K	INA[1]	DVMRP	✔	✔
Aggregate Performance	2.8M pps	5.3M Cells/Sec	CIDR	4	4
High Availability			**Protocol Support**		
Hot-swappable Power Supplies	✔	INA[1]	Frame Relay UNI	✔	No
Redundant Load-sharing Power Supplies	✔	INA[1]	Frame Relay NNI	✔	No
Hot-swappable Media Cards	✔	INA[1]	PPP	✔	No
Hot-swappable Fan Trays	✔	INA[1]	HDLC	✔	No
			ATM UNI 3.0/3.1	✔	Physical Layer
Media Cards			ATM PVC/SVC	✔	✔
Route Table Type	FRTLU	Cache	AAL 5	✔	✔
Max Route Table Size	150K	INA[1]	Classical IP over ATM - RFC 1577	✔	No
Type of Processors	Fujitsu Sparc 40 - 60 MHz	INA[1]	Traffic Shaping	✔	No
# of Processors per Card	1 to 2	N/A	Frame Relay over SONET	4th Qtr96	No
Active VCs ATM	512 per port	32,000	PPP over SONET	4th Qtr96	No
Active DLCI Frame Relay	932 per port	N/A	Management System		
Aggregate Throughput	10K pps-600K pps	5.3M Cells/Sec	Interface	Command Line	Web-based
Security			SNMP v1/v2	✔	✔
Basic Packet Filtering	✔	✔	SNMP - Read/Write	✔	✔
Packet Header Logging	✔	N/A	Remote Diagnostics	✔	✔
Radius Support - Admin. Authentication	✔	N/A	Remote Software Loading/System Reboot	✔	INA[1]

[1] Information not available.

ASCEND

Remote Networking Solutions That Work.™

Ascend Communications, Inc.
1275 Harbor Bay Parkway
Alameda, CA 94502, USA
TEL: 510.769.6001
FAX: 510.814.2300

E-mail: info@ascend.com
Toll Free: 800.621.9578
FAX Server: 415.688.4343
Internet Home Page:
http://www.ascend.com