Ascend

COMPETITIVE ANALYSIS

Selling Against the Ipsilon Story

The chart below depicts a typical Ipsilon configuration. The Ipsilon solution always includes an ATM switch, a PC router and special software that must reside in the existing "upstream" and "downstream" network nodes.

There are several problems that prevent the Ipsilon solution from becoming a viable choice of carriers and ISPs, which are marked by numbers in the diagram. Each area is discussed in the text below.

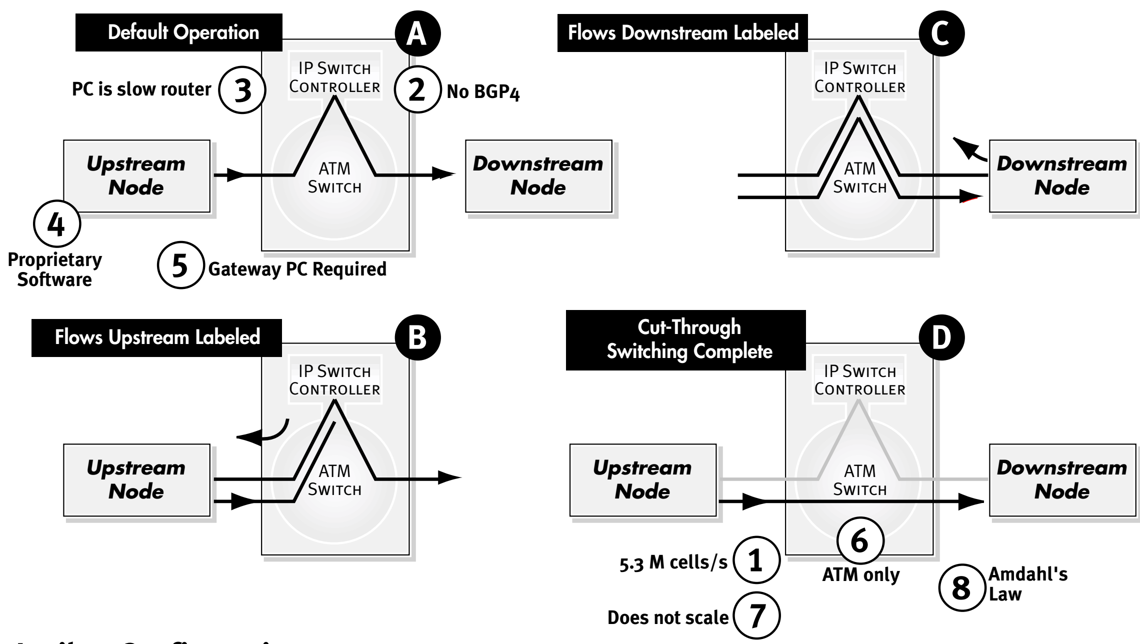

Ipsilon Configuration

(1) A Cell is Not an IP Packet

Ipsilon Claim: "Based on scalable, robust and proven technologies, the IP Switch ATM1600 provides up to 5.3 million IP packets per second (pps) throughput while maintaining full compatibility with existing IP networks, applications and network management tools."

Ascend Response: How does Ipsilon arrive at 5.3 Mpps?

- Ipsilon is using a 16-port Fore Systems ATM switch. Each port uses ATM running over SONET OC-3c at a line speed of 155 Mbps.
- Subtract the SONET framing overhead (about 9%): 155 - (0.09 * 155) = 141 Mbps line speed per line
- Multiply by 16 for total capacity: 16 lines * 141 Mb/s = 2257 Mbps total capacity.
- Convert total capacity to bytes: 2257 Mbps * 0.125 Bytes/bit = 282 MBps.
- Convert total capacity in bytes to ATM cells. (An ATM cell is 53 bytes long): 282 MBps * (1/53 cells/bytes) = 5.3 Mcells/s.

Note that Ipsilon's literature advertises 5.3 million PACKETS per SECOND.

An IP packet is not equivalent to an ATM cell! Even a tiny telnet packet is bigger than an ATM cell. Fact is, most traffic on the Internet is on the order of 5 to 10 cells per IP packet. At a minimum, you would have to reduce Ipsilon's claim by a factor of 5 to 10.

- 5.3 Mcells/s * 0.1 packets/cell=530 Kpps

This results in a rather mediocre (for all the headaches associated with a nonstandard ATM-based solution) total system performance of 530 Kpps for a 16-slot solution. A 4-slot solution would then claim about 133 Kpps for the entire system. This is about equal to the performance of a single GRF OC-3c ATM media card, or about 1/4 of the GRF 4-slot system solution.

② PC Routers are Slow and Vulnerable to Overload

Ipsilon Claim: "An Intel Pentium® microchip provides the processing power for the IP Switch Controller."

Ascend Response: Ipsilon says very little about their choice of the engine used to forward IP packets not destined for the ATM switch as a part of a flow. They, in fact, have chosen a PC to perform this task.

The problems with using a PC as a router are well known. As an inexpensive protocol research engine, the PC is ideal. But as a robust, carrier and ISP class scalable router, the PC is a bad choice for several reasons:

- Forwarding decisions are centralized and vulnerable to CPU saturation and lost packets.
- The route table is cached and vulnerable to reduced performance when IP addresses are not found in the cache.
- The route table lookup is software based; this is a slow process, especially with large route tables used in today's Internet.
- The PC's bus is shared, presenting a bottleneck to packet forwarding, and is not scalable.

③ BGP4 is Not Supported

Ipsilon Claim: "The IP Switch ATM1600 is compatible with standard Internet routing protocols, including the Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Classless Inter-Domain Routing (CIDR). It supports IP hardware multicast without any modification to standard IP multicast protocols, such as Distance Vector Multicast Protocol (DVMRP) and Internet Group Membership Protocol (IGMP)."

Ascend Response: Where is BGP4? Without BGP4 and its enhancements, the Ipsilon solution is relegated to the LAN! Without BGP4, Ipsilon is unable to address the NAP, Peering Center, Backbone or POP. The GRF supports BGP4 and all its enhancements, and it is specifically designed for NAPs, Peering Centers, Backbones, POPs and MegaPOPs.™

④ Ipsilon Requires Special Proprietary Software in Your Existing Network Devices

Ipsilon Claim *(Informational RFC 1953 - Ipsilon Flow Management Protocol Specification for IPv4 - May 1996):*

"This memo documents a private protocol for IPv4-based flows. This protocol is NOT the product of an IETF working group nor is it a standards track document. It has not necessarily benefited from the widespread and in depth community review that standards track documents receive."

RFC Abstract

"The Ipsilon Flow Management Protocol (IFMP) is a protocol for allowing a node to instruct an adjacent node to attach a Layer-2 label to a specified IP flow. The label allows more efficient access to cached routing information for that flow. The label can also enable a node to switch further packets belonging to the specified flow at Layer-2 rather than forwarding them at Layer-3."

Ascend Position: As is clearly noted in the above quoted RFC introductory text, the document is an RFC in name only and does not carry any weight in the IP community. It only appears to be official. The fact is that Ipsilon's solution requires special proprietary software to be placed in the ATM nodes attached to their ATM switch.

No special proprietary software is needed with the GRF. Notice also that the RFC introduction specifically mentions that no Layer-3 switching occurs. The GRF is based entirely on Layer-3 switching. **No other product can make that claim.**

Companion RFCs 1954 and 1987 are similar in nature and fail in the same way.

5 If You Don't Have Ipsilon's Special Software, Then a Special Gateway is Required

Ipsilon Claim: "The IP Switch Gateway is a Pentium-based hardware platform running the Ipsilon system software that converts incoming IP-based Ethernet, Fast Ethernet or FDDI packets into ATM cells and forwards the traffic to the IP Switch network. An ATM interface connects the IP Switch Gateway to the IP Switch ATM1600, which uses ATM hardware acceleration to achieve through-puts of up to 5.3 million packets per second. Multiple LANs can gain routed access to the high speed IP Switch mesh through the IP Switch Gateway."

"Hosts, routers and other edge devices that support Ipsilon's Flow Management Protocol (IFMP) can attach directly to IP Switches without the assistance of an IP Switch Gateway."

Ascend Response: Yet another PC-based platform in your network! Guaranteed to be a bottleneck. Conversion of FDDI or Ethernet-based frames to ATM is slow at best, and this is before they are routed or become a part of a flow sent across an ATM switch.

Why not just route IP packets directly at high-speed over a media-independent switch such as the GRF? The GRF stands alone, integrating routing and switching in one solution and one chassis.

6 Ipsilon Only Offers ATM, So Customers Can't Choose Other Media

Ipsilon Claim: ATM is all one needs.

Ascend Response: Ipsilon contradicts itself by providing gateways to support FDDI and Ethernet in its product offering. Customers need a choice of media. The Ascend GRF™ 400 offers customers a choice of HSSI, HIPPI, Ethernet, FDDI, IP over SONET and ATM. Since Ipsilon offers only ATM, customers are given no choice. The Ipsilon gateway converts Ethernet or FDDI to ATM. With the GRF 400, customers are given a choice and also have the added flexibility of mixing any of those choices in a single chassis.

7 No Realistic Scalability

Ipsilon Claim: "Scalable platform architecture."

Ascend Response: The Ipsilon solution is not scalable. Its solution requires everything to be connected to an ATM switch. This is not realistic for today's networks. The only scalability offered by Ipsilon already exists in the ATM switch itself, which has 16 slots. A fully configured ATM switch would completely overwhelm the PC router. The lack of BGP4 support limits any scaling to the LAN environment.

8 Amdahl's Law: The Truth About the Impact of Flows and Cut-through Routing on Overall Performance

Ipsilon Claim: Ipsilon claims that it has done studies of Internet traffic patterns and that 80 percent of the traffic is "flowable."

Ascend Response: Even if this is true, the effect on performance is negligible. Consider that 80 percent of the time you can use cut-through routing but 20 percent of time you must rely on conventional routing using a PC router.

Any time two processes are combined to complete a single task, the total time to complete the task can be used to determine an "average" rate of completion. This is especially important in competitive situations against Ispilon since its solution combines two processes, one fast and one slow. The lowest common denominator—the slow part of the process—drags down performance of its entire solution. In contrast, the Ascend GRF uses only one very fast switching process, by combining routing and switching into a single device.

For example, say you have planned a 1250 mile trip. You plan to travel 80 percent of the distance by airplane at 500 MPH and 20 percent of the distance by auto going 50 MPH. The airplane portion takes two hours, the auto portion takes five hours. Since the total time for the trip is seven hours, the "average" speed is 178.57 MPH (1250 miles divided by seven hours). That 179 rate represents 35.7 percent of speed possible when traveling by airplane.

This is an exact description of a weighted harmonic mean. When the two processes differ in speed, as in the Ipsilon solution, the formula used to determine this "average" is the weighted harmonic mean.* (For those interested, the formula is 1/[(.2/50)+(.8/500)].)

* Gene Amdahl of Amdahl Corporation fame used this same analysis to counter the media blitz in the mid-1980's from those who maintained that parallel processors using inexpensive PCs would soon outperform vector processing supercomputers. Ever since then this particular application of the weighted harmonic mean has been known as "Amdahl's Law."

The application to the Ipsilon solution is simple and direct. The packet per second rate during cut-through routing (using an SVC through the ATM switch) is equivalent to the speed of the airplane. The packet per second rate of the PC router is equivalent to the speed of the auto. At a ratio of fast to slow of 10 to 1, the effective rate for the system is about 36 percent of the high-speed path. A rather mediocre showing.

Let's put this into practical terms. Ipsilon is currently saying that the cut-through rate is 2 Mpps and the PC rate is 90 kpps. The ratio of fast to slow is 22.22 to 1. Using this ratio, the overall effective rate drops to an "average" 381 Kpps for the entire system of 16 slots. That's about 24 Kpps per slot. A single GRF400 OC-3c ATM media card will deliver 130 Kpps, five times the speed of Ipsilon's product.

The chart below summarizes this situation graphically for several different ratios. The Y-axis shows the percentage of the fast process achieved by the overall system for ratios of fast path to slow path as shown by the different lines plotted. The X-axis reflects the relative proportion of time spent in the fast and slow processes.

The main point is that the total time to complete the task is dominated by the slow process. Ipsilon's solution is therefore dominated by a PC router.

Since the GRF 400has no slow path it performs at 100% of its advertised rate at all times. The GRF is not subject to Amdahl's Law.

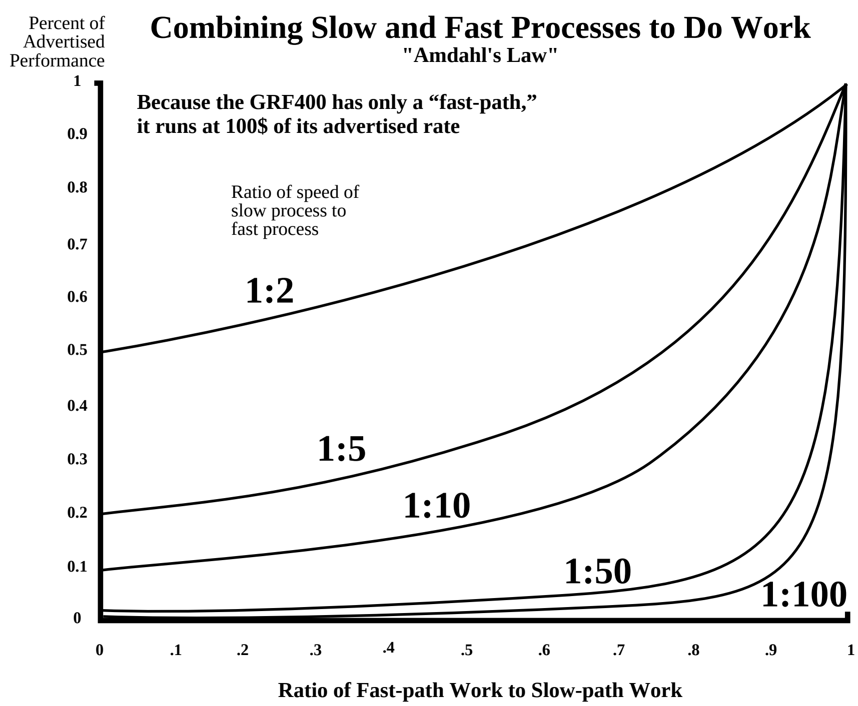

ASCEND
Remote Networking Solutions That Work.™

Ascend Communications, Inc.
One Ascend Plaza
1701 Harbor Bay Parkway
Alameda, CA 94502, USA
TEL: 510.769.6001
FAX: 510.747.2300

E-mail: info@ascend.com
Toll Free: 800.621.9578
FAX Server: 415.688.4343
Web site:
http://www.ascend.com