IP@ATM™

IP Networking in the WAN with the 1100 HSS®

ALCATEL

Introduction

Alcatel Data Networks is implementing IP networking capabilities in the 1100 HSS® multiservice switching platform. We have named our strategy for IP networking **IP@ATM™**. Having a single, robust WAN platform that will switch ATM, Frame Relay, and IP with network and service interworking, coupled with a comprehensive network management strategy, is the best solution for building scalable, high performance networks. Alcatel Data Networks is strongly committed to providing the best multiservice networking solutions to its customers that now will include integrated IP switching.

This paper addresses the issues and problems of IP networking in relationship to the Internet and intranets. It focuses on the key trends for implementing IP switching, debates the controversial issues, and tells how we will implement IP switching capability in the 1100 HSS.

The Issues and Problems in Today's IP Networks

The phenomenal growth of the Internet has created a need for a much larger and more robust IP networking infrastructure. The issues and problems plaguing the Internet are also now being seen in the evolving intranets that are growing larger and taking better advantage of Internet-related technologies and multimedia applications. The current state of the internet protocol and classical IP routers cannot cope with the burdening networking demands of **Scalability, Performance, Multiservice, and Quality of Service networking** that are very visible in networking today. While the internet protocol itself is evolving to try and address some of the issues, there is a realization that current day IP routing and forwarding techniques alone cannot solve the problems. In addition, it is also a fact that simply providing larger pipes (as is the case for gigabit routers) to carry more information is not the long-term answer to the problems.

In the data communications industry today, there are many IP networking issues being debated and, as a result, it can be very confusing. It is very easy to become absorbed in the vendor debates over which technology best solves the problems. The most common debate in the media and industry forums is the role of IP versus ATM as "the networking technology of choice," for example. Comparing IP with ATM is like comparing apples and oranges. And, the real debate is not about IP versus ATM at all nor routers versus switches. At the real heart of the issues are fundamental networking principles that are being revisited again to address the problems. The very basic principles being debated are really connectionless versus connection-oriented networking and flow-based versus topology-based connection establishment paradigms.

Connectionless Versus Connection-Oriented Networking

Connectionless networking is the basic networking principle used in most LAN technologies and in IP routers. This allows information to be sent into the network without knowing exactly where it will go or how (or when) it will get there. In short, there is no pre-defined connection path identified between a source and its destination that will be used repetitively to transport information. In the LAN, e.g., Ethernet, it is the concept of broadcast media (even in switched LANs). In current day IP networking, it is hop-by-hop routing and forwarding. Essentially both LANs and classical IP routers treat the network as common media and cannot effectively control how the overall network resources will be used. A startling example of how this principle is applied in reality is the fact that the only technique available for congestion control in IP connection-less networks is "packet discard." It is a well-known fact that IP routers have no particular form of traffic management nor congestion control other than very large buffers. When the buffers overflow data is discarded. Such techniques are the basis for some of the major problems in the Internet and intranets today.

Connection-oriented networking, on the other hand, offers the benefit of a dedicated connection, either real or virtual, between two endpoints in a network. The advantages are significant! A connection means that network resource utilization has been examined and used as a criteria for setting up connections to begin with. Network bandwidth allocation becomes an operator-controllable resource. Over-allocation of bandwidth becomes a conscious engineered network decision, not an arbitrary, non-controllable side effect. This is what bandwidth management is all about. In addition once connections are established, other forms of network controls can be effected on the connections in the form of real-time traffic management e.g., policing and traffic shaping, network-wide resource prioritization for Class of Service (CoS) provisioning and deterministic delay and delay variation measurement and control capability that can be used to set up multiple levels of Quality of Service (QoS). One very significant benefit of connection-oriented networking is the capability to perform real time congestion avoidance and control (simply put, intelligent flow control) to prevent traffic loss in the network. This can only be performed in networks where there is a fundamental concept of a known connection path so that path resource utilization information can be fed back to the network source of each connection where it can be throttled. It is no surprise to find connection-oriented networks, (X.25, Frame Relay, and ATM) in the heart of major service provider networks. Up until now, IP traffic has traversed these types of networks transparently. That is about to change.

Scaling Issues

The Internet today has already over 100,000,000 users with expectations that it will grow to well over 200,000,000 users by the year 2000. The number of hosts on the Internet has reached enormous proportions currently exceeding 30,000,000, and that is also expected to double by the year 2000. The router was not made to handle this kind of environment. The Internet is taking on the look and feel of the large telecommunications environment and it is expected that as time goes on there will be an equal number of Internet users as there are voice users, which now exceed 1 billion. The router has limited port and processing power and even as it evolves to gigabit speeds, it still cannot attain the right mix of networking capabilities that will be demanded in such a large scale environment. Public networks, in particular, require comprehensive, cost-effective, scalable access and backbone equipment that offers both low and high-speed connectivity with high availability and flexibility to support the growth of the network over many years. The router cannot do this. Switching systems using the fundamental principle of connection-oriented networking offer the better solution in terms of scalability. Routers pale by comparison.

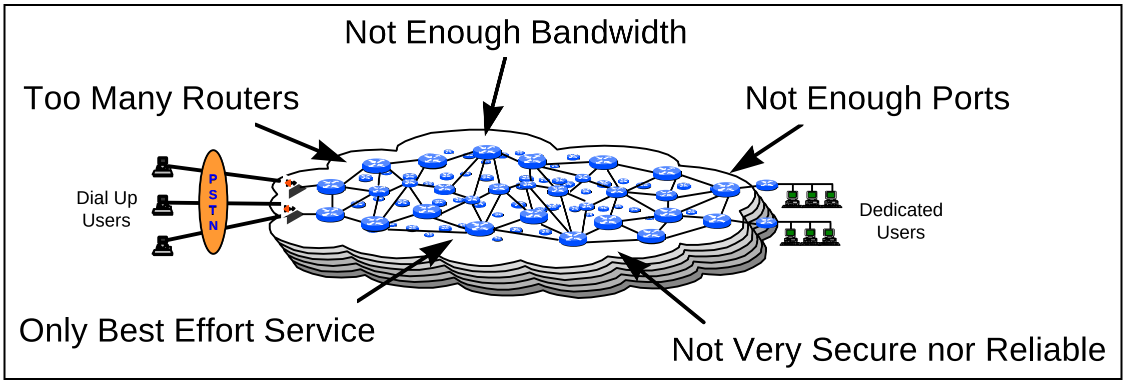

Performance Issues

LAN and WAN Requirements are Different

The primary application of connectionless networking can be found in the LAN. It was in the LAN that connection-less networking began and grew. In the LAN, bandwidth is relatively inexpensive as compared to the WAN. The price paid for retransmission of packets, when it occurs in the LAN, is so small that it is hardly even noticeable and no one worries about it too much unless it gets very large and starts to affect performance. In addition since LANs are typically deployed in campus environments, distances between network elements are short and hence the delay characteristics in the LAN are very low. The overall impact

of delay in connectionless LAN networking is negligible. No one talks about congestion management as a principle in LANs.
In the WAN however, it is an entirely different story. WAN bandwidth is the most expensive, recurring, operational resource aspect of a network. WAN links are typically long distance. Typically the price per bit of information transmitted across the WAN is a significant percentage of the overall networking cost base. The delay introduced in the WAN due to longer distance becomes a significant networking factor both from a QoS point of view and a network control point of view. When one considers the fact that if information were chronically lost in the WAN and had to be retransmitted, the cost of transporting information increases dramatically. It could even double. It is this reasoning that forces some important consideration as to how to deal with the subjects of bandwidth, traffic, and congestion management in the WAN. Taking a "reactive" approach to these issues, as is the case with IP routers, leaves no choice but to discard data. When the buffers fill up, the data is discarded. The problem is multiplied in the WAN because the delay aspect of the WAN increases the difference in the amount of time between the filling of the buffers and the movement of data through the network to the destination(s). If discard occurs in other IP routers on the way to the destination, the upstream routers have no way of knowing about the problem until it is too late. The result is more data discard.

The modern trend in connection-oriented networking, as we are implementing in the 1100 HSS with RateMaster™, is "proactive" congestion avoidance. With such a technique, the network connection utilization is constantly being measured and fed back to the ingress points of network connections where rate matching decisions can be made. Rate matching, as the name implies, attempts to match the rate of data submission into the network for each connection to that which the network can successfully pass to the other side of the network. If data discard has to be performed, as a matter of severe congestion or network problems, it is performed at network ingress points and thus protects the overall network from multiplying the problem. Connection-oriented networking in the WAN becomes necessary to introduce networking controls that ensure that once information is allowed into the WAN, it will traverse the WAN and be delivered to the destination side of the WAN and avoid repetitive network wide retransmission of data. The connection-oriented approach is also by far the best way in which network controls (CoS/QoS) can be used to effectively handle large numbers of highly deterministic, real time traffic sources such as voice and video. Additional benefits as described above add feature richness to the WAN and allow for other value-added services to be provisioned.

Some would argue that the additional overhead introduced by connection-oriented networking, in particular ATM, adds more cost to the WAN over that of connection-less native IP or even IP over Frame Relay. While it is true that ATM introduces additional overhead, it is for the reasons previously mentioned that make the difference. ATM offers all of the qualities for connection-oriented networking, and in particular, those associated with high speed bandwidth, traffic and congestion management that cannot be obtained with IP networking as it currently exists. If one considers the potential application and service impact of packet discard and cost of frame retransmission in the WAN versus ATM protocol overhead as a vehicle to have better network control, the choice becomes clearer. ATM with its multiple CoS/QoS and control architecture can actually improve the performance of the WAN.

A Movement Towards the Connection-Oriented Approach with IP Traffic

As it stands today, there is a *de facto* conclusion that IP networking, based solely on a connection-less networking philosophy, cannot scale to meet the overall networking requirements. The proof of this lies in the trends to introduce technologies like Multiprotocol Over ATM (MPOA), TAG switching and MPLS, and a host of other proprietary techniques to allow IP traffic to be transported over Frame Relay and ATM networks. While each proposed method has its place in the networking hierarchy and their own advantages and disadvantages, there is a realization that the connection-oriented approach in the WAN is a necessity.

The result is a marriage of convenience between connectionless and connection-oriented networking where the LAN meets the WAN.

Multiservice Networking

There is an increasing demand today for networking consolidation. The problem has manifested itself in that there has been so much growth and diversity in networking requirements over the past several years that service providers have had to install many different kinds of networks to meet their end user (customer) demands. Installing many different networks has created an enormous (and expensive) burden on network operations. The problem is complicated by the fact that many service providers simply cannot continuously change out their infrastructure to put in the latest networking products. The service providers who are constantly changing their equipment suffer from service discontinuities and inter-operability problems not to mention management problems.

The demand for Internet services has now become the driver by which major carriers see the need to implement large scale data networks. But what about voice? Now data networks have grown to become equally as important as voice. It is expected that voice will become a low bandwidth “add on” to data traffic in a multiservice networking environment. Network service providers can expect to be faced with even more diverse connectivity requirements that encompass not only different types of physical interfaces and speeds, but also multiple protocols and the need for more network and service interworking. Such an environment is depicted in the figure below. It is based upon using ATM as a core switching technology and surrounding that core with multiservice switches that accommodate all major data protocols such as ATM, Frame Relay, and IP as well as provisioning for handling voice and video. The addition of IP in multiservice switching offers complete desktop-to-desktop networking capability with router compatibility and significantly improved scalability and performance across the network.

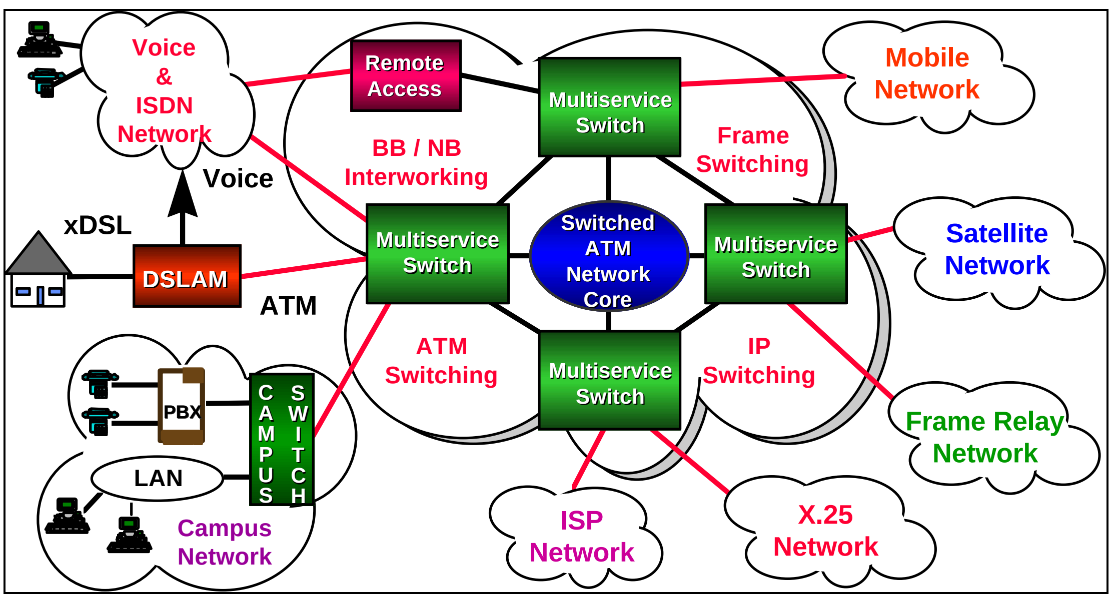

The traditional router has evolved to become a very sophisticated computing device. The fact that IP packet routing and forwarding as a networking technique is computational intensive brings the fundamental limitation of the router, as we currently know them, to the surface. The routing table in the Internet today is approaching 50,000 prefixes. The capability just to keep up with daily changes in the status of these routes is high enough to demand the full time attention of a processor in the router. Further, the performance requirements for packet forwarding in the router have grown enormously to meet the demand for higher speeds, now expected to grow to 622 Mbps and soon 2.4 Gbps. When you multiply these speeds by the number of physical interfaces required in an enterprise backbone networking hub, or even more so, a public network point of presence, the problems become significant. As always, the price one pays for networking equipment also plays an important role in the deployment of networks and, as the router has become more complex, the prices have gone up accordingly, particularly for the high-speed interfaces. The bottom line is that today's router has become a bottleneck in the Internet backbone. The fact is that as the router evolves, it takes on the flavors of switching systems.

The Need to Marry Connectionless and Connection Oriented Networks Using ATM

The key to success for IP switching lies in being capable of providing flexible IP routing with a switched packet forwarding mechanism that offers the highest possible performance. Today IP packet forwarding takes place in a connectionless environment. Essentially packets are sent from network element to network element without any advance knowledge of exactly how to get to the end connection point. Routing in IP networks identifies how to get a packet from one element to another based upon the IP address in each packet header. In the overall scheme of things, each IP router builds its routing tables based upon exchanging routing information with the other routers it connects to, eventually spanning the entire network. Connectionless networks are built primarily upon the concept of best effort service where packets are offered to the network regardless of the status of the intermediate network connections. This increases the potential for packet loss and/or delay and thus makes for a large degree of uncertainty in the ability to offer a guaranteed level of quality of service (QoS). Connectionless IP traffic has no traffic management mechanism to throttle intermediate network traffic and thus bottlenecks can form in transit routers. This bottlenecking is now a well known phenomenon in the Internet today.

If high quality service is to be provided in IP networks, whether it is provided as a higher grade of service to higher paying customers or to support the deployment of voice and/or video services in conjunction with data service, end-to-end QoS support is required. Some suggest that by merely throwing huge amounts of bandwidth into network connections that the problem will become a mute point. On the contrary, given the history of bandwidth utilization, all available bandwidth will ultimately be used up by applications. In short, there is a need for a QoS capability that can be used to prioritize and control the delay characteristics of diverse traffic types. This is where the best benefit of connection-oriented networking can be seen. In connection-oriented networks, such as ATM and Frame Relay, end-to-end connections are made in advance of data flow. The pre-determined data paths not only allow for a higher degree of control of network connection resources but also offer a more deterministic method for responding to changes in network status. ATM, in particular, because of its virtual path and channel capabilities can be used to build hierarchical networks that can scale to extremely large proportions. ATM was designed from its inception to accommodate multiple levels of QoS. As public networks grow, they will not only need these capabilities, but also controlled access capability. Here again, ATM was designed to provide this capability. It is no wonder then that the majority of large network service providers have chosen ATM as their choice for consolidating their networking backbone infrastructures.

TAG/MPLS

ATM's intrinsic capabilities can and will be exploited to produce a great marriage with IP. While the ATM Forum has moved forward and produced a much more slimmed down version of the Multi-Protocol over ATM (MPOA) specification, many believe that a better solution will come using shortcut signaling techniques that allow the IP packet forwarding mechanism the speed and flexibility to circumvent the more complex traditional ATM signaling protocols. The IETF is currently working on a standard shortcut technique that is called Multi-Protocol Label Switching (MPLS). While MPLS is ultimately to be applied to many different link layer protocols, the IETF is concentrating on ATM as the first link layer application. MPLS is expected to be published as a standard in mid-1999. The activity on MPLS is surrounded by the complexities of the IETF standards process, but, the original proposal was based on the "tag switching" (TAG) technique proposed by Cisco Systems. TAG is a technique that is well-suited to fit into the switched environment, and, in particular, ATM. One of the best benefits of TAG switching with ATM is the capability to scale from the Campus to the WAN.

A network of routers, typical of what exists today, requires that each node perform hop-by-hop routing and forwarding techniques, as depicted in the figure below. The result is complicated routing tables that are replicated throughout the network and the scaling and performance problems, as was pointed out earlier.

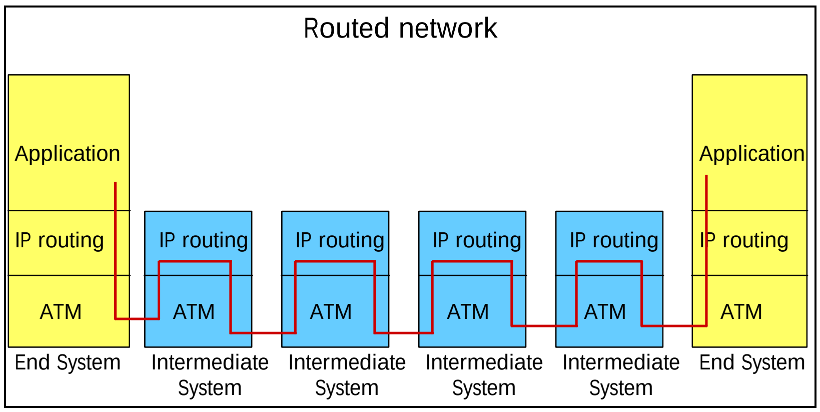

In a network that supports TAG switching, the intermediate switching nodes create a direct connection to the edge of the network which appears to the routers as a single hop. This is depicted in the diagram below. The TAG edge router and the TAG switch router could be the same device. Functionally, the TAG edge router creates the TAG switching connection through the network. It can perform both normal layer 3 forwarding functions and establish TAG connections. The TAG switch router only has to support the minimum protocol requirements to establish and maintain TAG connections. Because the intermediate network nodes are operating as an interior network, interior routing protocols can be used to build the routing tables and standard exterior routing protocols. For example, BGP can be used to interface to the end systems. The IP routing table in the backbone is simplified and the IP packet forwarding performance is boosted significantly.

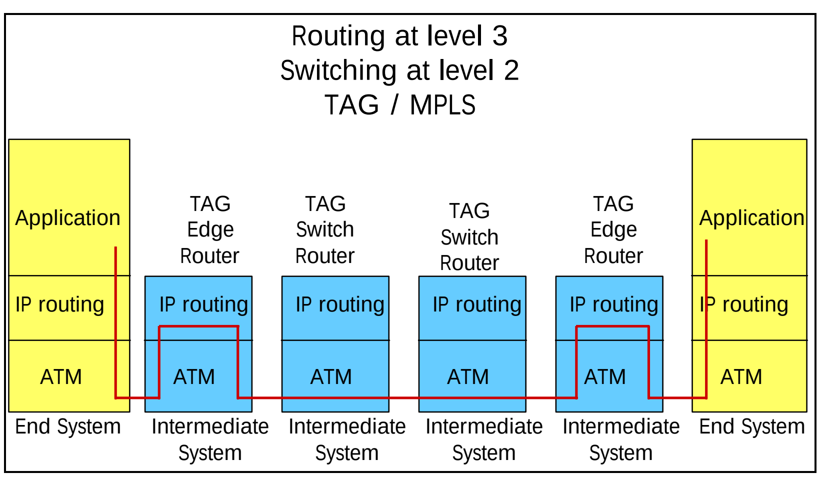

Adding New Functions to the Protocol Stack

In discussing the issues associated with building high performance, scalable IP networks it is clear that the network components themselves must be looked at from two points of view. They must be analyzed both from a network layer (layer 3) and link layer (layer 2) perspective in harmony to achieve the desired results.

The figure below depicts the functions required primarily at layers 2 and 3 in the networking hierarchy. The addition of L3 Characterization (L3C) and Shortcut Signaling (SCS) are key new functions that will reshape how IP traffic is handled. The new shortcut techniques such as TAG/MPLS allow for topology driven connections to be formed using new signaling protocol techniques that are streamlined to fit into the existing IP networking environment.

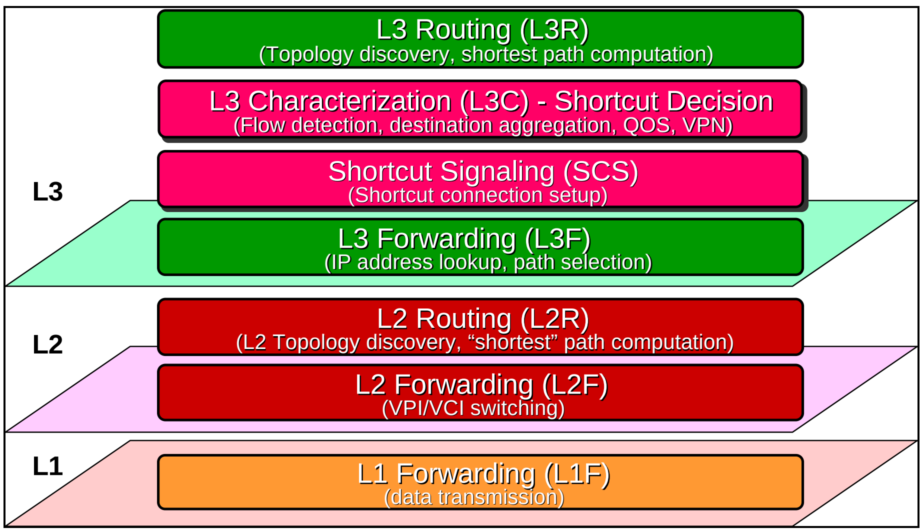

L3C performs an analysis of the type of traffic that is being handled and makes decisions as to how packet forwarding is to be accommodated. In the case of TAG/MPLS, the topology routing structure that was created in the L3R function is used to define a connection. SCS performs the streamlined connection setup functions. For TAG switching, it uses the Tag Distribution Protocol (TDP).

TAG/MPLS Versus MPOA

MPOA is not what it was supposed to be. The ATM Forum technical committee had no choice but to scale back the scope of the MPOA specification because the problems of adapting MPOA into the WAN environment proved to be too difficult. The MPOA specification today is clearly targeted as a LAN/Campus solution. The fact is that MPOA uses LAN Emulation (LANE) as a basis for its capabilities to have IP Routing and Forwarding functions. In basic terms, MPOA analyzes IP data flows and guesses at whether or not an ATM connection is needed as a short cut to reduce the burden of pure IP packet forwarding. If a flow is identified as qualifying for an ATM connection and if none already exists to the destination, MPOA uses ATM SVC signaling to set up an ATM connection to the other side of the network. MPOA uses the Next Hop Resolution Protocol (NHRP) to map ATM network addresses, e.g., E.164, to IP addresses and thus creates a very complex routing and addressing scheme. In effect, there are two addressing schemes (ATM and IP) and two routing schemes (RIP/OSPF and PNNI) that make network administration, operation, and troubleshooting very complex. MPOA, LANE, and NHRP require servers that must also be configured and maintained. In the WAN, the approach used by MPOA to create ATM connections per flow will have scaling problems as is the additional burden of the ATM connection signaling required to set up and tear down connections in the network. A significant amount of tuning will be required to optimize network performance.

	TAG/MPLS	MPOA	Advantage To:	Why?
Addressing and Routing	IP	Combined IP and ATM	TAG/MPLS	Best fit to IP Networking. Less Complex. Easier to Manage. No Addressing Conflicts
Connection Method	Shortcut PVC Topology Based	ATM SVC Flow Based	TAG/MPLS	Connections Established Before Traffic Moves. Faster Recovery from Problems.
Scalability Issues	Limited Only By Switch Connection and Traffic Capacity	SVC Capacity Limitation	TAG/MPLS	ATM SVC is a Limiting Scalability Factor in Any ATM Switch.
Performance Issues	Limited Only By Switch Traffic Handling Capacity	SVC Performance Limitations	TAG/MPLS	ATM SVC Performance in Large Scale Networks is Questionable. Connections are Always Re-established at the UNI
Performance Factors in 622M and 2.4G Trunks	Limited Only By Switch Connection and Traffic Handling Capacity	SVC Scaling and Performance Limitations	TAG/MPLS	622M Needs 16K Connections, Would require 600-800 CPS Per Trk. 2.4G Needs 64K Connections. Would Require 2800+ CPS Per Trk.

The more robust networking solution is based again on fundamental, solid networking principles. Everyone now knows that the key to optimal networking, combining the best for routing and forwarding with high scalability and performance, is based upon a topology-based, connection-oriented, switched network. Such networks have pre-defined routes calculated and connections established (including provisions for alternate routing) to eliminate the guess work and overhead associated with the flow based “connect as you go” networking philosophy. The flow-based approach will not scale. TAG/MPLS is designed with this in mind. MPLS is still under discussion, but TAG switching is here. TAG switching with ATM, for example, IP routing is the network routing. In addition, TDP is used after the network topology is discovered and routing table built to set up connections using the TAG as the connection identifier, **before** any data moves through the network, thus improving the probability for successful data delivery. An additional benefit is the consolidation of data traffic associated with different destination routes over single connections that reduces the overall number of connections and enhances the scalability of the solution. Multiple levels of QoS are handled by having more than one TAG to the same destination with different QoS attributes. Similar QoS profile connections can also be consolidated. The result is a more simplified and higher performing solution that can scale to meet the growth of the Internet and intranets. In this case, less is more.

Backbone IP Networking Using the Alcatel 1100 HSS

The 1100 HSS is a multiservice switching platform that already performs independent ATM and Frame Relay switching and is evolving to include IP packet routing and forwarding functionality. The most significant benefit of this product is the capability to perform both network and service interworking among the protocols. IP@ATM adds IP switching capability to the 1100 HSS so that backbone networks can be created taking advantage of the inherent of connection oriented switching using all of the benefits of using switches instead of routers.

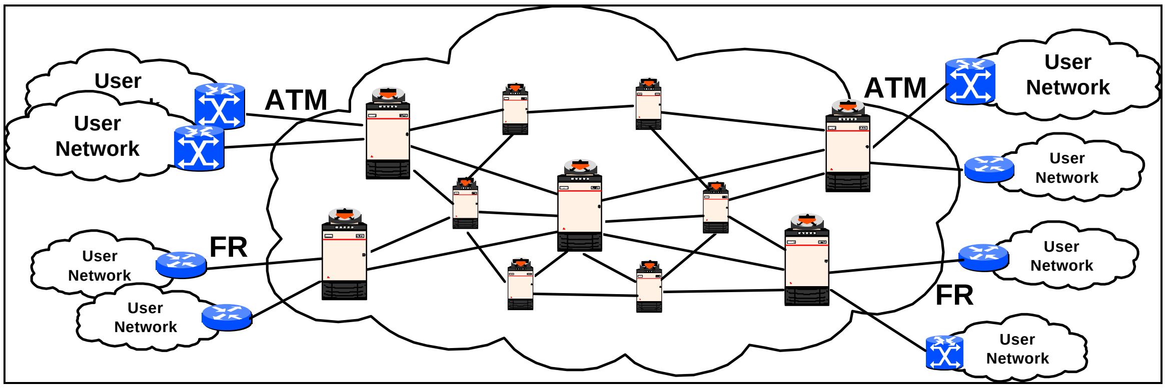

IP packet switching functions are implemented in the 1100 HSS using a distributed architecture and shortcut techniques based upon TAG switching and eventually MPLS. In the first phase of implementation, an external Cisco router will be connected to the 1100 HSS and managed in common through the Alcatel 1100 NMS. This is depicted in the diagram below.

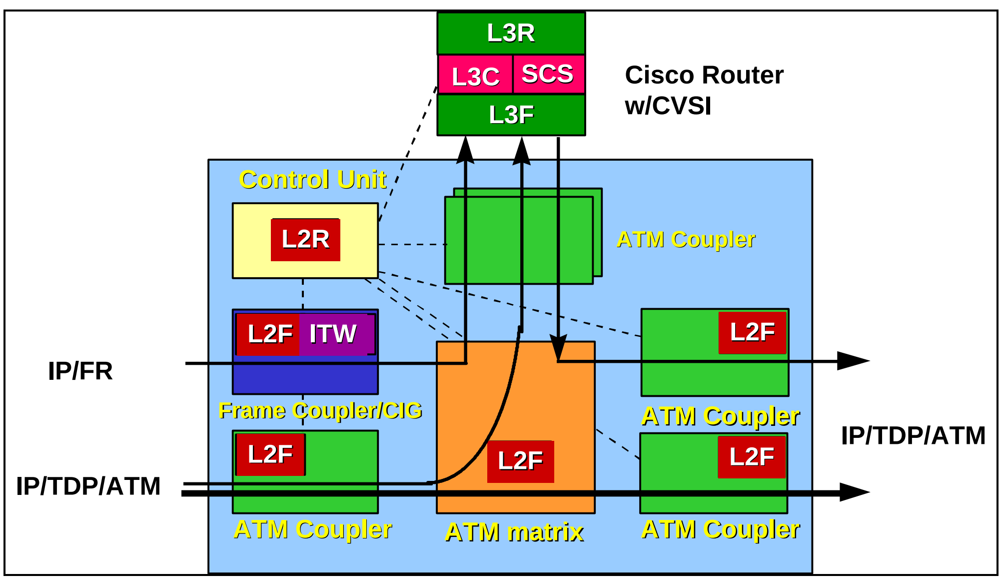

The router is connected to the 1100 HSS through ATM interfaces and uses the Cisco Virtual System Interface (CVSI) as the means to communicate with the 1100 HSS Control Unit. The router functions as an IP server for the 1100 HSS and performs standard IP routing and forwarding as well as TAG functions. The 1100 HSS will be capable of supporting native ATM and Frame Relay as well as IP over Frame Relay and IP over ATM. The advantages of this scenario are significant. The presence of an integrated Cisco router controlling IP traffic and TAG connections in the 1100 HSS ensures compatibility with any Cisco router that connects to the 1100 HSS. The benefit of direct compatibility with Cisco's IOS and Netflow enable an easy way to integrate connection-oriented networking into router networks.

The ultimate goal is to have an internal, distributed, plug-in IP Servers that will allow for routing and forwarding that can scale to multi-gigabit levels and support TAG/MPLS. The scenario is depicted in the diagram below. The 1100 HSS with its increased networking capabilities will add to any network to create a solid, robust backbone capability.

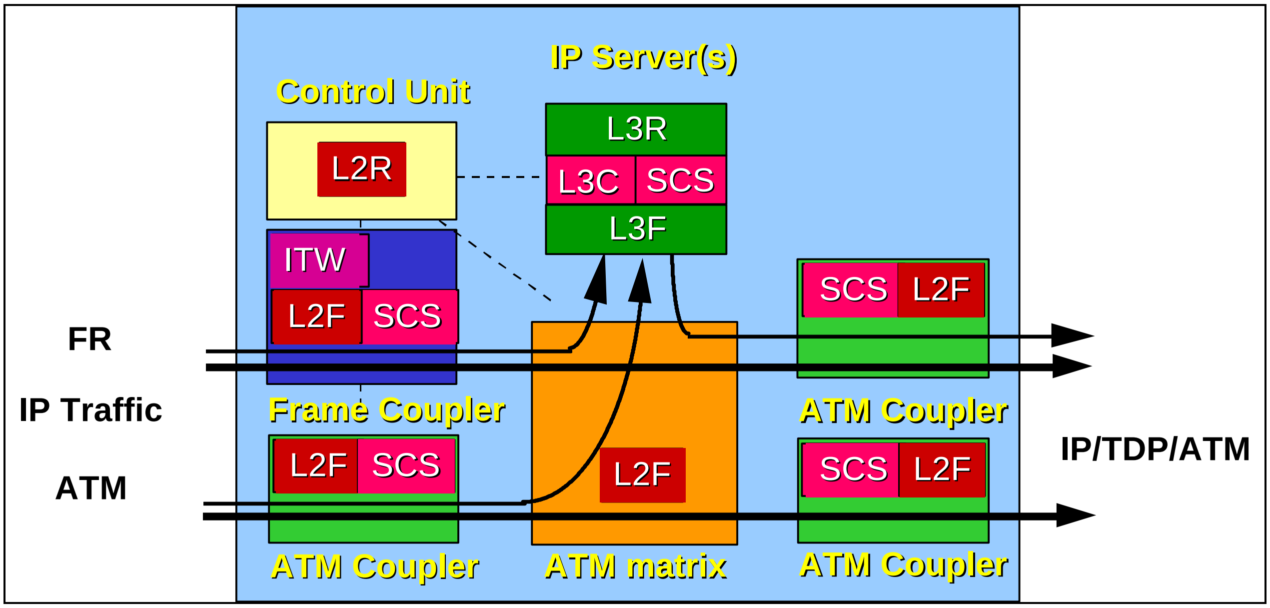

An 1100 HSS will be capable of operating as both a classical ATM and/or Frame Relay switch as well as an IP switch. The resulting product becomes a powerful Integrated Switch Router (ISR). There are many advantages to having a network based upon this kind of switch. Multiple routing topologies are created based individually upon ATM, FR, and IP. This allows for maximum flexibility to use the switch as a single platform for multiple independent network services. Management of the network is simplified. The advantages of incorporating IP switching eliminates the need for intermediate transit routers that form bottlenecks in large networks. Once shortcut path connections are in place, IP traffic moves through the switches without packet-by-packet routing. To the routers that access such a network, the intermediate switched connections looks like a single hop to the IP destination point.

The IP functionality is distributed and centers around an IP server integrated with the switch that performs IP route calculations on a continuous basis. The server can function as a full featured IP router as well as an engine for TAG/MPLS. The routing table that is generated is used to build the shortcut signaling relationship data that is then distributed to the individual external interface couplers. The IP server also has layer 3 traffic characterization and forwarding capability. Layer 3 forwarding can be performed as necessary. Arriving IP data flows over ATM or FR, that are characterized as shortcut, are forwarded at the ATM layer. The IP server itself is scalable with the capability to add additional servers as necessary.

The 1100 HSS IP switching implementation allows for the establishment of IP Virtual Private Networks. IP VPN capability will allow for the virtual partitioning of IP traffic handling functions to create a service platform that will support connectivity for multiple service providers, private Intranets and/or Extranets. Such capabilities can be combined with various types of access platforms, including xDSL access products such as the Alcatel 1000 ADSL, to offer maximum flexibility in service provisioning.

Class of Service and Quality of Service in Multiservice Networks

One fundamental reason for using ATM in the core of multiservice networking is to take advantage of the inherent multiple levels of CoS/QoS built into the ATM layer. The capability now exists to take advantage of this and apply the CoS/QoS philosophy to IP traffic. Today all IP traffic can be equated to the equivalent of ATM UBR traffic, i.e., best effort CoS. End user sources of IP traffic will eventually use the ReSerVation Protocol (RSVP) to request a certain level of QoS for a particular IP application that they want to use. The expectation is that the QoS level requested will be established on an end-to-end basis with the connection destination. A mapping of RSVP QoS level request to an ATM CoS/QoS level will take place in the access to the ATM portion of the network. In an 1100 HSS network the advantages will be considerable.

The 1100 HSS incorporates standard ATM CBR, VBR and UBR CoS levels and will also incorporate ABR CoS using RateMaster, a patented explicit, rate-based flow control mechanism that maximizes the use of the networks trunking resources under variable traffic conditions. The benefits will also apply to Frame Relay sources as well as there are also multiple CoS/QoS levels in the 1100 HSS for FR traffic, including the equivalent of ABR service using the Frame Relay version of RateMaster.

Network and Service Interworking Among ATM, Frame Relay, and IP

The 1100 HSS already incorporates standard ATM and FR network and service interworking. The addition of IP traffic sources in combination with ATM and FR allows for another level of interworking to make the process of ATM/FR interworking more efficient.

The standard ATM/FR interworking process today is configured statically, i.e., the relationship between ATM VPI/VCI and FR DLCI is predefined by a network operator. When IP traffic is involved, the process of interworking will become automatic using the shortcut techniques. IP traffic will be extracted from ATM cells or FR frames and directly routed to FR Frames and ATM cells respectively according to the shortcut connection paths established. The key advantage to this is the capability to move IP traffic over ATM and of easily interworking with FR to interface to service provider FR networks that are already established.

Summary

The addition of IP switching capability to the Alcatel 1100 HSS using TAG/MPLS technology offers significant advantages to network service providers. Having direct Cisco-compatible IP software eliminates any concerns over IP service interoperability with the vast majority of IP routers installed throughout the world. The ATM and Frame Relay multiservice switching aspects of the 1100 HSS are further enhanced with the capability to perform IP switching with network and service interworking among ATM and Frame Relay in a single powerful platform. IP VPN capability will provide a significant advantage to service providers, particularly those looking to establish wholesale network services, to offer controlled partitioned access to common subscriber sources such as xDSL access multiplexers, etc. As the next generation IP standards move to completion, the 1100 HSS will incorporate those functions that will guarantee continued interoperability in the network.