



**The ATM Forum**

[www.atmforum.com](http://www.atmforum.com)

© Copyright 2004

The ATM Forum

All Rights Reserved





# Table of Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. The Market Demand for Convergence</b>	<b>2</b>
<b>3. Enterprise Scenarios and the Impact of Pervasive Computing</b>	<b>5</b>
<b>4. Convergence in the Core: Technology Interworking</b>	<b>6</b>
<b>5. Convergence at the Edge: Integrated Platform Development</b>	<b>8</b>
<b>6. The ATM Forum</b>	<b>11</b>
<b>7. Internet Engineering Task Force</b>	<b>12</b>
<b>8. MPLS and Frame Relay Alliance</b>	<b>13</b>
<b>9. Case Study: Interoute – A New entrant IP Centric Carrier</b>	<b>13</b>
<b>10. Case Study: BT – the Incumbent Perspective</b>	<b>16</b>
<b>11. Summary</b>	<b>20</b>
<b>Appendix</b>	<b>21</b>
<b>Acknowledgements</b>	<b>21</b>
<b>About The ATM Forum</b>	<b>21</b>
<b>ATM Forum Specifications</b>	<b>22</b>
<b>The ATM Forum Broadband Roadmap</b>	<b>23</b>
<b>About the MPLS &amp; Frame Relay Alliance</b>	<b>23</b>
<b>Acronyms</b>	<b>24</b>

## 1. Introduction

The convergence vision for public telecommunications infrastructure is not new. The networking and telecom industry has always struggled to balance the desire for service-specific technology against the “one size fits all” approach. The need for innovation to deliver new services has led to the creation of new technologies, but has also resulted in a tendency to deploy service-specific networks – a commonplace feature of today’s network landscape. The creation of a single converged backbone capable of delivering multiple services is highly attractive for reasons of flexibility, economies of scale and operating efficiencies. Indeed technologies like ATM and ISDN were initially created specifically with convergence in mind.

Different technical requirements for the delivery of different services have logically led to the creation of multiple networks using multiple technologies. The challenge now is how to consolidate these multiple networks into a common infrastructure in order to reduce complexity and cost, but also to increase flexibility. Infrastructure technologies like ATM, Frame Relay, IP, and MPLS co-exist in the service provider infrastructure and are expected to continue to do so for some time yet, which makes the question of how best to interwork these technologies extremely important. Moving away from a “network per-service” model, the converged network of tomorrow will rely on weaving together multiple networking technologies and protocols into a single network. There are a number of possible solutions being proposed but the real question is – what is the optimal solution?

ATM and MPLS are likely candidates to feature in this next generation network vision, but there will be other technologies that must be integrated. Furthermore, the search for an optimal solution does not imply that there will be a single “all embracing” architectural model that solves all the needs of all service providers and enterprises alike. The networking and telecom industry is too diverse for this and the needs of the enterprise customers too widespread.


The challenge is to create a convergence blueprint that can be used to implement service provider specific networks and to identify and deliver the standards and specifications needed to do this. This allows for vendors and service providers to differentiate and ensures that, in a competitive market, there is still room for innovation and flexibility to make convergence a reality.

The real test of convergence is how service providers are using a combination of ATM and MPLS to deliver huge cost savings in their operational expenditure, while at the same time increasing their ability to offer new services.

## 2. The Market Demand for Convergence

Without pre-judging the need for specific technologies, it is clear that convergence means moving towards a common IP-based network capable of carrying voice, data, and video across some physical transmission/transport network. IP is key here as, quite simply, IP has won the battle for the networked desktop. That is not to say that IP can do it all and, as an application layer service, it will still require support from the lower layer networking technologies.

**Figure 1:**  
**Telecom Equipment Revenues (W. Europe)**



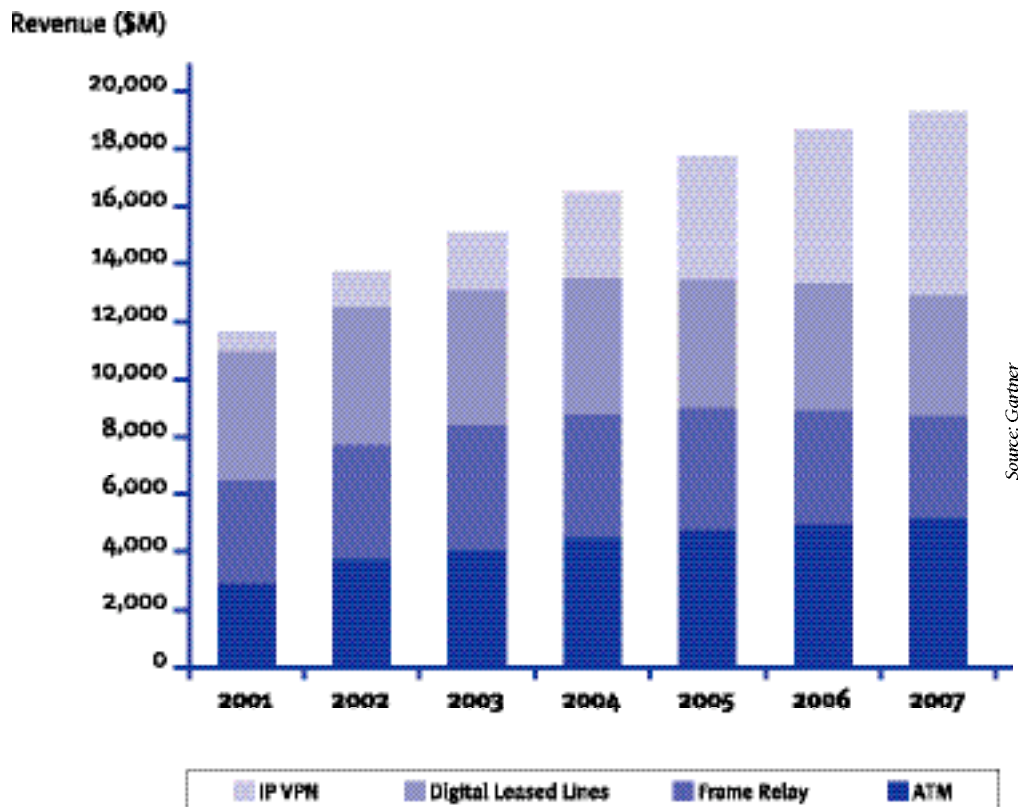
Of course, by some definitions we have convergence today with networks typically sharing the long haul transit at the optical layer. But above the “bit shifting” layer it is typical to have multiple service (or technology-specific) overlays. So, convergence is not just about convergence at a bit transport layer (Layer 1) – it is about convergence at Layer 2, Layer 3 and even above.

This means that service providers will want to sell/market an IP service rather than the underlying ATM or FR service that the IP service may depend on. What the customer gets may be FR or ATM based service but packaged as part of an IP service offering.

The market for convergence is actually getting more complex. It used to be that the main players were typically fixed operators, large enterprise users and the big vendors. Now there are more vested interests. Mobile service providers have entered the convergence arena and are, in some ways, the most fortunate as they have more of a green field opportunity for their GPRS and 3G networks. For them, the need for data infrastructure built on their voice centric networks is a new concept. Mobile networks are typically based on ATM (using 3GPP specification R4) and therefore able to use ATM's QoS. The incentive is strong to get it right, as GPRS will be the test bed for 3G. Compared to voice, revenues from mobile data are a drop in the fiscal ocean and, at least in Europe, must still prove itself, given that WAP has largely been seen as a failure. However, mobile service providers recognize that growth will be in data and multimedia services, although these new services must add value and not cannibalize the existing voice services.

Fixed operators are moving towards an IP infrastructure at a different pace, as they have more legacy migration concerns. So it's not a case of scrapping existing equipment – they must re-use equipment while they roll out their next generation multiservice infrastructure. Beyond the fixed and mobile service providers, there are also a number of xSPs seeing early demand for new services. Although they are typically dependent on other operators in the value chain for bandwidth services, they should not be excluded from the market picture.

**Figure 2:**  
**Service Provider Revenues (W. Europe)**



Then come the vendors, which have invested billions of dollars in the technology development to date and so have the experience and expertise to help drive the convergence needed at a technology level. But other significant players will play their part as well, including OSS/BSS vendors and those in content distribution and management. Lastly, it is important not to forget the enterprise users, whose need for private networks and VPNs are driving demand for the converged network.

Service providers are looking for ways to limit operational expenditures (OPEX) since capital expenditure (CAPEX) has been squeezed about as much as it can be. Demand for basic services is still growing but prices are dropping (voice, leased line) and customers now are more interested in a bundle of services. Likewise enterprise customers have cut back their expenditure, which has reduced services spending, but at the same time reductions in head count has made outsourcing more likely.

Why do service providers need convergence? Service providers need new revenue sources, which means looking for new markets or designing new services and applications. It is key for them to be more customer-focused and closer to the customer, both physically and in understanding terms. Competition is tough and each service provider has to be different from its competitors. Although there is a cost benefit to standardized services, there will still be a need for customized offerings to larger customers. Convergence can provide for flexible and quick deployment of these customer-focused offerings.

End users are becoming quite demanding. They want it now. Traditional lead times are not acceptable. Things are requested for today, tomorrow or even this afternoon – not in one month or three months. They need a bundled service. Although the margin on these services may be small, it makes them more likely to take up added value services like IP Virtual Private Networks (VPNs).

So, for the service provider, convergence means reduced cost and increased flexibility. But what is actually happening in the enterprise market will largely determine the success of service providers. How are computing trends in the enterprise market going to impact the end user demand that service providers must meet?



### 3. Enterprise Scenarios and the Impact of Pervasive Computing

Given that the enterprise is the core of the demand for services, we must examine what is happening here that could influence convergence. Most Enterprise IT managers face the challenge of how to mobilize the enterprise and create convergence for their fixed and mobile users. They need convergence of fixed and mobile as, quite simply, it becomes more difficult to categorize users without it. A desk-based employee, naturally assumed to be a classic “fixed” user, is also likely to have a mobile device (cell phone or PDA). Furthermore, the “always on” mentality means that this employee may need access to the enterprise resources from home, or when visiting a partner or customer site, or even in transit between the office and somewhere else.

Enterprise IT managers typically do something only when they realize that they have lost control, and pervasive computing is creating a more uncontrolled model. People buy what they want and expect the IT department to integrate it. Purchasing departments are contributing to this problem by specifying that people buy their own PDAs. It saves some capital cost in the purchase, but they will pay the price in operational and support costs. Compare this to a laptop, which comes with an enterprise-standard image build. Few companies would suggest employees simply go out and buy a laptop and install on it what they want. Yet this practice is becoming prevalent with other devices in ad hoc computing.

The need to mobilize the enterprise will impact the network infrastructure needed to support the users and the networking services that must be offered. For example, the drive for mobility is not just driving GPRS and 3G networks – it is driving the deployment of public wireless access points to the Internet.

So mobility will drive network providers to create a range of networks able to track the user throughout their day, which leads to a number of complexities. The user’s “profile” will change throughout the day depending on where he is, the network he is connected to and the nature of the access device. What can be done on a cell phone with a mobile data service using GPRS is quite different to what can be done using a PDA in a hotspot or a laptop back at the corporate headquarters.

Convergence to the enterprise customer is about creating a ubiquitous service that can offer:

- Adaptability - the service can intelligently adapt to the type of user device
- Utility - merging fixed and mobile networks into a single networked environment
- Control - reducing spend and avoiding redundancy

So, with a view of where the market is heading and what enterprise customers are demanding – what will the converged network look like and what are industry organizations like The ATM Forum, MPLS & Frame Relay Alliance, and others doing to develop the required specifications?

#### **4. Convergence in the Core: Technology Interworking**

As has been explained, the approach to convergence has traditionally been to implement a number of overlay networks over an optical transport, typically for TDM services, ATM & FR services and IP services (potentially a mix of public Internet and private VPN-style services). Bringing convergence to Layers 2 and 3 requires a different approach. The obvious choice, given the convergence on IP at the application layer, would be an IP-based convergence model at all layers. However, while IP is very good for “best effort”, connectionless data service, IP alone has some deficiencies both in offering QoS and in partitioning traffic from different customers – features normally offered by a connection-oriented model.

Some still argue that ATM, which was engineered around QoS using connection-based virtual circuits, is perfectly capable of meeting the needs of convergence as defined here. While many providers have yet to reach the limits of ATM, some are finding that, like FR before it, increased switching capacity will still be needed in tomorrow’s core networks.

A potential solution to the need for increased capacity is MPLS, which was designed to address the shortcomings of IP and deliver the benefits of a connection-oriented infrastructure, such as the optimization and traffic engineering. These benefits, coupled with features such as enhanced resilience using MPLS fast re-route and DiffServ coding to enhance the QoS of the IP network, make MPLS a natural choice for the converged core network and a popular delivery mechanism for new services at the edge (like IP VPN).

But how to approach convergence of different layers in the model, e.g. between data and optical? Generalized MPLS (GMPLS) provides a mechanism to extend the MPLS control plane to other technologies, like optical transport switches. Even more interesting is what is happening with service convergence, that is the ability to have different services delivered by a common network infrastructure on an end-to-end basis? In a converged network scenario, an access network is still required to reach to the customer site. The access network is connected to the network provider's MPLS core using a provider edge device. This means that the converged network could offer MPLS as the convergence layer in the core but still work with a number of other technologies at the aggregation points to the network and also in the access network. The aspects of this are covered in the next section.

Service convergence also means supporting multiple Layer 2 services (e.g. ATM, Frame Relay, Ethernet) over a common core, thereby extending the reach of existing services or allowing provision of today's standard services (like ATM and FR) over the new network. But can't ATM do all this? Potentially, but typically ATM and IP networks are quite separate in the service providers' infrastructure with the IP traffic being offloaded onto a separate IP network at the provider edge. So today, convergence tends to be only in the optical transport layer.

To suggest that a build-out of an MPLS infrastructure is the only solution would be an oversimplification. MPLS certainly holds some attractions for a green field site, but the challenges faced by an incumbent are much more complex. This can be seen in Sections 9 and 10, which feature case studies from a new entrant IP/MPLS carrier and an incumbent. Both are looking for convergence, but each is starting from a different point and potentially even finishing at a different solution.

So, what this means is that the core network must support the ability to do encapsulation – the transport of something across another. The basic form of encapsulation is network interworking, an example of which is the FR Forum's FREF5, which allowed the carrying of FR traffic over an ATM network. An example in the MPLS world is Pseudo Wire Emulation or some of the ATM/MPLS interworking specifications covered in Sections 6 and 7.

The more sophisticated requirement for encapsulation is to deliver service interworking, which allows disparate access methods to interwork. For example, ATM and Ethernet could be used to deliver an end-to-end Layer 2 service. An example today is the FREF8 specification for FR and ATM service interworking. Service interworking allows heterogeneous Layer 2 VPNs to be created.

So can this service consolidation take place in service provider networks? Are the necessary standards in place? Work has taken place in The ATM Forum's AIC and CS working groups since August 2001. The initial specification simply defined the user plane, then came the control plane using PNNI. The MPLS & Frame Relay Alliance is following up on this work with a generalized method of interworking ATM, Frame Relay, and MPLS control signalling protocols.

The Internet Engineering Task Force's Pseudo-Wire Emulation Edge-to-Edge (PWE3) Working Group is looking at a range of protocols over a range of networks (not just MPLS) and the final specifications are expected in the next three to six months. The ITU-T is following the lead of IETF and has approved various interworking specifications for ATM and MPLS for the transport of ATM cells, a Frame mode, and signalling interworking.

In addition, the MPLS and Frame Relay Alliance is specifying how to interwork various Layer 2 services at the edge when MPLS is used to interconnect them through the core. The end result is the choice of the local interface for each user depending on the availability of services in the local access network and which interface choice makes the most sense, both technically and economically, for the end customer.

So, with the developments in convergence in the core relatively well advanced, what does the picture look like as the network is extended out to the end user?

## **5. Convergence at the Edge: Integrated Platform Development**

The basic concept of an IP/MPLS core offering a high capacity packet transport is a simple one compared to the complexities at the access and provider edge. The technology to create the utility packet core is there – the main problem is CAPEX limitations. Internet traffic is still growing in the network, but the core has the capacity to cope. So some service providers are turning to the network edge, where the aggregation of customer traffic enters the core of the network, as the place where the biggest challenges to convergence now rest.

Reaching the customer always has the problem of needing to support multiple networks. Each country, each operator, each physical type of network is different, and many are rather outdated. Here the issue is both one of CAPEX and OPEX. At least 50% of all network CAPEX costs are in the access network and often with a long depreciation period. The technology tends to be put in place for a specific purpose and so it has traditionally been difficult to carry new services. Convergence to a single access technology simply won't happen.

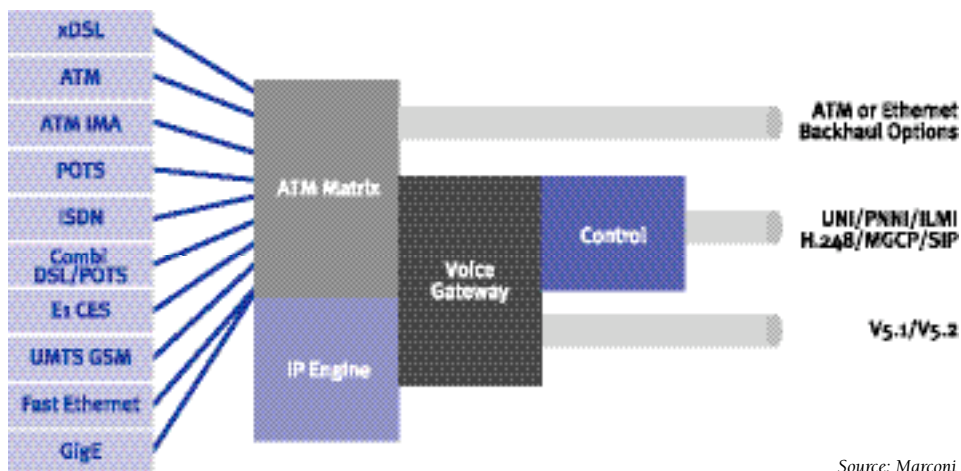
Because the access network is the service delivery path, it has a significant impact on the ability to offer new services. The emergence of DSL has opened up the old copper loop as an access path to deploy a new range of broadband services. But part of the problem is unpredictability. Much as any marketing department would wish to claim total vision, there is no escaping the fact that customer demand can lead to surprises. Few people predicted the growth in SMS text messaging in Europe or the way in which the Internet would become a part of everyday life. So how it is then possible to plan the network for the "As Yet Unknown Service"? Wherever possible one has to adapt existing technology to carry new services.

The model of proliferation of access technologies will prevail. There are various alternatives – from fixed wireless technologies like LMDS to the emerging free space optics. Other developments include HATP – High Altitude Telecom Platforms – hovering PoPs and even the existing utilities infrastructure.

Moving beyond the physical access, what transport technologies will be deployed in the access network? ATM is already there – used in the majority of DSLAMs today – and has many good characteristics for the access network. It has the major benefit of solid QoS for overcoming contention issues. Ethernet is also gaining ground due to its simplicity and ubiquity. Fiber becomes a natural choice in green field opportunities, such as the Dubai Marina development where 100Mbps Ethernet has been built into each home from ground up.

The need for QoS brings up the need to consider contention. The basic problem is that there is never enough bandwidth to the user and so there will always be a need to find the fairest way to manage the bandwidth. The mix could be data, voice and video – each with different bandwidth needs. One could use SDH – and rely on its time division multiplexing to separate traffic. However SDH lacks flexibility. ATM is clearly well positioned here, as it was designed with QoS in mind. On the other hand, Ethernet never really had to worry about QoS and so there have been various “add-ons” to bring QoS to the Ethernet world. MPLS also has good QoS support. The use of IP alone – designed around throwing bandwidth at the problem – will not be sufficient and so will also need QoS enhancers like DiffServ.

**Figure 3:**  
**Multi Service Access Network Architecture**



Source: Marconi

So, the access network will remain a bottleneck and there will be a number of different Layer 1 and Layer 2 technologies to support. Logically, the solution is a box that can do everything, although in reality it may end up being a selection of boxes with some commonality. This allows the service provider to cater for different user requirements and map the services onto the access loop using the most cost effective way to reach the user. Terminology emerging here is various. One service provider uses the term MSAN – Multi Service Access Node. Others refer to this as a gateway or MSPP – Multi Service Provisioning Platform.

So what does the MSAN have to do in a converged network? Simply as much as possible. It has to support a wide range of network interfaces and protocols and replace the functionality in individual access platforms (e.g. DSLAMs). There is also scope for these platforms to take on some of the functionality currently in the PE platforms, namely B-RAS.

With the greatest costs associated with the access network, the business case for the MSAN is about total cost of network ownership. One must consider CAPEX vs. OPEX. OPEX can be the bigger burden over the lifetime. CAPEX, although high initially, decreases over time. OPEX tends to be initially high to cover deployment and then drops back, only to rise again over time as the burden of support and maintenance kicks in. There is also a case for spending CAPEX to reduce the OPEX (in human terms) required to upgrade services.

By placing the intelligence in the MSAN, the “edge” is moving closer to the customer. One key issue will be how to handle voice in the MSAN. In the move to packet voice, there will need to be a media gateway to convert TDM voice to packet voice. The other issue, in creating an MSAN for a triple play proposition, is how to deploy video. Although in its early stages, the “Telco TV” model is gaining ground using DSL delivery. But in addition to simply supporting DSL, the MSAN will have to support some form of multicast technology and also be concerned with establishing that the network can support the video stream at the required QoS.

It does mean that the MSAN can, and will, support voice and then start to replace the existing TDM voice equipment. Much of this voice equipment is coming to end of life and the cost of maintaining it outweighs the cost considerations of moving to a packet voice approach. The converged network will support Softswitch functionality to replace the Class 4/5 switches. These will use MGCP/H.248 as signaling software with the voice traffic going over the MPLS core. This will require MPLS CoS to segment the voice and data traffic.

Another business case justification for the MSAN is that convergence actually reduces the footprint and saves space. So, space requirements will drop and this can reduce costs through releasing building space, the savings from which could then fund the equipment replacement!

So, having looked at the convergence issues for the core and for the edge, how are the technology forums helping the service providers and ultimately the end users accelerate their convergence plans?

## 6. The ATM Forum

ATM has become one of the most prevalent technologies in service provider networks over the last decade. Its role has been both as a core networking technology to offer a resilient and predictable connection oriented cell transport and as an aggregation and access technology where the QoS classes allow for service differentiation. ATM & FR services are still generating more revenue in today's networks than IP VPN. Over the years, The ATM Forum has had to place emphasis not just on developing ATM specifications but on tackling the interworking requirements. From the mid-90's there was a clear need to develop more intelligent IP/ATM solutions and since 2000, the emphasis has shifted to look at how ATM and MPLS will co-exist.

Whatever technology is used in the core, such as MPLS, transparency to the end customer is required, so the encapsulation of one traffic type across a different network must be user and application transparent. The ATM Forum produced some of the initial specifications in this area, with the key principle being that ATM/MPLS interworking allows ATM users to be interconnected via a tunnel across the MPLS network. Four specifications in this area have been approved since August 2001.

The goal is to aggregate multiple ATM VCCs and VPCs into a single MPLS Label Switched Path (LSP) so that only the LSP is visible to the core routers. This mechanism can support all ATM AAL types and the transport of OAM and RM cells, while also allowing support for single or multiple ATM cells in a single MPLS frame and offering complete transparency of ATM to ATM over MPLS.

Other work that has been done to enhance the MPLS support includes the addition of sequence numbers. ATM assumes FIFO sequencing, which is not necessarily the case with MPLS, for reasons of load balancing, for example. There it requires a mechanism to ensure that cells do not get incorrectly ordered during their transit of the MPLS network.

There are also extensions to the MPLS RSVP signalling protocol to set up the transport LSPs, although this work is still ongoing, and a specification to enable SVCs and SPVCs to be established that span the MPLS network through extensions to PNNI and AINI to support MPLS.

Looking at network and traffic management, one issue is how to map ATM QoS to MPLS service categories. There are two different ways of achieving this.

### **1. Class multiplexed LSPs**

Like CBR, UBR etc. Multiple traffic types using the same LSPs and so MPLS has to support most stringent requirement of the various traffic types. The mechanism to ensure MPLS can support the most stringent of the ATM service levels is network specific.

### **2. Class based LSP**

One LSP per ATM traffic class allowing each traffic class to be handled separately.

Note that a hybrid of both is possible, with perhaps CBR traffic being given its own LSP and other traffic classes being combined (multiplexed) into a different bit common LSP. These mappings can be used to enable PNNI to advertise the available resources on the MPLS LSP to the attached ATM networks, so allowing PNNI to make efficient use of the resources of the MPLS core network.

## **7. Internet Engineering Task Force**

MPLS standardization originated in the IETF, and, in its PWE3 Working Group, it continues to have an important role in the specification of the encapsulations necessary to support a converted IP-MPLS-based backbone. This work allows Layer 1 and 2 services to be run over MPLS, and forms the basis for a lot of work that is also taking place in other groups (like The ATM Forum's ATM/MPLS interworking activities mentioned in Section 6). PWE emulates Point-to-Point connections to tunnel traffic, such as TDM or ATM, across MPLS. Although not limited to MPLS tunnels (there is also another tunneling option using L2TP), much of the emphasis is indeed on MPLS given its role in the convergence model. The end result provides for such solutions as ATM over MPLS, FR over MPLS and Ethernet over MPLS. There is also active work on how to extend the LDP MPLS signalling protocol to set up the pseudo wires

The IETF also has two working groups standardizing VPN services over MPLS backbones. The first is the L2VPN WG (ATM, FR, Ethernet) and the other is the L3VPN WG. L2 is to standardize multipoint Layer 2 VPNs using techniques like VPLS (See section 10). It presents a Layer 2 service to the end user that emulates an Ethernet. There is also an IP only Layer 2 VPN specification called IPLS, which concentrates on building connectivity at Layer 2 for transporting IP-only frames.



## 8. MPLS and Frame Relay Alliance

In April 2003, the MPLS Forum and the Frame Relay Forum merged to form the MPLS and Frame Relay Alliance, with the combined vision of MPLS in the service provider backbone providing Frame Relay and IP services at the edge. The Alliance has also been working on a number of convergence-related specifications, including Layer 2 (ATM, Frame Relay and Ethernet) interworking over MPLS-based backbone networks, and generalized signalling interworking between ATM, Frame Relay, and MPLS. Both of these work items are meant to allow the preservation of existing edge services for both end customers and their service providers, while allowing the SPs to gracefully transition from multiple specialized backbone networks to a single converged MPLS-based backbone.

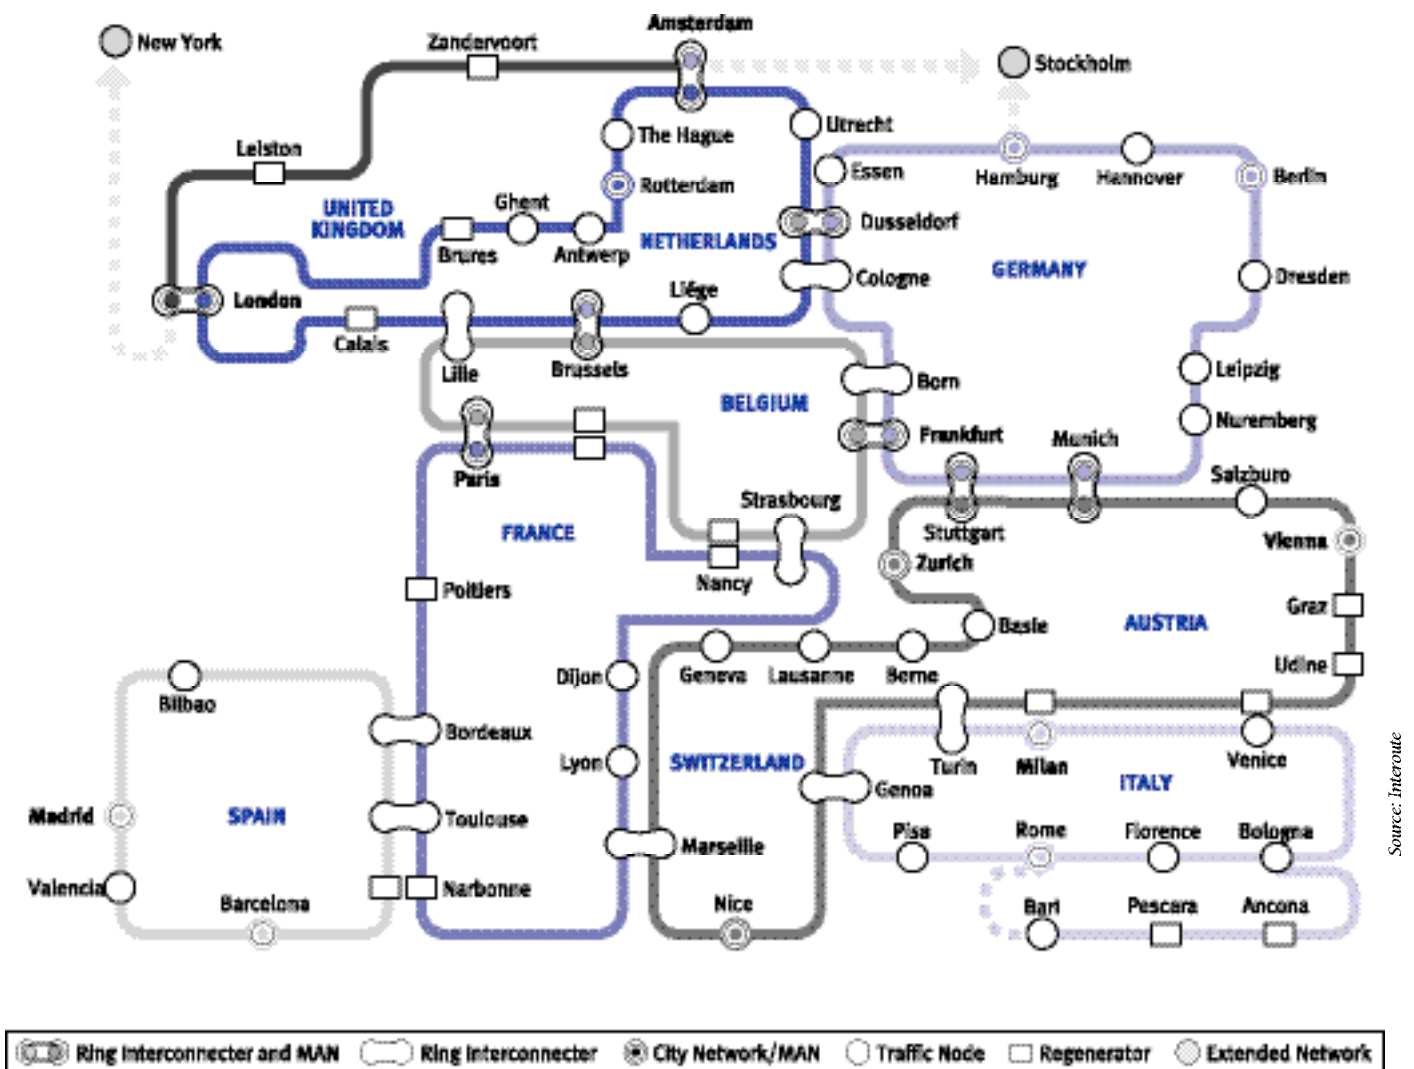
Another area of convergence that the Alliance has been working on is Voice over MPLS, for which there are a number of alternative solutions tailored to the specific needs of different existing service provider voice infrastructures. These alternatives include MPLS Forum 1.0, which specifies voice samples directly encapsulated over MPLS, carrying Voice over IP over MPLS, carrying ATM AAL1 and AAL2-encapsulated voice over MPLS pseudo-wires, and finally encapsulating TDM circuits over MPLS. All of these solutions are necessary due to the great diversity of ways voice is currently carried through different service provider infrastructures.

## 9. Case Study: Interoute – A New Entrant IP Centric Carrier

Interoute is a privately funded (part of the Sandoz Corporation) multinational IP-based carrier that has placed a strong emphasis on MPLS to deliver a range of services to its customers. The company has a presence in nine countries using its own network and covering 18,000 ducts of fiber and spanning 140 Metro/City PoPs. Having recently absorbed the KPN e-bone network, it is based around an IP infrastructure able to offer both Layer 2 and Layer 3 services. The Interoute philosophy has been not to put all their eggs in one basket and so saw a critical requirement to be able to provide a range of services. Their reasons for looking to MPLS to support their IP network were for reasons of scalability, simplicity and cost. With an array of vendors offering MPLS products it has created a very cost competitive market. While there is no dispute over the aspect of choice, simplicity has been less forthcoming, with a number of issues relating to IP VPNs that have led to complexity in the configuration process for the edge equipment.

As a service provider, their value proposition is to offload the burden of running a network from their customers. The model to which they operate is that the end customer doesn't want to have to worry about running their own Layer 3 routing network. Customers also need to be able to change their requirements quickly and have the ability to add or remove new sites and increase/decrease the bandwidth required on a per site basis. Service providers also need clear policies to allow for the deployment of new services, such as how to deploy VoIP using MPLS. The end goal is to create a bundle of services and use this bundle to win customers' loyalty so that they remain future customers.

**Figure 4:**  
**Interoute IP/MPLS Network**



Source: Interoute

One challenge is focusing on the applications -- something that Interoute perceive vendors sometimes overlook in favor of an emphasis on the next generation of “go faster” devices. Enterprise customers are looking to outsource, but they still need to be able to offer value-added services to their own customers. Therefore, Interoute has had to get closer to the customer, both in terms of reducing the local tail cost in the network (and therefore keep more revenue), but also in better understanding the customer’s needs. For example, class of service is only of interest if it can be put into the context of how it will solve a customer service problem. The end result is that Interoute have built intelligence into the network by making the PE function smarter. This allows them to differentiate types of traffic, offer multiple services and then sell the whole package of services. For Interoute, it is simply not about selling a box; it means a move to a total services mentality.

In order to offer a service bundle, business voice services must be supported. Typically, VoIP is added in to an existing network deploying data services, but Interoute took a different approach. The critical nature of voice requires that a network is designed to support voice and allows data to be added afterwards, since it is much more resilient.

There is a growing trend in the industry to offer low cost, simple Ethernet services. While Ethernet connectivity could be a good local access tail for a VPN service, Interoute is less convinced about Metro Ethernet as a service. Indeed, they could offer such a service as part of a portfolio of services, but they do not see a business case to invest in edge devices solely for deploying Ethernet services. Their multiservice philosophy requires multiservice equipment at the edge.

Although Interoute owns their own fiber network, they see a case for being more access agnostic. In some cases, a wireless local loop may be the most cost effective way to reach the customer. By populating as many buildings as possible in this way it brings the customers closer to the connectivity points. Also, by taking a wholesale DSL service they can provide an SME IP-VPN using wires only. This is a market segment ignored by many VPN providers and can offer a solution for five to seven sites sold through specialist SIs using unbundled Layer 3 services. As for capacity, they are currently upgrading their 2.5G backbone links to 10G, although this does have a CAPEX impact on the routers and the interface cards.

As for the acceptance of the Layer 3 PP VPN service and the whole outsourced network model, Interoute sees US customers as more likely to want to keep control of the CPE. In Europe and Asia, there is more acceptance of outsourcing, and so the bundled Layer 3 service is seen as a viable and attractive option.

## 10. Case Study: BT – The Incumbent Perspective

BT's network evolution has been the classic case of the incumbent – every time a new technology was introduced it led to a new service and a new network. Clearly, they need a single network that offers multiple services, which reduces the number of network elements, and hence reduces costs while making service deployment faster.

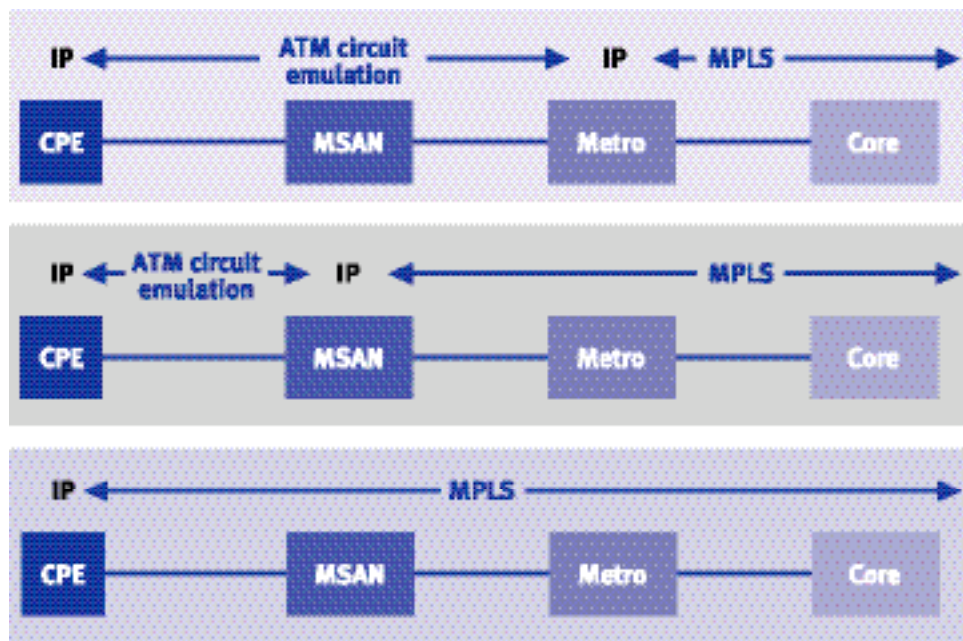
The network(s) of an incumbent like BT are multifaceted. Beyond the equipment in a home or business there is the aggregation layer. BT has 100,000 access devices (remote concentrators, DSLAMs, leased line SDH access points). At the PE there are PSTN voice switches, cross connects, and edge routers, which number about another 1,000 devices, plus about 170 network devices in the core.

Starting with the access, BT's goal was to bring the PSTN, SDH, and DSLAM functionality into one node under what BT terms the MSAN (Multiservice Access Node). The term node does not necessarily imply a single box; it may actually contain several discrete boxes. The core network was essentially made up of a large packet-switched network using (MPLS or IP/MPLS) and SDH cross connects delivering circuits at the VC4 level. This had the benefit of removing lower speed service drops (such as nx64 or E1) where the cross connect economics are poor. Beneath the core lay an optical transport switch offering a wavelength service providing wavelengths to both BT's customers as well to the packet-switched routers.

The extent of BT's problem to converge this network can be seen in the size of the network. There are 73 million copper pairs and over 400,000 fibers. Each network device typically has a different service management system. And it's not just about service management platforms – BT has six separate service management organizations. So convergence also means convergence at an organizational level. But what is the business goal from convergence? A 70% reduction in the number of access nodes delivering a £250 million savings (about \$475 million) just in access network OPEX.

The process of deploying the MSAN concept began in 2003 although plans have been in place for the last two years. Like the new entrants, BT needs to offer Layer 1, Layer 2, & Layer 3 VPNs all using the same network and organizational/management infrastructure.

**Figure 5:  
Choices for ATM/IP/MPLS Interworking**




Source: BT

At the provider edge, BT needed one multiservice management platform. In the core, BT has several IP cores (IP VPN, Internet, own IP network), plus an ATM network and the PSTN. It is easy to say, “let’s have a single MPLS network” but not so simple for what it means in practice. The MSAN devices, like the DSLAMs, are ATM-based and so the next convergence question is “where to interwork ATM and MPLS?” Typically the solution has been to do most things over ATM in the access network and then break it out and then encapsulate over MPLS. But this involves too many conversion steps. A slightly better solution is to directly interwork between ATM and MPLS using ATM/MPLS encapsulation. BT is also considering whether to push MPLS further out into the access network, but it’s simply too early for this development.

On converged management, the principal is simple – less wires, less boxes means fewer faults. Here OAM is key. With fewer points in the network to test it is easier to solve problems. Unified platforms mean more rapid service activation using web interfaces for customer to self-provision.

Figure 6:  
Cost Vs. Complexity Vs. Bandwidth Scenario



BT sees a trade-off between bandwidth and complexity. At one extreme there is a model of infinite bandwidth, which is simple but wasteful of network resources. Clearly this approach is not cost effective especially as BT has to account for bandwidth costs at market rates. Equally so they can use complexity to squeeze bandwidth down, such as PWE bandwidth header compression. But the danger is that this creates too much complexity just to save a bit of bandwidth. Multiple QoS classes are actually also a complexity that is used, amongst other things, to save bandwidth (when compared to say a CBR model). So, the challenge is to get the bandwidth and complexity balance right.

The best of breed approach is examined for each mode and a decision can be taken separately for each mode for functions such as addressing, routing, signaling, OAM, etc. At a signaling level, BT uses PNNI to set up SDH as well as ATM but acknowledges that this does not mean that one approach will work for everything and indeed doesn't believe in the G-MPLS "one size fits all" approach.

Convergence is the foundation of the BT's model, but their experience has shown that it is wise to fix potential problems before applying the concepts of convergence. The rationale being that converging on something that is flawed will require it to be fixed later and the impact of the problem is then multiplied. This implies more forward thinking to work out now what has to be fixed ahead of time. This means considering the legacy services like the PSTN, FR, low speed TDM circuits and for these to have a migration plan. The same applies to how BT will migrate its ATM customers over to the MPLS core. For this there is also a clear need for Service Interworking as it could be quite practical to have a customer with an ATM UNI at one site and an Ethernet Layer 2 service at another.

Finally the OSS environment is still key and for this BT requires a single OSS that does simply do everything! Security has become more of an issue, with the example stated by AT&T that they have seen more network intrusion attacks in last six months than in the previous ten years. One radical solution to this problem is to move the IP control plane out of band. The other aspect of security is keeping critical services operating when using shared infrastructure. BT quite simply can't have the PSTN go down due to a problem with the Internet traffic. This can be illustrated more graphically when considering that the UK's Air Traffic Control (ATC) network uses private circuits from BT. The implications of service disruption on those circuits due to interference from say their Internet traffic could mean lost planes and not lost packets!

Although there has to be a way to differentiate quality of service, one also has to consider how the customer is charged and indeed how the network should deliver QoS in a fair way. For many the obvious solution has been to introduce "medal tiers". But this leads to practical questions such as if a customer is not using its allocation of Gold bandwidth then should the Silver tier customer deserve to get this bandwidth just because it's there. So, many of the issues related to class of service become not ones of how to implement technically at a L2/L3 level but more about how to structure and price tiered services. This has led many to move away from the medal tiering to look simply at application-specific service level guarantees.

## 11. Summary

The value-add of the service provider is shifting from basic switching to managing the network as an intelligent information utility – automating and simplifying service delivery software and enhancing the BSS/CRM systems to bring the customer closer to the service provider.

Customers are becoming more aware of their networking needs and how to get them at the most cost effective levels. They want on-demand services and self provisioning (or at least the interface to do so) – and they want it all now. Customer friendly billing becomes even more important as the customer moves to a single bill for multiple services spanning a mix of fixed and usage-based tariffs.

Clearly no single vendor can deliver all this. This will drive increased collaboration and partnerships with a best of breed model bringing together multiple parties. Likewise SPs will need to form partnerships in the content and delivery supply chain (e.g. content and local access).

No single technology can do it all. There is a clear need to standardize on IP at the application to network interface, which will drive the use of IP-supporting technologies, like MPLS, starting in the core of the network and slowly spreading out towards the customer. Existing technologies won't go away and the more established the operator, the more the need to plan for equipment & service migration – squeezing every last drop of value out of existing CAPEX.

As for specific protocols, the advantages offered by the co-existence of MPLS and ATM in enhancing existing networks will naturally focus the attention around these two technology areas. Developments in the underlying transmission layer will simply make for more cost effective and faster transport of raw information – but the value will still remain in the differentiating and optimizing services offered to the end customer.

That is not about shifting bits – it is about convergence bringing the benefits of faster and more innovative service delivery to the user. It will be a long slow haul, but service providers will have to embrace convergence. This time it may well deliver on the promise that has been sought since the first mention of ISDN, some 20 years ago.



## Appendix

### Acknowledgements

The ATM Forum would like to thank the speakers who presented at the Broadband Exchange for providing content and their assistance in the creation of this white paper.

- Andrew G. Malis, MPLS & Frame Relay Alliance and Tellabs
- Jim Harford, The ATM Forum
- Maureen Coulter, Gartner
- Paula Richards, IBM
- Matthew Bocci, The ATM Forum and Alcatel
- Clive Deakin, Interoute
- Peter Willis, BT
- Alistair McBain, Marconi
- Mike Dell, Data Connection Ltd.

### About The ATM Forum

The ATM Forum is one of the most established names in the broadband arena.

It is partly for this reason that The ATM Forum felt well positioned to take on the ownership of the Broadband Roadmap.

The ATM Forum is an international non-profit organization formed with the objective of accelerating the use of ATM (Asynchronous Transfer Mode) products and services through a rapid convergence of interoperability specifications. In addition, the Forum promotes industry cooperation and awareness.

Since its formation in 1991, The ATM Forum has generated very strong interest within the communications industry. Currently, The ATM Forum consists of approximately 89 member companies, and it remains open to any organization that is interested in accelerating the availability of ATM-based solutions. The drive towards a more broadband related position shows the recognition that it is not just about the success of a single technology but the role that any one technology plays as part of the bigger broadband picture.

**ATM Forum Specifications**

The ATM Forum Technical Committee works with other worldwide standards bodies selecting appropriate standards, resolving differences among standards, and recommending new standards when existing ones are absent or inappropriate. Typically this process is entirely member contribution driven on the basis that its members are key players in the industry and therefore best positioned to know where this work must take place.

The Technical Committee was created as one, single worldwide committee in order to promote a single set of specifications, thereby ensuring interoperability among all vendors as ATM products and services become available.

The Technical Committee consists of a variety of multi-vendor working groups, which develop specifications in different areas of ATM technology. The ATM Forum is always interested to receive contributions for new work items to further the development of specifications.

To obtain a PDF document of any completed ATM Forum specification, please visit the following URL.

<http://www.atmforum.com/standards/approved.html>

## **The ATM Forum Broadband Roadmap**

The objectives of the Roadmap are to:

- Help The ATM Forum understand what new work is needed by industry and
- To be able to use this as a tool to communicate with other organizations to minimize/eliminate redundant work

The Roadmap includes definitions of the key emerging opportunities for Broadband and ATM including the market opportunity, the application definition, reference diagrams, the industry gaps and problems to be solved.

Additionally any key documents or terms of reference in the industry are identified and who else is working in this area.

Outputs are not only a time-sequenced list of future work items, but also the identification of potential collaborations with other organizations.

## **About the MPLS & Frame Relay Alliance**

The MPLS & Frame Relay Alliance is an international industry organization that is advancing the recognition and acceptance of MPLS and Frame Relay technologies in the global telecom industry. The Alliance is driving worldwide deployment of multi-vendor MPLS and Frame Relay networks, applications and services, through interoperability initiatives, implementation agreements, and educational and marketing resources and programs. The Alliance currently has more than 50 members. For a complete list of members, go to <http://www.mplsforum.org/about/members.shtml>. For Alliance membership information please contact Alexa Morris, executive director, at (510) 608-5914 or via e-mail at [amorris@mplsforum.org](mailto:amorris@mplsforum.org).

Additional information about the MPLS & Frame Relay is available online at <http://www.mplsforum.org/>.

## Acronyms

AAL	ATM Adaptation Layer
ATM	Asynchronous Transfer Mode
BICI	Broadband Inter-Carrier Interface
BOM	Beginning of Message
CAPEX	Capital Equipment Expenditure
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CER	Cell Error Rate
CLP	Cell Loss Priority
CLR	Cell Loss Ratio
CO	Central Office (Exchange)
COM	Continuation of Message
CPCS	Common Part Convergence Sublayer
CPE	Customer Premise Equipment
CRC	Cyclic Redundancy Code
CTD	Cell Transfer Delay
EFCI	Explicit Forward Congestion Indication
EOM	End of Message
EXP	Experimental Bits
GCRA	Generic Cell Rate Algorithm
ILMI	Interim Local Management Interface
INE	Interworking Network Element
ISH	Interworking Function
LEC	Local Exchange Carrier
LSP	Label Switched Path
LSR	Label Switching Router
MPLS	Multi-Protocol Label Switching
NMS	Network Management System
NNI	Network to Network Interface
NSAP	Network Service Access Point

OAM	Operation Administration and Management
OPEX	Operating Expenditure
PCI	Protocol Control Information
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PNNI	Private Network-Network Interface
PTI	Payload Type Identifier
QoS	Quality of Service
RCC	Routing Control Channel
RM	Resource Management
SAAL	Signaling ATM Adaptation Layer
SAR	Segmentation And Reassembly
S-bit	Stack bit
SDU	Service Data Unit
SNMP	Simple Network Management Protocol
SSM	Single Segment Message
TCP/IP	Transport Control Protocol / Internet Protocol
UDP/IP	User Datagram Protocol / Internet Protocol
UNI	User to Network Interface
UU	User to User
VC	Virtual Connection
VCC	Virtual Channel Connection
VPC	Virtual Path Connection
VPCI	Virtual Path Connection Identifier
VPI	Virtual Path Identifier
VPN	Virtual Private Network

