This is to certify that the
thesis entitled

EVALUATING TECHNOLOGY, A COMPARATIVE LOOK AT ATM
AND MPLS.

presented by

Michael P. Lowe

has been accepted towards fulfillment
of the requirements for the

M.A. degree in Telecommunication, Information
Studies, and Media

Major Professor's Signature

12/18/09

Date

EVALUATING TECHNOLOGY, A COMPARATIVE LOOK AT ATM AND MPLS.

By

Michael Paul Lowe

A THESIS

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

MASTERS OF ARTS

Telecommunication, Information Studies, and Media

2009

ABSTRACT

EVALUATING TECHNOLOGY, A COMPARATIVE LOOK AT ATM AND MPLS.

By

Michael P. Lowe

Technology is a fast moving industry with many vendors introducing new products constantly. Being a decision maker in a technical field often requires making sound choices on what new products and technologies to acquire. With so many products and technologies to choose from decisions makers need a framework to see beyond product marketing to evaluate if a product is a good fit for their organization. In this paper I illustrate three core factors to examine that will give a better understanding whether said technology is a proper fit. The three factors that I examine are history, technical workings, and market adoption. I use this framework to examine Asynchronous Transfer Mode and Multi-protocol Label Switching technologies. I chose to examine these two technologies because of their similarity of goals while highlighting they're stark differences relating to the three methods of history, technical workings and market adoption.

TABLE OF CONTENTS

LIST OF FIGURES

Introduction

In the world of technology new unveilings tend to be filled with huge fan fair and high aspirations. Networking technologies in particular are backed by buzz words and an avalanche of hype. ATM was one such technology. Originally envisioned as the next great networking solution for voice, data, and video; ATM was to be a giant leap forward. Almost a decade later another new networking technology is making many of the same claims, this newer technology is MPLS. MPLS promises to bring about a lot of the same functionality that ATM had been designed for but without the caveats of ATM's connection orientated nature. In the fast changing world of technology newer technologies come about to unseat older one's, rival technologies compete each claiming to be the best. How can decision makers see beyond the marketing and make confident decisions regarding their network when scenarios like these come into play? Decision makers need to take a look at the whole picture regarding said technology to make a confident and correct decision. I believe in a framework where researching a technologies history, it's technical workings and market adoption rates are the key factors a decision maker needs to consider. Evaluating a product in this framework will allow decision makers to make confident decisions in determining if a certain technology will deliver the expected results. I chose to evaluate Asynchronous Transfer Mode and Multi-protocol Label Switching technologies to highlight this evaluation framework. These two technologies have an interesting dynamic; one older, one newer but both trying to tackle many of the same goals through very different means.

History of ATM

Firstly let's take a look at the history of ATM. Discovering the history behind a technology can shed light on many important factors. Such as what the need was for this technology to come about, why certain decisions were made and if there were any disruptive processes that may have changed the course of this technology for better or worse. As a decision maker discovering the course of how a technology came to be will many times highlight if it is a good fit for your intended purposes.

In 1968 Bell Labs engineers began tinkering with a technology called cell switching. This was the start of what would ultimately become Asynchronous Transfer Mode (ATM) after spending more than two decades on the drawing board (Gould, Jeff 1994). In the late 1970's computers had become increasingly diffused throughout American and European societies. Their adoption as well as their processing power began increasing at impressive rates. This increase in computer volume and processing power ushered in an increased need for data networks to link them together. Using the public switched telephone networks (PSTN) started becoming insufficient as it provided limited bandwidth and was very susceptible to signal noise. In addition the constant circuit connections of the PSTN were not an efficient use of resources. Data transmissions tend to be busty with long periods of un-use; this constant connection spent large amounts of time idle. There also arose the desire to have a network that was capable of simultaneously handling both voice and data. Integrated services digital network (ISDN) was designed in the early 1980's and was the first network to achieve this goal. It used the PSTN infrastructure but transmitted a digital signal instead of an analog one. Since it ran over the PSTN it was still a circuit switched service but it utilized both of the two 64kbs channels that a standard telephone line carries. ISDN can use one channel for 64kbs data connection and the other for voice. Optionally it can also combine the two channels for a data only 128kbps connection (Becker, Ralph 2006). Time division multiplexing (TDM) is used for the co-ordination of sending and receiving on each channel utilizing a network clock. Essentially each channel has a designated time slot to send and receive data in; each channel goes one after another as their time slots expire (Stern, Mahmoud 2004). ISDN however was not a viable solution for the rate at which computers, and ultimately telecommunications in general were evolving. Using the PSTN as a wide area network (WAN) was not dependable and ISDN's limited bandwidth was hardly future proof.

By 1986 the International Telecommunications Union began to outline the successor to ISDN called broadband ISDN (B-ISDN). This was to be the next logical extension of ISDN with significantly more bandwidth to carry voice, data, and video services. This outline also contained recommendations to address the limitations of the PSTN, eventually phasing out PSTN's core to a more intelligent mixed digital service (Wood, Robert 2005). Telecommunications providers however still saw data networks as being similar to voice networks and envisioned circuit based networks for end-to-end connectivity. The initiative to create broadband ISDN grew into what would ultimately become Asynchronous Transfer Mode (ATM).

When the ITU (International Telecommunications Union) began work on the ATM standard there were many groups involved and each had their own agenda they wanted worked into the technology. Telecommunications providers who were at the forefront of ATM research had a grand vision. Similar to what they had previously done with the PSTN, they envisioned circuit switched networks covering the globe. This would make it more difficult for newer rivals such as coaxial cable systems to adopt the technology. In addition circuit networks were what telecom providers were comfortable with. Headaches also arrived in what was a split between proponents who wanted to use ATM mainly for data, and proponents who wanted to mainly use it for voice. At this time while T-1 lines just started taking hold in the US, most European regions didn't even offer 64kbps data lines. Interests could primarily be divided by saying the US wanted to use ATM for data and Europe wanted to use it for voice. While this might not seem like a big deal, it became one when deciding what type of packet transmission system to implement. At first the data lobbying groups were opposed to using a fixed length cell all together. Data is not normally sent in uniformly sized packets, its variable. Take Ethernet frames for example, they allow for a multitude of sizes up to 1500-bytes (Cisco Systems, Inc 1992-2008). So when a payload of data is larger than the uniform cell size, the data must be split up into pieces that fit the cell. With smaller cells more segmentation must take place and more reassembly must take place. This takes a lot of processing time and is not ideal, however many were eventually persuaded by the current work being done with fast switching. This was the argument; the trade off of smaller uniform sized cells would be faster to route versus variable sized packets. This would ultimately make up for any performance loss due to extra segmentation and reassembly of the data. However the debate over the cell format was not going to end there. As I mentioned Europe's main interests were to use ATM for voice communications and not data. However the size of the cell one way or the other affects which application it will be most suitable for. The data lobbying group came up with a 128-byte cell length they thought would be ideal for data transfers. However in voice communications the ITU standard concerning delay is any connection with over 20-milliseconds of delay requires echo cancellation equipment to be installed. The European telecoms figured that with a 128-byte cell even if voice was encoded at 16 kb/s it would take 64-milliseconds to fill the cell before transmitting it. In the United States voice carriers were use to transmitting over long distances and to a large degree already had echo cancellation equipment installed. European providers were not equipped for long distances unless they were out of country connections. This meant they would have to invest a large amount of money into a lot of new echo cancellation equipment. They lobbied for a very short 16-byte cell to be used which would leave plenty of padding space to avoid any echo problems. Even encoded at 64 kb/s a 16-byte cell would be filled in two-milliseconds, a far cry from the maximum allowance of 20-milliseconds. After much debate the closest compromise either side could reach was that the data group would go as small as 64-bytes while the voice group would go as large as 32-bytes (Gould, Jeff 1994). They officially hit a stalemate; the two sizes each group wanted were just too far apart to reach a real compromise. So the ITU made what would be a haunting decision and just simply picked the median between the two sizes. A 48-byte payload cell was ratified as the standard (Gould, Jeff 1994). After heated debates from the two sides neither of them walked away happy and ATM's cell structure was finalized by a political justification instead of a technological one.

Shortly after the specs for ATM were ratified a third interest group came into the picture to muddy-up the waters even more. As awkward as the cell size is, a group of vendors decided ATM would still work for high speed LAN networks. Before the technology even had time to mature and be tuned, vendors Adaptec and Fore had already raced LAN products to the market. In fact Fore's first generation of UNIX based products didn't even handle segmentation and reassembly; instead they dumped that workload onto the CPU. These early ATM LAN products also used proprietary technologies to fill in gaps that the initial ATM ratification didn't include; such as virtual circuit creation. These proprietary components made products from competing companies' incompatible (Gould, Jeff 1994). In addition, Ethernet continued to evolve heralding the 100Mbit fast Ethernet specification and continued improvements on simplicity. ATM LAN technologies simply could not compete against the ever evolving Ethernet due to higher cost, more complex design, and unproven track record. In the end ATM for the LAN faded, however the technology continued to be used in WAN implementations as originally intended.

Technical Workings of ATM

As I mentioned earlier an important piece of evaluating technology is to study how the technology itself works. This is a piece of the framework that more technical people will evaluate but it is also were a lot of marketing spin can come from. Taking a deeper and objective look into how a technology works will not only highlight it's abilities but perhaps even more important it's limitations. As mentioned earlier, ATM never caught on as a LAN technology. Ethernet equipment was already widely established in the market; it was cheaper and based on IP (Internet Protocol). Because of this the ATM networks that were deployed were mostly WAN implementations and any future reference to an ATM network in this writing will be in reference to a WAN implementation.

When standardizing ATM the hottest debate was over the cell structure. To understand how an ATM network functions the cell structure is a good place to start. Asynchronous Transfer Mode transfers a fixed size of data packets called cells. As mentioned this is different from many other network technologies. For example Ethernet transmits frames and can send data packets of variable size. Simplicity and low overhead were the main goals when designing the ATM cell. Engineers figured that using cells of the same size would allow the routing of data to be simpler and faster compared to networks with variable-length-packets. The small and uniform packet size was also intended to reduce jitter in voice transmissions. Each ATM cell is comprised of 53-bytes, 5-bytes of header information and 48-bytes of payload (IEC 2007).

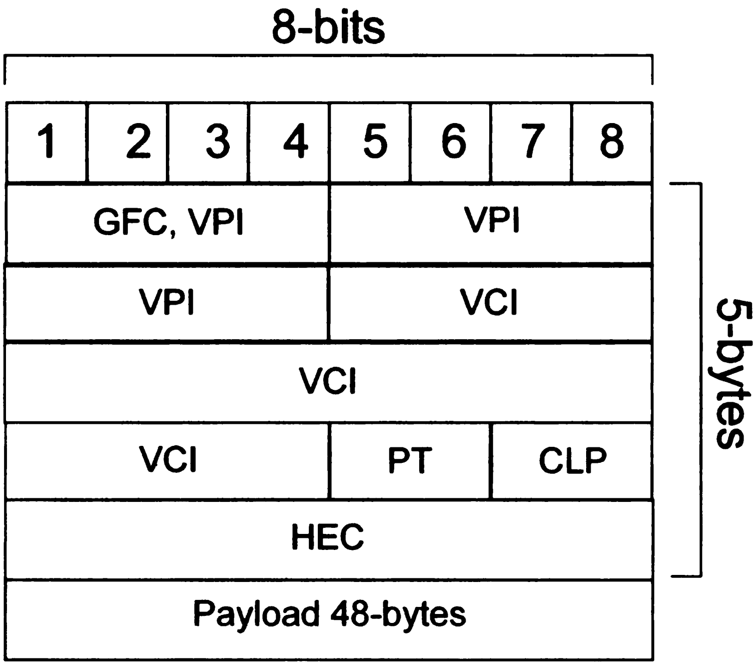

Figure 1. The separate components of an ATM cell.

In figure 1 the cell header is divided into octets to visually show how each component fits into the 5-byte header. You can see how adding each component of the five layers yields 40-pieces or bits which are the same as five-bytes (1-byte = 8-bits). Generic flow control occupies the first 4-bits of the header. Flow control negotiates the speed at which the sending and receiving devices transmit at to avoid cell loss. This would typically be switches and routers throughout the network. In the case of ATM, the flow control value is not constant from end-to-end. Each device can change the flow control value to correspond to the devices it's connected to, which is why it's called generic flow control (Cassidy, Kyle 2001). The next section of the header is the Virtual path identifier (VPI). This works in conjunction with the next section of the header called the Virtual channel identifier. VPI is comprised of 8-bits which signal an ATM switch which direction through the network the cell is suppose to take. Virtual path identification is only subject to each device in a network. Therefore this information is changed at each hop in a network (Cassidy, Kyle 2001). Working in conjunction with VPI is the next header section the Virtual circuit identifier. This section of the header holds the information for the virtual circuit that the cell must travel on. Virtual circuits will be explained later in greater detail. For now just know that a virtual path is the path through the network, and the virtual circuit is which circuit out of several the cell must travel through on that path. Virtual circuit information is the largest section of the cell's header at 16-bits; this is what allows ATM to carry multiple services over the same network. Payload type is the next component of the cell. This section is 3-bits in length and may contain information about the data as well as any traffic congestion the cell may have experienced (IEC 2007). Payload type works in conjunction with the next header section, the cell loss priority bit. This bit is flipped to either on or off to indicate whether the cell can be discarded by a switch in the network. If the payload type indicates that heavy congestion occurred the cell loss priority bit might flip from zero to one to tell the switch to discard the packet. Cell loss priority may also be used to give immunity to important packets where the bit is set to always stay zero. The last 8-bits of the ATM cell header are the header error control. This is the cyclic redundancy check (CRC) information used to run a checksum on the first 4-bytes of the header (Cassidy, Kyle 2001). The remaining 48-bytes in the cell are the payload, a chunk of whatever data is being transmitted.

After examining how the ATM cell is structured let's take a look at how these cells actually travel through an ATM network. Asynchronous Transfer Mode's development was heavily influenced by the telecom industry. Being such, it's actually a circuit based network, meaning each connection in the network is dedicated between nodes. You can think of the old telephone system as an analogy. There each call is established using a dedicated connection within a series of switches. Each caller has a physical connection with a dedicated line from end-to-end. Of course ATM is not that antiquated and uses virtual connections between end points, not physical ones. On the seven layer Open Systems Interconnection model (OSI), ATM is considered to be on the second layer, the data link layer (Xilink 2001).

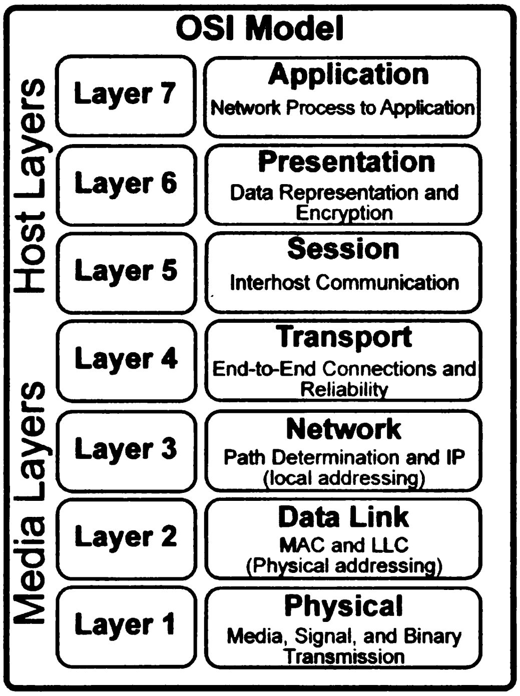

Figure 2. The seven layer OSI model.

The first word that jumps out at you from ATM is asynchronous which means; "*not occurring at the same time*" (Random House, Inc 2007). In the context of ATM this means that there is no set time slot for sources to have to send data in, which is typical of most circuit switched networks (Cassidy, Kyle 2001). Older circuit networks such as T-carrier used a standard clock to keep sender and receiver synchronized. This system of multiplexing is called Time Division Multiplexing (TDM). This is where ATM deviates from older circuit based networks and is sometimes referred to as a hybrid-technology. However, ATM is not a full-fledged asynchronous network. When the source is not sending data cells it sends a series of empty cells to keep the two nodes synchronized. This process is often referred to as bit stuffing (Cassidy, Kyle 2001). So ATM sends data cells asynchronously but keeps a synchronous connection between sending and receiving nodes.

Knowing how ATM's cells are structured and the manner of how ATM sends data we can now look at the connectivity structure of the network. ATM is capable of transmitting multiple streams of data concurrently, over one physical connection. For example a telephone call and a video stream can be run over the same physical cable. The way ATM accomplishes this is with virtual circuits.

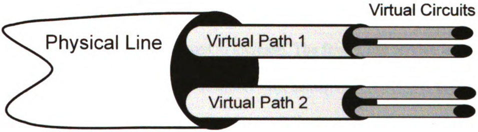

Figure 3. Diagram of virtual channels on one physical wire.

ATM can run multiple services at once by separating them onto separate virtual circuits; each circuit can handle one type of data stream. Referring back to the cell header information, data is assigned a virtual circuit and that information is encoded into each cells virtual circuit identifier (VCI). Don't confuse circuit identifiers with IP addresses; these are just in the context of keeping the individual paths separate. Virtual circuits are then bundled into virtual paths for organization. Virtual path information is also stored in each cells header in the virtual path identification section (VPI). Virtual paths can be used to bundle a series of virtual channels like a trunk, or in this case a virtual trunk. Using virtual paths in conjunction with a connection orientated link allows an ATM network to assign resources on a virtual connection basis. When the virtual connection is setup the proper amount of bandwidth is allocated to each virtual channel for each service; for example a data connection versus voice. Being able to more efficiently manage bandwidth in this manner has become increasingly important for companies. There's been an increasing market trend to more efficiently use the available resources to save cost (IEC 2007).

With an ability to control resource allocation through virtual channels, ATM defines five guaranteed service levels. This Quality of Service (QoS) is an aspect where ATM has always shined. Next generation networks have continued to followed suite as industry demand for it continues to increase. The first ATM service level is constant bit rate (CBR). CBR is a service meant for time sensitive applications such as voice and video (IEC 2007). These services are especially prone to cell delay. Any cells that arrive outside of a set window of time are considered no-longer useful and are discarded. The set value used to determine the maximum allowable delay at different network points is referred to as the cell transfer delay (CTD). Two of the service levels ATM provides are both variable bit-rate transfers; they are real time and non-real time. Variable bit rate non-real time (VBR-NT) is for transfers that are bursty or vary over time but are not as dependent on having a minimal CTD. An example would be certain types of non-interactive video or audio playback; anything where delays are not hampering the communication. Here statistical multiplexing is used where bandwidth is adjusted on the fly to save resources. The other variable bit rate service level is real-time variable bit rate (VBR-RT). Once again this is for traffic that is bursty, however the communication is less resilient to cell delay. Examples would be interactive video or voice. With VBR-RT a CTD value is specified and any late cells are discarded. Available bit rate (ABR) is a service level aimed at traditional computer communications such as email and file transfers. With ABR a CTD value is not specified since the communication is not sensitive to delay. However a minimum cell rate (MCR) can be specified to ensure delivery of a certain level of speed. The last service level is unspecified bit rate (UBR). Unspecified bit rate is a best effort service that contains no delay or bandwidth values. Basically it's whatever unused bandwidth the network can muster at that point in time. Like ABR this service level is used for traditional computer communications where delay is not of a high importance (Cassidy, Kyle 2001).

Working in conjunction with the various ATM service levels are the ATM adaptation layers. Not all networks are completely based on ATM; and thus do not use cells; there has to be adaptation standards to connect with other networks. There are five standardized adaptation layers. ATM adaptation layer one (AAL1) is for connection-orientated services that require a constant bit rate, are sensitive to cell delay, and missing cells. Each cell is given a sequence number, if the next sequence is unavailable when segmenting for transmission over the network a retransmission request goes out. Examples of networks that would interface using AAL1 would be DS1 or T1 connections. Adaptation layer two (AAL2) also places importance on cell loss and retransmission since it is mostly used for carrying voice and video. AAL2 however encapsulates variable sized packets within the ATM cell and uses variable bit-rate as opposed to constant-bit rate. The next adaptation level is actually two in one and is referred to as AAL3/4 and supports both connection and non-connection orientated networks. AAL3/4 is meant for variable bit-rate traffic that is sensitive to lost traffic but not necessarily delay. An example would be a Frame Relay network. The last adaptation layer is layer five (AAL5). AAL5 is for variable bit-rate data and has no built in error recovery or retransmission (Cassidy, Kyle 2001). It's essentially a cell with the whole 48-bytes of payload used for data. Higher layer protocols like TCP/IP can handle some of the error correction and detection that AAL5 is missing, the benefit is there are less processes taking place at segmentation and reassembly than with the other adaptation layers; simplifying the process to deploy.

History of MPLS

Now that we've seen the history and technical workings of ATM we must do the same for a rival technology in order to evaluate the two. In this case we'll look at the history behind Multi-protocol Label Switching (MPLS). As time went on from the advent of ATM, network needs changed, services became increasingly data centric and IP driven. Companies also wanted higher quality of service features built into the network. It had gotten pretty clear that ATM was not going to be the total solution it was meant to be. New approaches to building WAN networks were starting to be developed; the first standardized next generation network was MPLS.

MPLS's birth started when several companies had begun to experiment with what is now generally referred to as label switching. Most notable were Ipsilon's IP Switching, IBM's Aggregate Route-based IP Switching, and Cisco's Tag switching (Network World 2007). Each company came up with a proprietary approach to building a label switched network. It was now obvious that ATM's design proved to be a flawed approach. ATM has taken a lot of criticism for its cell based nature and its complexity. The main goal of label switching networks was to bring those connection orientated benefits into a non-connection orientated network; mainly IP. While simultaneously overcoming the complexity problems of many WAN based technologies, especially ATM. ATM's connection orientated paths allowed for many benefits but brought a lot of unwanted baggage along in the process. Resource efficiency also took center stage. Companies wanted to make better use of their bandwidth so dynamic traffic control was also a driving force for developing label switching.

As I stated MPLS began as a bunch of separate proprietary approaches incorporating a form of label switching. In 1996 Ipsilon was the first company to really start building hype around a label switched network. They hyped their IP Switching and Router Cut-through technologies. These were basically ATM routers reprogrammed to run IP protocol with an ability to skip the hop-by-hop nature of IP (Gair, Chris 2007). Ipsilon also concentrated heavily on homogenizing the enterprise network from what was predominantly Cisco seated hardware. They were able to generate a lot of hype for their IP switching but weren't able to generate a lot of customers. Gigabit Ethernet was also nearing the end of its drafting stage so a speed increase alone was not enough of a reason to adopt the Ipsilon technology. What they did generate though was an industry thirst for more intelligent IP routing. In 1997 Nokia bought the then struggling Ipsilon (Duffy, Jim 1997). Even though Ipsilon was the first company out the door with a form of label switching, it was Cisco and IBM's designs that most of the MPLS specification was drafted from. Both IBM and Cisco came to the market with proprietary label switching technologies in late 1996. IBM called their system ARIS (Aggregate Route-based IP Switching), and Cisco called theirs Tag Switching. Both systems used the same signaling technologies, and both used network topology information for the packets path determination. Where they were different was ARIS was mainly designed to run on top of an ATM network thus it focused on ATM specifics, such as virtual channels. Tag Switching was built to work generally with a mix of networks, it also allowed for labels to be assigned at any point in the network. ARIS only allowed labels to be assigned at the networks entry. The MPLS Working Group was formed on March 3rd 1997 to begin work creating a standardized specification for label switching, called Multi-protocol Label Switching. By the end of 1999 the MPLS Working Group had finalized specifications for both the signaling and encapsulation components across various layer-two transport technologies. Additional features such as VPN (Virtual Private Networking) and extra quality control functions became request for comment proposals (RFC) at nearly the same time.

The MPLS Working Group has shown a much more prudent approach to standardization than the ATM group had. An example would be in its signaling protocol decision. Both Reservation Protocol (RSVP) and Label Distribution Protocol (LDP) can be used as a signaling protocol with MPLS. The group favors RSVP; however they did not simply make it the standard and call it a day. Realizing that organizations have implemented LDP and some may even prefer it they continue to support LDP as a signaling protocol. While they won't continue standardizing LDP enhancements, they will support compatibility while focusing on RSVP for the future. This strikes an excellent compromise to not leave current LDP deployments out to dry while still keeping a future focused on their protocol of choice (Andersson L. & Swallow G., 2003). As you can see MPLS's history isn't quite as colorful as ATM's was and these types of differences may be important in the decision making process.

Ethernet

Before diving into the technical aspects of MPLS I think it's important to have a general understanding of Ethernet technologies since it has an impact on both MPLS and ATM. Ethernet started out as an experimental LAN technology by Xerox in the 1970's using the carrier sense multiple access collision detection (CSMA/CD) protocol. After much success used internally at Xerox, Intel and Digital Equipment Company formed a partnership to ratify the first 10mpbs Ethernet specification IEEE 802.3 in 1980. Ethernet is a frame based technology meaning that each data packet can be of variable length; in Ethernet's case the payload can be anywhere between 46 to 1500-bytes. The Ethernet encapsulation consists of a 24-byte header and a 4-byte footer which contains the error check information.

Encapsulation order left-to-right

PRE	SFD	DA	SA	Length/Type	Data Pad	FCS
7	1	6	6	4	46-1500	4

Length in bytes

PRE = Preamble
SFD = Start-of-frame delimiter
DA = Destination address
SA = Source address
FCS = Frame check sequence

Figure 4. The encapsulation components of an IEEE 802.3 Ethernet frame.

The first 7-bytes contain the preamble which tells the sending node that an incoming frame has arrived. The next byte is the start-of-frame delimiter which signals that the next portion of the header is the destination address and always ends with two consecutive 1-bits. The destination address is the next 6-bytes followed by the source address also 6-bytes in length. These two addresses are the MAC (Media Access Control) addresses of the destination and source nodes. The next section specifies the length of the payload and is four-bytes in length. The data payload is next which will consist between 46 and 1500-bytes. If the actual payload is less than 46-bytes then the remaining space is filled with random bits until the minimum 46-byte payload is reached; at the destination these bits are then discarded. The last 4-bytes are the frame check sequence which contains the CRC data to ensure that the payload was not corrupted in transmission. As mentioned, the recovery protocol CSMA/CD recognizes packet collisions on the network and handles the retransmission of those packets. If any nodes transmit onto the network at the same time they will then broadcast to the network that a collision took place. Each of the offending nodes then selects a randomly specified time to wait before resending the packet (Davis, Leroy 2008). Frames are sent along the network in a hop-by-hop fashion. They travel to each routing device in the network and perform a routing table lookup. The destination address in the Ethernet header is then updated to the next routing device's address and continues along the path in this fashion until the final destination is reached. Ethernet operates at the physical and data link layers of the OSI model. With TCP/IP being the dominant protocol stack to handle the network and transport layers. Ethernet has proven to be a very scalable technology; originally drafted to operate at 10mbps, it's seen increases to 100mbps, 1000mbps (Gigabit Ethernet) and 10000mbps (10Gigabit Ethernet). Where Ethernet has lacked more intelligent retransmission technology or QoS options it makes up for in raw throughput. With long roots in the networking world Ethernet is not only time tested and very entrenched in the market but its hardware costs tend to be lower than many newer technologies at their arrival. ATM was one such example and never could compete with Ethernet on price of hardware. It should also be noted that MPLS many times works in conjunction with Ethernet and doesn't necessarily look to completely replace it on a physical level.

Technical Workings of MPLS

Now that we have an understanding on how ATM works and some background on Ethernet we can start examining the technical workings of MPLS. When looking at how a Multi-Protocol Label Switched network functions, it's helpful to keep in mind a few main points to the technology. MPLS is considered to be one of the next generation networks in that it is aimed at being more efficient and more flexible in the way it routes traffic. Efficiency and flexibility were the two main goals in creating this network. As mentioned earlier speed increases were not the main driving force behind establishing label switching as a standard. This is an important aspect for anyone evaluating the technology to keep in mind.

Multi-protocol label switching attaches labels to data packets to route them. The path an MPLS packet takes through the network is called the label-switched path (LSP) where labels can be attached and detached at each switch in the specified network path. Label-switched paths are either control-driven where the path is determined before transmission or data-driven where the path is established according to a certain flow of data. Labels are distributed using one of three protocols, label distribution protocol (LDP), resource reservation protocol (RSVP), or are sent out using protocols that have been previously established for other networks such as open-shortest path first (OSPF) (IEC 2007). The switching of MPLS labels is very fast since the label is a short fixed length of 32-bits and is near the beginning of the packet.

MPLS is often referred to as a 2.5 layer protocol because it doesn't fit neatly within the seven layer OSI model. MPLS labels can generally be referenced as sitting between OSI layers-two (data-link layer) and three (network-link layer) or embedded in the header of layer-two (Cisco Systems 2007). Being such MPLS can work with almost any sort of traffic as labels can be adhered to any sort of packet, but is almost exclusively used for IP based traffic. Generally speaking labels identify the network path data takes through the network. As a data packet enters the network a layer-two label encapsulates the packet that will signal each router on the path the label must travel. An MPLS label is made up of 32-bits (4-bytes), so it adds a smaller amount of overhead data compared to other protocols, including ATM whose header is 48-bits long.

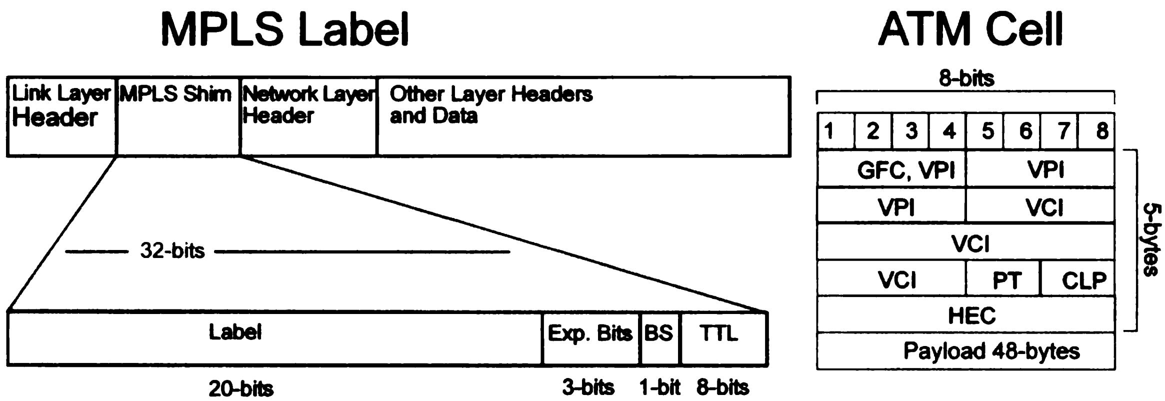

Figure 5. The structure of an MPLS label between OSI layers two and three.

At a glace the first thing you notice is that the structure of an MPLS label is a lot less complicated and takes fewer bits than an ATM cell. The first 20-bits make up the label information; this will contain the path the labeled packet will take through the network. The next three-bits are the quality of service identifiers which are also called the experimental field identifier. The end of stack bit is next, this may or may not be present in the label depending on if a label stack is being used. Label stacking is when a packet has multiple routing labels attached; the last label in the stack will contain the last stack bit identifier. By using multiple labels MPLS can build an MPLS domain hierarchy between connected networks. The last 8-bits are the time to live identifier, if a label does not reach it's destination within the set value it will be discarded at the next router. Which labels are attached to which packets are determined by a process known as forwarding equivalence class (FEC). This is a designated group of packets that have the same path requirements for traveling the network (IEC 2007). All packets in a class get the same treatment as they travel the network toward their destination. Forwarding equivalence class's importance is more apparent when you look at how MPLS packets travel through the network which I will explain in greater detail later on.

With any network an important consideration is how the physical layout of the network functions. This is especially true with connection orientated networks where the physical and logical topologies share a strong dynamic. Although MPLS is not a connection orientated protocol its physical layout is still very important in evaluating technologies. Within an MPLS network exists three types of label switched routers (LSR); ingress LSR, transit LSR, and egress LSR. The role each of these routers play is dependent on their location with regards to the data being sent. An ingress router is the first router a packet will encounter when it enters the MPLS network. The egress router is the last router encountered while exiting the network. Ingress and egress routers are also often referred to as edge routers since they comprise the outer boundaries of the network. Transit routers are simply label switched routers a packet will pass through as it moves through the core of the network (Riverstone Networks Inc., 2007).

Now that we're familiar with the components of an MPLS network we can look at how a packet travels through it and how these components all come together. When a packet enters the MPLS network it first arrives at the ingress label edge router. Here the packet is examined and is assigned a forwarding equivalence class. One of three methods is then used to send the route data for this new FEC to each routers label information base (LIB) to specify the path. Resource reservation protocol (RSVP) is the preferred protocol to distribute FEC information, however as mentioned label distribution protocol (LDP) can also be used or the route information may be piggybacked on an existing protocol such as border gateway protocol (BGP) or open shortest path first (OSPF). MPLS gains a large advantage over traditional IP routing with this method. Normally in an IP routed network as each packet travels over the network an IP table lookup must happen at each router along the way and a path determination must be calculated. Packets are then forwarded to the next router that satisfies a link towards the destination. This is referred to as hop-by-hop routing and is really a best effort service focused more on recovery than performance. With MPLS, once a path has been broadcast for that FEC each packet encapsulated with that FEC will take the same route through the network unless an alternative path is once again broadcast. MPLS spends a lot less time doing table look ups and path determinations. After the route data broadcasts the path, the ingress router will begin forwarding the packets to the specified label switched router. It may also be forwarded to a label edge router as those are also used for forwarding inside the network and not just as entry and exit routers. Each time a packet is received by a router it looks up the rule for that FEC in the LIB and one of three things happen. The router either imposes a swap, push, or pop function on the data packet. If a swap is the required action the router simply removes the packets label, attaches an updated one, and forwards the packet to the next destination. If the router imposes a push function than the existing packet and label have another label attached to form an MPLS label stack. Finally if the router is to use a pop function that means the label is removed and the data packet will exit the MPLS network (IEC 2007). However, if all data packets need to be routed through one egress router to pop the label and exit this would not only create a single point of failure but also create a performance bottleneck in the network. MPLS compensates for this shortcoming with an operation known as penultimate hop-popping. With Penultimate hop-popping the last label switch router can also remove the label prior to reaching any egress routers (Riverstone Networks Inc., 2007). The data packet then is able to just pass right through an egress router without any sort of lookup or label operations happening before it exits the network. This of course will greatly help load balance the network during times of heavy activity by allowing any last leg routers to perform the processing functions of the egress router.

Since MPLS can control a packets path through the network it allows for a much more refined control over quality of service and enables a depth of control over class of service. MPLS provides quality and class of service when packets enter a label switched path in two ways. The first method is called label inferred label switched paths (L-LSPs). Here the quality of service is determined by the forwarding equivalency class information. When the label switched path is created all packets entering that tunnel will be treated with the same class of service. The second method is to use the experimental bit information in the MPLS header. This field can be used to identify different classes of service to treat a packet; packets that are sharing the same tunnel can be treated differently depending on this set value. This method is called experimental bit inferred label switched paths (E-LSPs) because the quality of service information is being identified by the experimental bit section of the header (Bayle, Aibara, Nishimura, 2001). With this two-fold approach an MPLS network can institute a high quality of service level granularity.

MPLS can also be used to create what has been termed as network based VPN's. A VPN (Virtual Private Network) is a way to access remote network applications and services securely by using methods such as authentication, access control, and encryption many times over an unsecured network such as the internet (RSA 2001). Since MPLS can control a packets path, VPN's can be setup by allowing only that specified traffic to run over an established LSP residing on a completely secure infrastructure. MPLS can also use label stacks to hide all intermediate information about the network if a packet has to traverse a public network to create a VPN tunnel. Take for example two separate networks both comprised of ingress, egress and label switched routers. As the LSP is created the ingress router knows that label edge router three (LER3) required to enter network two is the destination and creates a LSP. Sitting between both ingress and egress routers may be a myriad of LSR's which a second LSP is created to traverse. When the packet arrives at the egress router or LER2, it will have its second LSP stripped since it's no longer needed. The packet is now left with the original LSP telling LER2 to route the packet to LER3 or the next ingress router of network two. When the packet arrives at LER3 the original LSP is removed and LER3 attaches a label for a LSP that will take the packet to LER4, the final destination. By using this label stack method not only can MPLS control preferred routes through the network to create a secure VPN, but the packet arrives at network two without any of the switching dynamics of network one still present.

Analysis

So far this paper has explained two of the three elements in the evaluation framework, history and technical workings. Now we need to take a look at what this knowledge actually means and we can determine both of these technologies strengths and weaknesses. Examining the first two framework elements in this way will highlight if any technology is in alignment with your needs. This is an important step, a technology might be superior in a greater number of categories but that however does not mean that it will necessarily be the right choice for you. So you have to be able to analyze what the data means. In addition the final frame work piece; market data, will help paint a complete picture. In the technology field a certain level of adoption and support must be maintained with any technical choice to be successful. We say how ATM for the LAN ended up quickly becoming vaporware, dead end product life cycles will always be a bad choice even for early adopters.

ATM setout to achieve pedestal status; created as the alternative to running a network over the PSTN its introduction was to be a turning point for new networks. ATM's main goals where to deliver a much higher bandwidth than ISDN, run converged voice and data, introduce quality of service features, and become the de-facto backbone network. While ATM was able to deliver on most of these goals it did so while introducing a lot of unwanted baggage into the process.

When introduced the 155Mbit/sec throughput ATM provided was much faster than ISDN BRI's 128kb/sec throughput. ATM was originally designed to be overwhelmingly faster than ISDN; it was also significantly faster than any of the T-carrier networks of the time. However ATM's cell design proved to be a performance bottleneck for transmitting data, mitigating its raw throughput. Since ATM transmits in fixed length cells, data must be continually broken up into 48-bytes; transmitted and then reassembled on the other side with a cyclic redundancy check (CRC) run on the data to be sure it was assembled properly. This process of splitting and recombining the data is called segmentation and reassembly (SAR). With a payload size of only 48-bytes there's a lot of segmentation and reassembly that needs to take place. Since ATM's introduction network throughputs have continued to expand at impressive rates. ATM today is able to reach a throughput of around 622.08 Mbit/sec, or roughly the speed of an OC-12 (Optical Carrier) network (Jaeger, Rob 2001). However the factor that limits ATM to this speed is the heavy amount of SAR processing that must take place. Currently the fastest SAR processing chips cannot operate at a speed any faster than around 622.08 Mbit/sec. This limitation of the SAR's processing not able to progress as fast as the physical line speed is one big factor that stopped ATM from becoming a champion for backbone applications. The Ethernet specification for example has scaled well beyond ATM's rate to 10Gbit/sec, nearly sixteen times faster (Eisenberg 2003).

Having the ability to carry both voice and data was another one of the main focuses behind creating ATM. While ATM was able to achieve this goal the solution simultaneously became its greatest weak point. ATM's design took two main approaches to ensure clean voice transmission; both however were based on short sighted principles. The first was to use fixed size cells to reduce jitter in voice transmissions. Not only was the 53-byte cell not ideal for voice but as I mentioned previously the amount of SAR that needs to take place wreaks havoc on data transmissions. The designers of ATM also never gave thought to what it would take for multi-sized packet networks to overcome these jitter hurdles. The original concerns with running voice over a multi-sized packet network such as Ethernet was delay. Using a network which tends to be bursty may take too long for a large packet to reach the receiver; this would end up as silence on the receiving end. However it didn't take long for multi-sized packet networks to evolve to the point where their increase in speed basically strong armed this problem. Providers can now not only run voice over multi-sized packet networks but can also implement the much more preferred Internet Protocol for addressing. The second approach ATM uses for voice communications are virtual channels. Establishing these dedicated links enables a quality voice connection to take place over a mixed network. Bandwidth needed to maintain the connection is allocated and the virtual channel is dedicated to that voice connection. ATM can guarantee a high quality voice service in a mixed network, however once again this approach brings a lot of unwanted complications. Establishing a myriad of dedicated connections is not the cleanest approach to network design. Again establishing a circuit like network was something the telecom engineers were comfortable with, but this method has significant problems when networks start scaling to large sizes. In an ATM environment every connection is dedicated end-to-end. As networks scale larger this constant increase in the number of dedicated links is a messy way to network. For example in a network consisting of four routers six VC links need to be created to have a complete network. As networks scale this problem is continually compounded and makes an ATM network overly complex in comparison to an IP network.

Quality of service is one area where ATM excels. Being that connections are dedicated virtual channels, the ability to guarantee a certain amount of bandwidth is one of ATM's strong points. ATM supports five quality of service levels; they're mostly characterized by how detrimental delay is to that type of activity. ATM does do a good job with maintaining quality of service levels over the network. However the downside is once again the extra complexity that virtual channels bring into the network architecture. In addition while five quality of service levels may be enough for certain applications by today's standards they're simply not that many. One of the main reasons MPLS was created was an industry thirst for more QoS. As a decision maker this is an important factor to consider when you look at your long term goals.

Asynchronous Transfer Mode never did reach the status its original standards committee envisioned it achieving. This was due to the combination of the extra baggage that it brought to achieve its goals and simultaneously other market factors. Ethernet based networks became so fast so quickly that their speed filled in for the shortcomings that ATM was meant to tackle. In addition these networks were not only easier to deploy but the equipment was cheaper than ATM's. In the wake of ATM's problems engineers started looking at tag switching technologies. Many of ATM's shortcomings were focal points in creating MPLS.

So how does MPLS perform; does this newer technology complete the gaps where ATM falls short? The biggest goal of MPLS was to bring the benefits of connection orientated networks such as ATM, to a connectionless IP network. Advanced quality of service features were also a major goal, as was the ability to establish VPN connections.

It's no secret that IP is the time tested and overall preferred network layer technology for most applications. Although the problem with mission critical applications such as voice has been that packets can take different routes through the network and arrive at the destination at different times. In the case of a low-latency dependent application such as voice, waiting for packets to catch up is disastrous. ATM approached this problem by having dedicated virtual channel connections; however this introduces other problems into the mix as I mentioned previously. MPLS however is aimed at having the packet priority of a connection based network while running over a connectionless based IP network. It's able to achieve this balance by use of label switching. A group of high priority packets such as voice can be grouped with the same FEC to travel the MPLS network through a fast available route; ensuring that packets arrive on time and in order. This is different than creating a virtual channel connection because these are a series of routes that exist in the network and once the transmission ends the path will be relinquished to the network to be reused. This also allows for a much more efficient use of network bandwidth since there won't be dedicated links sitting idle. Nick Kwiatkowski, a voice and video engineer for Michigan State University believes that to the end customer MPLS will feel no different than a leased circuit line, however to the Local Exchange Carrier MPLS is much more efficient. He states "Within the cloud, it is more efficient than ATM, by far. By only allowing TCP/IP encapsulated data, they are able to better share resources, as IP packets have a much more customizable payload than ATM datagram's (Kwiatkowski, Nick 2008)."

MPLS is also designed to be a much more flexible networking technology. MPLS is generally considered to sit between layers two and three of the OSI model and is commonly referred to as a layer 2.5 technology. This means that MPLS is not bound to one network type and can be incorporated into a variety of networks, even ATM. Any packet that enters an MPLS edge router simply gets encapsulated in a label; travels the network, and the label is removed by the egress router. MPLS can actually be used as a crutch for ATM networks, attaching labels to cells and sending them long haul over an MPLS network. There wouldn't even be any need to use any of ATM's adaptation layers. Of course this still would not solve the SAR problem, however it would mitigate the scaling problems created by virtual channels and it illustrates MPLS's ability to integrate with mixed networks.

Quality of service is something that new networks won't succeed without. This is one area where ATM shined with its connection orientated nature and was a corner stone in the MPLS design. MPLS takes a different approach however. Instead of outlining a set number of defined QoS levels like ATM does, MPLS can be more flexible potentially using a blend of QoS approaches. The most common is the L-LSP using label paths to group packets with a common priority. It's a simple concept although it allows for a wide range of control. Instead of having predefined QoS connection types, a baseline for priority is established in comparison to other traffic types in the network. In addition MPLS can use the experimental bit field in conjunction with L-LSP. This allows for packets inside each LSP to have three additional levels of priority. You can see how the MPLS forum took QoS into account when developing the standard. This two prong approach allows for a much more granular QoS control than ATM's five defined service levels. However as I previously mentioned QoS continues to be an important component engineers want to see in networking. Even with the several advances MPLS took over previous protocols the industry continues to want an even higher level of service in this respect. William Copeland an engineer for Verizon feels that the QoS level in MPLS is still not enough. When asked if MPLS brings a high level of QoS to IP networks he responded "no, the 3-bit EXP field should have been 6-bits to match DSCP" (Copeland, William 2008). DSCP stands for Differentiated Services Code Point and is a method of using flag bits in a standard IP header to enable service discrimination. Discrimination can happen on either a peak bandwidth level or on a service class level and has been in use for several years (RFC 2474, 1998). The issue that William Copeland points out is that DSCP flags consist of 6-bits, double the length of MPLS's EXP field and so it can't neatly fit into an MPLS label. Although DSCP can be integrated into MPLS by either putting the 6-bit flag into the EXP field using data transforms to reduce it to 3-bits or by copying it into the Label Information Base. Either way this takes some special configuring and operations to happen along the way. Had the EXP field been 6-bits it simply could be dropped right into that field (Welcher, Peter 2000). With DSCP already in use prior to MPLS this may have been an oversight by the MPLS forum when designing the protocol. As a decision maker your environment may already be using DSCP and this non-native compatibility may be a deal breaker for you. Robert McGowan who works for Converged Network Solutions has a different but related viewpoint on QoS over MPLS. He states "There is a significant lack of expertise among the carriers about QoS, how it works and how to implement, administer and troubleshoot WAN's with QoS. As the technology and carrier expertise align it has promise" (McGowan, Robert 2008). Such is the case with many new technologies there is usually a learning curve to adjust to, which is an important factor to remember when choosing technologies to implement. As William points out MPLS still has not fully satisfied QoS demands for the industry and I believe this will continue to be a strong focus for new networking technologies. However as Robert also pointed out some of the inadequacies of QoS with MPLS has to do with the newness of the technology. This may be overcome in the relatively short term as knowledge increases on deployment and management, or it may leave an opening for a competing technology with simpler integration to step in.

Connecting disparate offices and users has become an increasing concern with businesses. As the everywhere office grows companies have continued to equip branch offices and mobile users' access to the full suite of network tools. This is accomplished with Virtual Private Networks (VPN). VPN capabilities are making their way into next generation networks; and this is a large selling point for MPLS. In the past dedicated Frame Relay networks have been the most popular choice for running a VPN between offices. Frame Relay typically has a cheaper price tag and is relatively easy to configure.

Another method is to use a public infrastructure like the internet and encrypt data using the IPsec protocol. This is however a process intensive approach having to encrypt and decrypt packets at each end. MPLS can use LSP's to create VPN connections. In fact many people refer to MPLS as a network based VPN or netVPN. MPLS can build label switched paths within its network to only allow specified packets to travel across secure provider equipment. This allows a provider to setup a virtually dedicated connection similar to running a dedicated frame relay connection. However MPLS allows the provider to use label switched paths to run traffic through secure tunnels over more of its general infrastructure like an IPsec design. Providers can use their networks more efficiently reducing cost. In addition the VPN will be running over an IP based network but you will have the performance of a dedicated line without the delay of encrypting and decrypting packets. MPLS promises to provide VPN's that are cheaper, fast and scalable; so far MPLS's VPN capabilities have been one of its biggest selling points.

It's obvious that MPLS is taking a new approach to network design and can fulfill a lot of promises that ATM could not. How are these two technologies fairing in the market today? ATM has been around for a long time now and MPLS is still fairly new. Even with it's problems and it's failure to gain the adoption rates once envisioned, ATM is still in wide use with switch revenues of around five billion dollars in 2000 (Infonetics 2007). However with newer technologies able to fill ATM's role more effectively such as Gigabit Ethernet and MPLS arriving in 1999 we are seeing a sharp decrease in new ATM deployments. According to Xilink an ATM switch provider, worldwide ATM switch sales fell to $4.1 billion in 2001. ATM switch sales continued to decline by 6% in 2003 and by an additional 3% in 2004 (Moskalyuk, Alex 2004). Infonetics reports that ATM switch sales took an even steeper decline in 2005 with a decrease of 33% to $1.3 billion in sales for the year. It's obvious that new ATM deployments are almost a thing of the past and legacy systems will start to wane as tag switching networks and Gigabit Ethernet continues to replace them. In fact IT Jobs Watch tracks demand for technology based positions in the United Kingdom and is showing a steady decline in the number of ATM related career positions. Since October of 2006 to 2007 they measured a 26% decline in permanent jobs relating to ATM technologies.

However even with ATM on the decline there are a lot of legacy networks in place, and providers will continue to look for ways of keeping them useful until they absolutely need to be replaced. One area where ATM has been found to be advantageous is in use as Digital Subscriber Line (DSL) backbones. Here ATM's speed is enough to handle the load; even with SAR limitations on scaling there's enough headroom to be a viable solution for quite awhile. Usually ATM's circuit like design makes it difficult to work with, but DSL is also a circuit based technology and ATM integrates smoothly with it. Each end user is already connected by a dedicated line VIA the PSTN. ATM channels only need to be established between DSLAM's (Digital Subscriber Line Access Multiplexer), keeping the number of virtual channels much lower. ATM's Quality of Service features also make it a good choice for DSL. Network engineers can specify virtual channels with the required bandwidth running to multiple DSLAM's. One of the largest pitfalls in using ATM to feed DSL networks is once again DSL is an IP service and ATM is not, so adaptation layers must be used. Obviously using an all IP network for

DSL would make things simpler. However previous to label switching no IP based networks had the QoS capabilities that ATM had to ensure each DSLAM was partitioned with enough bandwidth for end users. MPLS may displace ATM as the de facto DSL backbone and keep the entire network IP based. For now though many of these ATM backbones are already in place so not only are they cost effective but ATM is integrating into DSL systems much smoother than it did into past roles. It seems ATM has found its niche serving not as a high speed LAN technology or a colossal multi-service backbone but as a feed for multi-node narrowband applications like DSL.

With MPLS being a fairly new technology it obviously hasn't reached a high saturation point and it is safe to generally say that MPLS is fairly expensive compared to other older technologies. However the time could not have been better for MPLS to arrive. IP has proven to be the all encompassing addressing protocol network engineers want to be using. However for the most part the low QoS connectionless networks provide hasn't allowed engineers to keep networks entirely IP based, until MPLS. IT Jobs Watch has shown a steady increase in MPLS related jobs since 2004. One of the biggest enablers of MPLS growth this early on has been its VPN features. As more and more businesses rely more heavily on their VPNs buying into a network that has VPN services built into the fabric is attractive. Especially with intensive applications like voice and video being increasingly run over these VPN's; an MPLS network will take heavy encryption loads off of servers and provide a better QoS. Business Wire reported that in 2004 network based IP VPN equipment sales totaled $347 million worldwide and estimated revenue of $658 million for 2009. That's nearly a ninety percent increase in five years.

MPLS certainly seems positioned to be the next big networking protocol and may even attain a status on par with Ethernet. At this point it's pretty obvious that IP based networks is the direction that everything is heading in. MPLS is essentially delivering the QoS levels for the first time that will allow engineers to design IP based networks from end-to-end. Add VPN features built into the network and MPLS might seem like the perfect network solution, however nothing is perfect. Some experts claim that MPLS's VPN features, one of it's strongest selling points thus far, is not all it's cracked up to be. One aspect being questioned is the fact that data is not encrypted running over the VPN. While the lack of having to process encryption makes for a speedier connection most security holes are related to human error. If the provisioning is done wrong with the VPN and the route is not secure the data being transmitted is vulnerable. The biggest criticism however is aimed at RFC 2547 which outlines a VPN method using Border Gateway Protocol (BGP) tables. BGP is a routing protocol meant to scale to very large sizes and it's the routing protocol the internet runs on. Normally an internet service provider would keep one BGP table that is associated with all links. However with RFC 2547 a BGP table would also have to be kept for each MPLS VPN connection. ISP's claim that running just one master BGP table is a daunting task, now they would have to run hundreds, possibly thousands of them. It's an argument mostly brought up by ISP engineers who bare the brunt of the administration. One approach to handle this problem proposed by Juniper Networks is to store the routing table on the customer premise equipment. This method is considered a layer-two VPN where as RFC 2547 is considered a layer-three VPN. The layer-two method has been received by engineers as a better option since it puts less duties on the ISP's and gives the customer more control while still being less demanding versus running an IPSec VPN. Although as a customer, if you wanted a completely hands off VPN, RFC 2547 would still be your preference. Neither VPN addresses the concern of non-encrypted data. However some engineers argue that an MPLS VPN is just as secure as a dedicated Frame Relay link which has generally been accepted by businesses. In the case of transmitting highly sensitive data as a small portion of your traffic you could still encrypt using a third party solution before transmitting.

Closing

Evaluating technologies can be a difficult thing for managers and engineers alike. However it is a very important process that will effect not only the day-to-day operations for years to come but also the abilities that your infrastructure is able to accomplish for several years. By examining potential solutions with this evaluation framework, decision makers will be able to determine a technologies strengths and weaknesses and decide if that product is in alignment with the organizations goals. While using ATM and MPLS as examples a lot of valuable information was uncovered that paints a whole picture that you're not likely to find in product marketing. While MPLS seems to be the newer more advanced solution compared to ATM we've found areas that both products thrive in. In addition there are several areas that MPLS might not be strong enough in to fit certain organizational needs. However, this research concludes that MPLS is on the rise while ATM simultaneously is in decline. The industry is embracing

MPLS and heading in a direction towards connectionless IP networks; while simultaneously moving away from connection based ATM. Joel Haist who works for Strategic Products and Services also agrees with this trend stating *"I definitely believe services like MPLS and SIP are on an upward slope overtaking older ways of connecting multiple locations together"* (Haist, Joel 2008). When examining new technologies to embrace, decisions makers will want to evaluate technologies through the three elements of this framework. Researching and evaluating a technologies history, technical workings and market adoption. This will ensure that decision makers choose confident and correct technologies that are in alignment with their organization.

Bibliography

Anderson, L. Swallow, G. "The Multiprotocol Label Switching (MPLS) Working Group decision on MPLS signaling protocols" Network Working Group, RFC 3468. February, 2003.

Becker, Ralph. "ISDN Tutorial" 5 December, 2006 <http://www.ralphb.net/ISDN/index.html>

Cassidy, Kyle. "Asynchronous Transfer Mode (ATM)." Inform IT. 11 May, 2001. < http://www.informit.com/articles/article.aspx?p=21316>

Cisco Systems. "MPLS FAQ for Beginners" Cisco Systems. 8, May 2006. <http://www.cisco.com/en/US/tech/tk436/tk428/technologies_q_and_a_item09186a0080 0949e5.shtml>

Cisco Systems. "Internetworking Technologies Handbook" Cisco Systems. 1992-2008. <http://www.cisco.com/en/US/docs/internetworking/technology/handbook/Ethernet.html >

Copeland, William "i030666 Survey Instrument" 2008.

Davis, Leroy. "Encoding Dictionary, Terms and Definitions" http://interfacebus.com 2 December, 2008. < http://www.interfacebus.com/Definitions.html >

Derbyshire, Arran "Why has communications evolved towards the ATM concept?" Surprise 96. 3 June, 1996. http://www.doc.ic.ac.uk/~nd/surprise_96/journal/vol2/arad/article2.html

Duffy, Jim. "Nokia catches a falling Ipsilon." Network World. 9 December, 1997. < http://www.networkworld.com/news/1997/1209ipsilon.html>

Gair, Chris. "The Rise and Maturity of MPLS" Technology Inside on the Web. 14 February, 2007. <http://technologyinside.com/2007/02/14/networks-part-3-the-rise-and-maturity-of-mpls/>

Gould, Jeff. "ATM's long, strange trip to the mainstream" Data Communications. June, 1994 pg. 120-130

Haist, Joel "i030666 Survey Instrument" 2008.

"Asynchronous Transfer Mode (ATM) Fundamentals" International Engineering Consortium. 2007 < http://www.iec.org/online/tutorials/atm_fund/>

"Multiprotocol Label Switching (MPLS)" International Engineering Consortium. 2007 < http://www.iec.org/online/tutorials/mpls/>

Jaeger, Rob "Transitioning from IP-over-LANE/ATM to IP/MPLS Networks" Juniper Networks. 2001. <http://www.juniper.net/solutions/literature/white_papers/200013.pdf>

Juniper Networks "How MPLS Works" Juniper Networks. 2007. <http://www.juniper.net/techpubs/software/erx/junose72/swconfig-bgp-mpls/html/mpls-config5.html>

Kwiatkowski, Nick "i030666 Survey Instrument" 2008.
Mahmoud, Sammy. Stern, Harold. "Theory of Time Division Multiplexing" National Instruments. 6 September, 2006 http://zone.ni.com/devzone/cda/ph/p/id/270

McGowan, Robert "i030666 Survey Instrument" 2008.

Moskalyuk, Alex. "ATM Switch Market to Delcine 3% in 2004." ZD Net Research. 13 March, 2004 http://blogs.zdnet.com/ITFacts/?p=5489

Network World "MPLS (Multi Protocol Label Switching)" Network World. 31, July 2007
< http://www.networkworld.com/community/node/18007>

"Asynchronous." *Dictionary.com Unabridged (v 1.1)*. Random House, Inc. 02 Apr. 2007. <Dictionary.com http://dictionary.reference.com/browse/Asynchronous>.

RFC 2474 "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers" International Engineering Taskforce. December 1998
< http://tools.ietf.org/html/rfc2474>

Riverstone Networks Inc. "Introduction to MPLS" Riverstone Networks Inc. 2007
http://www.riverstonenet.com/support/mpls/intro_to_mpls.htm

Welcher, Peter. "MPLS Part II" Chesapeake Netcraftsmen. September 9, 2000
< http://www.netcraftsmen.net/welcher/papers/mpls2.html >

Wood, Robert. "A Network Administrator's View of Multiservice Networks" Cisco Press.

9 December , 2005 <http://www.ciscopress.com/articles/article.asp?p=426645>