How Service Providers Can Provide Customers with Increased DSL Bandwidth and Reliability

Introduction

Service providers are anxious to differentiate their services by offering more reliable and higher speed broadband access services, regardless of the underlying infrastructure. ePipe's ML-IP solutions enable simple, low-cost provisioning of bonded broadband services with minimal infrastructure costs, when compared with IMUX, ML-PPP, ML-FR or load-balancing equipment.

Service Providers Today…

As a service provider, you aim to provide access to a service that an end user does not wish to provide themselves, or cannot afford to do so. Whether you are an Internet Service Provider (ISP) or an Application Service Provider (ASP), over the past decade your industry has been constrained by several factors, two of the most significant being:

Bandwidth, i.e. insufficient affordable bandwidth
Reliability, i.e. reliable communications at an affordable price

Of course, bandwidth and reliability are available for those customers willing to pay for it. The issue for service providers is how to provide an affordable product or service to the rest of the market.

The Bandwidth Challenge

The bandwidth challenge faced by many ISPs and ASPs is how to provide more bandwidth to or from their customers' networks within the limitations of the access technologies currently available. This is especially important for service providers that do not have control over the various access networks made available to customers by the telecommunications companies (telcos) or network access providers (NAPs). In practice this commonly means service providers are limited to deploying services using dial-up or DSL-based technologies as these can be supplied in a wholesale form to the service provider by the NAP. A telco on the other hand has a much wider choice of access technologies to offer the end customer and, therefore, the advantage of being able to offer products that scale in bandwidth and features, albeit at a cost. Such technologies include Frame Relay, T1/E1, ATM and higher data rate services.

The question for the service provider is:

How do I offer a product that provides bandwidth scalability using access technologies that are available to me, without having to involve the network access provider or telco?

Of course, this assumes the service provider is not also the telco or network provider who has access to the local loop. Even then, the question is still relevant and can be asked this way:

How do I offer more bandwidth to a cost sensitive market segment using access technologies the target customer can afford?

DSL (Digital Subscriber Line) access technologies, in particular Asymmetric DSL (ADSL), are proving increasingly popular to both home and business users as they are relatively inexpensive, easy to install (at least for the end user) and it runs over the existing copper phone network. DSL services are also popular because they offer speeds that were once only available to larger corporations; the market traditionally serviced by the telcos.

Service providers can offer ADSL services to their customers through the network access providers, who sell wholesale access to their own DSL networks. Thus, like dial-up Internet access, end users can gain access to the Internet using ADSL from a wide variety of service providers.

So, where is the 'fly in the ointment'?

Consider what happens when the customer requests more bandwidth than a single DSL link can provide. For many service providers the answer to the customer is 'we can't help', perhaps forcing the customer to go back to the telco. There are other options but these are usually significantly more expensive.

How do we solve this problem? Before we answer this, lets consider the issue of reliability.

The Reliability Challenge

Reliability, at its simplest, is all about ensuring continuous access to the information you need. For the end user connecting to an ISP, this means having always-on access to the Internet so that inbound or outbound traffic is always transmitted as fast as the service allows. For the ASP, this means ensuring their customers can access their applications when and how they want with bandwidth to ensure the application is useable.

A well understood way of achieving this in the network is via redundant paths, however this has the disadvantage of having expensive network infrastructure lying idle for most of the time. A better way would be to have multiple paths to the network (e.g. Internet or hosted network) and load balance the data across these paths. That way if a network path suffers an outage, there is another path for the data to take.

Solutions

As we have already mentioned, load balancing is one way of solving the problems of too little bandwidth and insufficient reliability.

The problem with most load balancing technologies is that they only work at the application level. In other words, they usually have to be application aware. A better way of ensuring increased bandwidth and reliability using multiple paths is via link aggregation or bonding. This has the advantage of working at the frame level and so is transparent to the application. An example is ML-PPP (Multilink Point to Point Protocol) which is commonly used to bond together dial-up (analog modem or ISDN) connections to increase bandwidth. Another example is ML-FR (Multilink Frame Relay). While these technologies are useful, they have the disadvantage of being tied to the infrastructure. ML-PPP only works point to point and requires the network access provider, not the service provider, to support it in their network. For example, in an ADSL network, the customer and network provider would both require compatible ML-PPP equipment to support bonding of multiple ADSL links. The same is true for ML-FR.

A better solution would enable multiple links to be aggregated together between the end user and the service provider itself. This has the advantage of providing the required bandwidth and reliability improvements while not requiring the network access provider to provision any equipment to support this. Only the end user and the ISP or ASP would need equipment to terminate each end of the aggregated connection.

But how can this be done across an unknown network infrastructure between the end user and the service provider, or at least an infrastructure over which the service provider has no control?

The answer is to aggregate the data logically or virtually at the IP layer and not the physical or link control layers. ePipe has called this technology Multilink IP™ (ML-IP). Figure 1 illustrates the different aggregation technologies and at which layers of the 7-layer OSI reference model they operate.
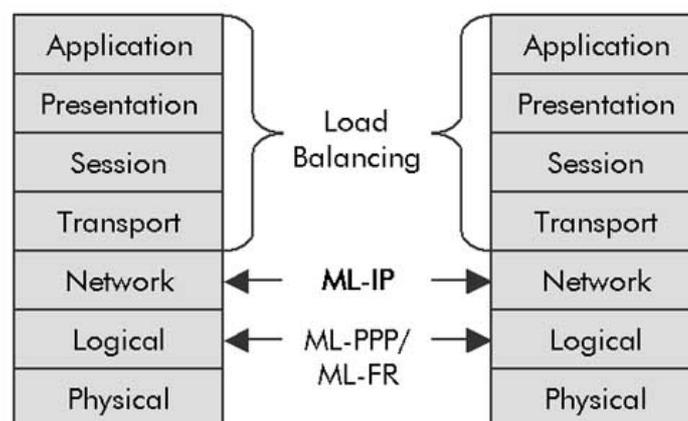

Figure 1 - Comparison of aggregation technologies using the OSI 7 layer model1

We will now explore the opportunities for service providers to use ML-IP in order to address customer demands for more (and more) bandwidth and increased reliability, thus allowing the service provider to broaden their market penetration. In particular we will look at how ML-IP can be used to address the bandwidth scalability issues that service providers have when developing products for their customers.

What is ML-IP?

Multilink IP (ML-IP) is multi-linking or aggregation of IP data streams, just as multilink PPP (ML-PPP) aggregates point-to-point data streams across analog modem, ISDN and ADSL connections. The key difference is that ML-PPP works at the PPP layer (data link layer or layer 2 of the OSI model) and is, by definition, a point-to-point technology; i.e. no intermediate networks/devices can be situated between the two ML-PPP capable devices. This limits ML-PPP to applications where distances between the ML-PPP devices are relatively small. In the case of ADSL, the network access provider would need to support ML-PPP in the DSLAM.

ML-IP allows traffic traveling across different network paths between two ML-IP gateways to be aggregated so that the benefits of increased bandwidth and link redundancy can be realized while being transparent to the applications using the network and the network infrastructure. As the aggregation is achieved end-to-end between the ML-IP gateways, the network access provider and/or other intermediate networks do not need to support the ML-IP technology. ML-IP can do this by aggregating at the network or IP layer, which is independent of the underlying physical and logical network transmission technologies.
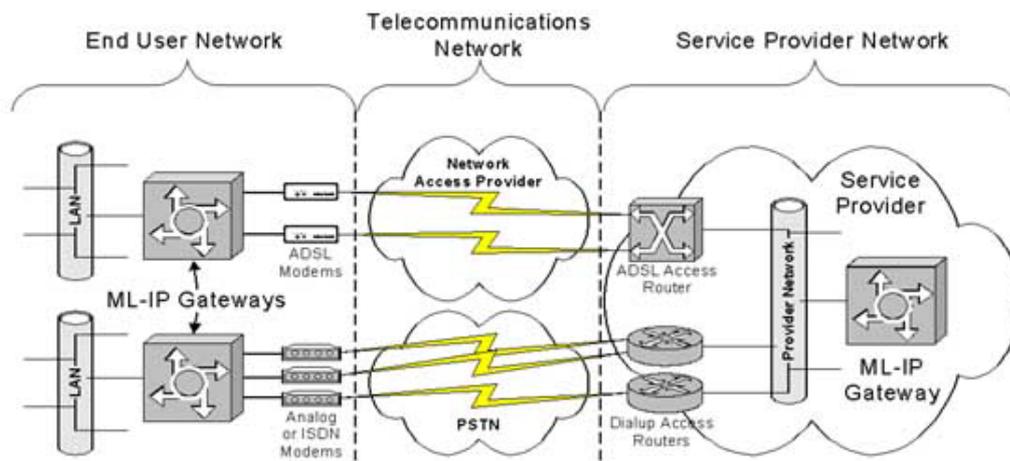

Figure 2 - Examples of ML-IP bandwidth aggregation

Each ML-IP gateway can be connected to the intermediate network using one or more connections. Each connection can be to the same or different access devices or routers. For example, you could connect one ML-IP gateway to the network access provider via a pair of ADSL services, which would most likely be terminated at the same DSLAM. Or, in the case of dial-up Internet access, the ML-IP gateway could have several dial-up analog modems dialing into different access routers within the ISP. Once again, this is possible as the aggregation occurs across the intermediate network to the other ML-IP gateway and not to the access router. Figure 2 shows these examples. A pair of end user networks is depicted connecting to a service provider (e.g. and ISP) via two ADSL or three analog/ISDN modems. Regardless of the intermediate network, ML-IP will aggregate the bandwidth to provide bandwidth scalability and increased reliability.

The initiating (or client) ML-IP gateway creates virtual channels to the receiving (or server) ML-IP gateway across each physical link connected to it. These channels are re-combined at the receiving ML-IP gateway. As these are virtual channels, each ML-IP gateway does not need to have the same number of links. For example, two networks could be connected over the Internet by establishing a ML-IP tunnel between one ML-IP gateway with 3 ADSL links and another with a single high speed connection.

The Benefits of ML-IP

Benefits for the Service Provider

There are a number of advantages for the service provider that supplies ML-IP based services to its customers. These include:

- A wider product range with broader market penetration
- Customer retention (decreased churn rates)
- Infrastructure independence (ML-IP not dependant on the network access provider)
- Customers see improved reliability

Benefits for the Customer

The customer also sees benefits when connecting to a service provider that supports ML-IP. These include:

- Increased bandwidth at an incremental cost rather than a significant cost increase through the migration to (usually) much faster but also much more expensive services.

- Improved reliability through the use of multiple connections means a service interruption with one service does not result in a complete service outage.

ML-IP in the Service Provider's Network

A service provider that supplies DSL services to its customers will be connected to one or more network access providers. The device that provides this connectivity is referred to as the Broadband Access Server (BAS).

The BAS connects the ADSL customers to the service provider's own network, allowing the customer to access the Internet, hosted servers or hosted applications.
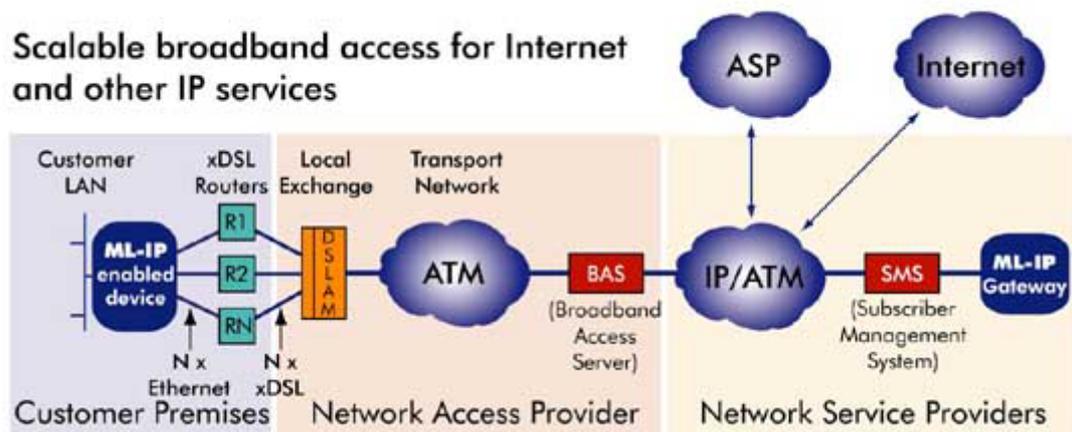
To enable bandwidth aggregation from the customer to the service provider, at least one ML-IP gateway will need to be located in the service providers network. This enables inbound virtual channels from the customer's ML-IP gateway to be bonded together. Exactly where the ML-IP gateway is located will depend on the design of the service providers network. However, in general terms, the ML-IP gateway needs to be located in between the BAS and the Internet router or hosted application server.

The ML-IP gateway within the service provider's network is typically a server class Linux-based PC with ML-IP Concentrator software installed. The system requires a 100 M Bit/sec Ethernet interface, which terminates fragments from ML-IP, bonded connections and forwards aggregate packets to the main Internet router. An additional 10/100 M Bit/sec Ethernet interface can be used for local administration.

Summary

ML-IP represents a huge opportunity for service providers to market reliable and higher speed broadband access services whether they are a broadband provider or not. ePipe Pty. Ltd. provides CPE devices and ML-IP software to enable simple, low-cost provisioning of ML-IP.

Providers can create 5 Megabit/sec, 10 Megabit/sec and 20 Megabit/sec broadband access services based around ePipe's CPE ML-IP gateways (complete ePipe 2344 gateways or custom Linux appliances) and high-end Linux servers as ML-IP concentrators.
Providers' investment in infrastructure compared with IMUX, ML-PPP, ML-FR or load-balancing equipment is miniscule. The ML-IP technology is simple to understand, easy to deploy and can be up and running very quickly.