Market Clarity

Australia's Source for Telecommunications Intelligence

Demystifying Layer 2 and Layer 3 VPNs

Version 1.0

15 June 2006

Market Clarity Pty Ltd
Suite 1404, Level 14, 33 Bligh Street
Sydney NSW 2000 Australia
Internet: www.marketclarity.com.au
Email: info@marketclarity.com.au
Phone: (02) 9221-9211
Fax: (02) 9221-9222
ABN 18 117 524 366

Publication Number: 06015

Market Clarity

Contents

Tables

Figures

Foreword

Leveraging international service and technology trends and recognising corporate demand for a national, Layer 2 VPN, Nextgen has launched a national VPLS (Virtual Private LAN service) with a market leading specification suite.

Written by Market Clarity, an independent research organisation, this paper discusses the various wide area networking (WAN) technology capabilities, including VPLS, and the business decisions that can influence an organisation's service and technology choices.

Demystifying the Difference between Layer 2 and Layer 3 VPNs compares and contrasts key WAN technologies such as Frame Relay, ATM, Ethernet, IP, IP/MPLS and VPLS service capabilities — and explains the key differences between these Layer 2 and Layer 3 network solutions. It also presents a set of decision-making guides to help organisations decide which service (or services) best fit their business and application needs.

Nextgen Networks trusts you will find this guide useful in your network planning activities.

Peter Harrison
General Manager, Nextgen Networks

1 Introduction

This White Paper is designed to give readers a pragmatic understanding of the technologies used to deliver WAN services in Australia.

With a growing number of telecommunications carriers and service providers now offering Layer 2 and Layer 3 network services, Australian businesses have an unprecedented opportunity to achieve a perfect fit between their needs and the services they purchase.

At the same time, however, customers need to have a good understanding of their business requirements, the applications their networks need to support, the costs associated with different services, and the ROI available from the right choice of service.

In this publication, Market Clarity presents a description of the behaviour and application of Layer 2 and Layer 3 VPN services, and analyses the technical and business features associated with these services.

To ensure readers gain a balanced understanding of the choices available to them, we have included legacy services such as ATM and Frame Relay (which still platy an important role in the delivery of high-reliability WAN services), as well as Ethernet, VPLS, IP+MPLS, Private IP, and IP-VPNs using the public Internet.

We also present examples of how the interactions between different business requirements, different application behaviours, and different legacy network requirements can impact the choice of new WAN services.

Finally, we advise readers to enter into a dialogue with prospective service providers. Just because a technology is capable of supporting a specific feature or topology, does not mean that a network operator has implemented any given feature.

2 Understanding VPNs

In the not too distant past, organisations that wished to establish always-on voice or data communications between their business sites used private lines (leased lines or ISDN) to connect their business locations across a "private" network infrastructure. For any given site to be connected to the network, a dedicated private line must be provisioned between a pair of sites. A simple example of this is shown below, in Figure 1, where three separate leased lines provide connectivity between the Headquarters and Sites 1, 2 and 3, respectively.

Figure 1. — Leased Line Private Network (Layer 1)

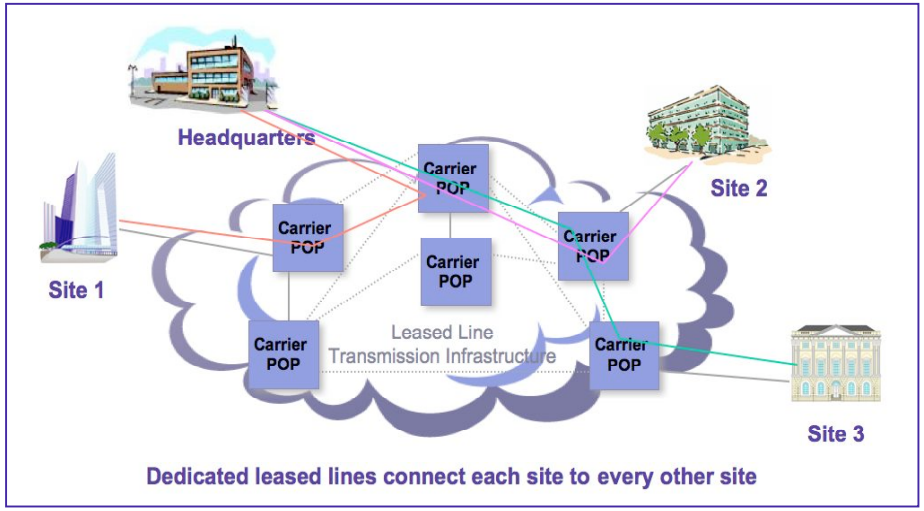

One of the disadvantages of this network topology is the requirement to route all remote site traffic through the "Headquarters" site, which is the only common link connecting Sites 1, 2 and 3, or to procure additional leased line connections (at additional expense) to connect Sites 1-2, 1-3, and 2-3.

In large networks the additional expense can be considerable.

Furthermore, this type of topology leads to the inefficient use of bandwidth between locations. Because each site is connected via a "dedicated" connection, the connection remains idle until such time as information is transmitted between a site pair, such as "Headquarters to Site 3."

As can be seen in Figure 1, there are actually three (3) physical connections going into Headquarters (the links connecting to Sites 1, 2 and 3), each of which is a dedicated point-to-point connection.

Clearly, the ability to aggregate bandwidth across all the available links going into a site (for instance, Headquarters) would allow information flows to take advantage of the total physical capacity going in and out of a site.

Even in "private" networks, transmission capacity inside a Carrier network is shared. Transmission equipment sits in each Carrier POP (point of presence), providing POP-POP high capacity connectivity. These high capacity connections are de-multiplexed (separated) into individual transmission streams, which can then be allocated to individual customers.

Where a network resource is either completely the property of a single user, or is sold for the exclusive use of a single user, it is considered a “private” network.

Figure 2, below, illustrates two kinds of private networks.

Figure 2. — Examples of Private Networks

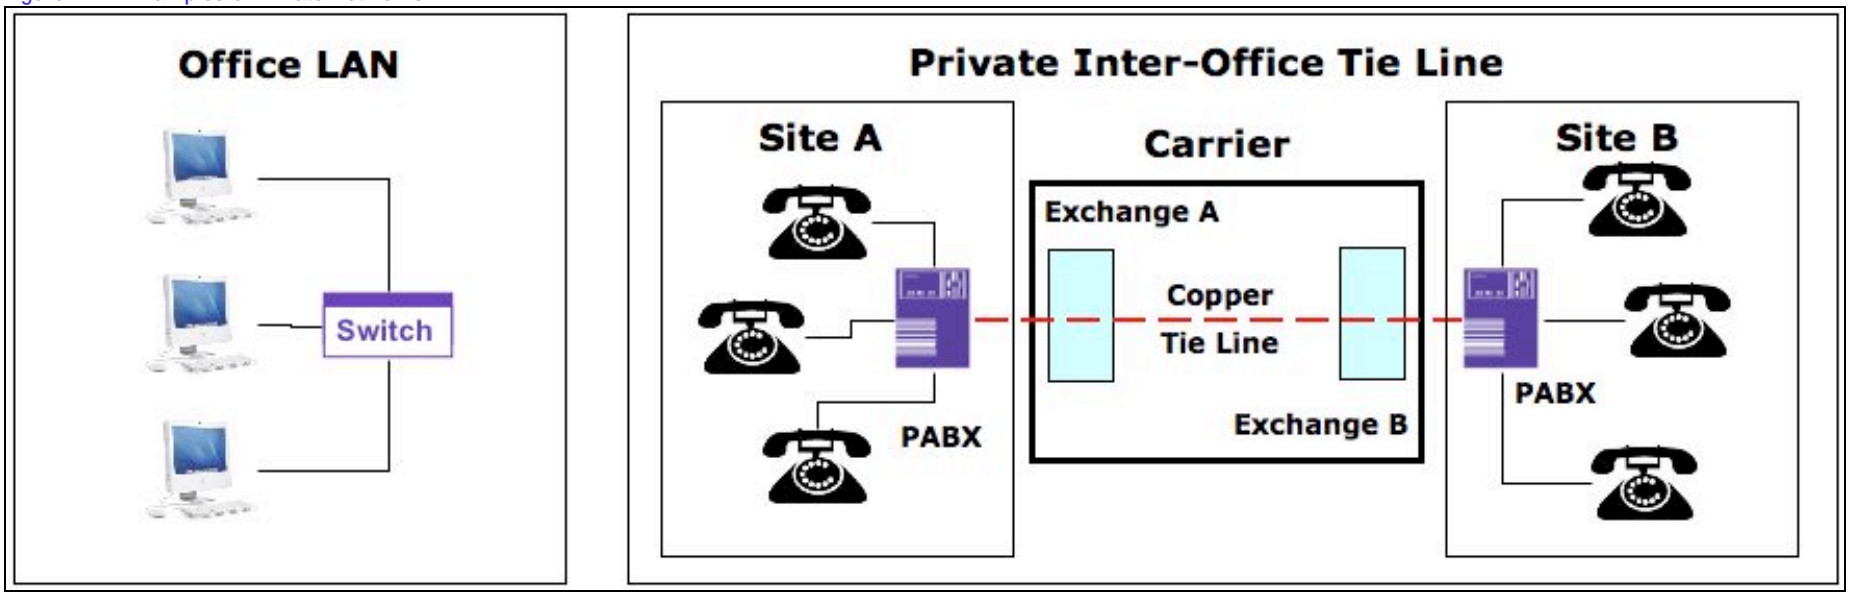

In the office LAN, all of the infrastructure — the Ethernet switch, the cabling, and the PCs — are owned by a single entity. In the case of the inter-office tie-line, the telecommunications carrier is providing a DC copper path between two points; the copper is not shared with any other customer.

However, leased lines (copper, fibre, radio) can span long or short distances. For longer distances, leased lines almost always refer to carrier-based TDM (time division multiplexing) circuits — a method of putting multiple data streams in a single signal by separating the signal into many segments, each having a very short duration. This type of leased line uses shared infrastructure, and each customer’s data stream has a dedicated segment of a transmission signal.

Similarly, a **Virtual Private Network** is a network in which a shared infrastructure is used to provide private services to its users.

Figure 3. — VPNs in the OSI Context

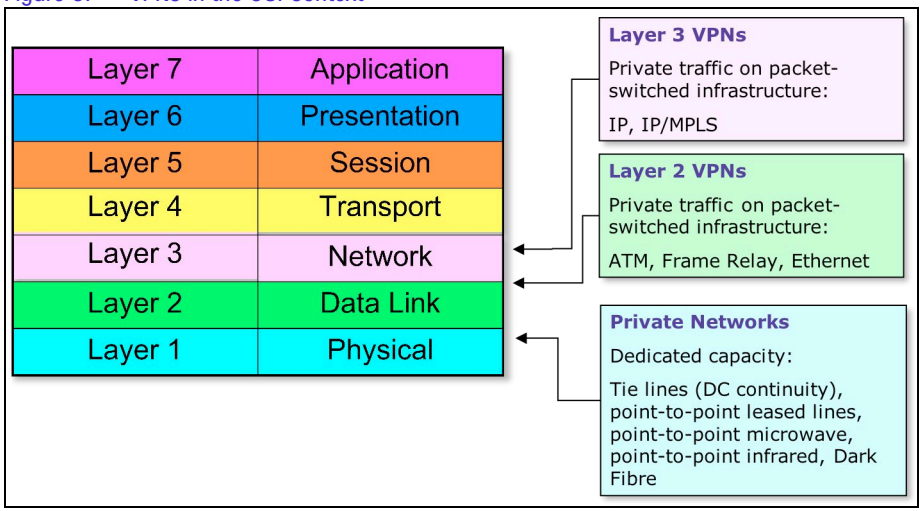

Shared infrastructure may exist in many ways:

- **Layer 1 —** At Layer 1 (the physical layer), users share the same physical medium for voice conversations or data sessions. In a radio network, different communities may use a single block of spectrum, with some kind of access control (such as frequency assignment or encryption) maintaining traffic privacy.
- **Layer 2 —** Customers share the switching infrastructure, with virtual paths or virtual circuits segregating customer traffic.
- **Layer 3 —** At Layer 3, all customer traffic uses the same internetworked routers. Traffic segregation is accomplished either by customers using encryption for traffic traversing the shared infrastructure (IP VPNs using the public Internet), or using IP+MPLS over a carrier infrastructure.

The remainder of this white paper will concentrate on Layer 2 and Layer 3 VPN services.

VPNs provide varying degrees of resource management, so as to meet customers' performance expectations (for example, to deliver the customer a 2 Mbps point-to-point link).

- **Layer 1 —** Management of a shared Layer 1 resource is highly dependent on the type of resource. A shared radio channel is very vulnerable to congestion, but a wavelength on a dark fibre network is immune to other customers' traffic.
- **Layer 2 —** Layer 2 VPNs can provide very good resource management within the capacity of the network.
- **Layer 3 —** IP+MPLS provides good resource management within a provider domain (that is, where a single service provider manages all of the resources associated with customer traffic). IP VPNs using the public Internet offer poor resource management in terms of the needs of a single customer.

3 Overview: Layer 2 and Layer 3 VPNs

3.1 Layer 2 VPNs

3.1.1 Layer 2 Virtual Circuits and Virtual Paths

Layer 2 Virtual Circuits

The Layer 2 technologies listed in the side box use a switched infrastructure to create an association between a customer's ingress port, and the network resources devoted to that customer.

Two approaches are used to achieve this: virtual circuits (VCs), and virtual paths (VPs).

Table 1, below, outlines characteristics of Layer 2 VCs.

Table 1. — Characteristics of Virtual Circuits (Source: Frame Relay Forum, 1998)

Characteristics	Private Lines	Virtual Circuits
Highly Secure - Used by Only One Company	X	X
Assigned Bandwidth	X	X
Automatically Route Around Network Failures		X
Able to Exceed Assigned Bandwidth		X
Efficient Bandwidth Utilisation		X

Layer 2 and Layer 3 Network Services

Layer 2	Layer 3
ATM	Internet-based IP VPNs
Frame Relay	"Private" Carrier IP Networks
Ethernet	IP+MPLS
Metro Ethernet	
VPLS	

Virtual Circuits are highly connection-oriented. The switching infrastructure used in a Layer 2 network:

- Receives requests for resources from the Network Termination Unit (NTU);
- Determines the availability of resources between source and destination; and
- Establishes the VC between source and destination.

Other characteristics of a Layer 2 virtual circuit include:

- **Multi-protocol networks —** Because the network exists at Layer 2, it does not constrain the customer's choice of Layer 3 protocols such as IP or IPX.
- **Multiple VCs —** A single physical connection into a customer site can be segregated into multiple VCs, each of which can be designed to carry a different type of application traffic.

Figure 4. — Layer 2 Services Can Support Multiple VCs

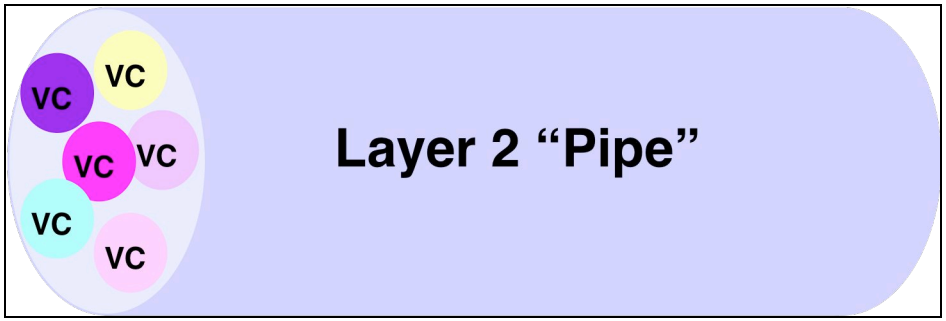

The use of multiple VCs on a single pipe has a variety of applications for customers. These include:

- **Application Segregation —** Traffic can be associated with different VCs configured for performance appropriate to different applications. One VC may be configured to provide "best effort" services for e-mail and Web browsing, while another VC may offer real-time performance appropriate to voice connections.
- **Topology Management —** On a high-capacity service entering a data centre, individual VCs may be associated with individual offices accessing the data centre.
- **User Group Segregation —** VCs may also be associated with different user groups to maintain internal traffic segregation (for example, to separate highly sensitive data from general office networks).

Figure 5. — VCs Over Carrier Networks

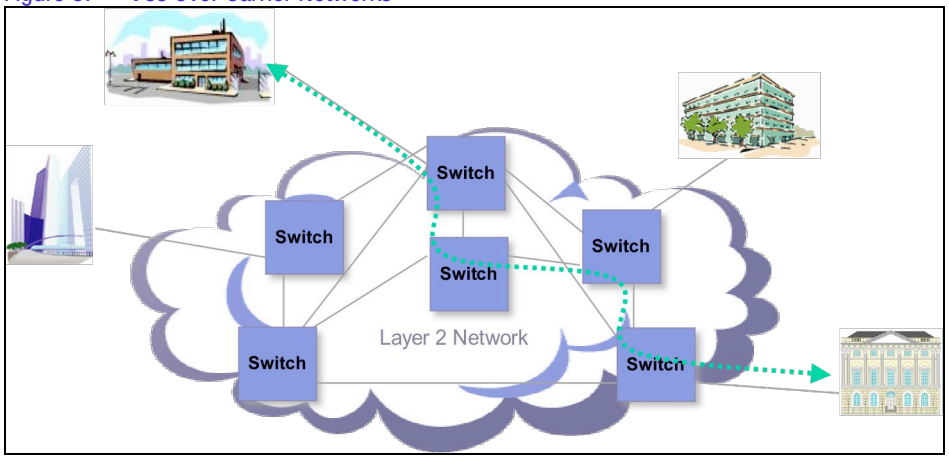

Deterministic Traffic Routing

Once switches are configured to use a particular path for a given VC, that VC's traffic will always follow the same path. This allows Layer 2 VCs to deliver deterministic traffic performance.

This capability allows carriers to provide service level guarantees (SLGs) to their customers.

Services such as Frame Relay and ATM utilise virtual circuits to transmit data between customer sites. Each virtual circuit within a customer network, called a permanent virtual circuit (PVC), typically attracts a separate charge.

Layer 2 Virtual Paths

ATM adds the concept of Virtual Paths (VPs) to the VC. VPs are established for group routing of VCs along a specific path, and each VC is able to have a different class of service.

The VP allows customers to create VCs with characteristics suitable for different traffic types, without requiring carrier provisioning for the newly created VC.

The concept of nesting Virtual Circuits (VCs) within Virtual Paths (VPs) is also used by the MPLS protocol. Here, the virtual circuits are called Label Switched Paths (LSPs).

And, the same concept has made its way into Metro Ethernet services with the concept of Stacked VLANs (see Section 3.1.3).

Figure 6. — Virtual Paths Within ATM VCs

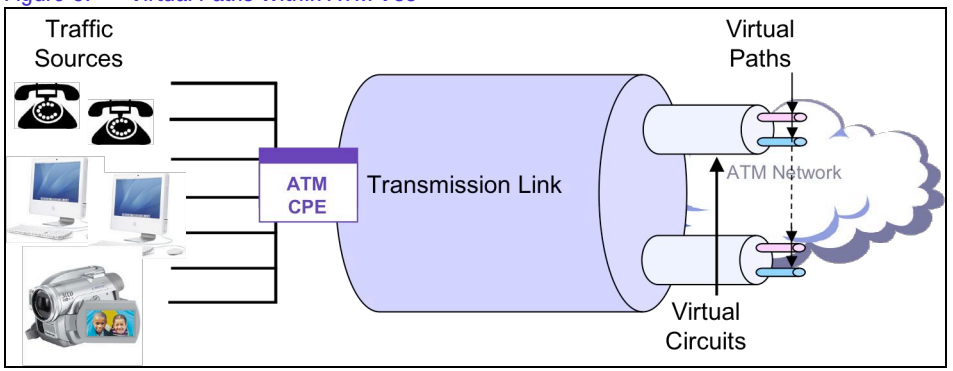

3.1.2 Layer 2 Ethernet Services

Figure 7. — Ethernet Services using VLANs for Traffic Segmentation

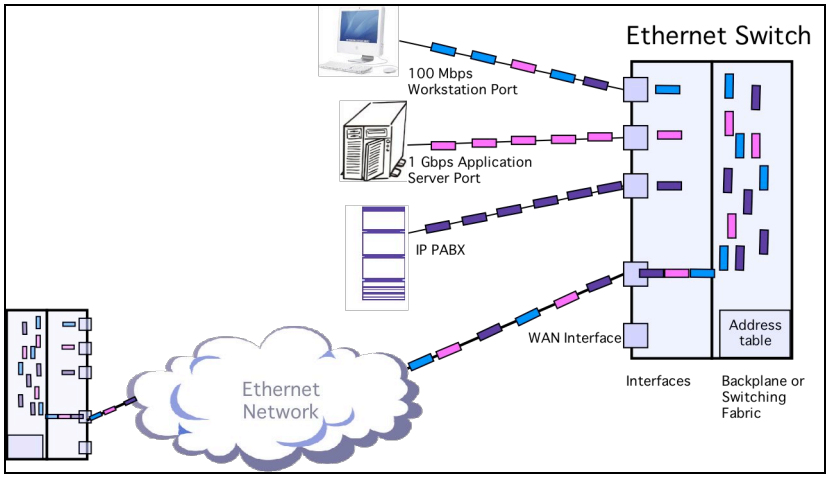

Ethernet-based WAN services use the same MAC (media access control) Layer-based addressing scheme as is used in the LAN.

A more recent development in Layer 2 carrier services is the use of Ethernet.

Originally developed as a LAN technology, Ethernet was for many years unsuitable for WAN applications due to its distance limitations.

The “reach” of a single Ethernet segment was a consequence of the collision detection protocols designed for Ethernet’s original shared media architecture. A single segment (in Ethernet’s early days, a single coaxial copper cable to which all stations were connected) needed to be short enough that signals would propagate quickly enough to be received within a collision detection cycle.

Switched Ethernet, in which a switch port is connected to a single transceiver, permitted the development of new protocols in the Ethernet standards stream to overcome this distance limitation. At the same time, the development of standards using optical fibre at the physical layer allows high-speed Ethernet-based communications over metropolitan distances.

The **Metropolitan Ethernet Forum** (http://www.metroethernetforum.org) provides information about the standards associated with deploying Ethernet as a WAN technology.

In operation, an Ethernet switch behaves very much like a bridge. It associates a MAC address, or a group of MAC addresses, with a particular port, with traffic forwarded using the outgoing port associated with that MAC address.

Virtual LANs (VLANs), a familiar concept within the enterprise, are also used in Ethernet-based WAN services as one technique to maintain segregation of customer traffic.

Further developments enabling the use of Ethernet as a metropolitan service are described in Sections 3.1.3 and 3.1.4, below.

3.1.3 Layer 2 Metro Ethernet Services

A number of developments in Ethernet standards have provided building blocks for the deployment of carrier grade Ethernet services. Key specifications include:

- **VLAN Stacking (IEEE 802.1ad, also called Q-in-Q):** VLAN stacking is a technique that allows a service provider to add an additional VLAN tag, called a S-VLAN Tag (Service VLAN Tag) to a subscriber's IEEE 802.1Q tagged Ethernet frame — allowing service provider VLAN tags to operate independently of customer VLAN tags. This feature prevents VLAN identifier conflicts between customer and service provider networks, as well as between customer networks. It also allows customers to specify multiple VLANs within a single service provider VLAN identifier.
- **Traffic Prioritisation (IEEE 802.1p):** The IEEE 802.1p standard specifies a mechanism for indicating Ethernet frame priority based on the User Priority bits. This three-bit supports the identification of up to 8 traffic classes (priorities). While 802.1p provides a mechanism to tag frames for prioritisation, there is no uniform approach to implementing underlying queuing mechanisms or bandwidth reservation.

In addition to work done by the IEEE, the **Metropolitan Ethernet Forum (MEF)** has produced a number of technical specifications, which define Ethernet service building blocks (service attributes) and a framework describing the building blocks for creating Ethernet services. Building blocks consist of Ethernet service attributes and parameters, allowing for concepts such as an Ethernet UNI (Ethernet User Network Interface).

Another important concept established by the Metro Ethernet Forum is an **Ethernet Virtual Connection (EVC)**. An EVC specifies connectivity between Ethernet UNIs. There are a number of EVC types of and a number of service attributes that an EVC can have, and the Forum's Ethernet Services Definitions define how to apply Ethernet service model building blocks to create services such as an Ethernet Line (E-Line) and Ethernet LAN (E-LAN).

- An **Ethernet Line (E-Line)** service provides point-to-point connectivity; and
- An **Ethernet LAN (E-LAN)** service provides multipoint-to-multipoint (any-to-any) connectivity.

These service types are illustrated in Figure 8.

Figure 8. — Metro Ethernet Services Types (Source: Metro Ethernet Forum)

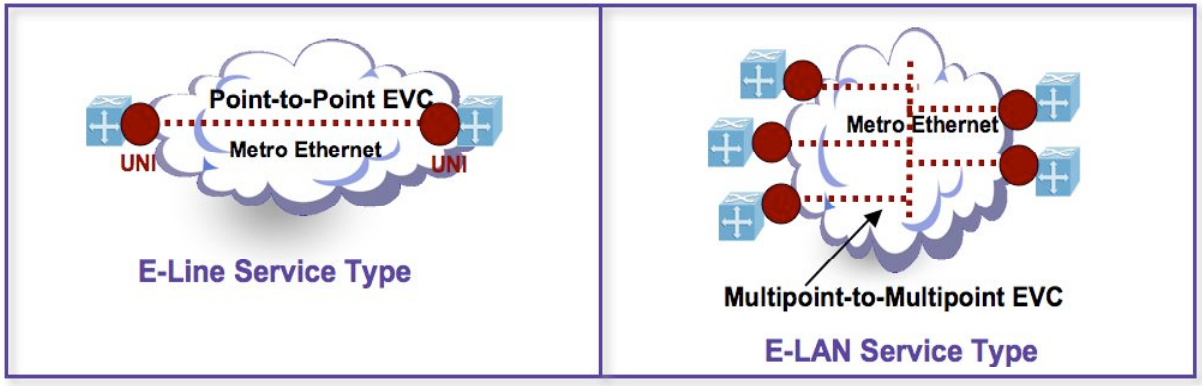

Additionally, the following service attributes can be defined on a UNI or Ethernet Virtual Connection basis:

- Ethernet Physical Interface;
- Traffic Parameters (Bandwidth Profiles);
- Service Performance Parameters;
- Class of Service;
- Service Frame Delivery;
- VLAN Tag Support;
- Service Multiplexing;
- Bundling; and
- Security Filters.

The Metro Ethernet Forum has also defined bandwidth profiles for Ethernet services consisting of the following traffic parameters:

- **CIR (Committed Information Rate)** — CIR is the average rate up to which service frames are delivered per the service performance objectives (such as delay or loss);
- **CBS (Committed Burst Size)** — CBS is the maximum number of bytes allowed for incoming service frames to be CIR-conformant;
- **EIR (Excess Information Rate)** — EIR is the average rate, greater than or equal to CIR, up to which service frames are delivered without any service performance objectives; and
- **EBS (Excess Burst Size)** — EBS is the maximum number of bytes allowed for incoming service frames to be EIR-conformant.

Keen readers will note that these parameters are very similar to those available with ATM and Frame Relay services. A key difference of Metro Ethernet over other Layer 2 technologies is its ability to support multipoint-to-multipoint, any-to-any connections. However, Metro Ethernet, by itself, has scalability issues due to its inherent reliance on VLAN tags to separate customer traffic. (VLAN tags utilise a 12-bit field in the Ethernet frame, and as a result can support a maximum of 4095 VLAN identifiers [excluding reserved values].)

3.1.4 Layer 2 VPLS Services

Although it is possible to build large-scale switched Ethernet networks to provide multipoint Ethernet Layer 2 VPN services using techniques such as Q-in-Q tunnelling, native Ethernet does not scale to very large, geographically distributed networks because of technical limitations inherent to Ethernet control protocols and bridges. These limitations have led to the development of mechanisms within the IETF whereby transport protocols, such as Ethernet, Frame Relay, ATM, and TDM can be transported over IP or MPLS networks utilising an emulation service called a pseudowire.

Building on the pseudowire concept, an Ethernet multipoint service known as **Virtual Private LAN Service (VPLS)** can be deployed using a combination of Ethernet and MPLS. Indeed, MPLS encapsulation of Ethernet is widely seen as a crucial step for VPN scaling, service provisioning and traffic engineering.

With VPLS, multiple customer sites can communicate as if they were connected via a private Ethernet LAN segment, but with superior robustness, reach and scalability.

Figure 9. — Layer 2 VPLS Services: Combining Metro Ethernet and MPLS

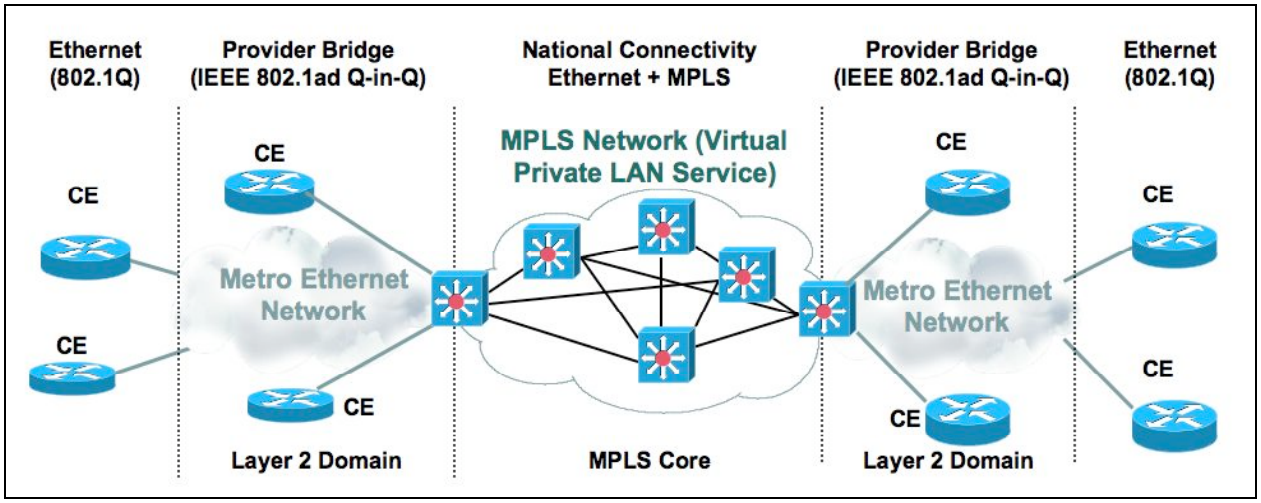

VPLS is a VPN technology that enables Ethernet multipoint services over an MPLS network infrastructure.

Customer edge (CE) equipment can connect directly to the MPLS/VPLS core network, or via a Metro Ethernet access network.

3.2 Layer 3 VPNs

This paper discusses three ways in which VPN capabilities can be delivered using Layer 3 IP routed networks:

- **Internet-based VPNs** — which use encryption at the boundary between the customer's private network and the public Internet to provide security and privacy of customer traffic while it traverses the public Internet;
- **Private IP Networks** — in which all of the infrastructure carrying customer traffic is owned by a single carrier or service provider; and
- **IP+MPLS Networks** — in which MPLS, combined with the DiffServ architecture, is used to add an overlay of QoS and CoS management to a routed IP Layer 3 network.

The key characteristics of Layer 3-based services include:

- **Flexible Any-to-Any Connectivity** — By default, every site in a Layer 3 network can connect to any other site on the network.
- **Voice, Data and Video** — IP divorces the content from the underlying network, eliminating the need for specialist products to handle different traffic types.
- **Very High Scalability** — The scalability of an IP-based network is based on the address space used in the network (which is not a constraint in private networks, since the address space is so much larger than the network), the communication links between routers, and the processing capacity of the routers themselves.

In Layer 3 technologies, the customer's traffic is not explicitly associated with a particular resource. The service provider architects a network according to the relationship between network capacity and performance, and customer expectations.

Hence the only guarantee associated with a public Internet connection is that of "best effort" routing of customer packets. Where a service exists solely within a single provider's domain (such as a carrier private IP network service), network infrastructure may be provisioned to provide sufficient bandwidth and routing capacity to meet customers' requirements. In the absence of any other technology, however, there is no QoS management in such an environment.

Technologies such as IP+MPLS are designed to replicate Layer 2-like service management in the Layer 3 environment.

3.2.1 Internet-based IP VPNs

The basis of an IP VPN is the use of encryption technology to provide a "virtual private network" for data traversing IP networks.
VPNs may be used to provide connections between remote clients and networks; or between two networks. In either case, the encryption device terminating the VPN to the network needs to support:

- **High throughput —** Where the VPN is terminated at a network gateway, the gateway device must be able to encrypt and decrypt data quickly enough to maintain overall network performance; and
- **Multiple sessions —** If it is used to provide remote access for a large number of remote users, or if it is terminating secure tunnels from a number of remote networks (such as a head office accepting VPN connections from branch offices), the network gateway must be able to process several tunnels simultaneously.

Internet-based VPNs are most useful in remote access for travelling or mobile users, since they enable access anywhere the user has Internet access. However, because the data must traverse the public Internet, they are not recommended for enterprise network-to-network connections if the applications using the service require deterministic QoS.

3.2.2 Private IP Networks

Private IP networks describe routed IP networks where the entire infrastructure is owned or managed by the service provider. These are traditional Layer 3 networks, as shown in Figure 10 below, but operate independently of the public Internet.

Figure 10. — Traditional Layer 3 Routing

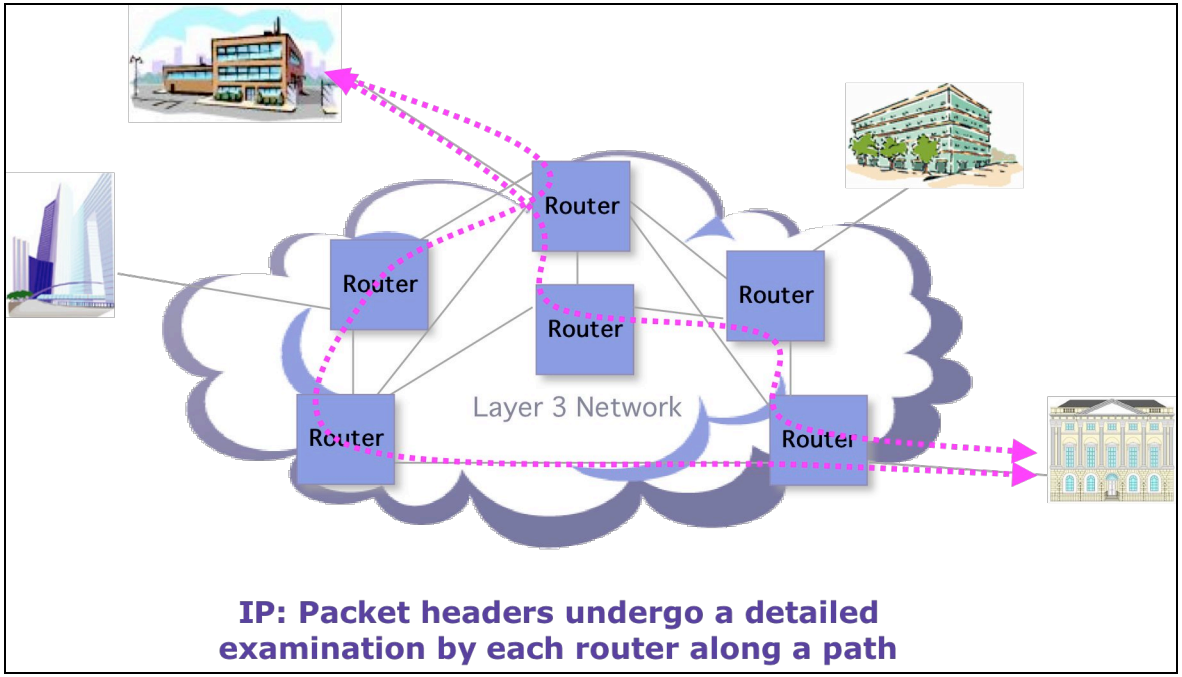

In the traditional Layer 3 network, every packet header undergoes detailed examination by all routers along a path. Moreover, as the state of the network changes (for example, if individual devices suffer congestion), routers will change their decisions about traffic paths, so that packets belonging to a single session will frequently traverse different devices between source and destination.

Since all customer traffic remains within a single provider domain, the network owner can adopt various strategies to deliver better performance than is available on the public Internet. These include:

- **Connection Provisioning —** The service provider can provision link capacity between its routers to support its customers' traffic requirements with lower oversubscription than is typical on Internet backbone networks;
- **Router Configuration —** The service provider will purchase backbone routers of sufficient capacity to serve its customers' requirements;
- **Network Topology —** Since the network is owned by a single entity, the provider can ensure a minimum of hops between its routers, minimising the risk of long on-network latency.

A pure routed IP network is unable to offer the deterministic QoS necessary for real-time traffic such as business IP Telephony.

One solution to the problem of delivering deterministic QoS on IP networks is the use of **MPLS (Multi-Protocol Label Switching)**. The MPLS protocol also allows routers to hold more information about the network links, which exist between two paths.

This gives IP+MPLS-based services much more deterministic behaviour than traditional "best-effort" IP networks, since the traffic between network nodes can be forced to follow a particular path through the network.

Figure 11. — MPLS Manages the Traffic Path

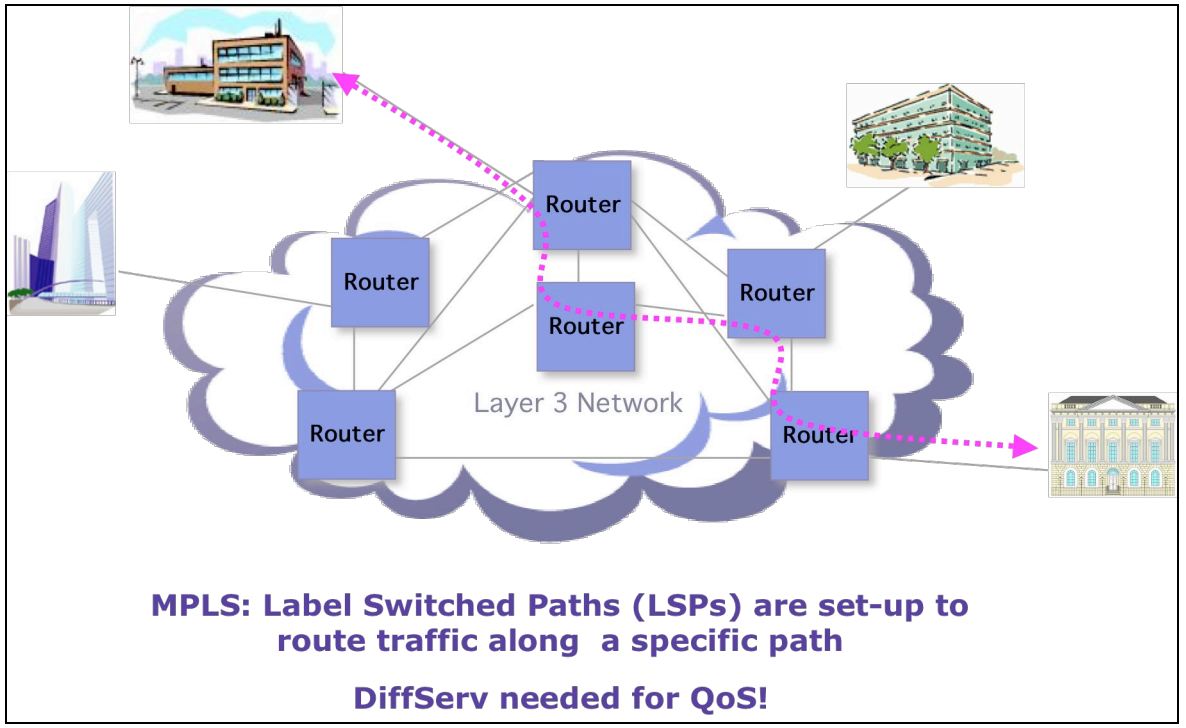

Routers that support IP+MPLS use labels to identify the routed path associated with particular traffic. Traffic without the MPLS labels is routed on a best-effort basis. In this way, IP+MPLS delivers connection-oriented behaviour on a connectionless routed network.

It also enables traffic engineering, since the network manager can design different paths for different traffic types. The most demanding traffic may take the shortest route, or it may be assigned to routes with the largest-scale routers (to avoid the risk of encountering a congested device).

MPLS can force packets into specific paths. However, MPLS by itself does not provide for QoS. To add QoS to an MPLS-based network, traffic engineering is required to differentiate IP packets by their QoS requirements, and to have the network honour that differentiation. IP+MPLS can achieve this using DiffServ, which provides for traffic classification, and allows the management of "Per Hop Behaviours" (PHBs).

MPLS typically specifies the behaviour of the network between "customer edge" routers. These are not CPE devices; rather, they are the customer-facing router ports at the edge of the provider network, to which a customer tail is connected. The IP+MPLS network provider will, however, either configure the CPE router (in a managed service) or provide information on configuration (if the customer is managing its own CPE router), so that appropriate markings can be applied to customer traffic before it enters the IP+MPLS network.

DiffServ provides a QoS treatment to traffic aggregates, providing a scalable and operationally simple solution. However, because it does not influence a packet path, it cannot guarantee QoS. By combining MPLS and DiffServ, network operators are provided with a scalable traffic classification schema that provides traffic engineered path selection and bandwidth guarantees, enabling true QoS.

3.3 Impact of Architecture on Tariff Structures

All of the Layer 2 and Layer 3 architectures outlined above also fall into two topological categories: hub-and-spoke topologies, and mesh topologies. These are outlined in Table 2, below.

Table 2. — Network Topologies

Network Type	Network Layer	Topology
Frame Relay	Layer 2	Hub-and-spoke
ATM	Layer 2	Hub-and-spoke
Metropolitan Ethernet	Layer 2	Mesh or Hub-and-spoke
VPLS	Layer 2	Mesh or Hub-and-spoke
IP+MPLS	Layer 3	Mesh (can also be designed for a hub-and-spoke topology)
IP (no MPLS)	Layer 3	Mesh
Public Internet (IP VPN)	Layer 3	Mesh

VC-based networks are typically based on a hub-and-spoke architecture (Figure 9, below) in which a central site (for example, a head office or a data centre) processes all traffic to and from branch offices, even if the communication is between two branch offices.

Figure 12. — Hub-and-Spoke Network

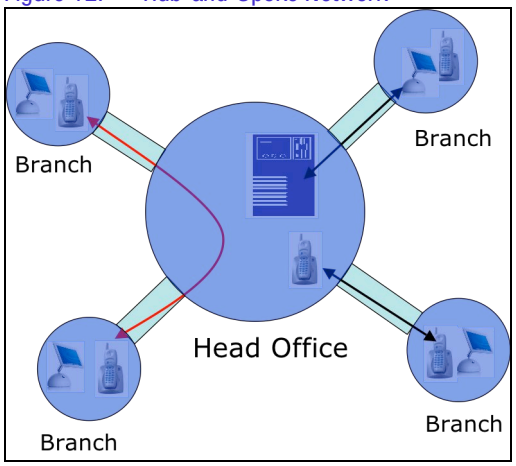

A **hub-and-spoke topology** is a simple and cost-effective way to support traditional client-server computing applications. However, applications such as IP Telephony and IP videoconferencing generate traffic that does not fit this pattern.

Where branch offices need to communicate directly, network traffic must still traverse the hub-and-spoke network. This means that all traffic must go in and out of the Head Office access links and switching infrastructure. This places an extra load on Head Office bandwidth, which is not required for efficient communications.

A meshed solution (Figure 14), wherein all sites can communicate with each other through a direct path, is a better fit for these types of applications.

As a VC-based network expands, the hub-and-spoke architecture may become several hubs serving localised spokes (Figure 13, below).

In this case, the growing distance between spokes may degrade the network's ability to support direct IP-based voice or video communications between remote spokes — due to the extra latency incurred with routing traffic in and out of hub sites, and the network owner may instead look to create a mesh topology (Figure 14).

The growth of a VC-base network also presents considerable challenges to the network designer in network dimensioning and QoS management.

Figure 13. — National Hub-and-Spoke Network

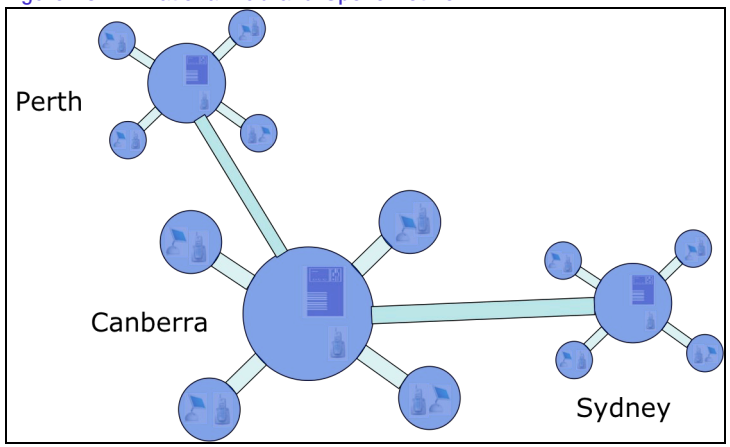

Figure 14. — Meshed Network

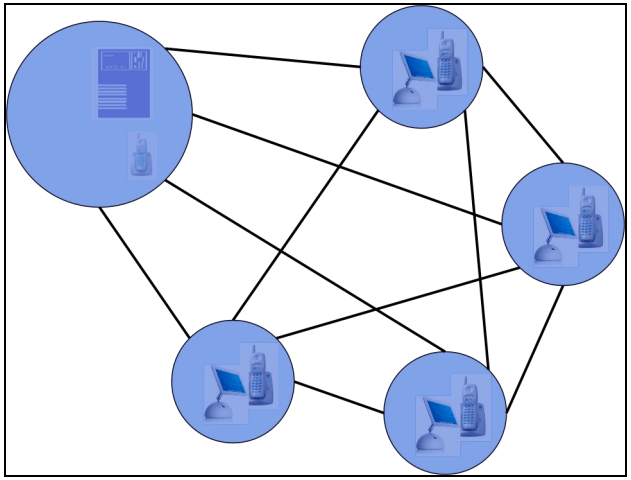

While virtual circuits can be used to create meshed networks, the cost of interconnectivity quickly escalates, because each VC is typically charged as a separate tariff item. Where customers require a high degree of any-to-any connectivity, they should consider implementing a carrier service that natively supports a meshed architecture.

Carrier services based on a meshed architecture eliminate the need for the customer to implement large numbers of point-to-point virtual connections, since the any-to-any connectivity is created within the carrier “cloud”.

4 Technology Comparisons

4.1 Overview

In this section, Market Clarity provides:

- **Technology Ratings —** Tables providing an at-a-glance indication of the technical aspects of various Layer 2 and Layer 3 technologies;
- **Business Features/Benefits Ratings —** Tables identifying the business characteristics of various Layer 2 and Layer 3 technologies in an at-a-glance format; and
- Detailed explanations of both the technology and business ratings.

Table 3, below, explains the rating system used to indicate technology and business ratings.

Table 3. — Colour-Code Key

+	Native to technology/architecture
+/-	Highly dependent upon configuration. Can be accomplished via the implementation of an optional feature, proprietary add-ons or special configuration.
-	Not supported by the technology/architecture

Market Clarity notes that while a technology or architecture may support a range of features and functions, service providers are free to implement, or not to implement, any given feature.

4.2 Technology Ratings

Table 4, below, summarises the technical features associated with Layer 2 and Layer 3 technologies.

Table 4. — Layer 2 and Layer 3 Services: Technical Features Summary

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Fast Network Recovery (Milliseconds)	+	+	+/-	+/-	+/-	-	-
Scalability (Number of Sites)	+	+	-	+	+	+	+/-
Scalability (Bandwidth)	+/-	+	+	+	+	+	+/-
Native Mesh Topology	-	-	+	+	+	+	+
Support for Traffic Engineering	+/-	+	+/-	+	+	-	-
Deterministic QoS	+/-	+	+/-	+	+	-	-
Guaranteed Bandwidth	+	+	+	+	+/-	-	-
Customer Specification of QoS Parameters	+/-	+	+/-	+	+	-	-
Low Network Complexity	+/-	+/-	+	+	+	+	-
Customer Control of Layer 3 Routing and IP Addressing	+	+	+	+	-	-	-
National VLANs	-	+/-	+/-	+	+/-	-	-
Customer Control of VLAN Numbering Across the WAN	-	+/-	+/-	+	+/-	-	-
Does not Require CE Router	+	+	+	+	-	-	-
Suitability for Small Sites	+	+/-	+	+	+	+	+
Support Mixed Physical Layer Technologies (Layer 1)	+	+	+	+	+	+	+
Support for Mixed Access Layer (Layer 2)	-	-	-	-	+	+	+
Native Support for Legacy Layer 3 Protocols	+	+	+	+	-	-	-
Less Vulnerable to Internet-Borne Security Threats	+	+	+	+	+/-	+/-	-
Multi-homing for WAN Redundancy	+	+	+	+	+/-	+/-	+/-

These features are explained in greater detail below.

4.2.1 Fast Network Recovery

All of the Layer 2 and Layer 3 technologies discussed in this paper are able to recover from network outages. The speed at which different technologies can recover from different kinds of outages, however, is highly dependent on the architecture and on the network layer at which that technology operates.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Fast Network Recovery (Milliseconds)	+	+	+/-	+/-	+/-	-	-

In general, Layer 2 technologies exhibit faster recovery from network outages. Wherever possible, carriers deploy redundant physical paths to support Layer 2 networks, and recovery is accomplished by switches designed to detect an outage, then switch to the redundant path as quickly as possible. The recovery of a Layer 3 network is more complex. First, routers detect the loss of a connection, discover alternative routes, and begin routing traffic along those new paths; then when the underlying connectivity is re-established, the Layer 3 network reverts to its original configuration.

One notable exception to the discussion of network recovery applies to MPLS and VPLS services that are designed using MPLS Fast-Reroute capabilities. This optional technique enables operators to set-up backup LSPs (either used in the network core, or for specific customers). MPLS Fast-Reroute enables traffic redirection in 10s of milliseconds.

We also note that a technique called Rapid Spanning Tree Protocol (RTSP), which is described in the IEEE 802.1w standard, provides for much faster recovery than a standard Ethernet service. However, even with RTSP, reconvergence occurs within sub-second to 3-second timeframes.

4.2.2 Scalability (Number of Sites)

With the exception of “native” Ethernet, all of the technologies shown above have moderate to high site scalability: that is, a single network can natively support a large number of end user locations.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Scalability (Number of Sites)	+	+	-	+	+	+	+/-

Native Ethernet services can quickly encounter MAC address learning limitations or VLAN numbering constraints, making it necessary for carriers to implement technologies, which quarantine user-side addressing from carrier-side addressing.

Internet VPNs are, in theory, almost infinitely scalable, since any limits on user-side addressing are far in excess of the requirements even of very large networks. The problem lies, however, in the encryption needed to ensure the network security, user authentication, traffic privacy, and provide activity audits over Internet-based VPNs. Maintaining the addressing of an encrypted Internet VPN adds an administrative overhead which rises as the number of sites grows.

4.2.3 Scalability (Bandwidth)

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Scalability (Bandwidth)	+/-	+	+	+	+	+	+/-

Nearly all of the private network solutions offer bandwidth scalability suitable for networks from the very small to the very large. ATM, Ethernet, VPLS, IP+MPLS and private IP networks are all offered by Australian carriers in speeds ranging from sub-Megabit up to the Gigabit range.

Frame Relay is the least scalable of the private network solutions now available. While there are standards for speeds of up to 622 Mbps, Australian carriers generally offer Frame Relay services at a maximum link speed of 2 Mbps. Where “Fast Frame Relay” services are available, these generally top out at 45 Mbps.

Public Internet VPNs can scale up to very high speeds on the access service. However, the lack of guaranteed bandwidth outside the domain of a single service provider make the public Internet unsuitable for Gigabit VPNs. Further, the likely cost of Internet traffic at very high throughput would make an Internet-based VPN uneconomic at high bandwidth.

4.2.4 Native Mesh Topology

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Native Mesh Topology	-	-	+	+	+	+	+

Where a corporate network supports centrally hosted data applications, a hub-and-spoke topology (as used by traditional Layer 2 services such as Frame Relay and ATM) is suitable.

However, where multimedia applications are required, and particularly where these applications need to provide communications between offices other than the central location, a hub-and-spoke topology can force traffic to take traverse long paths (for example, a VoIP call between two branch offices would have to traverse the head office network). By contrast, a mesh topology can support a session in which data travels directly between these two locations. Ethernet and VPLS services (at Layer 2) and all of the Layer 3 services are suitable for applications requiring mesh support. Where Layer 3 services are used, customers should ensure that the underlying Layer 2 topology is in line with their applications' requirements for meshing.

4.2.5 Support for Traffic Engineering

Quality of Service (QoS) refers to the capability of a network to provide better service to selected network traffic, and is often referred to as traffic engineering. QoS mechanisms can be deployed across various technologies, and at various network layers.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Support for Traffic Engineering	+/-	+	+/-	+	+	-	-

This table reflects the degree to which traffic engineering is available to meet customer requirements. It is, however, feasible that services that do not support "customer-side" traffic engineering are built on traffic-engineered networks. For example, a carrier may not be able to provide traffic engineered private IP links to its customers, but may traffic engineer the underlying links upon which it builds that private IP service.

The shared infrastructure of best-effort, routed IP networks prevents any significant customer-side traffic engineering. Frame Relay (using buffer and queue management, for example) and Ethernet (using VLANs) have a limited ability for traffic engineering. We do note, however, that Ethernet services, which utilise Metro Ethernet Forum's bandwidth profile specifications, can also support traffic engineering. Generally speaking, ATM, VPLS, and IP+MPLS services have the greatest degree of traffic engineering capabilities.

4.2.6 Deterministic QoS

This refers to the ability of a service to support predictable management of network QoS parameters. (Traffic engineering allows a network operator to establish a deterministic path. In other words, it is used to steer traffic to parts of the network where capacity is available.)

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Deterministic QoS	+/-	+	+/-	+	+	-	-

While best-effort IP services may offer high performance, it is not possible to engineer the behaviour of a best-effort network so that performance can be guaranteed. While customers may find latency, jitter and packet loss that is suitable for voice calls with IP or Public Internet based services, performance will not be guaranteed.

Frame Relay is able to behave in a deterministic fashion with respect to some parameters. For example, the customer's CIR (committed information rate) should always be available. However, Frame Relay lacks fine-grained control over QoS parameters. Similarly, standard Ethernet services can offer very high speed, but only limited management of QoS parameters unless Metro Ethernet Forum bandwidth profile specifications are implemented. The highest levels of QoS management are associated with ATM, VPLS, and IP+MPLS services. In the case of VPLS and IP+MPLS, DiffServ must be implemented to deliver optimum QoS control.

4.2.7 Guaranteed Bandwidth

While individual voice calls require only limited network bandwidth, a company operating a large network combining voice and data traffic will need a guarantee of a minimum available bandwidth. Both of the traditional Layer 2 carrier services presented here (Frame Relay and ATM) were designed

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Guaranteed Bandwidth	+	+	+	+	+/-	-	-

with bandwidth guarantees in mind. In the case of Ethernet services, various flow or rate limitation techniques allow carriers to offer services with bandwidth guarantees. VPLS can also deliver reliable bandwidth on an end-to-end basis.

Without techniques such as MPLS, DiffServ or resource reservation protocols, end-to-end bandwidth can't be guaranteed on best-effort routed IP networks.

4.2.8 Customer Specification of QoS Parameters

ATM, VPLS, IP+MPLS and Ethernet services that offer Metro Ethernet bandwidth profiles, all allow customers to associate particular network traffic types with the QoS parameters offered on the carrier side of the network boundary. In Frame Relay and standard Ethernet solutions, this is also possible but only within the more constrained QoS capabilities of these technologies. Pure best-effort routed IP solutions have extremely limited QoS capabilities, and in the public Internet QoS undertakings do not cross provider domains.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Customer Specification of QoS Parameters	+/-	+	+/-	+	+	-	-

4.2.9 Low Network Complexity

An Internet-based VPN quickly becomes unacceptably complex, due to the administrative burden associated with maintaining static address tables (which also limits its scalability, as discussed above).

While Frame Relay and ATM are familiar

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Low Network Complexity	+/-	+/-	+	+	+	+	-

to enterprise data managers, they are not LAN protocols. This creates a separation between LAN and WAN implementation, which adds complexity to network design. (While MPLS is not a LAN protocol, the customer side interface will typically utilise standard IP protocols.)

4.2.10 Customer Control of Layer 3 Routing and IP Addressing

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Customer Control of Layer 3 Routing and IP Addressing	+	+	+	+	-	-	-

In all Layer 2 solutions, Layer 3 routing remains completely under the customer's control. This is important where customers have large, established networks already using extensive private IP addressing schemes, particularly where the customer's addressing scheme includes a large number of devices with static addresses. Such devices may include boundary devices at individual sites (such as routers), as well as large numbers of devices that lack DHCP support.

Since these services operate at Layer 2, customers can if they choose make their own routing decisions. Here, the association between a traffic path and a set of IP addresses is made in the customer's router, rather than by the service provider.

While customers can maintain private addressing behind customer edge devices in Layer 3 networks, they have limited or no control over route selection for traffic once it enters the provider network. Further, where the Layer 3 solution is provided as a managed service, it may be necessary to harmonise customer addressing with provider requirements.

4.2.11 National VLANs

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
National VLANs	-	+/-	+/-	+	+/-	-	-

Increasingly, VLANs are used to implement traffic separation between different business units within an enterprise. For example, a company may wish to keep users from its sales

VLAN segmented from the finance VLAN.

Another common use of VLANs is to separate traffic with different quality of service requirements. Indeed, industry best practice for the implementation of IP Telephony across a LAN involves the use of VLANs to segment voice from other traffic types. Where IP-based video or videoconferencing is implemented, VLANs are commonly used to segment video streams, which have different QoS tolerances than voice.

The ability to deploy consistent VLANs on a national basis offers several advantages to enterprises. Users can be associated with the correct VLAN even when visiting an interstate office, and systems administrators' tasks may be more easily consolidated into a central location. Due to the scaling limitations of 802.1Q, and of Stacked VLANs (802.1ad), VPLS is the only solution, which natively supports VLAN deployment on a national scale.

4.2.12 Customer Control of VLAN Numbering Across the WAN

The national VLAN transparency referred to above also means that customers can easily apply their own VLAN numbering schemes between different sites using VPLS.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Customer Control of VLAN Numbering Across the WAN	-	+/-	+/-	+	+/-	-	-

Additionally, Ethernet services that support VLAN Stacking (IEEE 802.1 Q-in-Q) also provide for customer control of VLAN stacking; by hiding (encapsulating) the customer VLAN ID inside the provider's VLAN identifier, rather than requiring a customer to change VLAN numbering to suit a service provider's available VLAN identifiers. In order to scale these services to a national reach, additional transmission technologies such as ATM or IP+MPLS can be used.

In these cases, as with native ATM or IP+MPLS services, customer VLAN identifiers are encapsulated within the protocol's frame formats.

4.2.13 Does not Require CE Router

Layer 2 solutions present either an Ethernet port (in the case of Ethernet MAN and VPLS services), or a simple NTU port to the customer. Layer 3 solutions must terminate to a customer edge router. Layer 2 solutions therefore eliminate the administrative overheads associated with routers — an important consideration in networks with a large number of sites.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Does not Require CE Router	+	+	+	+	-	-	-

4.2.14 Suitability for Small Sites

Most of the solutions outlined above can be purchased at a bandwidth (and price) suitable for including smaller sites into the network. All of the Layer 3 solutions operate independently of the choice of lower-layer service, and can therefore be deployed down to dial-up speeds if necessary. Frame Relay's entry-level speed is typically 64 Kbps; while Ethernet services can be delivered over DSL services, generally at speeds of at least 512 Kbps.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Suitability for Small Sites	+	+/-	+	+	+	+	+

The exception is ATM, which in Australia is typically sold with entry-level speeds in the Mbps range, although new DSL-based ATM services are also available at lower speeds and prices. Nonetheless, the complexity of designing and managing ATM based services area beyond the requirements of smaller organisations.

Suitability for small sites should be considered in conjunction with the requirement for a customer edge router, since the presence of a router will add at least a small administrative burden to the small site.

4.2.15 Support for Mixed Physical Layer Technologies (Layer 1); Support for Mixed Access Layer Technologies (Layer 2); and Native Support for Legacy Layer 3 Protocols

Both Layer 2 and Layer 3 services are independent of the underlying physical layer (Layer 1) technology in use, although some services (such as ATM) are considered unsuitable for very low-speed links due to protocol overheads.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Support Mixed Physical Layer Technologies (Layer 1)	+	+	+	+	+	+	+
Support for Mixed Access Layer (Layer 2)	-	-	-	-	+	+	+
Native Support for Legacy Layer 3 Protocols	+	+	+	+	-	-	-

At Layer 2, the service is tied to the Layer 2 access technology used. Only Layer 3 technologies operate completely independently of both Layer 1 and Layer 2 choices.

This independence is only available, however, to users whose networks only need support for IP at Layer 3. Where customers need to support legacy Layer 3 protocols (such as SNA for mainframe interconnection, IPX for legacy Novell networks, and Fibre Channel-based storage area networks), the options available are to use a Layer 2 service which is independent of the Layer 3 protocol in use; or to translate all non-IP protocols to IP to use an IP-based WAN service.

4.2.16 Less Vulnerable to Internet-Borne Security Threats

Layer 2 services are not vulnerable to Internet-borne security threats, since they are transparent to what takes place at Layer 3. Furthermore, where the Layer 2 service supports traffic partitioning, a security event such as denial-of-service traffic overwhelming

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Less Vulnerable to Internet-Borne Security Threats	+	+	+	+	+/-	+/-	-

that part of the service devoted to “best effort” traffic should not impact the bandwidth set aside for other traffic, such as voice calls. However, where a Layer 2 service is also providing a connection to the Internet, any network or device exposed to that connection might still be visible to the Internet, and therefore vulnerable in the own right. Hence, the implementation of firewalls and other security mechanisms are a fundamental part of overall network design.

Carrier-provided services such as IP+MPLS and private IP networks offer a high degree of protection against Internet-borne threats, since these networks will use an Internet gateway hosted and managed in the carrier’s network to provide customer Internet access. However, where the same customer edge router is used for Internet traffic and private network traffic, the private network may be disabled by an attack against the router (for example, a denial-of-service attack originating from the Internet, or an attack exploiting a new router vulnerability). Furthermore, the implementation of firewalls, and possibly other security mechanisms is also recommended.

Although encryption provides good security for the traffic transiting Internet-based VPNs, the service is highly exposed to Internet-borne threats, particularly denial-of-service attacks, which can disable communications without compromising the content of those communications.

4.2.17 Multi-Homing for WAN Redundancy

Multi-homing refers to purchasing multiple connections into a network from different carriers, or services from a single carrier that utilise completely separate networks. In this way, critical communications can be maintained even if an outage affects one carrier’s network. This may be important to highly critical locations, such as data centres or disaster recovery centres, where a customer needs greater than the four-nines or five-nines availability available for most carrier services.

Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Technical Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Multi-homing for WAN Redundancy	+	+	+	+	+/-	+/-	+/-

At Layer 2, multi-homing requires only that the customer obtain connections from the carriers involved. The customer’s routers may then be configured to use both services simultaneously if desired, suffering a loss of bandwidth but not of connectivity should one carrier service fail. Alternatively, the customer may configure a single primary route, reverting to the alternative service only in the event of an outage.

It should be noted that multi-homing using a single Layer 2 architecture requires the presence of multiple carriers offering the appropriate service. However, at the cost of increased complexity, a mix of Layer 2 services (for example, ATM from one carrier and Ethernet from a second) may be deployed. Multi-homing is available, but more complex, in Layer 3 services. For example, incompatibility in providers' routing schemes may constrain a customer's ability to multi-home at Layer 3.

4.3 Business Features of Layer 2 and Layer 3 Services

Table 5, below, presents the business features of Layer 2 and Layer 3 services.

Table 5. — Layer 2 and Layer 3 Services: Business Features Summary

Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Multi-point Any-to-Any Connectivity (Native Meshed Network)	-	-	+	+	+	+	+
National Coverage	+	+	+/-	+	+	+	+
Bandwidth Scalability	+/-	+	+	+	+	+	+
Real-Time Application Support (Enabled by Deterministic QoS)	+/-	+	+/-	+	+	-	-
Service Level Agreements can be based on QoS Attributes	+/-	+	+/-	+	+	-	-
Private Carrier Network	+	+	+	+	+	+	-
Less Vulnerable to Interception	+	+	+	+	+	+	-
Less Vulnerable to Denial-of-Service (DOS) and Viral Attacks	+	+	+	+	+	+	-
Very High Speed Encryption Support	+	+	+	+	+	+	+
Single Connection for Enterprise and Internet access	+	+	+	+	+	+	+
Connectivity Support for TDM PABXs	+/-	+	+/-	+/-	+/-	-	-
Connectivity Support for IP Telephony	+/-	+	+	+	+	+/-	+/-

Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
QoS Support for IP Telephony	+/-	+	+	+	+	-	-
Logical Segregation of Applications and/or Users via use of National VLANs	-	-	+/-	+	+/-	-	-
Flexibility to Change VLAN Numbering without Carrier Coordination	+	+	+/-	+	+	+	+
Flexibility to Change IP Addressing without Carrier Coordination	+	+	+	+	-	-	+/-

4.3.1 Multi-point Any-to-Any Connectivity

The natively-meshed architecture of Ethernet, VPLS, and Layer 3 services makes them highly suitable for converged multimedia business applications, which are more likely to require direct connection between offices without traversing a central site. This, for example, provides optimal performance for services such as IP Telephony and IP Videoconferencing.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Multi-point Any-to-Any Connectivity (Native Meshed Network)	-	-	+	+	+	+	+

4.3.2 National Coverage

With the exception of Ethernet MAN services all of the Layer 2 and Layer 3 services are available for nation-wide deployment. And, it is possible to

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
National Coverage	+	+	+/-	+	+	+	+

purchase WAN interconnectivity for Ethernet MANs.

4.3.3 Bandwidth Scalability

With the exception of Frame Relay, all of the Layer 2 and Layer 3 services offer very high bandwidth scalability, with maximum access speeds in the Gbps range.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Bandwidth Scalability	+/-	+	+	+	+	+	+

4.3.4 Real-Time Application Support (Enabled by Deterministic QoS)

ATM, VPLS, and IP+MPLS are the services which best support real-time applications such as IP Telephony and IP Videoconferencing. Ethernet (unless Metro Ethernet Forum bandwidth profiles are used) and Frame Relay services have more limited QoS capability, but can be managed in such a way as to provide adequate real-time application support. Best-effort IP and Internet-based VPN services should not be used where real-time applications require business-class performance.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Real-Time Application Support (Enabled by Deterministic QoS)	+/-	+	+/-	+	+	-	-

4.3.5 Service Level Agreements Based on QoS Attributes

Carriers offering services such as ATM, VPLS, Ethernet (where Metro Ethernet Forum bandwidth profiles are used) and IP+MPLS will typically offer either standard or negotiated SLA commitments associated with different traffic classes. This allows customers to plan their network traffic requirements and, more importantly, to seek correction and penalties where services fail to meet their SLA targets.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Service Level Agreements can be based on QoS Attributes	+/-	+	+/-	+	+	-	-

These are available to a lesser extent in Frame Relay and standard Ethernet services, within the more limited range of QoS management available under these services. Best-effort IP and Internet services do not offer SLA commitments associated with QoS parameters, except in the unusual case that the customer is able to secure a commitment to service availability.

4.3.6 Private Carrier Network

Internet VPNs do not operate over a private carrier network. While they are suitable for some applications (particularly connecting remote or travelling users connecting to the enterprise network using public services), they should be avoided for permanent connections such as inter-office networks.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Private Carrier Network	+	+	+	+	+	+	-

4.3.7 Less Vulnerable to Interception, Denial-of-Service (DOS) and Viral Attacks

Deployment over private carrier networks makes most of the Layer 2 and Layer 3 options highly immune to in-transit traffic interception.

Interception of traffic over the public Internet is more likely, however it should be remembered that most attacks and data theft are directed towards computers attached to the Internet, rather than the Internet connections themselves.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Less Vulnerable to Interception	+	+	+	+	+	+	-
Less Vulnerable to Denial-of-Service (DOS) and Viral Attacks	+	+	+	+	+	+	-

4.3.8 Very High Speed Encryption Support

Where encryption is necessary to protect traffic, it is now available at "wire speed" for most technologies and at data rates up to Gbps. Encryption provides marginally less protection for Layer 3 networks, because route addresses need to be transmitted "in the clear." Layer 2 technologies can encrypt all of the traffic travelling between two network ports.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Very High Speed Encryption Support	+	+	+	+	+	+	+

4.3.9 Single Connection for Enterprise and Internet Access

All of the technologies described are able to carry private and Internet traffic if required, so long as the carrier providing the connection offers Internet gateway services.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Single Connection for Enterprise and Internet access	+	+	+	+	+	+	+

4.3.10 Connectivity Support for TDM PABXs

All of the Layer 2 technologies, and IP+MPLS, provide varying degrees of support for traditional TDM-based PABXs.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Connectivity Support for TDM PABXs	+/-	+	+/-	+/-	+/-	-	-

The best support is provided by ATM, which is well-established as the basis of large-scale voice/data networks. Frame Relay and Ethernet interfaces are also widely available for TDM PABXs. Connecting TDM PABXs to IP networks requires a media gateway, which may be implemented either as a standalone unit, or as an expansion unit for the PABX. However, the IP network should also support voice-grade QoS, which is only available where the IP network is based on technologies such as MPLS.

4.3.11 Connectivity and QoS Support for IP Telephony

All of the technologies listed above are able to provide connectivity for business IP Telephony systems.

IP Telephony is best supported by those technologies which provide good QoS management, so that the voice network is not subject to conditions on the data network.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Connectivity Support for IP Telephony	+/-	+	+	+	+	+/-	+/-
QoS Support for IP Telephony	+/-	+	+	+	+	-	-

4.3.12 Logical Segregation of Applications and/or Users via use of National VLANs; Flexibility to Change VLAN Numbering without Carrier Coordination

Where companies are using VLANs to create closed user groups, VPLS or Ethernet services that support Stacked VLANs allow the best cross-site transparency of VLANs. This enables two important features: the VLANs can be centrally administered, and users can easily connect to their "home" VLANs from any office.

VPLS provides the best support for National VLANs. And, Ethernet and IP+MPLS services can be configured to provide similar services, at the cost of greater complexity.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Logical Segregation of Applications and/or Users via use of National VLANs	-	-	+/-	+	+/-	-	-
Flexibility to Change VLAN Numbering without Carrier Coordination	+	+	+/-	+	+	+	+

Frame Relay, ATM, private IP networks without MPLS, and Internet-based IP services do not directly support national VLANs.

Most of the technologies outlined above are not sensitive to the VLAN numbering scheme used by the customer. However, a carrier-provide Ethernet service may cause conflict between customer and provider VLAN numbering, in which case technologies such as Q-in-Q (Stacked VLANs) are required to maintain separation between provider and customer VLANs.

4.3.13 Flexibility to Change IP Addressing without Carrier Coordination

Layer 2 technologies operate independently of the customer's IP addressing scheme, allowing customers to implement new services or IP addressing without serious addressing considerations.

Network Layer	Layer 2	Layer 2	Layer 2	Layer 2	Layer 3	Layer 3	Layer 3
Business Features Summary	Frame Relay	ATM	Ethernet (Point-to-Point and VLAN)	VPLS (Ethernet + MPLS with DiffServ)	IP+MPLS (with DiffServ)	IP (No MPLS)	Public Internet (IP VPN)
Flexibility to Change IP Addressing without Carrier Coordination	+	+	+	+	-	-	+/-

Coordination between provider and customer addressing is more likely to pose problems where Layer 3 carrier services are used. Even where NAT (network address translation) is used to separate provider and customer address schemes, the customer's addressing choices are constrained by the need to maintain carrier-facing addresses that fit the carrier's address schemes.

5 Business Considerations Impacting the Selection of Layer 2 and Layer 3 Services

5.1 Introduction

In this section we present the impact of business considerations on the choice of network services. Scenarios discussed include:

- **Network Consolidation** — What are the considerations for an enterprise is seeking to unify multiple networks for reasons of cost or administrative overhead?
- **Convergence** — What decisions influence the choice of carrier services for a company looking to create converged voice-data networks?
- **Legacy Services** — What impact does the presence of legacy networks have on the choice of carrier WAN services?
- **QoS and Site Connectivity** — How do QoS requirements influence network topology, and how does this impact the choice of carrier WAN services?

These are discussed in the following sections in the form of decision trees for each business case, with accompanying explanatory text.

5.2 Network Consolidation

Figures 15-17 show a simplified example of a network currently using a Layer 2 data service to connect each site, along with Internet access at each site. This topology can be consolidated down to a single network service. After consolidation, each site would have a single connection and use a carrier-hosted gateway for Internet access.

Examples are shown with Layer 2 (Figure 16) and Layer 3 (Figure 17) services.

Figure 15. — Before Network Consolidation

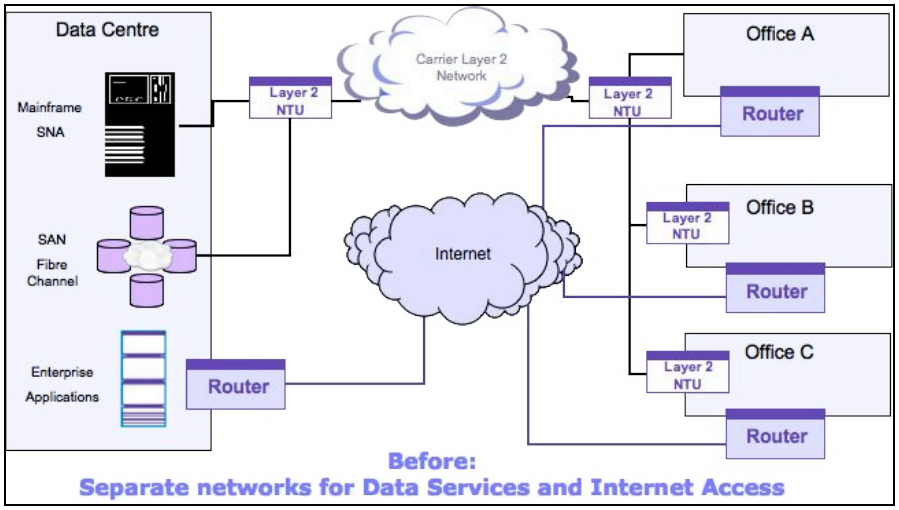

Operating multiple Internet gateways multiplies service costs, and also causes both administrative and security headaches for the enterprise.

While the existing Layer 2 service is providing good service to the organisation, it's expensive to maintain two separate links at each site; one for Internet access, and the other for the corporate data network.

By consolidating all data services onto a single infrastructure, the customer can realise greater efficiency, lower cost, and a simpler network.

While this Figure shows a Layer 2 VPLS service being used as the basis of the consolidated network, Layer 3 services may also be suitable for such a project (see Figure 17).

Figure 16. — After Network Consolidation: VPLS

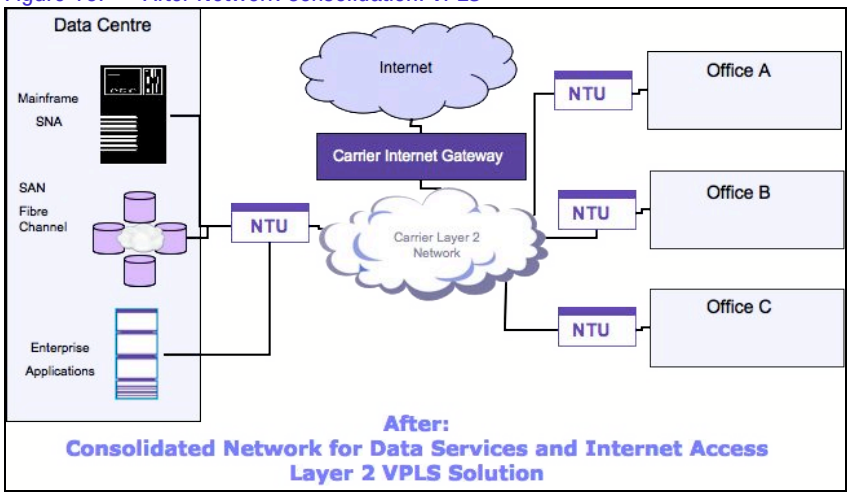

Figure 17. — After Network Consolidation: IP or IP+MPLS

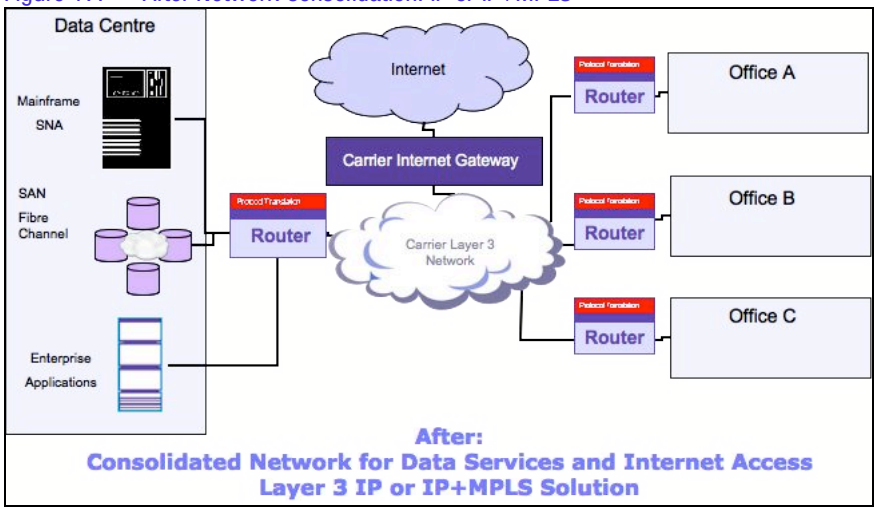

Figure 17 shows a Layer 3 network service providing data network connectivity as well as Internet access.

One important difference between the Layer 2 and Layer 3 solutions is the requirement for customer premise equipment to provide protocol translation services for legacy applications that utilise SNA and Fibre Channel protocols to traverse the Layer 3 IP-based data network.

5.2.1 Decision Points

As the Decision Tree in Figure 18 shows, the first question faced by a company which decides upon a network consolidation project is whether it has a preference for maintaining a Layer 2 switched network, or a Layer 3 routed network.

Factors influencing this choice may include:

- Whether the company has already invested in an extensive private routed network scheme;
- Whether it needs to support multiple Layer 3 protocols (in which case an all-IP WAN would require protocol translation); and
- Whether it plans to implement Ethernet-based WAN services.

These are discussed in further detail in later sections.

Once these factors are assessed, the enterprise can proceed with decisions based on factors relevant to Layer 2 and Layer 3 services.

If the enterprise decides to implement a **Layer 3** service, its next key question is whether it has a requirement to support **QoS** for real-time applications across the WAN connection.

If QoS is not required, the customer can choose either a carrier-based private IP network, or a carrier-based IP+MPLS solution.

If QoS is required in a Layer 3 environment, the customer should examine carrier-based IP+MPLS network solutions.

Figure 18. — Network Consolidation Decision Tree

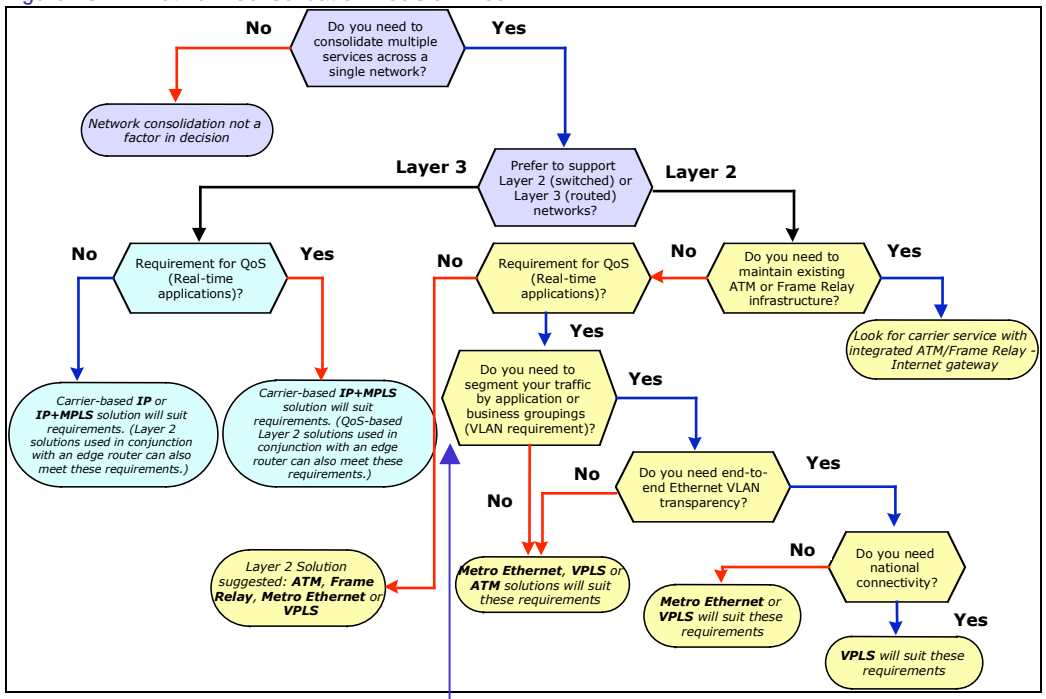

If **VLANs** are not required, ATM, Metropolitan Ethernet or VPLS may be suitable for the customer's requirements. However, if VLANs are required, the customer should consider whether or not it requires end-to-end VLAN transparency. (Typically, VLANs have been implemented on a site-by-site basis. However, with the advent of services such as Metropolitan Ethernet and VPLS, it has become feasible to make VLANs in different sites behave as a single VLAN.)

If end-to-end **VLAN transparency** is not required, ATM, Metropolitan Ethernet, or VPLS may suit the network's requirements. However, if VLAN transparency is required one further question must be posed: Is national connectivity required? A customer whose sites are located in a single city can achieve end-to-end VLAN transparency using either Metropolitan Ethernet or VPLS services. If **national connectivity** is required, a customer with multiple Metropolitan Ethernet services would need to purchase a separate WAN interconnect service to pass traffic between cities. Therefore, where national connectivity is required, a VPLS service may best suit the customer's requirements.

If the business requires a **Layer 2** network, it has several other considerations to inform its choice.

Does the network need to maintain support for existing ATM or Frame Relay network infrastructure?
Such requirements may arise because existing network services remain under long-term contract, or because some sites are in locations that are not well served by newer technologies.

If ATM and/or Frame Relay support is required, the consolidated network should be based on these technologies.

If ATM and/or Frame Relay support is not required, the customer can proceed to the next stage in selecting a Layer 2 service, the determination as to whether **QoS support** is required.

If QoS support is not required, the customer can select whichever of the available Layer 2 services (ATM, Frame Relay, Metropolitan Ethernet or VPLS) best suits its bandwidth, geographical, and cost requirements.

Nation-wide VLAN numbering provides various advantages:

- **Consistent Security** — A VLAN created to segregate a sensitive group of users can be made seamlessly available to those users as they move between offices.
- **Consistent Network Administration and Behaviour** — VLAN characteristics can be maintained from a single location, and VLANs in different offices will exhibit similar behaviour.
- **Simplified WAN Connectivity** — If VLANs are not visible in multiple offices, some other mechanism is required to connect users in one location to a VLAN in another location. For example, the business may assign specific WAN routes to particular VLANs. This requires more complex WAN routers, and also makes route administration more complex.

However, the use of a ubiquitous VLAN scheme across multiple network locations may have the following adverse impacts:

- **Network Upgrades** — If a customer has a large number of sites where VLANs have not been implemented (for example, small remote networks), it may need to upgrade those sites to create a consistent view of the enterprise-wide VLANs.
- **VLAN inconsistency** — If VLANs have been implemented as silos in different locations, creating a consistent VLAN scheme may require an unacceptable administrative effort. Furthermore, different sites may use equipment, which while supporting interoperable VLANs uses proprietary configuration and administration tools.

5.3 Convergence

The growing maturity of business IP Telephony environments are driving demand for networks able to support converged applications. Suitable QoS is available in both Layer 2 and Layer 3 carrier services, so a service choice decision should be driven by the technology that best suits the customer's other network requirements.

Figure 19, below, shows the decision tree associated with the selection of a carrier service to support converged applications.

VLANs may not be required in small-scale networks, where a single 100 Mbps Ethernet LAN has sufficient capacity to support the business' data requirements without impacting the quality of voice calls. However, at even moderate network scale, VLAN support is recommended to ensure that unexpected traffic on the data network (for example, large file transfers) does not prevent the operation of the voice system.

The next decision then involves an organisations' preference for either **Layer 2 or Layer 3 services**.

Where **Layer 3** services are preferred, an IP+MPLS best suits the requirements of converged applications.

Where **Layer 2** services are preferred, a further question pertaining to **national connectivity** must be posed.

Organisations with **national connectivity** requirements are best served by VPLS or ATM solutions, whereas organisations with metropolitan connectivity requirements are best served by VPLS or Metro Ethernet services.

Figure 19. — Convergence Decision Tree

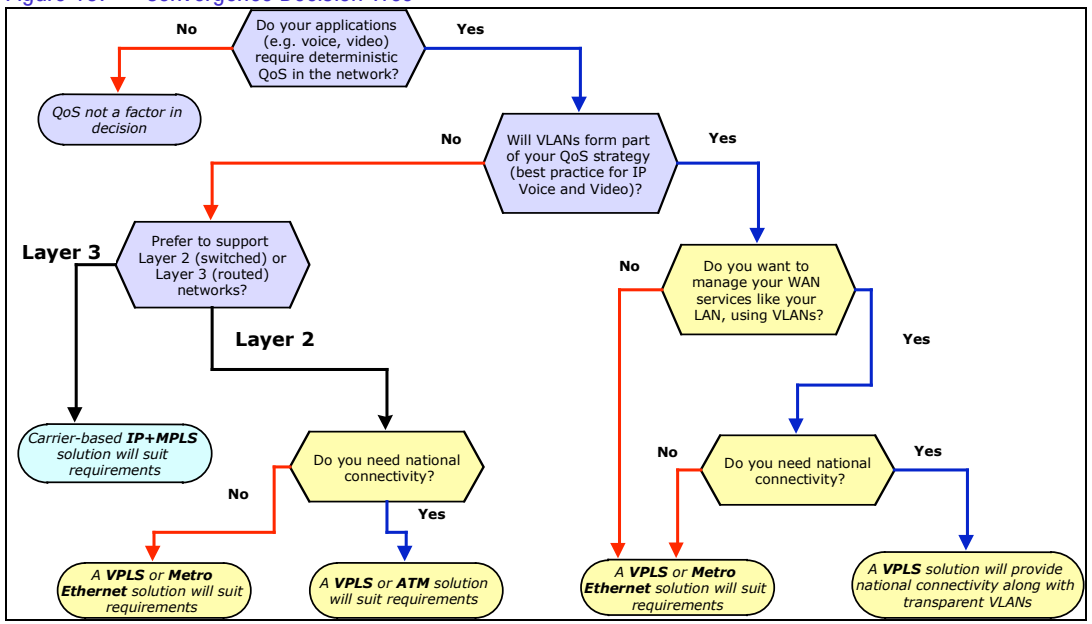

Once the decision to implement converged applications has been made, and **QoS support** therefore becomes a requirement for the network, the customer then needs to consider whether **VLANs** will be required in the enterprise network.

Where **VLANs** will form an integral part of an organisation's QoS strategy, a further consideration is whether the organisation prefers to manage its WAN in the same manner as its LAN. In circumstances where this is not a requirement, the organisation can choose from VPLS or Metro Ethernet services.

If **national connectivity** is required, a VPLS service may best suit the customer's requirements.

Organisations with metropolitan connectivity requirements best served by VPLS or Metro Ethernet services.

5.4 QoS and Site Connectivity

Where networks require QoS support for real-time applications, customers should also take into account the relationship between network topology and QoS.

In some businesses, a hub-and-spoke architecture may be an accurate reflection of business communications. The business model may rely on the ability of branch offices to contact head office staff. On the other hand, a business whose operations depend on communications between branch offices, may be best served by a mesh architecture. The decision tree for QoS and Site Connectivity is shown in Figure 20, below.

Figure 20. — QoS and Site Connectivity Decision Tree

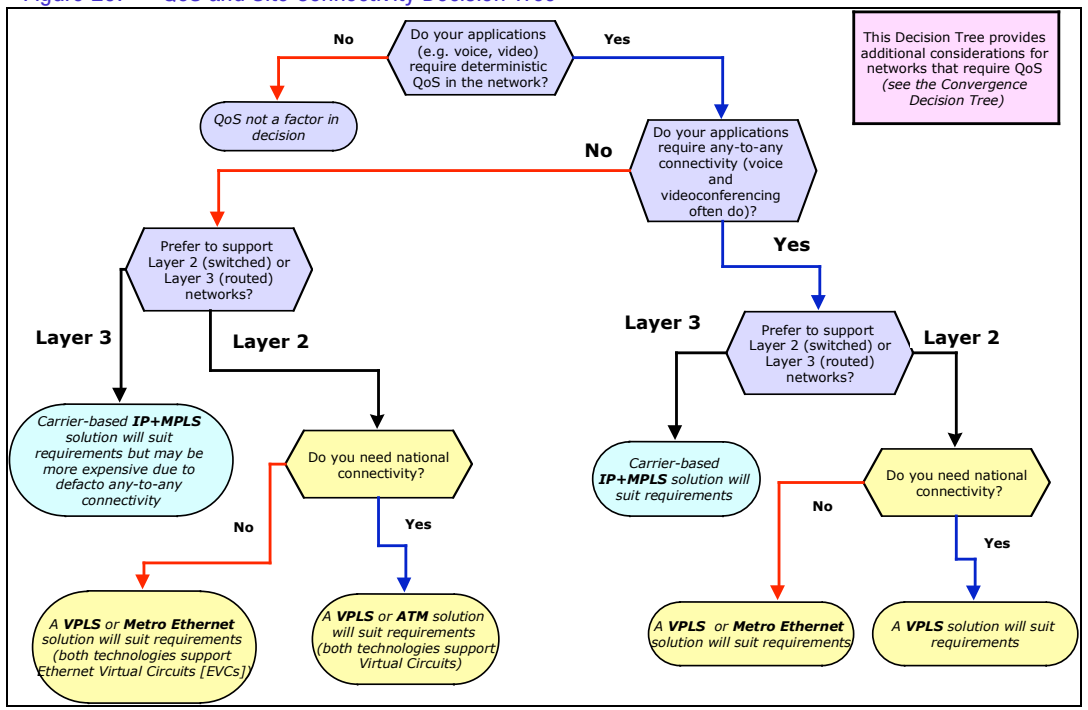

If **any-to-any connectivity** is not required, the next question to consider is the organisations' preference for either **Layer 2 or Layer 3 services**.

Where **Layer 3** services are preferred, an IP+MPLS best suits the requirements of converged applications.

Where **Layer 2** services are preferred, a further question pertaining to **national connectivity** must be posed.

Organisations with **national connectivity** requirements are best served by VPLS or ATM solutions, whereas organisations with metropolitan connectivity requirements are best served by VPLS or Metro Ethernet services.

If **any-to-any connectivity** is required, the next question to consider is the organisations' preference for either **Layer 2 or Layer 3 services**.

Where **Layer 3** services are preferred, an IP+MPLS best suits the requirements of converged applications.

Where **Layer 2** services are preferred, a further question pertaining to **national connectivity** must be posed.

Organisations with **national connectivity** requirements are best served by VPLS solutions, whereas organisations with metropolitan connectivity requirements are best served by VPLS or Metro Ethernet services.

5.5 Legacy Network Support

The existence of, and ability to support, legacy network protocols may have a significant impact on the customer's choice between Layer 2 and Layer 3 services.

Layer 2 services operate independently of the Layer 3 protocol in use. Frame Relay, ATM, Ethernet, or VPLS services can, with suitable interfaces, carry SNA mainframe traffic or Novell IPX traffic alongside newer TCP/IP traffic without requiring protocol translation. Legacy protocols can only be carried across Layer 3 networks if the traffic undergoes protocol translation, or if it is encapsulated in IP packets. Figure 21, below, shows the Decision Tree applicable where the network must support legacy protocols.

If native support for **legacy network protocols** is not required, the customer can choose between Layer 2 and Layer 3 services.

Where **Layer 2** services are preferred, readers should follow the decision tree described in the box on the right-hand side of this page.

If the enterprise decides to implement a **Layer 3** service, its next key question is whether it has a requirement to support **QoS** for real-time applications across the WAN connection.

If QoS is not required, the customer can choose either a carrier-based private IP network, or a carrier-based IP+MPLS solution.

If QoS is required in a Layer 3 environment, the customer should examine carrier-based IP+MPLS network solutions.

Figure 21. — Legacy Networks Decision Tree

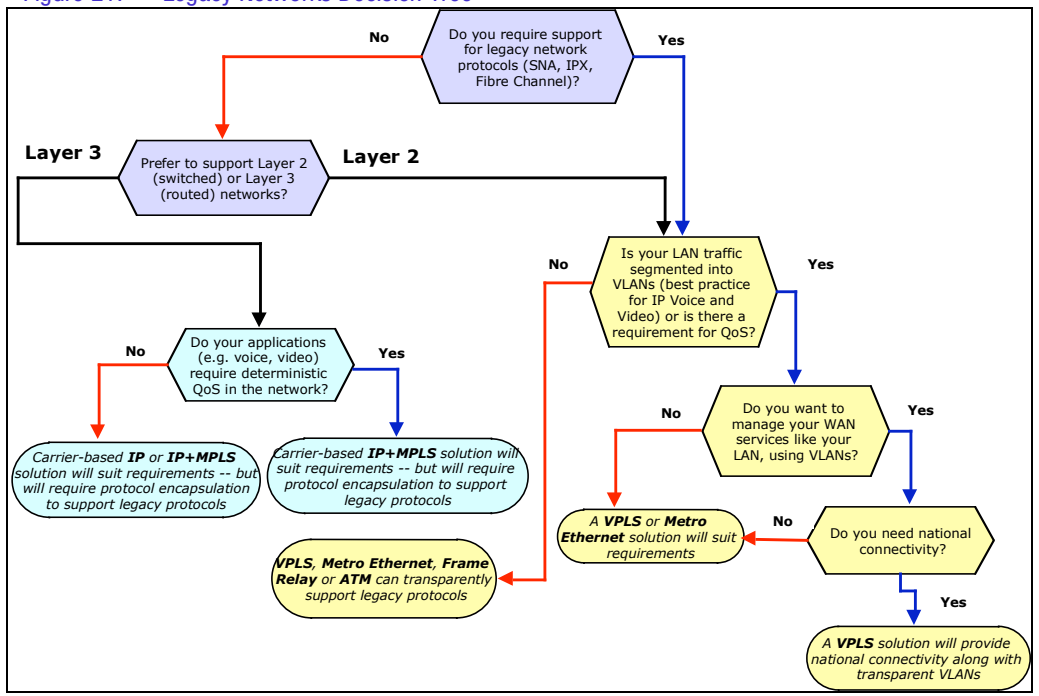

If native support for **legacy network protocols** is required, the next determination is whether **VLAN and QoS support** is required. Where not required organisations can transparently support legacy protocols with VPLS, Metro Ethernet, Frame Relay or ATM services.

Where **VLAN and QoS support** is required, a further consideration is whether the organisation prefers to manage its WAN in the same manner as its LAN. In circumstances where this is not a requirement, the organisation can choose from VPLS or Metro Ethernet services.

If **national connectivity** is required, a VPLS service may best suit the customer's requirements. Where national connectivity is not required, either VPLS or Metro Ethernet solutions will provide an appropriate service.

6 Concluding Remarks

This white paper has presented a range of information designed to assist organisations' WAN service decision-making processes. New to the technology mix is VPLS — the latest in a long line of data network technology enhancements to be offered as a service, and one that offers a number of key features that warrant its consideration in an organisation's network planning processes:

- Any-to-any or virtual circuit based connectivity;
- Customer control of routing;
- Customer control of VLANs;
- VLAN transparency;
- National connectivity;
- Traffic engineering for true QoS and high reliability;
- Scalability;
- Security;
- The cost advantages of Ethernet interfaces;
- A variety of physical layer solutions — fibre, copper, wireless;
- Bandwidth flexibility from Kbps to 10 Gbps; and
- Service pricing that takes advantage of commodity Ethernet equipment.

VPLS combines the advantages and features of Layer 2 and 3 networks, into a simple, yet flexible solution that draws upon established Ethernet and MPLS standards.

This paper provides a general discussion of WAN technologies and services. However, decisions should not be made without specific consideration of the following factors:

- An organisation’s business requirements;

- The network characteristics required by an organisation's applications;
- Near-term requirements for new applications such as IP Telephony or IP Video;
- Existing investment in network equipment and software;
- Existing contractual commitments for networks services, support and maintenance;
- The skill sets of an organisation's IT staff;
- The inherent capabilities of various WAN technologies;
- Service providers' implementation of various WAN technologies — just because a technology can support a feature or characteristic, do not assume that a service provider has implemented these features;
- Cost comparisons of various network solutions — including installation costs, migration costs, costs for new or upgraded equipment, staff training, and recurring service fees; and
- The total cost of ownership (TCO) of various network solutions calculated over the proposed contract term.

As demonstrated throughout this paper, different technologies (and the services built upon these technologies) are best suited for different combinations of the above considerations. While most of the technologies presented in this paper are familiar to technology readers, we have also provided information on a new service type — Virtual Private LAN Service (VPLS), which is a best-of-breed architecture that builds on key aspects of Frame Relay and ATM, using IEEE, Metro Ethernet Forum (MEF) and IETF MPLS standards.