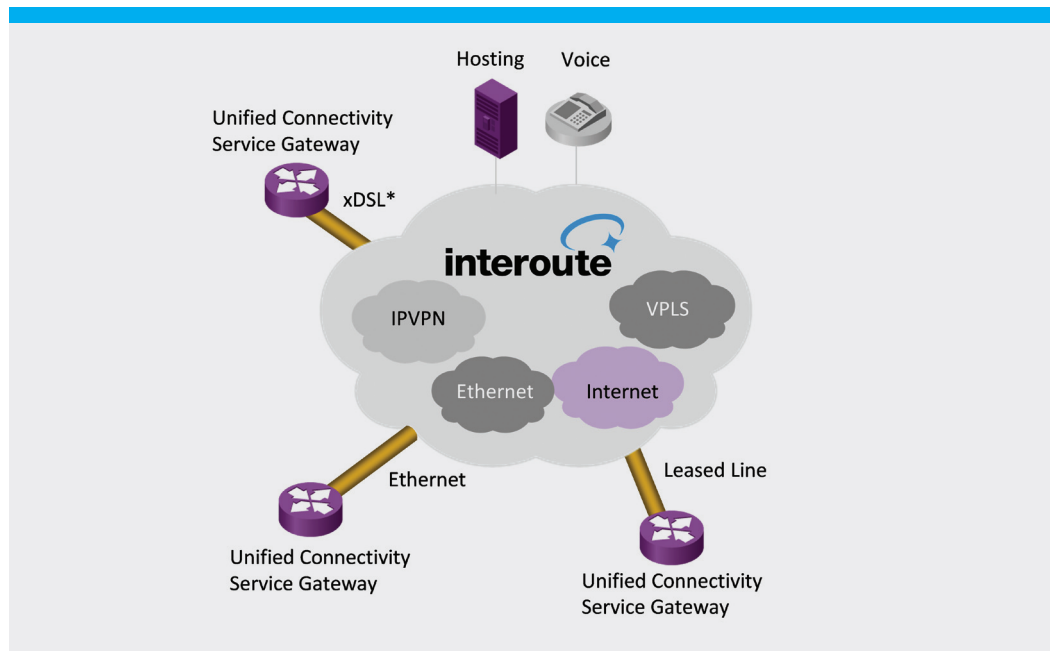


Unified Connectivity: The birth of the enterprise network as a dynamic ICT services platform

FIGURE 1
Unified Connectivity
Service Gateway



INTRODUCTION

Unified Connectivity (UConn) is a revolution in WAN outsourcing that eliminates the dependence on access technology and creates a dynamic connectivity environment for the enterprise. UConn is an integrated connectivity platform that allows an IT manager to securely deliver any voice, computing and data service(s), over any access type. It redefines the approach to ICT service provisioning by giving the IT manager a scalable, secure and unconstrained connectivity platform on which they can build their services. This dramatically reduces the complexity at the network edge, smooths migrations and enables IT departments to implement network changes to answer the immediate business needs, rather than when the 3 to 5 year procurement cycle allows it.

Using the new UConn Service Gateway, Interoute extends the core network out to the enterprise premises and can deliver any connectivity service be it MPLS VPN, VPLS based Ethernet, Internet or dedicated service capability such as a SIP trunk or data extension over any access technology, as illustrated above. These services can be delivered individually or in multiples, enabling IT departments to select a different service mix on a site-by-site basis. Taking into account the depth of coverage, range of services offered and the way that services are abstracted from physical access Unified Connectivity is unique in the market place today.

Abstract

Choosing a managed wide-area network solution can be difficult; this document aims to help ICT professionals and business managers in making the right choice for their data networks.

Today's service providers offer a plethora of different technologies to achieve similar ends. The technologies vary, promoted by services providers and vendors alike, and promise numerous advantages, but often impose punitive constraints on the network topology and design.

Worse, the technologies, configuration options, and even terminology are often incompatible with one another, forcing enterprises into choosing the most suitable technology of compromise for their entire network and preventing them from fairly comparing service providers. And changing technologies and configurations can be a very costly exercise.

Should an IT manager choose a conventional IP VPN based on layer-3 IP routing, or a more transparent multi-site Ethernet LAN? Alternatively, perhaps it is better to procure ordinary ubiquitous Internet access services and build an overlay "do-it-yourself" topology with IPSec tunnel technology? How exactly does the choice of technology, and topology, affect the performance of applications that really matter to the business?

Today service providers offer enterprises little in the way of answering these questions and often mandate a fixed technology – and thus topology – path. The nature of WAN outsourcing means that typically these decisions are ones that the IT manager will have to live with for 3 years.

Interoute simplifies the act of selecting managed WAN services through the use of a new development called Unified Connectivity: a fully-featured network capability and services platform that allows a enterprises to make use of multiple managed Ethernet and IP services with single cost-effective customer premise equipment (CPE) presence that depends only on local bandwidth requirements.

Unified Connectivity is set to revolutionise the way IT departments design, develop and procure connectivity for their user community in the broadest sense of the definition. For the first time the IT department will be able take complete control of the network topology free from the limitations of scale and in house skills. Unified Connectivity not only simplifies the complexity of WAN outsourcing through eliminating the constraint of access technology it offers both management options (layer 2 and layer 3) combined. Because the actual connectivity service is separated from the physical constraint of the access service and physical demarcation device, IT departments can modify topology and capability without the need for a site visit or any physical intervention from their network service provider. This separation of the logical and physical elements of the network creates a dynamic, scalable and future proofed connectivity platform servicing the entire enterprise from the remote worker to the branch office, headquarters or core data centre.

With Unified Connectivity change becomes an intrinsic part of the service and the proposition.



Benefits

UConn delivers true flexibility. By separating the physical connectivity from the services that are available over that connection, enterprises have the ultimate flexibility to change service mix in line with business needs and not in line with a 3-5 year telecommunications procurement cycle. To enable this flexibility, all of the UConn enabled products have a harmonised contract & SLA structure, so changing services does not involve a major legal exercise.

Reduction in complexity and cost

- UConn reduces the amount of circuits, equipment, power, and people required to deliver and maintain multiple connectivity services in an IT organisation, ultimately reducing costs. It also provides access to other services that enable similar efficiencies in voice, hosting, computing and CDN services.

Pricing is irrespective of technology

- Because the UConn platform delivers a range of options, Interoute is committed to delivering the one that is right for a business need. Our pricing policy also reflects this; for example, charging exactly the same price for a layer 2 (VPLS) and a layer 3 (MPLS) service that uses the same access circuit and service gateway.

Embrace change, don't penalise it

- Modifying a service on UConn enabled sites is now completed by a software change, meaning it can be undertaken mid-contract. Interoute does not make excessive charges to change services, and in cases where physical network elements or locations do not change, there is no requirement to extend your contract.

Rich feature set

- Supporting numerous access options from 3G mobile, DSL and Leased Line to Gigabit Ethernet, UConn is a single ubiquitous platform that delivers the right connectivity based upon your requirements at each individual location.

Managed WAN technology options

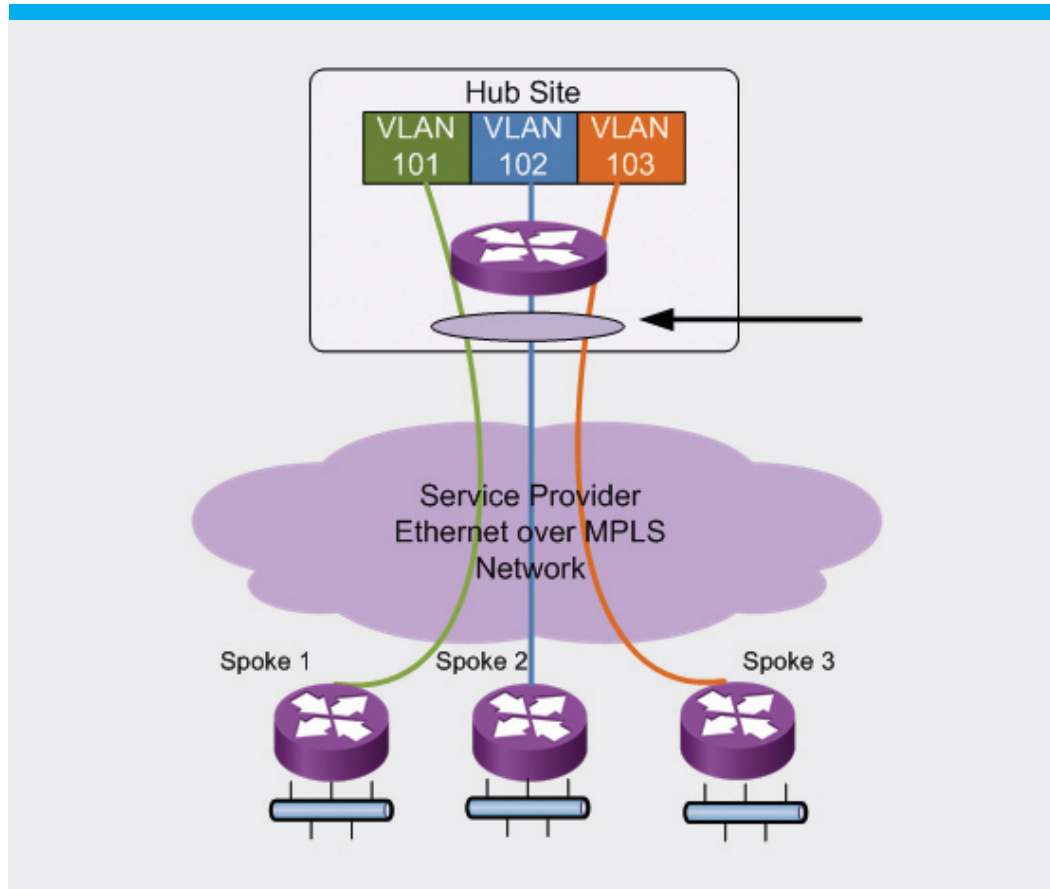
There are many options available for enterprises seeking a service provider to operate their wide-area network infrastructure.

At the most basic level, enterprises can opt for traditional "leased-line" services which provide point-to-point connectivity between two sites. The enterprise IT department operates LAN-based routers and switches and uses the appropriate interface cards on the router or switch devices in order to connect the leased-line services into the LAN infrastructure. Bandwidth on the leased-line service is fixed across the whole route, specified in advance and cannot be changed. The IT department, in operating the equipment at either endpoint, has a large degree of control and responsibility.

The advantage with these services is that they are extremely well-understood and standardised and, for the CTO who still wants to have a high-touch level of involvement with the operation of his WAN infrastructure, they represent a sound investment.

The disadvantage is that the services are typically not flexible – bandwidth is fixed, and any changes need expensive physical upgrades – and the services are often expensive because of the fact that the service provider must "nail-up" bandwidth resources for the service irrespective of whether they are used or not. This leads to situations where the server infrastructure location is a compromise between the availability of access circuits with appropriate bandwidth, acceptable physical environmental capabilities – such as power, space, airflow – and end-user demands.



FIGURE 2 Virtual
Circuit VPN
(using EPLs)

Building virtual topologies with virtual circuits

Building on from this basic level of functionality is virtual circuit technology. Rather than establishing point-to-point circuits between enterprise sites, virtual circuit technology requires only the establishment of physical bandwidth to a service provider's closest POP. From there enterprise traffic is aggregated efficiently on an optimally-operated and multiplexed packet backbone network along with similar traffic from other customers. Virtual circuits are then established between enterprise sites and bandwidth resources are logically assigned as appropriate to each virtual circuit to the sum of the underlying access bearer circuit.

Since bandwidth is no longer "nailed-up" across the service provider network, but shared between all enterprises, this allows an amount of flexibility in changing the topology: it is possible to reconfigure the destination and bandwidth of virtual circuits to adapt to new network requirements with little or no physical impact so long as the requirements don't exceed the capacity of the local access bearer.

There are various generations of virtual circuit technology, including:

- **Frame Relay** – for access bandwidth to 2Mbps.
- **ATM virtual circuits** – for higher access bandwidths, typically 45Mbps, 140Mbps.
- **Ethernet Virtual Private Line** – with flexible bandwidth options.

Variations on the virtual circuit technology also employ encrypted IPSec tunnels in order to use public Internet networks for private packet transport. Functionally these technologies are very similar in terms of what they deliver to the enterprise, but the significant advantage of the Ethernet-based technologies is the implicit compatibility with common LAN infrastructure, such as routers or switches, without needing expensive SDH or TDM line cards.

These mechanics allow the service provider an economic way to provide long-haul bandwidth to multiple enterprises efficiently.

The advantage to the enterprise of this technology is that it is presented with a cost-effective interface to the WAN service. The service provider is able to provide long distance bandwidth to multiple enterprises economically and efficiently by taking advantage of the fact that different enterprise traffic profiles can be mixed in order to "over-book" certain routes. An additional advantage is that the actual logical topology – the connectivity of sites – is not physically bound, and can be changed with logical configuration work if new network demands dictate.

The virtual circuit technology is still based upon point-to-point links, however, which means that network designers need to define explicit connections between sites using knowledge of the traffic flows within the enterprise network in order to determine optimal virtual circuit configuration. In addition, complex provisioning software is often needed by the service provider network in order to auto-provision virtual circuits which can otherwise be laborious.

MPLS IP VPNs put routing in the Cloud

To address this limit in flexibility and overhead in configuration, service providers and vendors developed MPLS technology to capitalise on the simplicity and privacy of virtual circuits but with the added flexibility of automatic meshed setup to connect all sites together.

MPLS VPNs make use of this to provide private any-to-any IP-layer connectivity on a shared packet backbone. The any-to-any connectivity through automatic full-mesh routing negates the need to rigidly pre-determine site-to-site traffic demands for a virtual circuit traffic matrix, and the fact that all enterprises make use of a shared packet backbone means that the service provider gains huge economies of scale by operating a single network, and the maintenance procedures that goes with it – shortest-path routing policy for minimum latency, capacity planning for bandwidth guarantees, software upgrades for security – only once.

The service provider typically provides a managed IP router on the enterprise site as part of the service which maintains the private IP routing information necessary to direct traffic within the enterprise network. On the service provider network, the enterprise private IP routing information is maintained within a virtual routing table and remains isolated from all other enterprises, permitting multiple enterprises with overlapping address space to happily co-exist on a single packet backbone.



For enterprises with pure IP-layer topologies – and the amount of legacy devices and protocols that demand exception to this is ever declining – managed MPLS VPNs represent a productive opportunity to outsource the operation of wide-area networks with little risk or compromise. Enterprises have only the need to correctly identify the aggregate amount of bandwidth that each site demands, and don't need to possess an intricate knowledge of site-to-site traffic matrices. Enterprise sites are automatically assumed to be divided into logical IP subnets, in line with best common practices, and most service providers provide network supporting services such as DHCP to assist with IP address allocation, or routing protocol integration in order to interconnect with IP networks at each site.

From Ethernet WANs to Ethernet VPNs using VPLS

Building on this capability, the most recent technology addition – VPLS – allows enterprises to do exactly the same thing with private Ethernet networks, taking away any dependency on IP topology information.

VPLS technology offers significant advantages by making, for the first time, large scale wide-area networks possible at the basic Ethernet level. Some of the features that made Ethernet networks simple and convenient for use in enterprise LANs have also hampered scalable large-scale deployments across WANs.

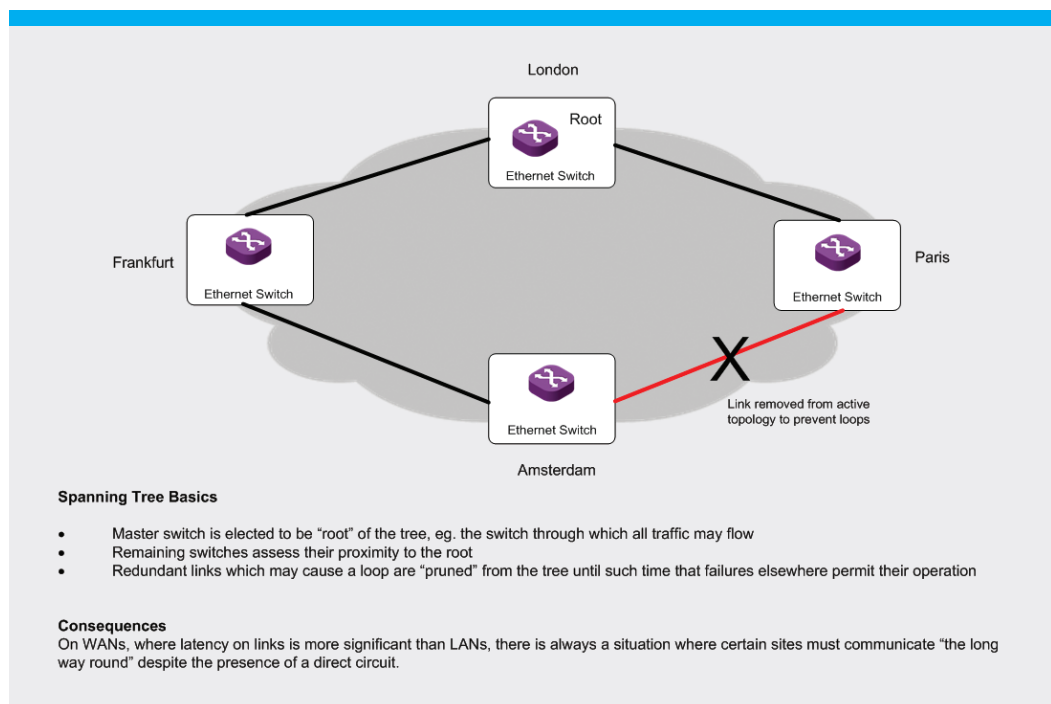
Specifically;

- The default forwarding decision for switched Ethernet networks is to flood to all ports in order to ensure a packet reaches its destination,
- There is no implicit facility within the Ethernet frame that limits the distribution of the frame, in contrast with the IP "time-to-live" header that is fundamental to limiting the impact of IP routing loops.

As a consequence of this, WAN networks based on Ethernet have traditionally been dependent upon the IEEE 802.1d Spanning Tree Protocol (STP) in order to operate optimally loop-free networks with the underlying redundancy required to tolerate the link failure and repair situations commonplace on circuits of large distance. STP provides an assurance of a working topology that will service all nodes in the network in a loop-free manner, but it doesn't provide any fine-grained control over how the traffic is routed. In local area networks, this matters little, but in wide-area networks, where bandwidth is more expensive, and links incur delay, there is much more interest.

The other main obstacle to building large-scale Ethernet networks is centred on the fact that Ethernet frame forwarding, as implemented by Ethernet switches, is based upon a destination MAC addresses. This MAC address is a "burned-in" entity which is assigned once to a network interface card at manufacture, rather than being logically assigned based upon department, network or location by a network administrator and as a result there is no opportunity to aggregate routing information, which is what allows IP routing its scalability.



FIGURE 3 Spanning
Tree Protocol
Operation

As a result, Ethernet switch devices must maintain a view of every single end-device on the network, rather than simply tracking logical groups of devices as they would do in the IP world.

VPLS negates the need for the traditional Spanning Tree Protocol by grouping all WAN interfaces together and changing the forwarding rules to establish a WAN/LAN boundary. Frames are only ever forwarded across this boundary, never within it. Not only does this mean that network-wide loops are avoided, it also means that the same automatic full-meshing that is possible with MPLS for IP VPNs can easily be applied to Ethernet VPNs, and so packet delivery is more direct. In addition to this, "Hierarchical" VPLS further encapsulate Ethernet MAC frames within a site-level address in order to provide the aggregation and grouping affect that allows scalability.

Even with these advances, most network design experts agree that Ethernet WANs, whether implemented with VPLS or other technology, should really be limited in size since Ethernet-layer protocols still require the ability to broadcast (talk to all stations) which places a limit on scaling. In addition, establishing a logical addressing layer that has significant meaning allows finer control in routing traffic and can also be useful for other parts of the enterprise network, eg. security policies, capacity planning.

Typically VPLS-based Ethernet WANs lend themselves very well to supporting:

- small-to-mid-sized LANs with no established logical addressing policy
- service-provider demarcation zones for larger networks.

In the first case, the enterprise is essentially not performing any logical-layer routing, while in the second case, the enterprise is providing his own routing function across the service provider facility.

Choosing a Service Provider and a Technology

From the many myriad of options available, enterprises need to choose not only a service provider they can trust, not only one who provide a cost-effective option and reliable SLA, but also one who operates a technology that supports the topology that they want with the features and benefits that they need.

While enterprises are interested in the different technology types and the benefits to be gained, they don't necessarily have the expertise, and this can be a difficult decision to make.

If that weren't enough, it is hard for a IT manager to change his or her mind at a later date if the original technology choice was an inappropriate one since different service providers are often tied to different technologies.

The ones that can offer different technologies can rarely do so without effecting physical or contractual change, which may mean the enterprise is forced to live with an inappropriate resource and topology model for up to three years, based upon current typical contract lengths.

Introducing Unified Connectivity

Interoute's Unified Connectivity service revolutionises the business of selecting managed WAN services by removing the need for the enterprise choose between what are, often deeply religious, technology issues.

Modern packet networks provide the advanced services that they offer to enterprises through the use of MPLS technology. Underneath the services offered to enterprises, it is MPLS that provides the private virtual circuits (or label-switched paths in MPLS terminology) that link together enterprise endpoints as and when topology requires.

The MPLS "provider edge" router on the service provider network is the place where service aggregation occurs, and it is this point where it is most convenient to offer multiple services to an enterprise. Unfortunately, since most enterprises often don't share physical sites with service providers, it is not always possible to extend several services to an enterprise site without considering the specific capabilities of the local access technology. While some access technologies can lend themselves to natively delivering multiple logical services that remain separated, others – such as common xDSL deployments – cannot.

By extending the MPLS technology to the very edge of the network – the enterprise site – Interoute's Unified Connectivity service offers the enterprise its own virtual MPLS backbone upon which services dedicated to it can run in the same identical fashion to how they would run on a shared IP backbone. Essentially, the enterprise gets its own "virtual service provider" to deliver its services using a combination of the main service provider's backbone and whatever access technologies are appropriate for the site.

By pushing the MPLS capability to the edge of the network, the access technology is significant only for its bearer capacity rather than its capability. For example, VPLS VPNs are traditionally only available on Ethernet access interfaces, yet Unified Connectivity can offer these services independent of the WAN media used for access to the enterprise site.



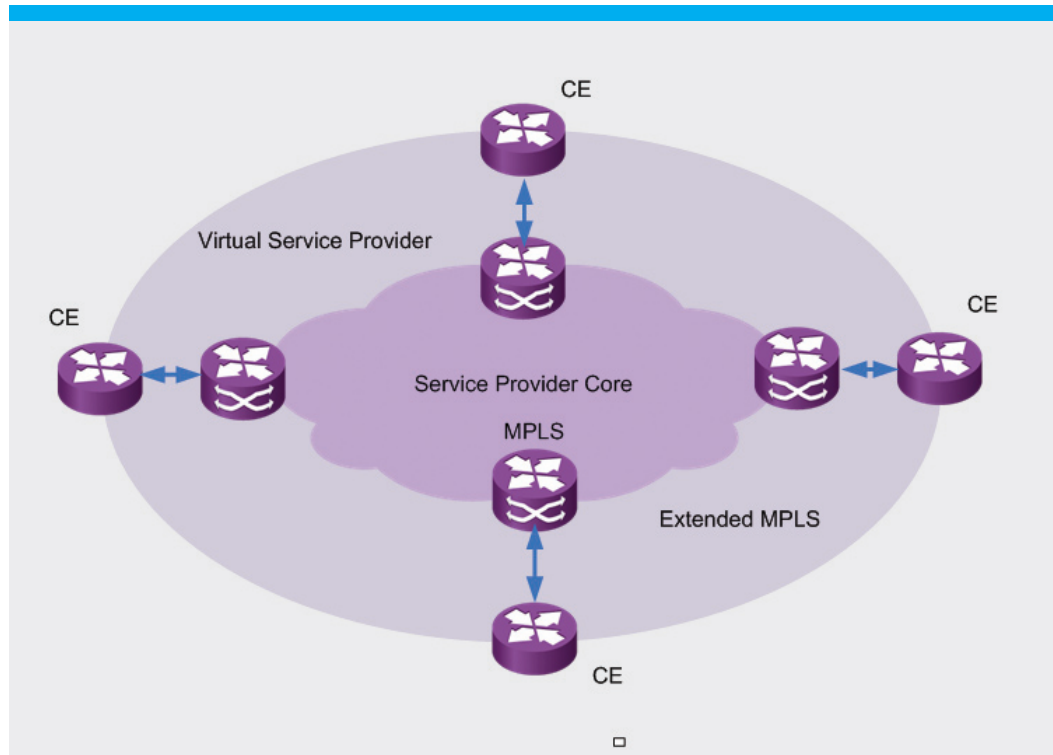
How does Unified Connectivity help organisations manage their WAN?

Rather than making a technology and topology decision straight away, all an enterprise needs to do is to consider the following elements of all its sites:

- **Required network bandwidth:** from entry-level DSL/2Mbps to high capacity Gigabit
- **Availability:** standard single-router service, or high-availability fully-redundant
- **Number of connectivity services:** the number of different services that might be required

This initial information is all that is required for an enterprise to assess the cost of an Interoute managed WAN solution. It doesn't matter whether an enterprise wants a virtual-circuit service, an MPLS-based IP VPN, an Ethernet VPN, or even domestic Internet access for DIY IPsec VPN, Interoute Unified Connectivity can support any or all topologies throughout the lifetime of the service as a simple logical configuration change, typically actionable through the intuitive web-based Hub interface. By making use of a simple and concise product code scheme, an enterprise selects the WAN connectivity option required on a site-by-site basis. A simple bandwidth calculator can help decide the amount of bandwidth required at a given site dependent upon number of users.

FIGURE 4 Virtual Service Provider Network



Complex technology and topology decisions such as whether to use IP VPN routing or private wire Ethernet virtual circuits, or meshed Ethernet VPN can be deferred until late in the deployment process for handling by administrators and engineers closer to the actual network.

Once the actual sites are specified and understood, IT department can select one or many logical services which can be offered at the site. All Interoute Unified Connectivity CPE devices offer multiple services: at least three LAN interfaces are available, and sometimes as many as eight. It doesn't matter if an enterprise wants to trial both IP VPN and Meshed Ethernet VPLS VPN services at a site, or simply configure three distinct IP VPNs. Since Interoute adopt an advanced template-based configuration system, the production of deterministic, predictable and quality-assured configurations to support customer requirements is possible with ease.

Any of the following services can be provided with Unified Connectivity:

- **SVC-IPVPN:** Managed layer-3 IP VPN, providing managed WAN routing for the customer's private IP network.
- **SVC-EVPN:** Managed layer-2 Ethernet VPN, providing scalable multi-site layer-2 switching within the customer Ethernet network.
- **SVC-IPETHX:** Managed layer-2 Ethernet virtual-circuits, providing point-to-point links with commit/burst tariffs.
- **SVC-INET:** Local access to the public Internet, with or without managed Internet firewall functionality.

One of the challenges when offering multiple logical services is the recognition that different services and different traffic types might have different significance to business operations and while local site bandwidth

access remains the significant cost component in any enterprise VPN, it is to be expected that local access WAN links may sometimes get congested.

In these cases, Interoute's Unified Connectivity Congestion Management feature provides a simple way for an enterprise to apportion sections of its WAN bandwidth, on a site-by-site basis to different applications, traffic types or services using Quality-of-Service technology.

Unified Connectivity; a Procurement revolution

The separation of the physical and logical elements of the service not only offers the freedom to change the technical topology of the network, but also radically changes how the network is procured.

Unlike traditional network procurement Unified Connectivity makes no distinction between the services be it Ethernet or a Managed MPLS VPN which leads to transparent and simple pricing structure. This freedom to switch regardless of technology means that changing services becomes easier to manage and faster to deliver completely changing the approach to WAN outsourcing. Rather than being required to predict technology shifts and choose a particular technology or approach to WAN management the IT department is left with a very simple set of criteria by which to assess their WAN requirements.

In simple terms the WAN configuration could be decided by identifying the amount of capacity required and the level of resilience. Alternatively this could be further simplified by stating a level of acceptable budget and the capacity and resilience derived from that. This completely de-risks the approach to the network as it can be easily modified during the life of the contract and can adapt as business needs change.



Conclusion

Despite the fact that the network requirements of most enterprises are extremely similar, there still remains significant complexity and confusion in the arena of managed WAN services.

Interoute's Unified Connectivity services aim to help remove some of this complexity by allowing organisations to focus on the real physical requirements of their connectivity, without compromising the logical topology options available to them.

This advance is then complemented with commercial innovation that removes the constraints of traditional procurement methods by allowing the customer to exploit the capability to change services at any time during the contract within a simple, predictable and transparent pricing structure.

Interoute through an innovative approach to utilising MPLS technologies has created a completely different way of thinking about the Wide Area Network. The WAN is no longer a network where the way in which it is managed influences enterprise IT planning. The WAN is now the platform for the enterprise ICT services dynamically changing to adapt to business needs, enabling a leap in terms of service delivery for all ICT services. The network has never been so free.

