

# Scalability and Functionality Challenges for MPLS VPN Networks

**This article explains the principal network engineering developments that have underpinned significant growth in provider-provisioned virtual private network (PP-VPN) services based on border gateway protocol/multi-protocol label switching (BGP/MPLS) technology. The key advances can be broadly categorised into two important areas – scalability and functionality. This article assesses in detail the underlying facets of scalability and functionality that continue to drive engineering advances in BGP/MPLS VPN networks. This assessment is carried out from the perspective of real-world deployment and operation of a very large-scale BGP/MPLS VPN network. As well as stressing the obvious advances that have been made in such networks since their inception, the article also highlights future challenges that will be faced by service providers.**

**VPN service providers are under continuous pressure to develop their core network platforms to support an array of service features**

## Introduction

Virtual private networks (VPNs) are an attractive solution to serve the enterprise networking requirements of a wide range of businesses from small-to-medium enterprises (SMEs) to multi-national 'blue-chip' corporate organisations. Essentially, VPNs provide a seamless network infrastructure that allows multiple customer sites to communicate over a shared backbone network, as though they are using their own private network, regardless of geographical location. Typical applications that run across an organisation's virtual private network include corporate intranet, mail services and voice over IP (VoIP) telephony.

Three clearly distinct categories of VPN networking technology exist:

- traditional VPNs based on layer 2 frame relay (FR) and asynchronous transfer mode (ATM) technology;
- CPE-based VPNs based on protocols such as layer 2 tunnelling protocol (L2TP) and IPsec;
- provider-provisioned VPNs (PP-VPNs) based on layer 2 switching such as

Ethernet, or layer 3 IP-based paradigms such as BGP/MPLS VPNs.

A good overview of these main types of VPN can be found in Knight et al<sup>1</sup>. This article focuses exclusively on BGP/MPLS VPNs, as defined in RFC4364<sup>2</sup> and other related Internet Drafts.

Generally speaking, since the early part of the decade, there has been a notable shift in customer demand from traditional layer 2 VPNs to layer 3 IP-based VPNs. In the USA alone, dedicated IP VPN revenues are projected to reach \$34.6 billion between 2004 and 2009. According to the Vertical Systems Group, network-based services delivered over carrier-based MPLS or IP infrastructures will account for almost half of the revenues, totalling \$17.1 billion by 2009<sup>3</sup>. This significant market growth has been driven mainly by perceived cost benefits, as well as the major attraction of being able to combine all IP-based networking requirements (data, voice-over-IP, etc) into a single integrated VPN service capability.

The ability to support IP VPN services in a scalable manner is a fundamental requirement for service providers who wish to compete in a global market-place where a large proportion of corporate clients have a sizeable geographic spread and encompass large volumes of access connections. But scalability is nothing without functionality. In other words, as well as meeting the demands to support increasing numbers of access circuit volumes, VPN service providers are under continuous pressure to develop their core network platforms to support an array of service features. This article considers in detail the underlying facets of scalability and functionality that continue to drive advances in BGP/MPLS VPN networks. This assessment will be carried out from the perspective of real-world deployment.

## BGP/MPLS VPN Network Advances

This section provides an overview of BGP/MPLS VPN network technology and presents an insight of a real-world

deployment, from its inception to the present day.

### Overview of BGP/MPLS VPN networks

The key core network elements of a provider-provisioned BGP/MPLS VPN network are provider edge (PE) and provider core (P) routers as shown in Figure 1. PE routers terminate customer access circuits, whereas P routers perform packet forwarding and typically do not have directly connected customer access circuits. All PE and P routers run label switching so that they can build MPLS label switched paths (LSPs) from each PE to each other PE. This is achieved through use of the label distribution protocol (LDP) in conjunction with the interior gateway protocol, such as open shortest path first (OSPF). When a PE forwards a VPN-addressed packet across the core, it adds an inner MPLS label to identify the VPN of which the packet is a member and then an outer MPLS label to identify the egress PE router. Any intermediate P or PE routers switch the packet to the egress PE using the outer label only. The inner label is used by the egress PE to determine the VPN/port to which the packet should be forwarded.

The customer edge (CE) router is not considered part of the provider's core network. It acts as a peer of the PE router, but not a peer to other CE routers. Each PE router supports multiple routing/forwarding tables, called virtual route forwarding (VRF) tables. VRFs are logically separate and may

contain IP addresses received from the CE router which overlap with addresses in other VRFs (e.g. in Figure 1, VPN\_A, Site 2 has the same private routes as VPN\_B, Site 2). VPNs are formed by defining individual customer accesses to be members of a specific VRF, with several sites formed on one PE by defining all sites to use the same VRF.

The PE routers use an extended variant of the border gateway protocol (BGP) for signalling between themselves and propagating information about the actual routes of each VPN, as well as the inner MPLS label. The extended BGP carries each VPN route together with two new fields, the route distinguisher (RD) and the extended community. The RD is added to each VPN route to ensure that routes from different customers are unique; extended BGP only treats VPN routes as equal if both the RD and the IP prefix mask are equal. Extended communities are used by BGP to indicate a group of routes, thus defining VPN membership information for exchange between PEs. Only a very rudimentary overview of BGP/MPLS VPN principles has been afforded here for brevity – a much more detailed treatment of BGP/MPLS VPNs can be found in Rosen and Rekhter<sup>2</sup> and Pepelnjak and Guichard<sup>4</sup>.

### Real-world deployment of BGP/MPLS VPN

In the UK, BT has offered IP VPN services based on BGP/MPLS technology since 2001. Over the past six years, this network

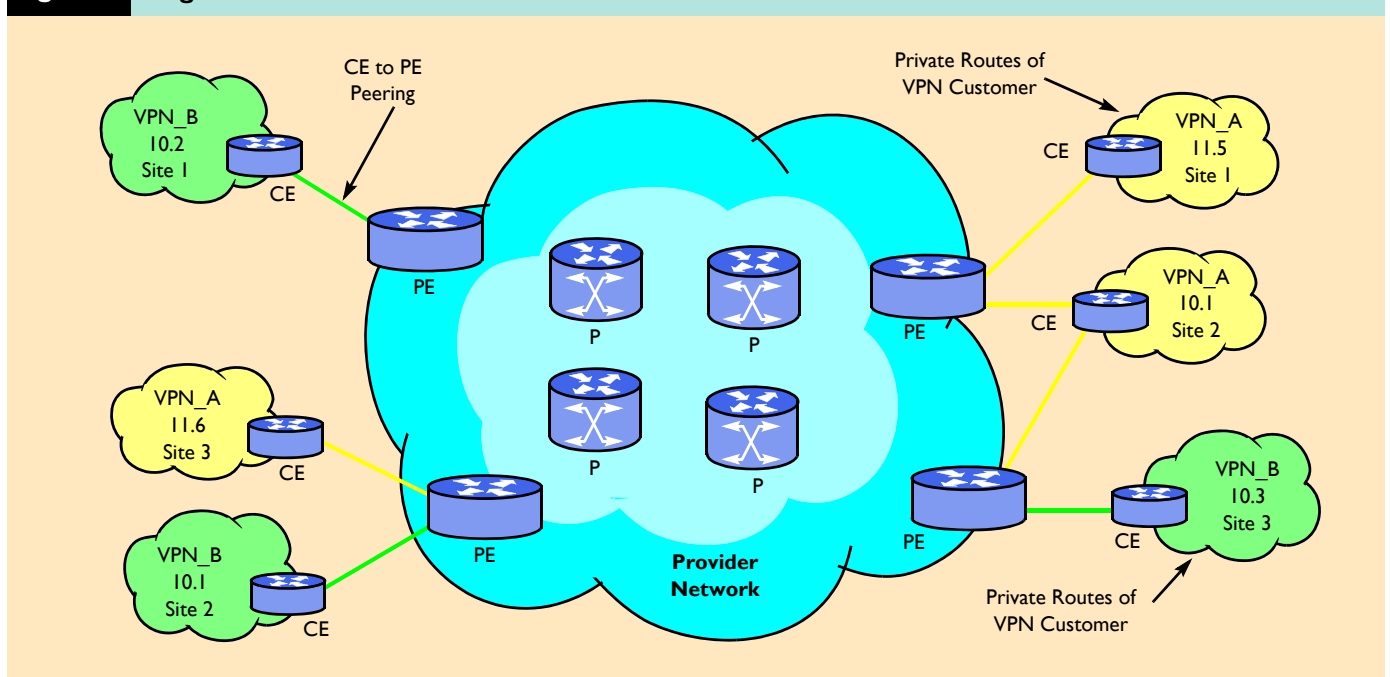
infrastructure has undergone a major transformation in terms of scale and functionality. By way of illustration, Table 1 provides a snapshot of the scale of the network around the time of IP VPN service launch in 2001, compared with the current network characteristics. As can be seen, the growth rate has been phenomenal in almost every respect. As well as network scaling, many significant advances have been made in functionality that have supported key service feature developments, also illustrated in Table 1.

To sustain such network growth, service providers must maintain the ability to support VPN services in a scalable manner, as well as continuing to develop an array of service features. These two aspects of engineering development are therefore very closely related as shown in Figure 2. The diagram also illustrates factors that commonly influence and effectively bind together all aspects of scaling and service functionality. In other words, it is not possible to make advances in VPN network developments to improve either scalability or feature-richness, without due consideration of cost, security and vendor interoperability.

### Scalability Challenges of BGP/MPLS VPNs

This section examines the key facets of network scalability that underpin platform growth and development of BGP/MPLS

**Figure 1** Diagrammatic overview of BGP/MPLS VPN network



**Table 1** BT's BGP/MPLS VPN platform growth and feature development in UK over 6 years

'Growth' Metric	Circa 2001	Circa 2007
Number of PE routers	<50	>1000
Maximum WAN link speed	OC-3/STM-1 (155 Mbit/s)	OC-48/STM-16 (2.44 Gbit/s)
Maximum customer access speed	E3/T3 (34 Mbit/s, 45 Mbit/s)	Gigabit Ethernet (1 Gbit/s)
Number of PoP locations	<10	~100
Number of distinct customer VPNs supported	Tens	Thousands
Number of customer ports	<1000	Tens of thousands
Number of customer routes	Thousands	Hundreds of thousands
Service Feature	Circa 2001	Circa 2007
Quality-of-Service and traffic categorisation	Best-effort only, supporting data traffic	6 DiffServ Code Point (DSCP) service classes, supporting voice, video, premium data and best-effort data
Access connectivity options	Leased line only	Leased line, layer 2 (ATM, Frame Relay), Ethernet, xDSL
Customer routing	Static only	Static, default or BGP
Access resilience options	Fixed line only	Fixed line, dial, xDSL as back-up

VPNs. These aspects of scalability are indicated in Figure 2, namely PE density, logical scaling, network footprint and core capacity.

**PE density and logical scaling**

One of the critical aspects of physical network scalability is the ability to support a rapidly growing number of customer accesses without having to deploy a concomitant number of PE access routers. Deployment of dense aggregation PE devices allows network scaling in terms of the number of access connections supported, and allows the cost per port to be driven down. Additional economic benefits accrue with the deployment of dense aggregation PEs, mainly due to reduced accommodation requirements such as racking, power and ventilation.

Keeping the number of PE routers in check by densely aggregating large numbers of customer access circuits is important not just from a cost-efficiency perspective, but also as a means of maintaining logical scalability. In a BGP/MPLS VPN network, logical state is required to maintain

connectivity and reachability between PEs in the form of LSPs and BGP peerings. LSPs are created in the control plane by label distribution protocol (LDP) and once established, facilitate MPLS packet forwarding in the data plane. BGP peerings are created in the control plane via TCP sessions which in turn facilitate the exchange of customer-specific VPN routing information stored in PE routers. It is evident therefore, that a direct relationship exists between the number of PE routers in the network used to terminate customer access connections and the accumulation of logical state.

Technology advances in the PE router space mean that many vendors now offer dense aggregation edge devices which can potentially handle thousands of access connections; however, it is no longer just a question of how many access connections can be supported per box. As shown in Figure 3, the key to dense aggregation is to deliver a PE device that can support a large volume of accesses as well as a wide range of required service features (e.g. QoS, multicast, etc) and logical scaling (e.g. BGP

peerings) without degrading the performance of the device.

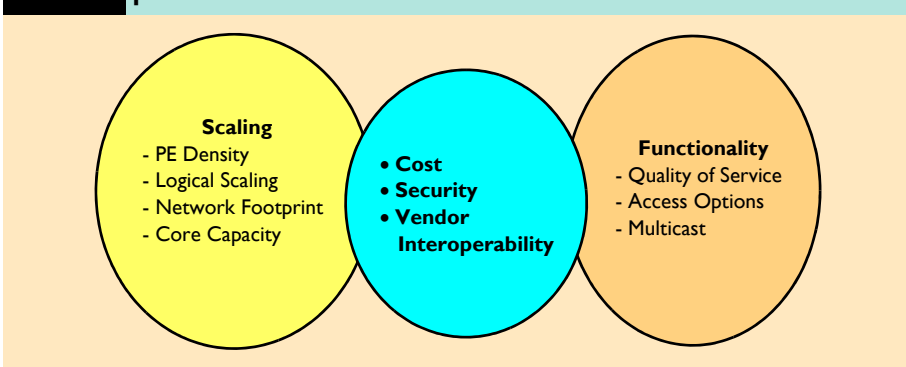
**Network footprint**

A critical aspect of physical scaling is the ability to support a sizeable network footprint providing wide-ranging geographical coverage, preferably on a global scale. Building network points of presence (PoPs), and the associated wide area network (WAN) connectivity between such PoPs, comes at a significant cost, however, and many providers of VPN services selectively target obvious geographical locations such as major cities and finance centres.


It is not always possible to invest in network infrastructure in all countries, however, due to prohibitive build costs. An alternative approach is to reach a commercial agreement with other service providers to form strategic partner arrangements. From a BGP/MPLS VPN perspective, this is made technically possible by exploiting one of several interconnect options between autonomous systems (ASs) as outlined in RFC4364<sup>2</sup>. These optional means of creating 'multi-AS backbones' differ mainly according to the means of transferring routing information between autonomous systems and in turn there is a significant emphasis on network security aspects.

In other words, the selected means of interconnecting one BGP/MPLS VPN network with another depends largely on the concept of 'trust' between such networks. The physical scaling challenge of extending a network footprint is therefore a very good example where the combined constraints of cost and security provide a key influence over the preferred expansion strategy.

**Figure 2** Main engineering challenges for BGP/MPLS VPN service providers



**Figure 3** Slow-down in required number of PEs achieved with dense aggregation technology



### Core capacity management

A final consideration of physical scalability relates to the increasing challenge of handling large volumes of customer VPN traffic in the core network. In other words, although there may be a large amount of 'sold' access bandwidth in a VPN network, the actual aggregate traffic levels presented to the network (i.e. the 'used' bandwidth) from customer access connections tends to be a fraction of the sold amount.

Nevertheless, there are a growing number of requirements for very high access bandwidths such as Gigabit Ethernet. This is an increasingly common requirement of customers in the finance sector who require a hub-spoke VPN whereby the hub site is a data centre which serves a large number of client 'spoke' sites in a fan arrangement.

The fundamental aim of effective core capacity management is to deploy the necessary core WAN links to support the traffic requirements and meet the necessary performance criteria such as packet throughput, delay and jitter. This must be done in sympathy with cost constraints, however. Due to increasing commercial and competitive pressures on service providers, the days of deploying 'gluts' of backbone capacity regardless of cost are long gone, and there is increasing emphasis on 'asset-sweating'. Expensive WAN link capacity upgrades are consequently avoided or deferred wherever possible.

Deferral of WAN link upgrades is possible by optimising the flow of traffic in the core network such that traffic loadings are spread as evenly as possible around available WAN links in the network topology. MPLS traffic engineering (MPLS-TE) is one such technique that can be used

as a means of optimising the available capacity<sup>5</sup>. MPLS-TE operates on the principle of explicitly routed tunnels to be established using an extended version of the resource reservation protocol (RSVP)<sup>6</sup> based on suitable link costs. This allows, for example, high-cost (i.e. heavily loaded) links to be avoided, and lower-cost (i.e. lightly loaded) links to be used in preference, thus ensuring that traffic is spread around the network in an optimal fashion.

### Functionality Challenges of BGP/MPLS VPNs

This section assesses some of the key VPN service features that have driven major developments of BGP/MPLS VPN network platforms. These service features were indicated in Figure 2, namely quality-of-service, access types and multicast. Although not an exhaustive list *per se*, these have been key areas of service development in IP VPN services since their inception.

#### Quality of service

Very early deployments of BGP/MPLS VPNs had either no, or at best a very basic, ability to apply IP quality of service (QoS) as a means of differentiating the classification and treatment of IP traffic entering the network. Over time, however, existing and potential new VPN customers demanded more sophisticated support for QoS features on BGP/MPLS VPNs which would allow prioritisation of certain critical data applications, and in the case of converged application networks, the support of voice. Initial QoS implementations were largely

vendor-proprietary and generally supported three or four QoS classes. Such early QoS implementations used the IP 'precedence' value to classify the traffic, comprising the first three bits of the type-of-service (ToS) byte in an IPv4 packet (Figure 4). The ToS byte is the second byte in the IP header.


More recently, BGP/MPLS VPN service providers have developed a more extensive and standardised means of supporting IP QoS with six service levels. The six service classes are based on IETF 'per-hop behaviours' as defined by the DiffServ working group<sup>7, 8</sup> and the recommended DiffServ code point (DSCP) values for these. They can be broadly described as follows.

- Expedited forwarding (EF)  
This is equivalent to class 1 in the vendor-proprietary QoS models, designed and optimised for the delivery of jitter- and delay-sensitive applications such as VoIP.
- Assured forwarding (AF)  
This is equivalent to class 2 in the vendor-proprietary QoS models, intended to support priority data applications. The AF class is split into four equivalent sub-classes, AF1–AF4 which can be used to segregate data or video traffic applications with priority being maintained over the default class.
- Default (DE)  
This is equivalent to class 3 in the vendor-proprietary QoS models, to support 'best-effort' (i.e. non-prioritised) data traffic.

The QoS treatment on the access link of a customer VPN connection will be based on the DSCP value in the IP header, which corresponds to the first six bits of the type-of-service (ToS) byte, as shown in Figure 4.

The primary motivation for the additional classes provided by the DSCP design is to extend the scope for customers to partition bandwidth between applications across their own access link, thus making as efficient use of the available bandwidth as possible. In addition to optimising traffic sent across the access links from a customer

**Figure 4** DiffServ code point mapping into IPv4 ToS byte



site into the VPN core network, the DSCP markings will dictate the way in which such traffic is placed into queues and conveyed across the core network. This means that under heavy traffic load and congestion situations, the QoS policies will dictate how packets are treated in terms of scheduling, queuing, discard eligibility, and so on.

It is worth remarking that design and architecture development of a QoS architecture for MPLS VPNs on an end-to-end basis is very complex as it involves the CE to PE access link, the PE to P router links, and the core WAN P to P router links. A more detailed exposition of QoS principles in a generic sense, and QoS in a specific BGP/MPLS VPN context, can be found in Willis<sup>9</sup> and Carter<sup>10</sup>, respectively.

**Access options**

As discussed earlier in the article, VPN customers encompass a wide and diverse range in terms of number of access circuits, geographic coverage, access bandwidths and so on. In order to keep their own costs down, customers of IP VPN services continually seek the best possible value in ‘cost-per-megabit’ as part of their overall service delivery. In the early days of commercially operated BGP/MPLS VPNs, the principal access technology supported was leased line (Table 1), such as E1/T1 circuits. More recently, Ethernet and xDSL access technologies have become very popular choices for virtual private networks. Ethernet access tends to plug a gap as being a more cost-effective means of delivering high bandwidth access (10M, 100M, Gigabit), while xDSL technology provides a low-cost option for sub-2 Mbit/s speeds.

Ethernet and xDSL are interesting ‘new-wave’ access technologies as they provide a stark contrast in terms of load and traffic presentation on to the network core. This can be summarised simplistically by observing that Ethernet access types tend to be relatively low volume but high bandwidth, whereas xDSL access types are

high volume but low bandwidth. Figure 5 illustrates this point by demonstrating the breakdown per access type and bandwidth for the installed base of VPN customers on BT’s BGP/MPLS VPN platform, circa 2006. ADSL access connections account for around one third of the installed customer base, while presenting only 6% of the total access bandwidth to the network. Ethernet on the other hand accounts for only 2% of installed access connections but represents 61% of the ‘sold’ access bandwidth.

The unique characteristics of a wide range of access types are important for the service provider because it introduces a wider range of technical requirements of the PE router technology. In other words, a PE device that is optimised in terms of performance in dealing with very large-volume, low-bandwidth accesses such as xDSL, is not necessarily the appropriate device to terminate Ethernet and leased line accesses. As such, it is not uncommon to consider deploying different PE types to support different access types. The service provider must then meet the challenge of ensuring that the different PE devices deployed for different access types are aligned in terms of the other service features that may be supported as part of the VPN service (QoS, multicast, etc), irrespective of access technology. This problem would be further compounded in a multi-vendor environment.


**Multicast VPNs**

IP multicast involves efficient transmission of a single datagram to a selected number of receiving hosts by way of suitable packet replication and forwarding in the core network. This results in a single multicast data stream replacing multiple point-to-point ‘unicast’ data flows. The collection of hosts is called a multicast group, and each group is uniquely identified as a Class D IP address. Potential requirements and applications for IP multicast in a WAN environment, such as BGP/MPLS VPNs, are

**IP multicast involves efficient transmission of a single datagram to a selected number of receiving hosts by way of suitable packet replication and forwarding in the core network**

manifold. For example, in-house IT departments can exploit multicast for distributing regular software updates (e.g. anti-virus upgrades and patches) to company employees. Company CEOs can exploit multicast technology to facilitate ‘IPTV’ intranet broadcasts (e.g. to report on quarterly financial results) to company employees. There are also specific requirements for multicast in the finance sector, such as distribution of real-time market data (e.g. stocks and shares) between financial service providers and clients.

Multicast support in a BGP/MPLS VPN network environment is detailed in an Internet Draft which is currently under development in the IETF<sup>11</sup>. This builds on, and broadens the scope of, a previous Internet Draft specification known as the ‘Draft-Rosen’ implementation of multicast VPNs, and involves generic routing encapsulation (GRE) tunnels between PE routers for the conveyance of multicast packets across a BGP/MPLS VPN core. This has enjoyed widespread implementation by a number of router vendors, and is consequently used as the basis of multicast support in real-world BGP/MPLS VPN



## using a data MDT has the benefit of reducing the amount of multicast traffic on the backbone, as well reducing the load on some of the PEs

networks. This is a fairly detailed technical specification which consequently cannot be repeated in this article – hence only a brief summary of the salient aspects of the ‘Draft-Rosen’ multicast VPN architecture will be described.

In a BGP/MPLS VPN network supporting multicast, a service provider will determine whether a particular VPN is multicast-enabled. If it is, it corresponds to a ‘multicast domain’ (MD), and a PE router that attaches to a particular multicast-enabled VPN is said to belong to the corresponding multicast domain. For each MD, there is a default ‘multicast distribution tree (MDT)’ through the backbone, connecting all of the PEs that belong to that MD. Construction of the default MDT does not depend on the existence of multicast traffic in the domain, meaning that it will be established before any such multicast traffic is seen.

Multicast data is flooded to all PEs within a customer’s VPN irrespective of whether there are ‘interested’ downstream receivers, which can be very wasteful of network capacity, especially if the multicast application requires a large bandwidth traffic source. For this reason, there is a method for establishing individual MDTs for specific multicast groups, namely ‘data MDTs’. A data MDT delivers VPN data traffic for a particular multicast group only to those PE routers which are on the path to

receivers of that multicast group. As illustrated in Figure 6, the high-bandwidth source connected to CE B3, is required by CE B1, but not CE B2. The data MDT therefore does not build a branch of its tree to the PE router to which CE B2 is connected. Using a data MDT therefore has the benefit of reducing the amount of multicast traffic on the backbone, as well reducing the load on some of the PEs. Construction of a data MDT to serve the requirements of particular high-bandwidth multicast groups, is triggered by a bandwidth threshold being exceeded on the default MDT.

It is worth noting that the ‘broadening’ of the Draft-Rosen multicast VPN specification, that is currently under way in the IETF, includes the development of ‘label-switched’ multicast, which involves point-to-multipoint MPLS label switched paths (LSPs)<sup>11</sup>. In the label-based multicast domain, the generic terminology P-multicast service interfaces (PMSIs), will replace the notion of MDTs as outlined in the preceding part of this section.

In summary, of all the service features detailed in this article, multicast is probably the most challenging from the perspective of bringing to market, mainly due to the all-pervading nature of the required engineering developments. It impacts on almost all aspects of the network including QoS, traffic management, billing, network

dimensioning (including PE and P router planning rules), network diagnostics and performance reporting. As such, this feature has only very recently been included as part of VPN service portfolios supported on BGP/MPLS network platforms.

## Conclusions

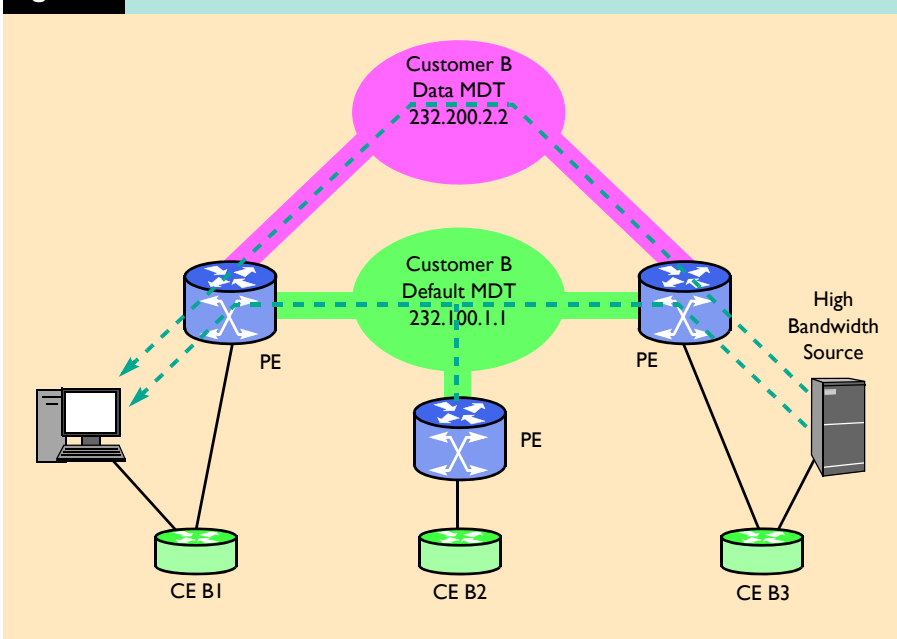
This article has outlined the key advances in BGP/MPLS VPN technology that have facilitated network scalability and support of an increasing range of service features. It has been demonstrated how these two strands of engineering development are closely related, and are often influenced by overriding design challenges such as cost-efficiency, security and vendor interoperability. Although significant advances have been made in BGP/MPLS VPNs, there remain a significant number of emerging challenges that would benefit both ongoing standards activity and academic research, including the following areas.

- Increased core network convergence**  
 BGP/MPLS VPNs are typically supported on self-contained IP-network infrastructures. Increasingly, service providers aim to drive down capital and operational costs by converging all their ‘legacy’ service platforms (VPNs, Internet, voice) on to a common converged core<sup>12</sup>.
- Performance monitoring and service-level guarantees**  
 There are increasingly demanding requirements to support very stringent performance guarantees for some VPN customers, e.g. millisecond-level traffic reporting.
- Tools and systems support**  
 There is a growing emphasis on improved automation and systems intelligence to help manage the complexity, feature richness, increasing demands from customers, etc. This is especially pertinent in the performance monitoring, reporting and fault diagnostic spaces.  
 Developments in such spaces will ensure that BGP/MPLS VPNs continue to evolve and adapt to support a rich set of service features in a scalable fashion.

## References


- 1 Knight, P. and Lewis, C. Layer 2 and 3 Virtual Private Networks: Taxonomy, Technology, and Standardization Efforts. *IEEE Communications Magazine*, June 2004, pp124–131.

**Figure 6** Default and data MDTs in a multicast VPN network



- 2 Rosen, E. and Rekhter, Y. IETF RFC4364: BGP/MPLS IP Virtual Private Networks (VPNs). February 2006 – <http://tools.ietf.org/rfc/rfc4364.txt>
- 3 IP-Based VPN Market To Hit \$34.6 Billion In US. Information Week, December 2005 – <http://informationweek.com/story/showArticle.jhtml?articleID=174403430>
- 4 Pepelnjak, I. and Guichard, J. MPLS and VPN Architectures. Volume 1, Cisco Press, 2003.
- 5 Osborne, E. and Simha, A. Traffic Engineering with MPLS. Cisco Press, 2002.
- 6 Braden, R. et al. IETF RFC2205: Resource Reservation Protocol – Version 1 Functional Specification. September 1997 – <http://tools.ietf.org/rfc/rfc2205.txt>
- 7 Davey, B. et al. IETF RFC3246: An Expedited Forwarding PHB. March 2002 – <http://tools.ietf.org/rfc/rfc3246.txt>
- 8 Heinanen, J. et al. IETF RFC2597: Assured Forwarding PHB Group. June 1999 – <http://tools.ietf.org/rfc/rfc2597.txt>
- 9 Willis, P. J. An introduction to quality of service. *BT Technology Journal*, April 2005, **23**(2), pp 13–27.
- 10 Carter, S. F. Quality of service in BT's MPLS-VPN platform. *BT Technology Journal*, April 2005, **23**(2), pp 61–72.
- 11 Rosen, E. and Aggarwal, R. Internet Draft: Multicast in MPLS/BGP IPVPNs. April 2007 – <http://www.ietf.org/internet-drafts/draft-ietf-13vpn-2547bis-mcast-04.txt>
- 12 Reeve, M. H. et al. Networks and systems for BT in the 21st century. *BT Technology Journal*, January 2005, **23**(1), pp 11–14.

### Biography



**Paul Veitch**  
BT Design

Paul Veitch received MEng and PhD degrees in Electrical & Electronic Engineering from the University of Strathclyde in 1993 and 1996, respectively. He joined BT in September 1996, working on IP, ATM, SDH and 3G network architectural design. In 2000, he joined UUNET (now Verizon Business) and led a number of projects on IP backbone network design. In 2003, he returned to BT and is currently the infrastructure solution design authority for BT's UK IPVPN platform, based at Adastral Park, Suffolk, UK.  
[paul.veitch@bt.com](mailto:paul.veitch@bt.com)