

Cisco Metro Ethernet Access Services Directions in Control-Plane Concepts



Abstract

This document surveys the control-plane technologies and architectures being evaluated for use in deployment of Metropolitan Area Network (MAN). Ethernet access services by numerous data-communication service providers. It identifies the relevant selection criteria for these control-plane architectures, and then describes various control-plane technologies and how they fit these selection criteria.

This document should assist both service providers in their understanding of these architectures and corporate end users in their understanding of the kind of service definitions they can expect to find available as part of the new category of metro Ethernet access services.

Market Overview

The promise of flexibly provisioned bandwidth provided at dramatically reduced service rates based on cost-effective, ubiquitous Ethernet technology has jump-started a new category of data-communication services referred to as metro Ethernet access services. The growing market for metro Ethernet access services has motivated every class of data-communication service providers to examine their service strategies to determine how these metro Ethernet access services will integrate with their existing data-communication service offerings. In this paper these service providers are categorized by the following groups:

- *Ethernet service providers (ESPs)*—These are startup service providers, generally founded within the last three years, that have established their business strategy around providing data-communication services based entirely on Ethernet technologies.
- *Traditional service providers*—These are the established service providers, including independent local exchange carriers (ILECs), inter-exchange carriers (IXCs) and competitive local exchange carriers (CLECs), that have a variety of existing service offerings based on more traditional WAN technologies, such as ATM, Frame Relay, and leased-line services.

Between these two categories of service providers, the business strategy decisions that must be made to accommodate the new metro Ethernet access services are very different.

The upstart ESPs must define metro service offerings based entirely on Ethernet that will encourage the most rapid adoption of these metro Ethernet access services that is possible. To attract customers, the ESP may offer metro Ethernet access services that resemble traditional point-to-point data-communication services, or it may create new service definitions that take advantage of the unique characteristics of Ethernet technologies.

The traditional service providers must define metro Ethernet access service offerings that will establish their position in this new service market, while continuing to provide the traditional data-communication services, including ATM, Frame Relay, and leased line. To retain existing customers, the traditional service providers will seek to deploy creative service definitions that combine the traditional data-communication services with the new metro Ethernet access services.

For both of these service-provider categories, the technical decisions they make about the control-plane architecture of their networks will be a major determinant of the service definition they are able to deliver to their end users, and in turn, whether they are successful in executing their business strategy. This makes the selection of the metro Ethernet control plane all the more critical.

The criteria that the service providers use in selecting their metro Ethernet control-plane architecture will force consideration of new technologies and new concepts. Several of the most critical of these criteria are described in the next section.

Control-Plane Architecture Evaluation Criteria

The following criteria are used as the basis for evaluating and comparing the various control-plane architectures. Other evaluation criteria, not included below, may be of interest to specific service providers, but the evaluation criteria included here must be considered by any service provider planning to deploy a metro Ethernet access service in order to deliver a marketable service:

- *Cost-effectiveness*—In order for metro Ethernet access services to live up to their promise of reduced-cost network services, they must be delivered over a network architecture that can be established with relatively low capital expenditure, and maintained with relatively low operation costs. The operating cost of a metro Ethernet access service can be significantly affected by the control-plane architecture that is used to support a particular network topology and provisioning system. In fact, the reduction in operating costs that can be achieved by certain control-plane architectures can in some cases justify a higher level of capital expenditure needed to deploy those control-plane architectures.
- *Service level*—End users have become accustomed to receiving a predefined service level from their service providers for the traditional data-communication services. For this reason, design of the control-plane architecture of the metro Ethernet access service needs to account for the characteristics of the service levels that can be provided to the end user. These service-level definitions can be as simple as defining the level of service availability on the metro Ethernet connection, or as complex as defining different classes of service for different types of traffic transmitted by the end user over the metro Ethernet connection, including such parameters as latency, committed information rate (CIR), and so on.
- *Point-to-point versus multipoint*—The concept of multipoint connectivity has always existed with traditional Layer 2 data-communication services. Frame Relay and ATM both allow a single physical interface on the customer premises equipment (CPE) to support multiple virtual circuits, each connected to a different remote end point. However, with this model of multipoint connectivity, there was always a requirement that a Layer 3 routing function exist in the CPE to provide communication between the Layer 2 point-to-point circuits, so that a multipoint logical topology could be established. With metro Ethernet access services, the concept of multipoint connectivity is slightly different. Instead of the CPE supporting multiple virtual circuits over a single physical interface, there is only a single logical interface from the CPE to the provider edge. But, at the provider edge, this single logical interface can be mapped into multiple point-to-point circuits to establish a single Layer 2 broadcast domain. With this model of multipoint connectivity, it is possible—but not necessarily desirable—to eliminate the Layer 3 function on the CPE. For larger multipoint topologies, a Layer 3 control plane, such as Multiprotocol Label Switching virtual private network (MPLS VPN), should also be considered.

- *Transparency*—Metro Ethernet access services, based on Ethernet technologies, are introducing a new element into the control-plane architecture decision that previously did not exist with other traditional Layer 2 service definitions. Because Ethernet is by far the predominant technology used in enterprise networks, it is possible for the metro Ethernet control plane to accommodate a transparent interface to the end customer so that the end customers do not need to make any changes to the Layer 2 virtual-LAN (VLAN) configurations within their own networks, even if the connection to the metro service provider is based on Layer 2 instead of Layer 3 protocols. This concept never existed with traditional data-communication services, such as ATM, Frame Relay, or leased lines. These Layer 2 transport protocols were different from the Ethernet Layer 2 protocol used in the end user's LAN, so that a Layer 3 interface from the end user to the service provider was a technical requirement.
- *Scalability*—Scalability refers to both the number of end users and the geographic distribution of those customers. As the name suggests, metro Ethernet access services originated to provide high-speed connectivity within a metro service area. However, as metro Ethernet access services have evolved, end users have expressed the requirement for the Ethernet service interface connecting the end user to the service provider to provide general-purpose WAN connectivity to either private or public (that is, Internet) networks. For some service providers, it may be acceptable from the perspective of their customers for the metro service definition to be limited to a specific metropolitan area. For other service providers, it may be a requirement that the metro Ethernet access service definition scale to inter-metro distances. In either case, the control plane must allow new metro access and aggregation network elements to connect into the core metro network without major disruption or reconfiguration.
- *Interoperability*—For traditional service providers, the issue of interoperability is of paramount importance. One of the most obvious strategies for these service providers to retain their existing base of customers using ATM, Frame Relay, and leased-line services is to demonstrate the ability to support interconnectivity between legacy services and Ethernet-based services. At a minimum, the ability to transport the new metro Ethernet access services over the same network infrastructure as the legacy services must be considered by traditional service providers in order to take greatest advantage of their network infrastructure. However, ESPs must also consider interoperability, particularly when they must interconnect to other service providers to deliver geographically scalable services.

A variety of control-plane architectures is evolving, based on both Layer 2 and Layer 3 network service concepts, that are being considered by service providers for metro Ethernet access services. Service providers are being forced to examine closely exactly what expectations their end users will have for such metro Ethernet access services.

Layer 2 Control Plane Supporting Layer 2 Services

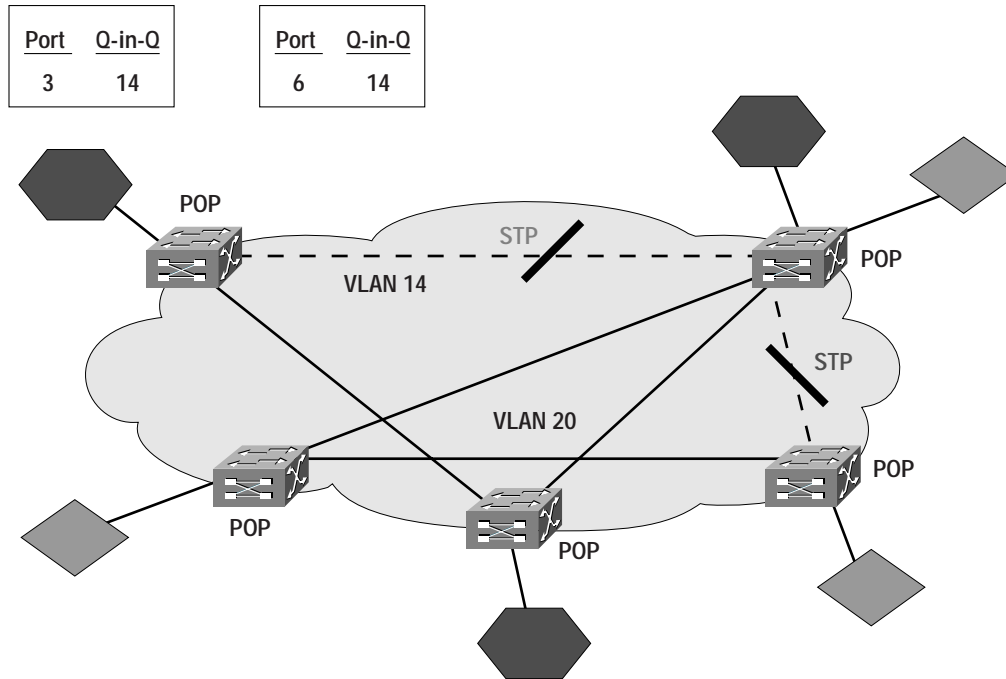
Service providers have traditionally delivered Layer 2 point-to-point data-communication services. Services such as Frame Relay, ATM, and leased-line circuits are the most widely used Layer 2 services for which traditional service providers have an established customer base. These Layer 2 services have been attractive to service providers not only because of the revenue stream they generate, but also because such services do not require the service provider to participate in any end-user Layer 3 network designs.

Metro Ethernet access services can be designed to deliver Layer 2 services in a manner that is very similar to the traditional point-to-point data-communication services. However, metro Ethernet access services can also be designed to take advantage of the specific characteristics of Ethernet technologies, which can support Layer 2 services that go beyond traditional point-to-point service definitions.

Although many service providers are migrating to Layer 3-based core networks to take advantage of the scalability of the Layer 3 control plane (as provided by either IP- or MPLS-based network technologies), they will continue to support a significant base of users that want only Layer 2 services. For that reason, the metro Ethernet access services must be able to support Layer 2 service definitions and technologies.

This section reviews the control-plane architectures that are being considered for the deployment of Layer 2 metro Ethernet access services based on Layer 2 control-plane architectures.

Figure 1 Layer 2 Control-Plane Architecture-802.1Q or Q-in-Q



VLANs Based on 802.1Q

This control-plane architecture as shown in Figure 1, above, is based on the standardized IEEE 802.1Q VLAN control plane deployed in the Ethernet LANs of many major enterprise corporations. As such, it assumes a virtual broadcast domain between all network elements that are assigned to a specific 802.1Q VLAN. The number of 802.1Q VLAN domains that can be supported within a single Layer 2 Ethernet topology is 4096. Network redundancy and availability are based on IEEE 802.1d Spanning Tree Protocol for bridged Ethernet networks.

- *Cost-effectiveness*—Upon first analysis, a metro Ethernet access service based entirely on an 802.1Q control plane would seem to be the most cost-effective strategy. The costs of high-performance Ethernet switch network elements, with standards-based VLAN and spanning-tree protocols, are very attractive when compared to traditional Synchronous Optical Network (SONET)- or ATM-based network elements. However, the relatively low cost of these network elements is at least partially offset by the control-plane provisioning that must occur with each new customer VLAN added to metro Ethernet network. With an 802.1Q control plane, it is necessary to configure VLAN trunking across all network elements on both the edge and core of the network, including the 802.1d spanning-tree hierarchy for each new VLAN added to the network, and this scenario becomes increasingly complex and, therefore, costly, as the size of the network increases.
- *Service level*—Defining service levels in a metro Ethernet access service definition can be very difficult, unless the network is designed to guarantee that no oversubscription can exist at any point in the network. This strategy has been used with reasonable effectiveness by some of the new ESPs by simply overprovisioning the bandwidth available at all points in the network. However, the larger the metro Ethernet network becomes, the more difficult it becomes to guarantee no oversubscription. Furthermore, there is a limit to the customer service levels that can be effectively supported in large metro Ethernet networks based on 802.1Q. For example, it can be very difficult to support low-latency service levels,

even when no oversubscription exists, because of the transient queuing delays that can occur on the high-speed Ethernet interfaces in the core of the network when the network is heavily utilized. The ability to prioritize traffic forwarding based on the 802.1p value associated with the 802.1Q VLAN tag is a requirement for all network elements. In addition, even the simplest service-level guarantees for availability can be difficult to deliver in an 802.1d-controlled metro Ethernet network, because the network failure recovery is based on the Spanning-Tree Protocol (STP), which can have unacceptably long reconvergence delays (30 to 90 seconds is not uncommon) in large network topologies. This issue of STP reconvergence delays is being addressed by the new IEEE standards of 802.1s and 802.1w, which will achieve subsecond reconvergence times, but not quite the 50-ms reconvergence times delivered by SONET and Route Processor Redundancy (RPR).

- *Point-to-point versus multipoint*—Metro services based on an 802.1Q control-plane architecture do support a multipoint service definition. The metro Ethernet access service provider is able to offer a true meshed multipoint Layer 2 service with an 802.1Q control plane.
- *Transparency*—Metro services based on an 802.1Q control-plane architecture are not transparent to the end users' VLAN configurations. The end users must present Ethernet traffic to the service-provider network either with no VLAN tag or with a preassigned VLAN tag. If the end users used their own VLAN tag configurations, the service-provider network would not be able to prevent different end users from using the same VLAN tag value, a scenario that would result in those customers receiving one another's data transmissions. Obviously, this would not be an acceptable service definition.
- *Scalability*—The primary limitation of a metro Ethernet access service based on an 802.1Q control plane is its scalability. It is essentially impossible to scale a metro Ethernet access service beyond a single metropolitan vicinity without introducing a secondary control plane to support inter-metro service requirements. The reasons for this limitation on scalability are twofold: First, the limitation of 4096 VLAN IDs places a strict limit on the number of end users that can be supported. This limit is probably acceptable in a metropolitan geographic vicinity, but is prohibitive for larger-scale networks. Second, the larger a metro Ethernet network based on 802.1d spanning-tree reconvergence becomes, the more complex and cumbersome becomes the provisioning of such a network. Furthermore, the larger the network, the longer can be the delays in the 802.1d network reconvergence. As previously mentioned, the 802.1s and 802.1w standards have been adopted to improve the efficiency of reconvergence in a spanning-tree control plane. However, in spite of these improvements, they have still not simplified the configuration of spanning tree, which only becomes more complex as more VLANs are added and the number of network elements increases.
- *Interoperability*—Another critical limitation of the 802.1Q control plane is its inherent lack of interoperability with traditional data-communication services. A metro Ethernet access service based solely on an 802.1Q control plane cannot be used to transport data-communication traffic generated by other types of legacy services, such as ATM or Frame Relay. It may be possible, on the edge of an 802.1Q network, to support a protocol translation to another service type, such as using RFC 1483 bridged encapsulation to communicate to an ATM network or Ethernet over MPLS (EoMPLS) to communicate through an MPLS network. But such hybrid architectures require a secondary control plane. Several of these hybrid architectures are discussed later in this document.

VLANs Based on Q-in-Q

The issue of Layer 2 Ethernet transparency has resulted in the evolution of the 802.1Q standard to a new control-plane model, sometimes referred to as Q-in-Q. The concept of Q-in-Q is quite simple: In order to enable the metro Ethernet access service provider to provide a service that is completely transparent to the Layer 2 VLAN configuration of the end user, when the service provider's edge device receives an Ethernet frame from the end user, a second-level 802.1Q tag is placed in the Ethernet frame immediately preceding the 802.1Q tag that has been inserted by the end user's network. The service-provider network then uses this "outer" 802.1Q tag as the control-plane information as the end user's Ethernet frame transits the

service-provider network, and then removes this “outer” tag as the end-user Ethernet frame exits the service-provider network. Although several Ethernet switch vendors offer their own versions of the Q-in-Q control plane, none of these versions is interoperable with other vendors’ versions, so the Q-in-Q model remains a strictly proprietary control-plane architecture. It should be noted that in almost every respect other than transparency, the control-plane architecture of Q-in-Q is essentially the same as the 802.1Q VLAN control plane.

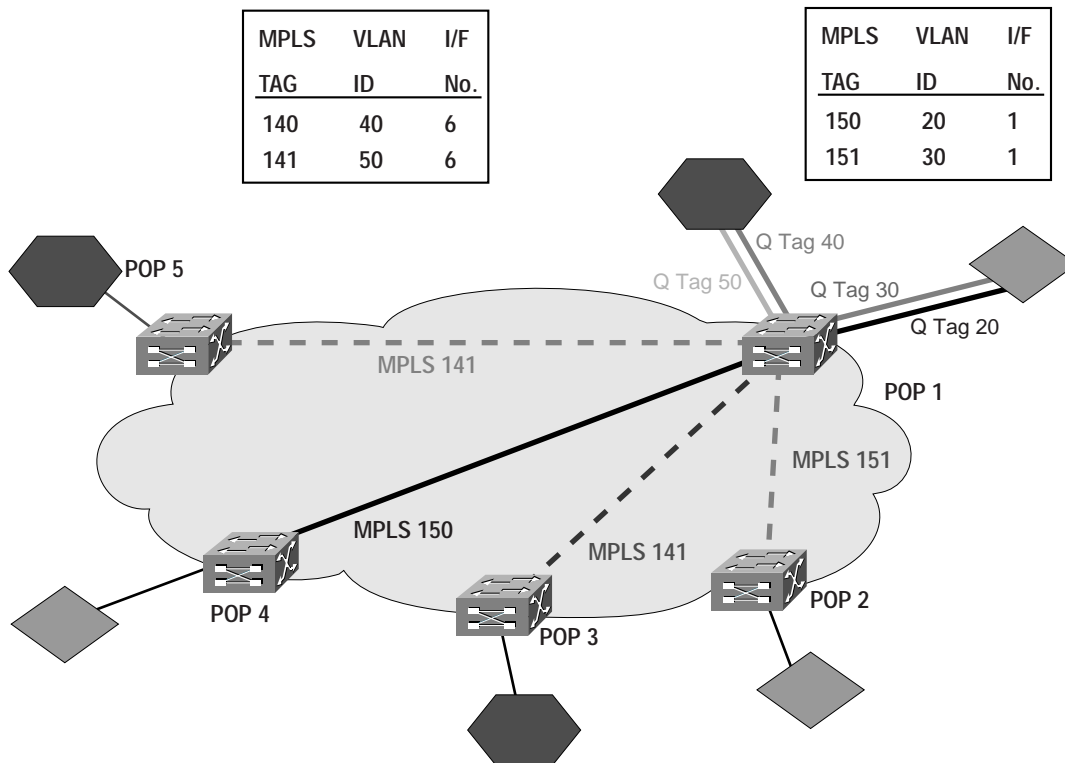
- *Cost-effectiveness*—See the previous comments from the 802.1Q VLAN control-plane discussion.
- *Service level*—The service-level characteristics of the Q-in-Q control-plane architecture are similar to those of 802.1Q. As with 802.1Q, the larger the network that uses the Q-in-Q control plane, the more difficult it becomes to guarantee any kind of service-level guarantee. Another relevant concern is whether the Layer 2 Ethernet class-of-service (CoS) bits normally associated with 802.1P standardized Ethernet switches are or are not supported in each vendor’s proprietary implementation of Q-in-Q. At the point of access, it will be necessary for the service-provider access device to apply a preprovisioned CoS value to the second-level Q-tag.
- *Point-to-point versus multipoint*—See the previous comments from the 802.1Q VLAN control-plane discussion.
- *Transparency*—As previously explained, the primary reason for the Q-in-Q control-plane architecture is to support complete Layer 2 Ethernet transparency to the end users’ Ethernet network. Q-in-Q is specifically designed with the intent of supporting transparency for end users’ VLAN configurations. At this point, the Q-in-Q features supported by most vendors do not support the ability to assign each end-user Ethernet frame to a different Q-in-Q domain, depending on the value of the 802.1Q tag associated with that frame. Future implementations of Q-in-Q may support such functionality, but it will require a more complex provisioning capability by the service provider in order to support such functionality.
- *Scalability*—Q-in-Q has significant limitations on its scalability that are essentially identical to the limitations on scalability for the 802.1Q VLAN control plane, as previously discussed.
- *Interoperability*—If anything, a metro service based on a Q-in-Q control plane is less interoperable than that of the 802.1Q control plane, which, as described above, has limited interoperability. The primary reason for the poor interoperability for a Q-in-Q network is the fact that it is an entirely proprietary, vendor-specific, control plane. As with 802.1Q, efforts are under way to develop a hybrid control plane between Q-in-Q and EoMPLS (see the following section).

Layer 3 Control Plane Supporting Layer 2 Services

Although many service providers are migrating to Layer 3-based core networks to take advantage of the scalability of the Layer 3 control plane (as provided by either IP- or MPLS-based network technologies), they will continue to support a significant base of users who want only Layer 2 services. For that reason, the metro Ethernet access services must be able to support Layer 2 service definitions and technologies. Figure 2 provides a sample network topology of this type of control-plane architecture.

This section reviews the control-plane architectures that are being considered for the deployment of Layer 2 metro Ethernet access services based on Layer 3 control-plane architectures.

Figure 2 Layer 3 Control Plane for Point-to-Point Services: EoMPLS Point-to-Point or L2TPv3



EoMPLS—Point to Point

One of the most important recent technological innovations to address the metro Ethernet control-plane requirement is Ethernet over MPLS, or EoMPLS. As the name suggests, EoMPLS performs an encapsulation of an Ethernet frame into an MPLS label switch path, allowing an MPLS core network to provide transport of native Ethernet frames. The original definition of EoMPLS was developed by a joint submission to the Internet Engineering Task Force (IETF) of a draft RFC, referred to as the Martini draft, which defined a point-to-point transport of native Ethernet frames (see “References” section for more details). The Martini draft RFC appears to be gaining acceptance, has been implemented in numerous vendors’ products, and is likely to become an IETF standard soon. In this section, it is assumed that the EoMPLS control-plane encompasses all service-provider network elements, including access devices. A related control-plane architecture is a hybrid EoMPLS/802.1q solution, which assumes the EoMPLS control plane originates only from the service-provider point of presence (POP), but the connection from the POP to the metro access device at customer premises uses 802.1Q control plane.

- *Cost-effectiveness*—Network elements that support the EoMPLS encapsulation function are inherently more expensive than network elements based strictly on traditional Ethernet switching. As such, a metro Ethernet network based on EoMPLS extended to the service access devices may require a larger capital investment than a metro Ethernet network based entirely on 802.1Q. This increased capital investment can be justified to some extent by the increased network scalability and interoperability (see discussions below) provided by the EoMPLS control plane. In addition, the standard MPLS control-plane features, including Label Distribution Protocol (LDP) and MPLS traffic engineering, can make the provisioning and management characteristics more efficient, thereby offering the potential of lowering network operation expenses, or at least limiting their rate of increase as the size of the service provider’s network increases.

- *Service level*—Because EoMPLS is designed to transport native Ethernet frames over an MPLS core network, EoMPLS can take advantage of the quality-of-service (QoS) and traffic-engineering capabilities that are designed into MPLS. As a result, EoMPLS will be able to accommodate service-level agreements (SLAs) that are at least equivalent to existing Frame Relay SLAs. However, using the MPLS experimental bits (MPLS exp bits) to identify a specific CoS, such as for packetized voice, EoMPLS is expected to be able to deliver enhanced SLAs as well.
- *Point-to-point versus multipoint*—As mentioned, the EoMPLS functionality, as defined by the Martini IETF draft, is inherently a point-to-point Layer 2 service definition. The intention of the Martini draft was to define a metro Ethernet access service that offered similar network topologies that are already deployed by end users based on Frame Relay services. One of the differences between EoMPLS and Frame Relay services is that the metro Ethernet access service allows the end user to connect using an Ethernet interface, and (perhaps) provides more granular bandwidth provisioning than Frame Relay. Of course, as with Frame Relay, EoMPLS supports a multilink capability, so that a single physical connection between the end user and service provider can support multiple EoMPLS Layer 2 circuits. As such, the end user can achieve a multipoint network topology by using this multilink EoMPLS capability, and then performing a Layer 3 routing function between the multiple EoMPLS Layer 2 circuits.
- *Transparency*—Many service providers are likely to offer EoMPLS services as a Frame Relay alternative. As such, that service definition will require the end user to use a VLAN tag assigned by the service provider, much as the Frame Relay data-link connection identifier (DLCI) is assigned to the end user by the service provider. However, the functionality of EoMPLS, as defined by the Martini draft, does include support for end-user VLAN transparency, but such transparency will restrict use of the multilink features of EoMPLS, as previously discussed. In other words, if VLAN transparency is required between two locations—and only two locations—then EoMPLS provides a solution. However, if it is necessary to support both VLAN transparency and multiple Layer 2 virtual circuits on a single end-user Ethernet interface, then the EoMPLS definition based on the Martini draft is probably not the appropriate control-plane architecture. This limitation is addressed in some of the hybrid control-plane architectures discussed in subsequent sections.
- *Scalability*—EoMPLS is inherently a highly scalable control-plane architecture for metro Ethernet access services because of the underlying characteristics of MPLS technology, which is capable of supporting very large, geographically dispersed network service environments. Additionally, the limitations of VLAN tagging, which are so restrictive in metro Ethernet networks based on 802.1Q control planes, do not constitute a significant issue when EoMPLS is used as the control-plane architecture, even though the service provider will usually require the end user's traffic to include a service provider-assigned VLAN tag, although the reason is not intuitively obvious. With the 802.1Q control plane, the 802.1Q VLAN tag has global significance within the metro Ethernet access service domain. As such, each 802.1Q VLAN tag can be associated with one and only one customer throughout the entire network, meaning that the maximum number of 802.1Q VLANs is also the maximum number of customers within that domain. However, with EoMPLS, the 802.1Q VLAN tags are terminated at the metro access device (for hybrid architectures, at the metro aggregation device), meaning that VLAN tag values assigned to one customer at one metro access point can simultaneously be assigned to another customer connected to another metro access point.
- *Interoperability*—EoMPLS, based on the Martini draft, offers excellent potential as a control plane that will provide a high degree of interoperability with traditional WAN services, such as ATM, Frame Relay, and leased-line connections. Once again, the primary reason for this is the underlying MPLS technology. MPLS technology enables the service provider to take advantage of the core network to deliver virtually any combination of WAN services, from the simple leased-line circuit, to Frame Relay or ATM, to EoMPLS and on up to Layer 3 VPNs. At the next level of interoperability, where an EoMPLS point-to-point circuit can be provisioned on one side of the connection but terminate the other side using ATM or Frame Relay, significant technology research is being performed. Existing technologies, such as RFC 1483 (Multiprotocol Encapsulation over AAL5), are being evaluated in conjunction with MPLS as a strategy for interconnecting dissimilar Layer 2 service types over an MPLS network. This effort is sometimes referred to as ATOM (Any Transport over MPLS). Deployable ATOM solutions should be feasible within the next product development cycle. EoMPLS will play an important role as part of an overall ATOM deployment strategy.

L2TPv3 (including Cisco's UTI)

MPLS-based control-plane architectures provide many important benefits, such as EoMPLS, to service providers that are planning to deploy metro Ethernet access services. This paper has already discussed the scalability and operational advantage that can be achieved as a result of deploying an MPLS core network. The ability to simultaneously deliver legacy Layer 2 services, such as Frame Relay and ATM, along with Layer 2 metro Ethernet access services over an MPLS core network is also a major benefit. However, as important as these benefits are, the fact remains that many service providers, for various reasons, will choose not to base their core networks on MPLS technology. However, this does not mean that these service providers do not need the benefits that MPLS can deliver, or that they see no value in a Layer 3 control plane. As an alternative, a new high-performance tunneling protocol based on IP technologies is being developed to provide service providers that deploy IP-based core networks with many of the same capabilities as those provided from MPLS-based core networks. This new tunneling protocol, which is currently under standards review in the IETF, is referred to as Layer 2 Tunneling Protocol Version 3 (L2TPv3). Cisco's prestandard implementation of this new tunneling protocol is referred to as the universal transport interface (UTI). L2TPv3 is not just another IP tunneling protocol. It is specifically designed to support high-speed broadband interfaces operating at 1-Gbps speeds or faster. Generic routing encapsulation (GRE), another well-known IP tunneling protocol, could never operate at such high speeds. For service providers planning to deploy metro Ethernet access services but are reluctant to deploy MPLS technology, the L2TPv3 IP-based tunneling protocol should not be overlooked.

- *Cost-effectiveness*—Network elements that support the L2TPv3 Ethernet encapsulation function are, like EoMPLS network elements, inherently more expensive than network elements based strictly on traditional Ethernet switching. As such, a metro Ethernet network based on L2TPv3 extended to the service access devices will require a larger capital investment than a metro Ethernet network based entirely on an 802.1Q control plane. This increased capital investment can be justified to some extent by the increased network scalability and interoperability (see discussions below) provided by the L2TPv3 control plane. Unlike MPLS, which supports LDP (Label Distribution Protocol), the L2TPv3 protocols base their control plane on IP routing protocols, such as Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP), over which this tunneling protocol is implemented.
- *Service level*—The L2TPv3 tunneling protocols do not support some of the traffic-engineering features provided by MPLS, such as fast reroute (FRR) and MPLS-traffic engineering (MPLS-TE). As a result, in order to satisfy SLAs offered to end users, it may be necessary to overprovision the network, in much the same way as is necessary with an 802.1Q control-plane architecture. However, L2TPv3 can take advantage of traditional differentiated-services (DiffServ) QoS mechanisms available from IP protocols, such as Class-Based Weighted Fair Queuing (CBWFQ) and low-latency queuing, to deliver service-level guarantees for different classes of end-user traffic.
- *Point-to-point versus multipoint*—At the time of this writing, the L2TPv3 protocols provide a point-to-point service definition only. This service definition, like the point-to-point EoMPLS based on the Martini draft, can provide a metro Ethernet access service that is similar to existing Frame Relay services. The IETF subcommittees that are evaluating the L2TPv3 protocols recognize that a multipoint Layer 2 service definition may be required, and they are developing proposals to enable L2TPv3 to support this functionality.
- *Transparency*—L2TPv3 does provide support for VLAN transparency for the end user, but only for point-to-point configurations. However, L2TPv3 protocols can also use VLANs assigned by the service provider as a means of supporting multiple L2TPv3 tunnels from a single physical interface, just like EoMPLS. In this configuration, it is up to the end user to perform Layer 3 routing between the different VLAN tags to support a multipoint architecture. As with EoMPLS, the mapping of VLAN tags on end-user Ethernet frames into L2TPv3 tunnels means that the VLAN tags are only locally significant, a situation that is also beneficial to the scalability of L2TPv3.

- *Scalability*—L2TPv3, being based on IP protocols, is highly scalable, in terms of both the number of users it can service and the size and reach of the network that it can support.
- *Interoperability*—L2TPv3 tunneling protocols provide a very high level of interoperability between different Layer 2 service types. It is possible for service providers to take advantage of their network infrastructure to support all the traditional WAN services, including ATM, Frame Relay, and leased lines, as well as the new metro Ethernet access services. L2TPv3 is also intended to provide interoperability between different Layer 2 services, enabling a point-to-point circuit based on L2TPv3 to be terminated at one end by ATM and at the other end by Ethernet.

EoMPLS—Multipoint

The market for metro Ethernet access services is evolving rapidly, and is fostering the development of new technologies for control-plane architectures. Since the original IETF submission of the Martini draft of EoMPLS in late 2000, several modified and enhanced versions of EoMPLS have been proposed in the form of new IETF draft RFCs (see the “References” section for more details).

One of the most intriguing aspects of these enhancements is the definition of a Layer 2 Ethernet broadcast domain between multiple EoMPLS Cable Switch Paths (LSPs), which would allow the service provider to support a Layer 2 multipoint metro Ethernet access service. This Layer 2 service would use a modified form of Ethernet bridging between multiple EoMPLS LSPs, and, thereby, would not require a Layer 3 connection to support a multipoint end-user topology. These new IETF draft RFCs have just recently been submitted, and they have generated significant controversy. As a result, they are not likely to become IETF standards for a while, and what form they take when they do finally become standard RFCs is still open to speculation. This means that the development of products that support an EoMPLS Layer 2 multipoint service will be based on prestandard (that is, proprietary) implementations for the foreseeable future.

Some of the service providers that are interested in this control-plane solution have already offered their end users a multipoint metro Ethernet access service using a core ATM network. Unfortunately, the inefficiencies of the ATM segmentation-and-reassembly (SAR) function have made it difficult to attain the performance and bandwidth that their end users are seeking from a metro Ethernet access service. Multipoint EoMPLS has the advantage that it supports similar services as the metro Ethernet access services based on a core ATM network, but it does not depend on ATM transport.

- *Cost-effectiveness*—As with point-to-point EoMPLS, the network elements that support a multipoint EoMPLS control plane will be inherently more expensive than network elements based on traditional Ethernet switching. And again, as with EoMPLS, one of the benefits of using MPLS will be reduced operational expenses (or at least a limit to their rate of increase as the size of the network increases).
- *Service level*—The service-level characteristics for the multi-point EoMPLS control plane will be similar to those of point-to-point EoMPLS. However, service levels in a multipoint EoMPLS control plane could be adversely affected by how much Layer 2 broadcast traffic exists on the network and on how the EoMPLS network element handles the broadcast replication.
- *Point-to-point versus multipoint*—As the name indicates, multipoint EoMPLS is designed with the intent of supporting a multipoint Layer 2 network that could also support point-to-point configurations. In order to support a Layer 2 multipoint configuration without needing to use spanning-tree protocols for failure recovery as part of the control plane, the new IETF drafts are proposing a “split-horizon” for the Ethernet bridging functionality, which, instead, will rely on the underlying MPLS functionality for failure recovery. This will require that the multipoint EoMPLS configuration will support only full mesh multipoint topologies. This is not really a limitation, but it is important that it be understood as part of the metro Ethernet access service definition given to the end user.

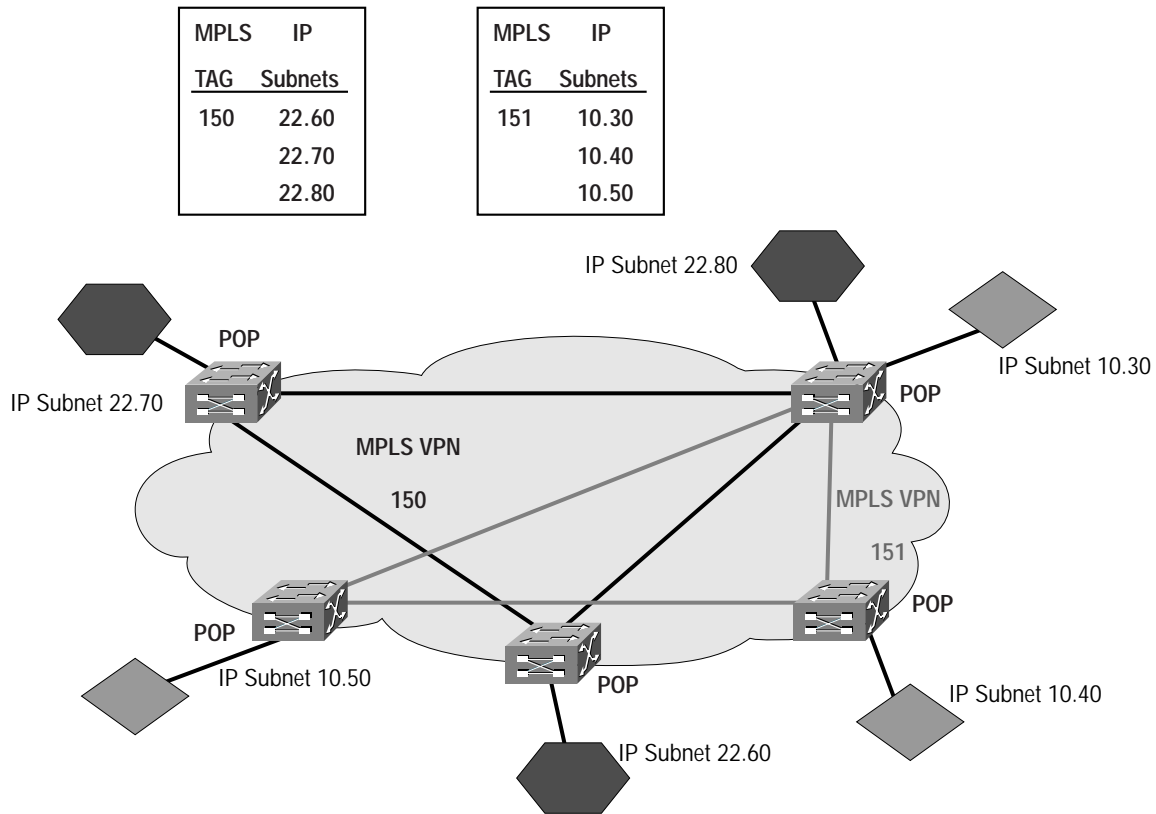
- *Transparency*—The multipoint EoMPLS control plane will support transparency of end users' VLANs. However, the end users' VLANs will not be included as part of the Ethernet bridging forwarding decision performed by the access device in a multipoint EoMPLS control-plane architecture. As a result, the end user will be able to configure network topologies that carry their private VLAN trunks to each node on the multipoint EoMPLS VPN. As the multipoint EoMPLS standards evolve, they may define a service wherein the VLAN tag on each Ethernet packet received from the end user will actually determine the multipoint VPN over which that packet will be transmitted. But the standards are not yet mature enough to say that this will be supported.
- *Scalability*—As with point-to-point EoMPLS, multipoint EoMPLS, being based on MPLS technologies, is designed to be highly scalable for supporting large networks and large numbers of users. However, certain aspects of the multipoint EoMPLS control plane are likely to inhibit scalability of this control-plane architecture, particularly in initial deployments. Specifically, the multipoint EoMPLS control plane will need to support Ethernet broadcast traffic, meaning that packet replication will need to occur for all MPLS LSPs that are part of a multipoint EoMPLS broadcast domain. Because the first vendor implementations of multipoint EoMPLS will be supported using standard microprocessors or network processors, but not implemented in application-specific integrated circuits (ASICs), the packet-replication function could potentially consume a significant amount of processing capacity, a scenario that could limit the overall packet-forwarding rate on the network access device supporting the multipoint EoMPLS control-plane function. As a result, there are still many unknowns with respect to the scalability of the multipoint EoMPLS control plane.
- *Interoperability*—As with the point-to-point EoMPLS control plane, multipoint EoMPLS should provide significant benefits to the service providers that need to consolidate their traditional WAN services (ATM, Frame Relay, leased lines) over the same core network as their new metro Ethernet access services. On the other hand, the intercommunication between different types of MPLS Layer 2 services with multipoint EoMPLS is likely to be more complex than with point-to-point EoMPLS. None of the existing proposals for ATOM is based on a multipoint EoMPLS control plane.

Layer 3 Control Plane Supporting Layer 3 Services

As service providers continue to deploy core networks based on Layer 3 control planes, such as MPLS or IP, it is only a matter of time before they offer Layer 3 services as an alternative to the traditional point-to-point Layer 2 services, such as ATM and Frame Relay. The continuing rapid growth of the Internet will significantly influence service providers (other than Internet service providers [ISPs]), that have previously been reluctant to offer Layer 3 services. Metro Ethernet access services are likely to be among the first services that will be offered based on a Layer 3 service paradigm, although Layer 2 metro Ethernet access services will remain an important service offering for many years to come.

This section reviews the control-plane architectures that are being considered for the deployment of Layer 3 metro Ethernet access services.

Figure 3 Layer 3 Control Plane with Layer 3 Services—MPLS VPN



MPLS VPN

MPLS VPN is the principal technology for delivering virtual-private-network services based on both a Layer 3 control plane and a Layer 3 data plane. This service will provide significant flexibility to end users with respect to modifying their WAN topology over time. However, it will also require the service provider to support a new provisioning paradigm that accounts for the Layer 3 addressing structure of the end users' network, a scenario that service providers have been able to avoid by offering just Layer 2 services. Metro Ethernet access services will be one mechanism for delivering the MPLS VPN service to the customer.

- *Cost-effectiveness*—MPLS VPN services will require very capable Layer 3 network elements to support this service. As such, these network elements will tend to be more expensive than Layer 2 Ethernet switches that were used in metro Ethernet configurations based on the 802.1Q control-plane architecture. The higher cost of an MPLS VPN network element will be offset by the incremental service fees that the service provider will be able to charge for the MPLS VPN service because of the outstanding flexibility in terms of interoperability and scalability provided by these services to the end user.
- *Service levels*—MPLS VPN is able to take advantage of all the QoS, traffic engineering, and resiliency capabilities that are intrinsic to the MPLS technology. As a result, MPLS VPN can support a wide range of service-level definitions.
- *Point-to-point versus multipoint*—MPLS VPN will inherently support any logical topology, both point to point and multipoint.
- *Transparency*—As a Layer 3 network service, MPLS VPN will not be concerned with Layer 2 VLAN transparency for the end customer, so this criteria is not relevant to an evaluation of MPLS VPN.

- *Scalability*—As with all MPLS-based services, MPLS VPN is highly scalable in almost every respect. The only possible restriction on this scalability will be the size of the routing table to be supported by the service provider for each end user's MPLS VPN service. Any MPLS VPN network element will have a limitation as to the total number of route paths it can support, so as the number of end users connected to a particular MPLS VPN network element grows, the number of routes that can be supported per customer by the MPLS VPN network element will decrease. This limitation will be vendor specific.
- *Interoperability*—MPLS VPN services will provide a very high level of interoperability, because all Layer 2 service types (that is, Frame Relay, ATM, metro Ethernet) can simultaneously connect into an MPLS VPN network. MPLS VPN services can also interoperate and coexist with other types of MPLS services, including EoMPLS.

Hybrid Control-Plane Architectures

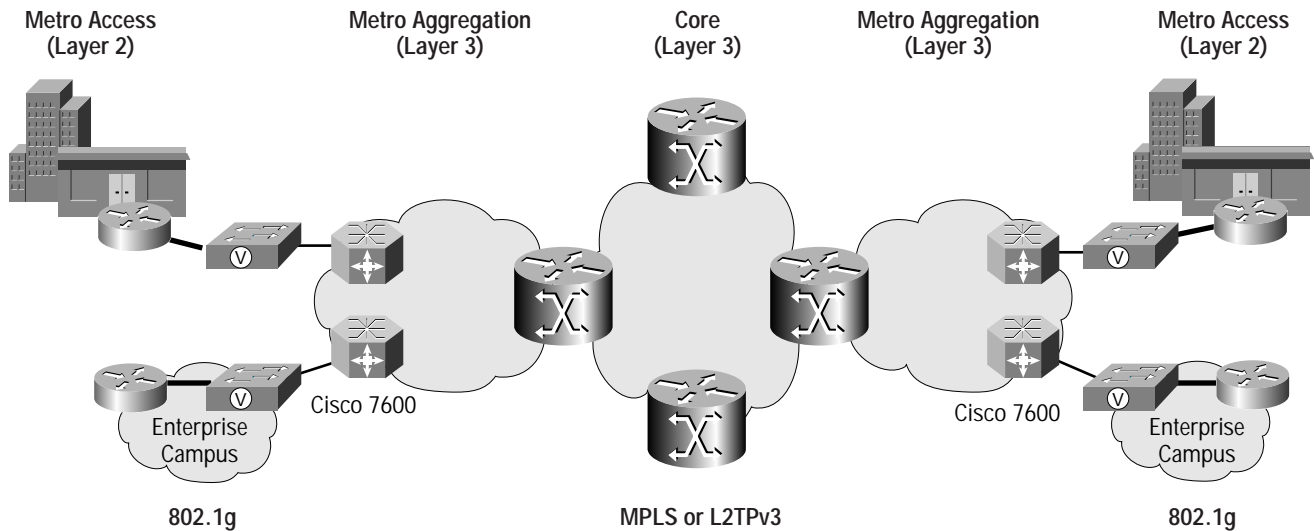
As the preceding discussion makes clear, none of the metro Ethernet control-plane architectures—Layer 2 or Layer 3—is completely satisfactory under all evaluation criteria. Table 1 summarizes ratings for each of the control-plane architectures across all the evaluation criteria.

Table 1 Summary Comparison of Layer 2 and Layer 3 Control Planes

		Layer 2		Layer 3			
		802.1Q	Q-in-Q	EoMPLS Point-to-Point	EoMPLS Multipoint	L2TPv3	MPLS VPN
Cost-Effective	CAPEX	Satisfactory	Satisfactory	Acceptable in certain designs	Acceptable in certain designs	Acceptable in certain designs	Acceptable in certain designs
Cost-Effective	OPEX	Acceptable in certain designs	Acceptable alternative	Satisfactory	Acceptable alternative	Satisfactory	Acceptable alternative
Service Level		Acceptable alternative	Acceptable alternative	Outstanding	Acceptable alternative	Satisfactory	Satisfactory
Point-to-Point vs. Multi-point		Outstanding	Outstanding	Acceptable alternative	Outstanding	Acceptable alternative	Outstanding
VLAN Transparency		Unacceptable	Outstanding	Satisfactory	Satisfactory	Satisfactory	Acceptable alternative
Service Scalability		Acceptable in certain designs	Acceptable in certain designs	Outstanding	Satisfactory	Outstanding	Satisfactory
Service Interoperability		Acceptable	Unacceptable	Outstanding	Satisfactory	Outstanding	Outstanding

This table quickly reveals a pattern wherein the most scalable and interoperable Layer 3 control planes may be less cost-effective than the Layer 2 control planes, which, in turn, do not offer satisfactory scalability or interoperability. An obvious inference of this table is that it may be possible to achieve the optimal metro Ethernet control plane by deploying a hybrid architecture that combines a Layer 2 control plane at the edge with a Layer 3 control plane for the core of the network. There are so many different combinations of Layer 2 and Layer 3 control planes that we will not examine each of these combinations separately. Instead, we will attempt to make general conclusions that can be drawn for all these hybrid Layer 2 /Layer 3 control-plane architectures.

Figure 4 Hybrid Layer 2/Layer 3 Control Plane with Either Layer 2 or Layer 3 Services



- **Cost-effectiveness**—Hybrid control planes allow for relatively low-cost metro Ethernet access devices based purely on Layer 2 control planes. Because the Ethernet access device will constitute the largest number of network elements in the metro Ethernet network, this can greatly reduce the total capital expenditure required, as compared to a pure Layer 3 control-plane architecture. With respect to operating expenses, the hybrid control-plane architecture is not significantly different from either a pure Layer 2 or pure Layer 3 control-plane architecture. There are already provisioning systems that map Layer 2 Ethernet service connections into Layer 3 service connections in a simple fashion.
- **Service levels**—The ability to deliver service-level guarantees in a metro Ethernet access service based on a hybrid control-plane architecture can be equivalent to any of the pure Layer 3 control-plane architectures, particularly when the metro access network, based on Ethernet transport with 802.1Q or Q-in-Q control plane, is not oversubscribed. If the metro access network is oversubscribed, then the metro access network element must provide more sophisticated queuing functionality to ensure that high-priority queues are not starved. As for reconvergence, although spanning tree may be used in the access network, it controls only a very small number of devices, so the spanning-tree structure is simple, and, provided 802.1s and 802.1w are supported, the reconvergence, when needed, can occur in subsecond time frames.
- **Point-to-point versus multipoint**—Hybrid control planes offer tremendous flexibility for supporting both Layer 2 and Layer 3 multipoint services. By combining either 802.1Q or Q-in-Q with multipoint EoMPLS, the service provider can establish a Layer 2 multipoint solution. If the service provider wants to deliver a Layer 3 service, 802.1Q can be combined with MPLS VPN. In either case, the 802.1Q control plane at the edge provides a Layer 2 circuit from the CPE to the MPLS provider-edge device.

- *Transparency*—If Q-in-Q is used as the Layer 2 control plane for the metro access network, it is possible to support VLAN transparency in any of the hybrid control-plane combinations. The only exception is MPLS VPN, but VLAN transparency is not a requirement with MPLS VPN.
- *Scalability*—Hybrid control planes can be just as scalable as any of the pure Layer 3 control-plane architectures discussed. For each of the possible hybrid control-plane combinations, the Layer 2 metro access control plane, based on 802.1Q or Q-in-Q, allows for the 802.1Q tags to be locally significant only, allowing the number of user connections to scale to the limit of any of the Layer 3 control planes, and well beyond the 4096 limit of 802.1Q VLAN tags. The use of the Layer 3 control plane in the metro aggregation and core networks further enhances the scalability of the hybrid control plane. Figure 4 illustrates the relation between the Layer 2 and Layer 3 control planes in a hybrid control-plane architecture.
- *Interoperability*—Hybrid control planes can support as high a level of interoperability as a pure Layer 3 control plane. In fact, the hybrid control plane may offer advantages for interoperability between metro Ethernet access services and traditional WAN services, such as ATM and Frame Relay, as long as the network element performing network aggregation can interface to each of the various Layer 2 service types.

Metro Ethernet Control-Plane Selection Strategies

So, with the variety of control-plane solutions that are being proposed for supporting the new metro Ethernet access services, how should a service provider proceed in selecting the appropriate one for the metro Ethernet access services it intends to offer?

As a starting point, the service provider must clearly understand the relation between the evaluation criteria presented in this document and the type of metro Ethernet access service that the service provider intends to deliver. The fact that one of these control-plane technologies might be rated as unacceptable for a particular evaluation criterion does not mean that control-plane technology should not be deployed.

For example, the Layer 2 control planes, 802.1Q and Q-in-Q, both have significant limitations for scalability. But this may not be important to the service provider if it does not intend to offer a nationwide metro Ethernet access service, but instead will deliver services that span only specific smaller-scale metropolitan deployments, and the interconnection between these separate metro networks is intended to be performed by separate long-haul service providers that would connect to the metro Ethernet access service.

Similarly, the fact that the Layer 3 control planes require higher capital investment than the Layer 2 control planes does not mean that a pure Layer 3 control plane should never be deployed. In markets where the potential customers are densely populated, such as New York City, it might make the most economic sense to deploy the metro aggregation network element at the maximum-transmission-unit (MTU) location, rather than the central office or POP, and then have customers connect directly from their CPE device into the service provider's metro aggregation network element, completely bypassing a metro access network element. This might also apply in cases where each customer is requesting a large amount of bandwidth, such as a full Gigabit Ethernet. In such cases, it could make sense for a Layer 3 control plane to extend all the way to the CPE.

So, the first step for a service provider in selecting a control-plane architecture for a metro Ethernet access service is to be as thorough as possible in defining the nature of the metro Ethernet access services that it intends to offer to customers. When that service definition is established, then, using the criteria set forth in this document, the service provider can make an assessment of which control-plane architecture will be most suitable for delivering that metro Ethernet access service.

Summary

As this document has indicated, the deployment of metro Ethernet access services will require both service providers and data-communication equipment vendors to develop new control-plane technologies to deliver the services that will meet the end users' requirements. Although Ethernet switching technology has a well-established track record for delivering cost-effective network solutions, the control-plane architectures that will be used to deploy metro Ethernet access services are still evolving. Both service providers and vendors will need to track the evolution of these control-plane technologies as their product development and network deployments proceed.

Because these control-plane technologies are still evolving, perhaps the most important objective for a service provider in selecting a metro Ethernet control-plane architecture is to maintain the greatest degree of flexibility. This flexibility in selecting a metro Ethernet control plane must account for both the kinds of services that can be delivered and the types of transport technologies (that is, SONET, Gigabit Ethernet, RPR) those services can be delivered on. With this in mind, it makes sense to deploy metro Ethernet access services over network elements that offer the most flexibility in supporting any combination of the control-plane architectures that have been discussed in this document. Furthermore, those network elements should also allow deployment of these control-plane architectures over any combination of transport technologies that will be used in a metro Ethernet network topology.

As the market leader in Ethernet switch technologies, Cisco Systems recognizes the metro service-provider market as one of the most important new markets for Ethernet switching technologies and products. For this reason, Cisco has made significant investments in developing the control-plane architectures that will support these metro Ethernet access services. For example, Cisco was the first vendor to both announce and ship a functional implementation of Ethernet over MPLS based on the IETF Martini draft. Furthermore, Cisco is deeply involved in the standards bodies that are defining these control-plane architectures.

Cisco intends to maintain its leadership in Ethernet switch technologies by delivering to the market each of the control-plane architectures discussed in this document, in such a way that any of these architectures can be deployed over any of the common metro Ethernet transport technologies. Cisco's control-plane modularity is one example of how Cisco delivers the most complete modular network systems for deployment flexibility with operational consistency through multiuse products and systems. In addition, Cisco intends to deliver the management and provisioning capabilities that will combine these technologies into a complete solution that will allow service providers to aggressively deploy metro Ethernet access services in order to achieve the highest possible end-user satisfaction from the service-provider customers that are, ultimately, driving the demand for these metro Ethernet access services.

References

- L2TPv3-Introduced at IETF 51 (London) August 2001-draft-ietf-l2tpext-l2tp-base-00.txt
- Lassere, Transparent VLAN Services over MPLS-draft-lasserre-tls-mpls-00.txt
- IETF Martini Draft: "Encapsulation Methods for Transport of Layer 2 Frames over MPLS" -draft-martini-l2circuit-encap-mpls-02.txt (work in progress)
- IETF Martini Draft: "Transport of Layer 2 Frames over MPLS" -draft-martini-l2circuit-trans-mpls-06.txt



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy Les Moulineaux Cedex 9
France
www.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Inc
Capital Tower
168 Robinson Road
22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE
Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia
Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2001 Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0110R)