MPLS and Traffic Engineering in IP Networks

Daniel O. Awduche, UUNET (MCI Worldcom)

ABSTRACT Rapid growth and increasing requirements for service quality, reliability, and efficiency have made traffic engineering an essential consideration in the design and operation of large public Internet backbone networks. Internet traffic engineering addresses the issue of performance optimization of operational networks. A paramount objective of Internet traffic engineering is to facilitate the transport of IP traffic through a given network in the most efficient, reliable, and expeditious manner possible. Historically, traffic engineering in the Internet has been hampered by the limited functional capabilities of conventional IP technologies. Recent developments in multiprotocol label switching and differentiated services have opened up new possibilities to address some of the limitations of the conventional technologies. This article discusses the applications of MPLS to traffic engineering in IP networks.

Internet growth in recent times has been very impressive. A report from the U.S. Department of Commerce [1] suggests that the rate at which the Internet has been adopted has surpassed all other technologies preceding it, including radio, television, and the personal computer. Today, the Internet has become a convenient and cost-effective medium for collaboration, education, electronic commerce, and entertainment. A common consensus is that the Internet will metamorphose into a medium for the convergence of voice, video, and data communications. Although the long-term market behavior of the Internet is difficult to forecast, Internet traffic is clearly growing in a geometric progression. Reported compounded traffic growth rates range from two to ten times per annum.

Large Internet service providers (ISPs) have responded to the challenge of Internet growth by employing three complementary technical instruments:

- Network architecture
- Capacity expansion
- Traffic engineering

Network architecture deals with the abstract structure of networks, the components or object classes of the network, their functions, and the relationships between them. A good, scalable network architecture, premised on sound architectural principles, is imperative in the quickly evolving Internet environment.

The second instrument employed by large ISPs to respond to traffic growth is rapid expansion of capacity and network infrastructure. In 1996 most large ISPs in the United States operated backbones with DS3 (44.736 Mb/s) links. In 1997 and 1998, OC-12c (622 Mb/s) links became pervasive. In 1999 a number of major ISPs upgraded to OC-48c (2.488 Gb/s) links. By the year 2000, some ISPs expect to begin deployment of IP backbones with OC-192c (9.953 Gb/s) links, provisioned directly over dense wavelength-division multiplexing (DWDM) facilities.

The third instrument employed by service providers to address the Internet growth challenge is traffic engineering. This subject has attracted significant attention in recent times [2–9]. A motivation for Internet traffic engineering is the realization that architectural paradigms and simple capacity expansion are necessary, but not sufficient, to deliver high-quality Internet service under all circumstances. Internet traffic engineering is the aspect of Internet network engineering that addresses the issue of performance optimization of operational networks. It encompasses the application of technology and scientific principles to the measurement, modeling, characterization, and control of Internet traffic [2]. It also includes the application of knowledge and techniques to achieve specific performance objectives, including reliable and expeditious movement of traffic through the network, efficient utilization of network resources, and planning of network capacity. Ultimately, good traffic engineering increases the value of a network to both the service provider and the Internet user community.

Historically, effective traffic engineering has been difficult to achieve in public IP networks. The reason for this is the limited functional capabilities of conventional IP technologies. One particular shortcoming of conventional IP systems is the inadequacy of measurement functions. For example, a traffic matrix, which is a basic data set needed for traffic engineering, is difficult to estimate from interface statistics on IP routers. The limitations of intradomain routing control functions are another issue with conventional IP systems. Interior gateway protocols (IGPs), such as Intermediate System–Intermediate System (IS-IS) and Open Shortest Path First (OSPF), commonly used to route traffic within autonomous systems in the Internet, are topology-driven and employ per-packet progressive connection control. Each router makes independent routing decisions using a local instantiation of a synchronized routing area link state database. Route selection is based on shortest path computations using simple additive link metrics. This approach is highly distributed and scalable, but flawed. The flaw is that these protocols do not consider the characteristics of offered traffic and network capacity constraints when making routing decisions. This results in subsets of network resources becoming congested, while other resources along alternate paths remain underutilized [2]. This type of congestion problem is a symptom of poor resource allocation, and is an issue that traffic engineering specifically attempts to redress.

Recent developments in multiprotocol label switching (MPLS) [2–8] open new possibilities to address some of the limitations of IP systems concerning traffic engineering. A framework for MPLS is presented in [5] and an architecture for it described in [8]. The requirements for traffic engineering over MPLS were articulated in [2]. Although MPLS is a relatively simple technology (based on the classical label swapping paradigm), it enables the introduction of sophisticated control capabilities that advance the traffic engineering function in IP networks [2–4, 6, 7]. A particularly interesting aspect of MPLS is that it efficiently supports origination connection control through explicit label-switched paths. When MPLS is com-

0163-6804/99/$10.00 © 1999 IEEE

bined with differentiated services and constraint-based routing, they become powerful and complementary abstractions for quality of service (QoS) provisioning in IP networks.

This article discusses the applications of MPLS to traffic engineering in IP networks, focusing specifically on service provider networks. The basic concepts and challenges of traffic engineering in the Internet are introduced first. These concepts and challenges are followed by the capabilities that make MPLS applicable to traffic engineering in such environments. A review of the overlay methodology that was used for traffic engineering in classical IP networks (prior to the advent of MPLS) is also provided. This article covers intradomain traffic engineering, that is, traffic engineering within a single autonomous system in the Internet.

The remainder of this article is organized as follows. The following section introduces the basic concepts and challenges of traffic engineering in the Internet. We then describe the functional capabilities making MPLS applicable to traffic engineering in IP networks. The last section contains the concluding remarks.

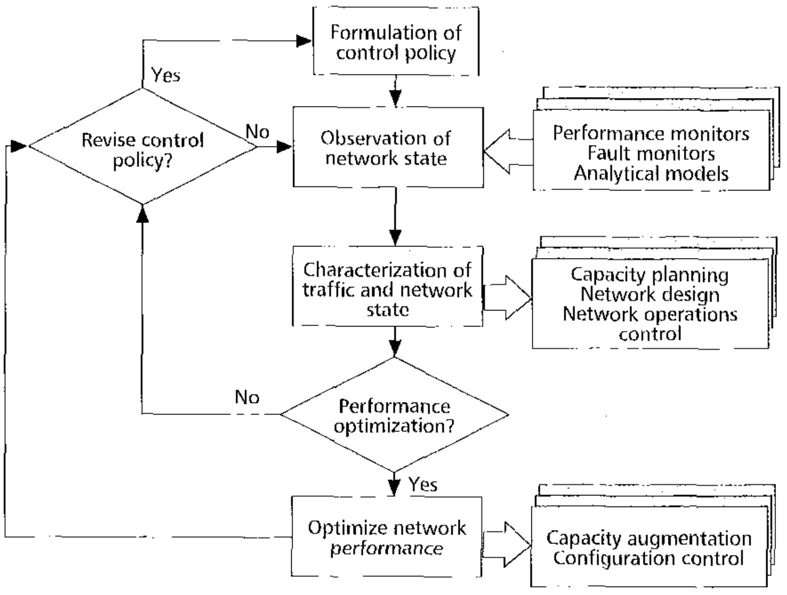
Formulation of control policy
Yes
Revise control policy?
No
Observation of network state
Performance monitors
Fault monitors
Analytical models
Characterization of traffic and network state
Capacity planning
Network design
Network operations control
No
Performance optimization?
Yes
Optimize network performance
Capacity augmentation
Configuration control

Figure 1. *The traffic engineering process model.*

The Concepts and Challenge of Traffic Engineering in IP Networks

This section introduces the concepts and practical functions of traffic engineering in operational IP networks. The challenge of traffic engineering in autonomous systems within the Internet is highlighted, and an overview of the classical IP over ATM overlay model is provided.

Internet Traffic Engineering Concepts

In concept, a network consists of a demand system (traffic), a constraint system (interconnected network elements), and a response system (network protocols and processes). Traffic engineering establishes the parameters and operating points for all three aspects of the network in an operational context. Consequently, Internet traffic engineering is fundamentally a control problem [2].

The Traffic Engineering Process Model — A number of stages can be identified in the Internet traffic engineering process model. The first stage is the formulation of a control policy. The control policy depends on the network context, cost structure, revenue or utility model, operating constraints, and success criteria. The second stage is the observation of the network state through a set of monitoring functions. This is the feedback component of the traffic engineering process model. It may include preprocessing activities such as data reduction and data transformation. The third stage is the characterization of traffic and analysis of the network state. Various qualitative and quantitative techniques can be applied in the characterization and analysis stage. Bottlenecks and pathologies that impede (or potentially impede) network performance are identified. The results are used for network performance optimization, network operations control, network design, and capacity planning. The fourth stage is the optimization of network performance. This is accomplished by applying control actions, if necessary, to drive the network to a desired state according to the control policy. Control actions may involve modifying or relaxing network resource constraints (e.g., augmenting capacity), manipulating traffic management parameters, or modifying the parameters associated with routing through a configuration control system.

Traffic engineering is an adaptive process. The four stages of the process model defined above are iterated. In an operational context, it is best to minimize the level of manual intervention involved in traffic engineering by automating the tasks whenever possible. The traffic engineering process model is illustrated in Fig. 1.

Traffic Engineering Objectives — A practical function of traffic engineering in IP networks is the mapping of traffic onto the network infrastructure to achieve specific performance objectives. High service quality, efficiency, survivability, and economy are crucial objectives in today's commercial, competitive, and mission-critical Internet. Traffic engineering requires precise control over the routing function to achieve the objectives. Indeed, an essential requirement for traffic engineering in IP networks is the capability to compute and establish a forwarding path from one node to another. This path must fulfill some requirements, while also satisfying network capacity and policy constraints. Generally, performance objectives can be *traffic-oriented* and/or *resource-oriented.*

Traffic-oriented performance objectives relate to the improvement of the QoS provisioned to Internet traffic. Traffic-oriented performance metrics include packet loss, delay, delay variation, and goodput. The effectiveness of traffic-oriented policies can also be measured in terms of the relative proportion of offered traffic achieving their performance requirements. When service level agreements (SLAs) are involved, protecting traffic streams that comply with their SLAs from those that are noncompliant becomes an important factor in the attainment of traffic-oriented performance objectives. Resource-oriented performance objectives relate to the optimization of the utilization of network assets. Efficient resource allocation is the basic approach to secure resource-oriented performance objectives. A traffic engineering system

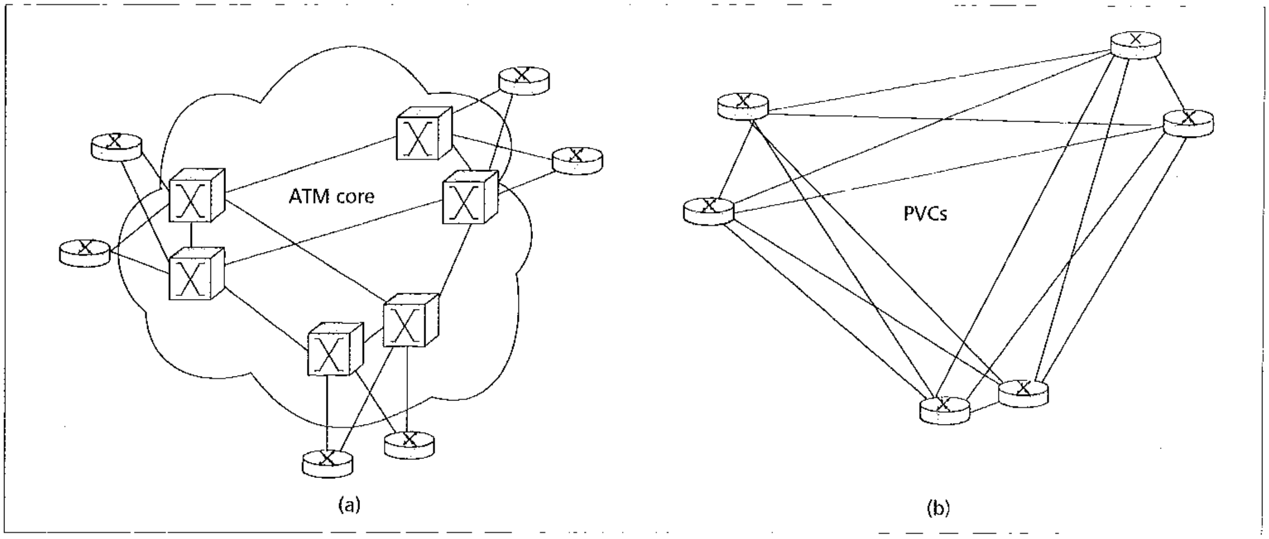

Figure 2. *The overlay model, IP over ATM: a) physical network topology; b) logical network topology.*

is said to be "rational" if it addresses traffic-oriented performance problems while simultaneously utilizing network resources efficiently.

Minimizing congestion is a central goal of traffic engineering. Congestion typically manifests under the following scenarios:

- When network resources are insufficient or inadequate to handle offered load
- When traffic is inefficiently mapped onto resources, causing subsets of resources to become overutilized while others remain underutilized [2]

Congestion problems resulting from insufficient or inadequate resources can be addressed by: augmenting network capacity, or modulating, conditioning, or throttling the demand so that the traffic fits onto the available capacity (e.g., using policing, flow control, rate shaping, link scheduling, queue management, tariffs, *et al.*). Congestion problems resulting from inefficient mapping of traffic onto resources can be addressed by increasing the efficiency of resource allocation. An example of this increased efficiency of resource allocation would be to route some traffic away from congested resources to relatively underutilized ones.

Reliable network operation is another important objective of Internet traffic engineering. Multiple failure recovery scenarios must be devised to ensure continuity of service following network impairments. Adequate capacity for service restoration must be provisioned, therefore, and the operational capability must exist to expeditiously reroute traffic through the redundant capacity when faults occur. Reoptimization may be required following restoration to make more effective use of the residual post-fault capacity. It may be advantageous to utilize subsets of the redundant capacity to improve network performance and efficiency when the network is fault-free.

Traffic engineering becomes even more critical in a multiclass service environment like the emerging differentiated services Internet, where traffic streams with different service requirements are in contention for network resources. In these environments, traffic engineering establishes the resource sharing parameters so that the network provides preferential treatment to some service classes in accordance with a utility model.

The Challenge of Internet Traffic Engineering

Traffic engineering in conventional IP networks is a challenging problem. Singularities and discontinuities characterize Internet growth. Very rapid growth occurs over relatively short intervals of time. This rapid growth is then followed by modest growth over relatively longer intervals of time. Accurate forecasting is therefore quite difficult. Furthermore, Internet traffic exhibits very dynamic behavior with characteristics that are not yet well understood. Traffic also tends to be highly asymmetric.

The operating environment is also in a continual state of flux. New resources are added constantly. Resources also fail regularly. New Internet applications with bandwidth requirements which may have significant global impact are introduced all the time. Facility location is also an issue. Sometimes network resources are sited in less than ideal locations due to facility constraints. Additional complications are introduced by interdomain traffic traversing autonomous systems' boundaries. These environmental factors result in the network topology not usually correlating with the traffic matrix. Addressing these issues requires continual monitoring and performance optimization of public IP networks.

Traffic Engineering with the Classical Overlay Model

The overlay model is a technique that was applied, prior to MPLS, to circumvent some of the limitations of IP systems regarding traffic engineering. The basic idea is to introduce a secondary technology, with virtual circuit and traffic management capabilities (e.g., asynchronous transfer mode, ATM), into the IP infrastructure in an overlay configuration. The virtual circuits of the secondary technology serve as point-to-point links between IP routers. Figure 2 illustrates the overlay

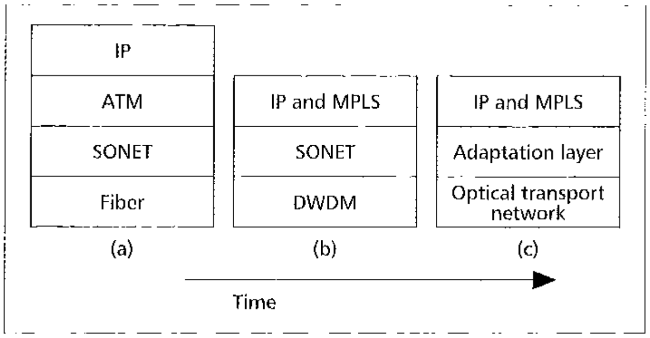

Figure 3. *Technology layer evolution.*

model with ATM switches at the core surrounded by an epidermis of IP routers logically interconnected by ATM permanent virtual connections (PVCs).

The overlay approach extends the design space and allows arbitrary virtual topologies to be defined and superimposed onto the physical network topology. The overlay model also permits estimation of a rudimentary traffic matrix from statistics on the PVCs that interconnect routers. Traffic engineering can also move traffic from overloaded links to relatively underutilized links by changing the designated transit lists (DTLs) of a subset of PVCs.

There are fundamental drawbacks to the IP over ATM overlay model. Perhaps the most significant problem is the need to build and manage two networks with dissimilar technologies. The overlay model also increases the complexity of network architecture and network design. Reliability is also a concern because more network elements now exist in series on the routed path. Scalability is another issue because the number of adjacencies in the overlay graph generally increases quadratically with the number of routers, thereby increasing the CPU and network resource consumption associated with routing. Other issues include the quantization and encapsulation overhead associated with ATM, the technical difficulties inherent in developing microelectronics that perform segmentation and reassembly (SAR) at very high speeds (OC-48c and above), and the possibility of routing instability in the IP domain induced by multiple PVCs failures following a single interswitch link impairment in the ATM core.

The trend, therefore, is to evolve core IP networks away from the overlay model and toward more integrated solutions. This evolution is now possible because of developments in MPLS and recent advances in high-performance gigabit/terabit routers and optical internetworking systems. Figure 3 presents the expected technology layer evolution of core IP networks from IP over ATM over synchronous optical network (SONET) over fiber to IP with MPLS over SONET over DWDM; and, finally, to IP with MPLS over an adaptation layer interfacing with a versatile optical transport network (OTN).

MPLS and Traffic Engineering in IP Networks

The functional capabilities making MPLS attractive for traffic engineering in IP networks are described in this section. General discussions of MPLS technology itself are detailed in a number of documents from the IETF MPLS working group [2, 5, 8].

MPLS allows sophisticated routing control capabilities to be introduced into IP networks. These capabilities are buttressed on the fact that MPLS efficiently supports origination connection control through explicit label-switched paths (LSPs). An explicit LSP is one whose route is determined at the origination node. Origination connection control permits explicit routes to be established which are independent of the destination-based IP shortest path routing model. Once an explicit route is determined, a signaling protocol is then used to install the LSP. Through explicit LSPs, a quasi circuit switching capability is superimposed on the IP routing model [2]. When deployed in IP over SONET or IP over DWDM configurations, the traditional L3 and L2 functions are virtualized in one network element called the label switching router (LSR). Fewer network elements are required than with overlay alternatives, reliability is increased, and operating costs and queuing delays are reduced. Additionally, MPLS simplifies network architecture and network design relative to the overlay model. MPLS coupled with differentiated services and constraint-based routing fundamentally changes the way core IP networks are designed and traffic engineered. Figure 4 depicts a simplified topology that results when LSRs that are equipped with packet over SONET interfaces replace the overlay network previously shown in Fig. 2.

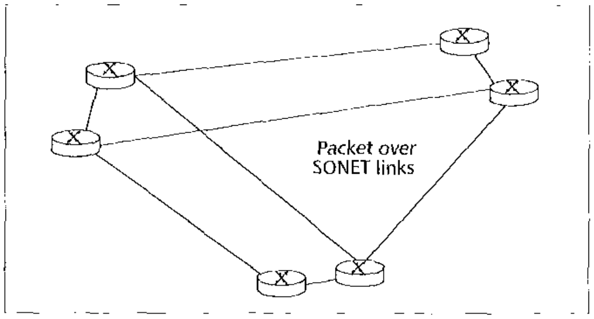

Figure 4. *Simplified IP network topology with MPLS.*

One feature of traffic engineering in IP networks is the traffic trunk concept. A traffic trunk is an aggregation of traffic belonging to the same class [2]. It is essentially an abstract description of traffic that allows certain attributes of the traffic to be parameterized. It is independent of the underlying technologies. The problem of mapping traffic trunks onto a given network topology is a central issue of traffic engineering. In MPLS networks, traffic trunks are mapped onto the network topology through the selection of routes for explicit LSPs. The terms *LSP tunnel* [3] and *traffic engineering tunnel* (te-tunnel) [9] are commonly used to refer to the combination of traffic trunk and explicit LSPs in MPLS.

LSP tunnels allow the performance of an operational network to be optimized in various ways. For example, if congestion problems caused by suboptimal routing are detected, LSP tunnels can be rerouted to alleviate the problem. LSP tunnels can be parameterized, and network resources can be allocated to them based on those parameters. Multiple LSP tunnels can be created between two nodes, and the traffic between the nodes divided among the tunnels according to some local policy. LSP tunnels permit the introduction of flexible and cost-effective survivability options. Statistics derived from LSP tunnels can be used to construct a

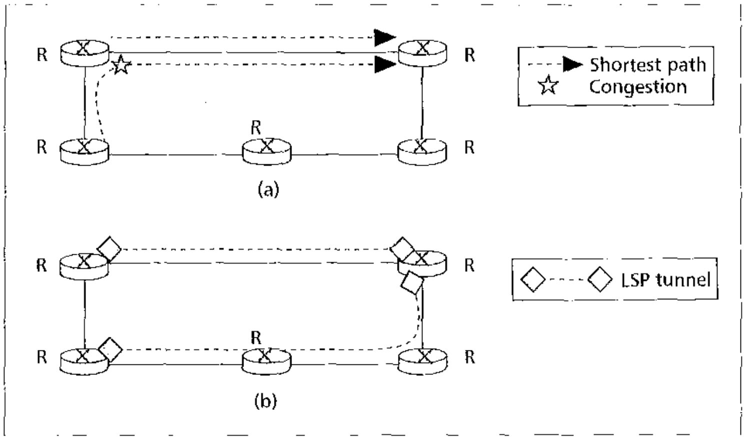

Figure 5. *Addressing congestion problems with LSP tunnels: a) congestion caused by intersecting shortest paths; b) traffic redistribution using LSP tunnels.*

rudimentary traffic matrix. Figure 5 shows how LSP tunnels can be used to redistribute traffic to address congestion problems caused by shortest path IGPs.

COMPONENTS OF THE MPLS TRAFFIC ENGINEERING MODEL

An MPLS traffic engineering model consists of the following basic functional components:

- Path management
- Traffic assignment
- Network state information dissemination
- Network management

These are elements of the MPLS control plane and are distinct from the forwarding plane.

Path Management — Path management concerns all aspects related to the selection of explicit routes, and the instantiatiation and maintenance of LSP tunnels. A path management policy defines the path selection criteria as well as rules for sustaining already established LSP tunnels [2]. Path management consists of three primary functions: path selection, path placement, and path maintenance.

The path selection function specifies the explicit route for an LSP tunnel at the origination node of the tunnel. An explicit route can be represented as a sequence of hops or a sequence of abstract nodes. It may contain both strict and loose subsets. An abstract node is a group of nodes whose internal topology is opaque to the origination node. Explicit routes can be defined administratively or computed automatically by a constraint-based routing entity.

Constraint-based routing is a generalization of QoS routing. It is used to compute routes that satisfy a set of requirements, subject to constraints imposed by the network and administrative policies. Constraint-based routing reduces the level of manual intervention involved in traffic engineering.

The second component of path management is the path placement function. This is used to instantiate LSP tunnels using a signaling protocol, which also serves as a label distribution protocol. Two MPLS signaling protocols are currently defined: Resource Reservation Protocol (RSVP) extensions [3, 4] and constraint-based routed LDP (CR-LDP) [6]. The third component of path management is the path maintenance function, which sustains and terminates already established LSP tunnels.

A set of attributes can be associated with LSP tunnels and network resources to guide the path management functions and to provide controls over constraint-based routing. An important operational requirement is the capability to manipulate the attributes of active LSP tunnels to cause certain transitions (e.g., explicit route changes) to occur gracefully without adversely impacting network operations.

LSP tunnel attributes include traffic parameters, adaptivity attributes, priority attributes, preemption attributes, resilience attributes, resource class affinity attributes, and other policy options such as policing attributes [2]. Traffic parameters specify the bandwidth characteristics of the LSP tunnel, and may include peak rates, mean rates and burst sizes — or the parameters may simply specify an effective bandwidth. The adaptivity attributes indicate the sensitivity of an LSP tunnel to the dynamics of the network state. Adaptive LSP tunnels can be rerouted automatically when better routes become available. Nonadaptive LSP tunnels are pinned to their established routes except under faults. Priority attributes impose a partial order on multiple LSP tunnels, according to which path selection and path placement are sequenced. Currently, eight setup priority levels are specified [3]. The preemption attributes determine whether a new LSP tunnel can acquire the resources allocated to an existing tunnel. Preemption is implemented using a combination of setup and reservation priorities [3]. Various prioritized restoration schemes can be implemented in a multiclass environment using preemption. Resilience attributes specify the response of an LSP tunnel to impairments that impact its route. A basic resilience attribute specifies whether an LSP tunnel is to be automatically rerouted following faults along its established path. Extended resilience attributes incorporate more sophisticated recovery policies, including policies that instantiate multiple parallel LSP tunnels together with rules to determine their relative preference under faults. Resource class affinity attributes impose additional policy restrictions on the qualification of sets of resources for LSP tunnel path selection. An affinity relationship between an LSP tunnel and a resource class indicates whether the resource class is to be included or excluded from the path of the LSP tunnel.

Resource attributes define additional properties of network resources that further constrain the routing of LSP tunnels through them. Resource attributes include the maximum allocation multiplier (MAM), the default traffic engineering metric, and resource class attributes. The MAM concept is analogous to subscription and overbooking factors in frame relay and ATM networks. The default traffic engineering metrics can be used to establish route optimization criteria for LSP tunnels independent of IGP metrics.

Resource class attributes are used to categorize resources, primarily links, into different classes. Uniform policies, such as inclusion and exclusion, can then be applied to each resource class with respect to LSP tunnel path selection. A link can belong to more than one resource class. The resource class attribute is part of the link state parameters. Resource class attributes can be used to contain traffic within specific topological regions of a network.

Traffic Assignment — Traffic must be assigned to an LSP tunnel once the tunnel is established. Traffic assignment concerns all aspects related to the allocation of traffic to established LSP tunnels. It consists of a partitioning function and an apportionment function. The partitioning function partitions ingress traffic according to some principle of division. The apportionment function allots the partitioned traffic to established LSP tunnels according to some principle of allocation. The potential flexibility in traffic assignment fundamentally distinguishes MPLS from ATM.

One way to automate the traffic assignment problem is to view LSP tunnels as shortcuts through the IGP domain [9]. Additional attributes may be introduced to control the assignment function when there are multiple paths to a given node. Filtration rules may be applied to restrict the class of traffic mapped onto a given LSP tunnel. Filtration rules may, for example, be used to define the way differentiated services behavior aggregates are mapped onto LSP tunnels.

Load distribution across multiple LSP tunnels between two nodes is an important traffic assignment issue. The load distribution problem can be addressed by implicitly or explicitly assigning weights to each LSP tunnel and apportioning traffic in relative proportion to the weights. Load distribution across parallel LSP tunnels can also be implemented as a feedback function of the state of the network.

Network State Information Dissemination — Network state information dissemination concerns the distribution of relevant topology state information throughout the MPLS domain. This is accomplished by extending conventional IGPs to propagate additional information about the state of the network in link state advertisements [7]. The additional information distributed includes maximum link bandwidth, maximum allocation multiplier, default traffic engineering metric, reserved bandwidth per priority class, and resource class attributes. The topology state information is used by constraint-based routing entities to select feasible routes for LSP tunnels.

Network Management — Network management is an important aspect of traffic engineering over MPLS. The success of the MPLS approach to traffic engineering eventually depends on the ease with which the network can be observed and controlled. Generally, an MPLS network management system includes a set of configuration management functions, performance and accounting management functions, and fault management functions. Collectively, these functions allow the state of managed MPLS objects (e.g., LSP tunnels) to be acquired and their characteristics (and ultimately network performance) controlled. Point-to-point traffic flows can be characterized by monitoring traffic statistics on LSP tunnels. Path loss characteristics can be estimated by monitoring ingress and egress traffic statistics at both endpoints of an LSP tunnel and noting discrepancies. Path delay characteristics can be estimated by sending probe packets through LSP tunnels and measuring the transit times. Event notifications can be issued when the state of a managed MPLS object exceeds prescribed thresholds. Bulk retrieval of LSP tunnel traffic statistics can be used for time series analysis and capacity planning purposes. An operational requirement is the capability to list, at any given point in time, all the nodes traversed by an LSP tunnel, and for each node to list all of the LSP tunnels originating from it, terminating on it, and traversing through it.

Because optimizing the performance of large-scale networks is an intractable problem, offline traffic engineering support tools may be required to augment the online capabilities of MPLS. Such offline tools may be interfaced with the MPLS network management system to provide external feedback control.

Conclusion

This article discusses the applications of multiprotocol label switching to traffic engineering in IP networks. The concepts and challenges of traffic engineering in the Internet are reviewed. The overlay model is described. This model is based on IP over ATM and is an alternative to MPLS. Finally, the functional capabilities making MPLS useful for traffic engineering in IP networks are highlighted.

References

[1] USDC "The Emerging Digital Economy" U.S. Dept. of Commerce, Apr. 1998, http://www.ecommerce.gov/emerging.htm

[2] D. Awduche *et al.*, "Requirements for Traffic Engineering Over MPLS," RFC 2702, Sept. 1999.

[3] D. Awduche *et al.*, "Extensions to RSVP for Traffic Engineering," IETF Internet draft, work in progress, Feb. 1999.

[4] D. Awduche, A. Hannan, and X. Xiao, "Applicability Statement for Extensions to RSVP for LSP-Tunnels," IETF Internet draft, work in progress, July 1999.

[5] R. Callon *et al.*, "A Framework for Multiprotocol Label Switching," IETF Internet draft, work in progress, Nov. 1997.

[6] B. Jamoussi *et al.*, "Constraint-Based LSP Setup Using LDP," IETF Internet draft, work in progress, Feb. 1999.

[7] T. Li, G. Swallow, and D. Awduche, "IGP Requirements for Traffic Engineering with MPLS," IETF Internet draft, work in progress, Feb. 1999.

[8] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," IETF Internet draft, work in progress, July 1998.

[9] N. Shen and H. Smit, "Calculating IGP Routes over Traffic Engineering Tunnels," IETF Internet draft, work in progress, June 1999.

Biography

Daniel O. Awduche [M] (awduche@uu.net) is the manager of advanced technology at UUNET, an MCI Worldcom Company, a global provider of Internet communications services. In this role he leads a team of engineers engaged in architecture, design, and development activities aimed at UUNET's next-generation network. He orchestrated UUNET's MPLS strategy from concept to the first large-scale deployment of MPLS in the world. He is very active in the IETF where he serves as co-chair of the Internet Traffic Engineering Working Group. He is a co-author of several IETF documents on MPLS. He is also active in the Optical Internetworking Forum, where he serves as editor of the "Optical Internetwork Architecture: Functions and Reference Models" document. He is a member of the ACM. He did graduate studies in computer systems engineering at the University of Massachusetts, Amherst.