

Traffic Engineering with MPLS in the Internet

Xipeng Xiao, Alan Hannan, Brook Bailey

GlobalCenter Inc, A Global Crossing Company
141 Caspian Court
Sunnyvale, CA 94089

[xipeng, alan, bbailey}@globalcenter.net](mailto:{xipeng, alan, bbailey}@globalcenter.net)

Lionel M. Ni

Department of Computer Science
3115 Engineering Building
Michigan State University
East Lansing, MI 48824-1226

ni@cse.msu.edu

Abstract

This paper discusses *Traffic Engineering with Multi-Protocol Label Switching (MPLS)* in an *Internet Service Provider's (ISP)* network. In this paper, we first briefly review *MPLS*, *Constraint-based Routing* and enhanced link state *Interior Gateway Protocols (IGPs)* to provide a background for Traffic Engineering. We then discuss the general issues of designing an MPLS system for Traffic Engineering. The design of GlobalCenter's MPLS system is presented. Based on our experiences, a generic procedure for deploying MPLS system is proposed. We also discuss how to provide *Quality of Service (QoS)* in a network with MPLS. Putting these together, we present the readers the practical issues of Traffic Engineering and a working solution for Traffic Engineering with MPLS in the Internet.

Keywords: Traffic Engineering, Multi-Protocol Label Switching (MPLS), Constraint-based Routing, Link State IGP, Quality of Service (QoS)

1. Introduction

Traffic Engineering is the process of controlling how traffic flows through one's network so as to optimize resource utilization and network performance [1, 2, 3]. Traffic Engineering is needed in the Internet mainly because current IGPs always use the shortest paths to forward traffic. Using shortest paths conserves network resources, but it may also cause the following problems.

1. The shortest paths from different sources overlap at some links, causing congestion on those links.
2. The traffic from a source to a destination exceeds the capacity of the shortest path, while a longer path between these two routers is under-utilized.

There is a debate of whether network capacity will one day become so cheap and abundant that these two problems will be eliminated. This debate is beyond the scope of this paper. Here we simply note that currently all ISPs have the above problems. By performing Traffic Engineering in their networks, ISPs can greatly opti-

mize resource utilization and network performance. Revenue can be increased without large investment in upgrading network infrastructure. Therefore, Traffic Engineering is definitely useful for ISPs now.

Traffic Engineering is difficult to do with IGP in large networks for the following reasons:

1. Among the *Equal-Cost Multi-Paths* (ECMPs) from a source, every path will have an equal share of load. This equal ratio cannot be changed. Therefore, one of the paths may end up carrying significantly more traffic than other paths because it also carries traffic from other sources.
2. Load sharing cannot be done among multiple paths of different cost.
3. Modifying IGP metric to trigger some traffic shift tends to have side effects, and undesirable traffic shifts may also be triggered.

In order to do Traffic Engineering effectively, the *Internet Engineering Task Force* (IETF) introduces MPLS [4], Constraint-based Routing [6] and an enhanced link state IGP [7]. They are briefly reviewed in this section.

1.1 MPLS

MPLS is an advanced forwarding scheme. It extends routing with respect to packet forwarding and path controlling.

Each MPLS packet has a header. In a non-ATM environment, the header contains a 20-bit label, a 3-bit *Experimental* field (formerly known as *Class of Service*, or CoS, field), a 1-bit label stack indicator and an 8-bit TTL field. In an ATM environment, the header contains only a label encoded in the VCI/VPI field. An MPLS capable router, termed *Label Switching Router* (LSR), examines the label and possibly the experimental field in forwarding the packet.




Figure 1. MPLS

At the ingress LSRs of an MPLS-capable domain IP packets are classified and routed based on a combination of the information carried in the IP header of the packets and the local routing information maintained by the LSRs. An MPLS header is then inserted for each packet. Within an MPLS-capable domain, an LSR will use the label as the index to look up the forwarding table of the LSR. The packet is processed as specified by the forwarding table entry. The incoming label is replaced by the outgoing label and the packet is switched to

the next LSR. This label-switching process is very similar to ATM's VCI/VPI processing. Before a packet leaves a MPLS domain, its MPLS header is removed. This whole process is showed in Fig. 1. The paths between the ingress LSRs and the egress LSRs are called *Label Switched Paths* (LSPs). MPLS uses some signaling protocol like RSVP [2] or LDP [5] to set up LSPs.

In order to control the path of LSPs effectively, each LSP can be assigned one or more attributes. These attributes will be considered in computing the path for the LSP. Such attributes are summarized in Table 1.

Table 1. LSP Attributes

Attribute Name	Meaning of the attribute
Bandwidth	The minimum requirement on the reservable bandwidth of a path for the LSP to be set up along that path
Path attribute	An attribute that decides whether the path of the LSP should be manually specified or dynamically computed by Constraint-based Routing
Setup Priority	The attribute that decides which LSP will get the resource when multiple LSPs compete for it
Holding Priority	The attribute that decides whether an established LSP should be preempted the resource it holds by a new LSP
Affinity (color)	An administratively specified property of an LSP (see Section 2.2 for detail)
Adaptability	Whether to switch the LSP to a more optimal path when one becomes available
Resilience	The attribute that decides whether to reroute the LSP when the current path is affected by failure

1.2 Constraint-based Routing

Constraint-based Routing (CBR) computes routes that are subject to constraints such as bandwidth and administrative policy. Because Constraint-based Routing considers more than network topology in computing routes, it may find a longer but lightly loaded path better than the heavily loaded shortest path. Network traffic is hence distributed more evenly.

For example in Fig. 2, the shortest path between router A and router C is through link A-C with IGP metric $m=1$. But because the reservable bandwidth on the shortest path is only $(622-600)=22$ Mbps, when Constraint-based Routing tries to find a path for an LSP of 40 Mbps, it will select path A-B-C instead, because the shortest path does not meet the bandwidth constraint.




Figure 2. Constraint-based Routing

It should be noted that the reservable bandwidth of a link is equal to the maximum reservable bandwidth set by network administrators minus the total bandwidth reserved by LSPs traversing the link. It does not depend

on the actual amount of available bandwidth on that link. For example, if the maximum reservable bandwidth of a link is 155 Mbps, and the total bandwidth reserved by LSPs is 100 Mbps, then the reservable bandwidth of the link is 55 Mbps, regardless of whether the link is actually carrying 100 Mbps of traffic or more or less. In other words, Constraint-based Routing does not compute LSP paths based on instantaneous residual bandwidth of links. This reduces the probability of routing instability [6].

Constraint-based Routing can be online or offline. With online Constraint-based Routing, routers may compute paths for LSPs at any time. With offline Constraint-based Routing, an offline server computes paths for LSPs periodically (hourly/daily). LSPs are then configured to take the computed paths.

1.3 Enhanced Link State IGPs

In order for Constraint-based Routing to compute LSP paths subject to constraints, an enhanced link state IGP must be used to propagate link attributes in addition to normal link state information. Common link attributes include:

1. reservable bandwidth, and
2. link affinity (color), that is, an administratively specified property of the link.

Enhanced link state IGPs will flood information more frequently than normal IGP. This is because, even without topology changes and hence no normal IGP flooding, changes in reservable bandwidth or link color can trigger the enhanced IGP to flood information. Therefore, a tradeoff must be made between the need for accurate information and for avoiding excessive flooding. Change in reservable bandwidth is flooded only when it is significant. And a timer should be used to set an upper bound on flooding frequency.

When the enhanced IGP builds LSR's forwarding table, it will take into account LSPs originated by the LSR, so that the LSPs can actually be used to carry traffic. One of such approaches is described in [9]. In a network with MPLS, Constraint-based Routing and enhanced IGP, the process of forwarding table construction is showed in Fig. 3.




Figure 3. Forwarding Table Construction

1.4 Summary

With MPLS, Constraint-based Routing and an enhanced IGP, Traffic Engineering can be done much more effectively. The two problems discussed in the introduction can be solved.

First, by setting the maximum reservable bandwidth of each link (e.g., to the physical bandwidth of the link), and by setting the bandwidth requirement for each LSP, Constraint-based Routing will automatically avoid placing too many LSPs on any link. This solves the first problem. For example in Fig. 4, Constraint-based Routing will automatically place LSP B→E on a longer path to avoid congestion in link C→E.

Second, if traffic from router C1 to router B1 exceeds the capacity of any single path from C1 to B1 (Fig.




Figure 4. Congestion avoidance

5), then multiple LSPs can be configured from C1 to B1, and load ratio of these two LSPs can be specified as desired, so that load can be distributed optimally. This solves the second problem. For example in Fig. 5, if the total traffic from router C1 to router B1 is 160 Mbps, and the traffic from C2 to B1 is 90 Mbps, then two LSPs can be configured from C1 to B1. Their load ratio is automatically derived from their bandwidth specification. Here we clearly see that with MPLS, Constraint-based Routing, and enhanced IGP, not only can load sharing be done among multiple paths of different cost, but also load ratio can be specified as desired. With IGP shortest path routing and ECMPs, the problem described in the example cannot be solved.

In addition to the advantages mentioned above, MPLS also provides the followings:




Figure 5. Load sharing

1. Explicit routes (ERs) can be specified for LSPs. Network administrators can use ERs to control traffic trajectory precisely.
2. Per-LSP statistics can provide accurate end-to-end traffic matrix.

3. Backup LSPs can be used to provide graceful degradation in the case of router or link failure.

In particular, end-to-end traffic matrix makes network planning possible in an IP network without a dedicated connection-oriented network layer such as ATM or Frame Relay.

The organization of the rest of the paper is as follows. In Section 2, we discuss the general issues of designing an MPLS system for Traffic Engineering. The national MPLS system of GlobalCenter, a Global Crossing Company, is also presented. In Section 3, an approach is proposed for deploying MPLS system in an evolutionary way. Providing *Quality of Service* (QoS) in a network with MPLS is briefly discussed in Section 4. Finally, the paper is concluded in Section 5.

2. Designing an MPLS System in a National Network

Before the design issues of MPLS systems can be discussed, readers must have a conceptual understanding of the architecture of ISP networks.

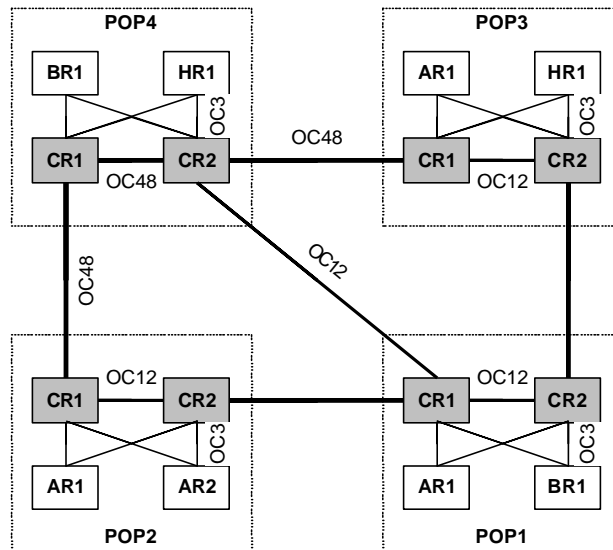


Figure 6. A sample part of an ISP network

2.1 An Overview of ISP Networks

ISP networks consists of *Points-of-presence* (POPs) and links interconnecting the POPs. A sample part of an ISP network is shown in Fig. 6. From a simplified perspective, a POP consists of a combination of one or more *access routers* (ARs) connected to customers, *border routers* (BRs) connected to other ISPs, *hosting routers* (HRs) connected to Web servers of companies such as Yahoo, and *core routers* (CRs) connected to other POPs. Packets from an AR or BR or HR to other POPs must first be sent to a CR. The CR will then relay those packets to the CRs in other POPs. The network architecture of POPs is usually highly symmetric, as showed in Fig. 6. Different POPs are normally connected in a ring topology to increase reliability. A large ISP may have more than 30 POPs.

2.2 Generic Issues of Designing an MPLS System for Traffic Engineering

To build an MPLS system for Traffic Engineering, the following design parameters must be determined:

1. the geographical scope of the MPLS system;
2. the participating routers;
3. the hierarchy of MPLS system;
4. the bandwidth requirement of the LSPs;
5. the path attribute of the LSPs;
6. the priority of the LSPs;
7. the number of parallel LSPs between each endpoint pair;
8. the affinity of the LSPs and the links;
9. the adaptability and resilience attributes of the LSPs.

The process of deciding the scope of an MPLS system is driven by administrative policy. Specifically, if the network architecture of a region is irregular (as opposed to the regular architecture showed in Fig. 6), or the capacity of a region is tight, then the region should be included in the MPLS system.

The second step is to decide the participating routers in the MPLS system, i.e., the ingress LSRs, the transit LSRs and the egress LSRs. This should also be guided by the administrative policy. Network administrators may want to forbid some routers from participating in the MPLS system for some reason, for example, because those routers cannot be trusted or because those routers do not have enough processing power and/or memory capacity. Another factor for consideration is the tradeoff between the number of LSPs and efficiency of the links. More ingress and egress LSRs mean more LSPs and thus higher LSP-routing complexity. But because the average size (bandwidth requirement) of the LSPs is smaller, Constraint-based Routing has more flexibility in routing the LSPs. Higher link efficiency may be achieved.

After the LSRs are decided, network administrators need to decide the hierarchy of the MPLS system. One alternative is to fully mesh all LSRs, resulting in a single layer of LSPs. For large ISPs, there can be hundreds of LSRs. A full mesh will result in a huge MPLS system. Another alternative is to divide one's network into multiple regions. LSRs in each region are meshed. This forms the first layer of LSPs. Some selected LSRs from each region, for example the core routers, are also fully meshed to form the second layer of the LSPs. This hierarchical design can significantly reduce the number of LSPs in the network, and hence the associated processing and managing overhead.

Unless an end-to-end traffic matrix is available beforehand, the bandwidth requirement of the LSPs is usually unknown and has to be guessed for the first time LSPs are deployed. Later, the measured rate of the LSPs can be used as the bandwidth requirement of the LSPs. This process is detailed in the next Section.

LSP paths can be manually specified or dynamically computed. Unless offline Constraint-based Routing is used to compute the paths, manually specifying paths for LSPs is difficult. Therefore, LSPs are usually dynamically computed by an online Constraint-based Routing algorithm in the routers.

Important LSPs, such as those carrying large amount of traffic, can be given a higher priority than other LSPs. In this way, these LSPs are more likely to take the optimal paths. This will result in higher routing stability and better resource utilization from a global perspective.

Multiple parallel LSPs can be configured between an ingress-egress pair. These LSPs can be placed on different physical paths, so that the traffic load from the source to the destination can be distributed more evenly. By using multiple parallel LSPs, the size of each LSP is also smaller. These LSPs can be routed more flexibly. These are the primary motivations for parallel LSPs. It is recommended that parallel LSPs be used to keep the size of each LSP below 25 Mbps.

Affinity, or color, can be assigned to LSPs and links to achieve some desired LSP placement. For example, if network administrators want to prevent a regional LSP from traversing routers or links outside the region, color can be used to achieve the goal. All regional links can be colored *green*, and all inter-region links can be colored *red*. Regional LSPs are constrained to take only *green* links. In this way, regional LSPs can never traverse any inter-region link. The process of assigning color to LSPs and links is again guided by administrative policy.

Depending on the stability of the network, when better paths become available, network administrators may or may not want to switch LSPs to the more optimal paths. The switching of LSPs to better paths is called *LSP reoptimization*. Reoptimization is not always desirable because it may introduce routing instability. In the case that reoptimization is allowed, it should not occur too frequently. Performing reoptimization once per hour may be a good choice. As to the resilience attribute, LSPs are generally allowed to be rerouted when failure occurs along their paths. In the cases of failure, it may even be desirable to reroute LSPs regardless of their bandwidth and other constraints.

2.3 The MPLS System in GlobalCenter's US National Network

GlobalCenter is one of the 10 largest ISPs in the US. GlobalCenter's US national network consists of approximately 50 POPs of more than 300 routers. An enhanced IS-IS is used as the IGP. All routers are in the same IS-IS level. All routers also run BGP.

In GlobalCenter, it is decided that MPLS will be deployed nation-wide, for both Traffic Engineering and *Virtual Private Networks* (VPNs). Approximately 200 routers with interfaces ranging from DS3 through OC-48c participate in the MPLS system. Fully meshing these 200 routers would result in an MPLS system of about 40,000 LSPs. Although it is not a problem for the network to support this many LSPs, it is decided that better performance and manageability can be achieved by deploying a hierarchical MPLS system of 2 layers of LSPs.

The nation-wide network is divided into 9 regions. All routers in each region are meshed to form the first layer of LSPs. Core routers in all regions are also fully meshed to form the second layer of LSPs. The number of LSRs in a region ranges from 10 to 50. Therefore, the number of LSPs in a region ranges from 100 to 2500. Intra-region traffic, ranging from several Mbps to several Gbps, is transported by these regional LSPs. The core LSP mesh consists of approximately 2500 LSPs. Inter-region traffic, in the magnitude of several Gbps, traverses first the regional LSPs in the source region, then the core LSPs, and last the regional LSPs in the destination region.

3. Deploying an MPLS System for Traffic Engineering

A procedure for deploying an MPLS system for Traffic Engineering is proposed. It is based on our experiences of deploying GlobalCenter’s MPLS system. It is recommended that a non-critical area be selected for deployment to gain operational experience before a critical area is selected.

3.1 Step 1: Statistics Collecting using MPLS LSPs

The first step is to deploy LSPs without specifying their bandwidth requirement. The purpose of this step is to collect traffic statistics in the interested area, particularly the traffic amount between each pair of routers.

The reason why LSPs are used to collect statistics is illustrated below. Current statistics collectors can only report the inbound and outbound traffic of a network interface. It cannot report the final destination of the traffic. For example in Fig. 7, even if it is known that the traffic from router A to router B and from router B to router C are both 1 Mbps, it is still unknown how much traffic from A is destined to B and how much is destined to C. Therefore, the end-to-end traffic matrix has to be guessed. This makes Traffic Engineering difficult and less effective.



Figure 7. Statistics Collecting

However, if two LSPs are configured from A to B and from A to C, then the traffic from A to C will not enter the LSP from A to B. Therefore, the statistics of these two LSPs will tell network administrators precisely how much traffic from A is destined to B and how much is destined to C.

3.2 Step 2: Deploy LSPs with Bandwidth Constraint

The second step is to deploy LSPs with their bandwidth requirement specified. Usually, the measured rate of an LSP is used as the bandwidth requirement of that LSP. Because the measured rate of an LSP can vary sig-

nificantly at different time, it must be decided which measured rate to use. One common choice is the 95-percentile of all measured rates over a period of time, e.g., a week. This value is usually close to the real peak value as opposed to a traffic spike.

Online Constraint-based Routing will compute paths for the LSPs such that for every link, the maximum reservable bandwidth of the link is greater than or equal to the sum of the specified bandwidth of all LSPs traversing the link. Under this premise, whether high utilization will occur in a link depends on how close the total actual traffic rate of all the LSPs matches the total specified bandwidth. If the total rate is greater than or equal to the total specified bandwidth, high utilization may occur in the link.

There are two approaches to avoid the above problem. The first approach is to under-subscribe the links. For example, the maximum reservable bandwidth of an OC3 link that tends to be highly utilized can be set to 60% of the link capacity, i.e. 100Mbps. This approach is useful for controlling the utilization of a specific link. The second approach is to inflate the bandwidth requirement of the LSPs by a factor. For example, if the desired utilization of all links is 60%, then a factor of $1/60\%=1.67$ can be used. More specifically, the required bandwidth of every LSP is set to the 95-percentile of the measured rate multiplied by the factor. The philosophy of this approach is to reserve some headroom for each LSP to grow. It is useful for controlling the utilization of all links.

With either approach, a tradeoff must be made between the need to avoid congestion and the need for link efficiency. Under-subscription or a large inflation factor can both cause some LSPs to take sub-optimal paths, even though the optimal paths have enough physical bandwidth (but not enough reservable bandwidth) for these LSPs. This reduces efficiency of the links.

No matter what has been done, it is still possible that the utilization of a specific link become too high. In that case, network administrators can either manually move some LSPs away from the congested link by using explicit routes, or lower the reservable bandwidth of the congested link to force some LSPs to reroute, or add parallel LSPs along some lightly loaded alternative paths to offload some traffic from the congested link. Load sharing ratio of the parallel LSPs is determined by their specified bandwidth. In order do these effectively, tools for showing the paths of the LSPs, the reservable bandwidth of the links, and for suggesting the alternative paths are very useful.

A simulation tool like the WANDL simulator [13] can be used before the actual deployment to make sure that the Constraint-based Routing algorithm can find paths for all the LSPs. Otherwise, more capacity has to be added to the network. If the simulator shows that some link will have very high utilization, the two approaches described above can be used to solve the problem.

3.3 Step 3: Periodic Update of LSP Bandwidth

The third step is to further collect statistics and adjust the bandwidth of the LSPs if necessary. This is needed because network traffic is growing and traffic distribution is changing all the time. This step should be done periodically, e.g. daily or weekly, by some scripts (as *cron* jobs). It is wise not to adjust the bandwidth requirement of all LSPs at the same time. Otherwise, significant change may be observed during the bandwidth adjustment process.

3.4 Step 4. Offline Constraint-based Routing

If optimal link efficiency is desired, offline Constraint-based Routing can be used to compute the paths for the LSPs. By taking all the LSP requirements, link attributes and network topology information into consideration, an offline Constraint-based Routing server may be able to find better LSP placement than online Constraint-based Routing, where every router in the network finds paths for its LSPs separately based on its own information. Therefore, links can be used more efficiently. Offline Constraint-based Routing will compute paths for LSPs periodically, e.g. daily. Then the path information of the LSPs is downloaded into the routers. Online Constraint-based Routing is still used in the routers, so that if some routers or links fail, new paths will be computed for those affected LSPs immediately. A simple offline Constraint-based Routing algorithm is described below.

Compute paths for LSPs one by one, starting from high priority LSPs. For LSPs with the same priority, compute paths for them in the order of decreasing bandwidth requirement. The goal of this algorithm is to optimize the overall *bandwidth-routing metric* product. That is, the goal is to let the largest LSP take the best possible path, the second largest LSP take the next best possible path, and so on, inside each priority class. The algorithm is briefly described below.

- 1) Sort the LSPs in decreasing order of importance as described above;
- 2) For a particular LSP, first prune all the unusable links;

A link can be unusable for an LSP because of some reasons such as:

- the reservable bandwidth of the link is not sufficient for the LSP or the delay of the link is too high (e.g., satellite links);
- the link is administratively forbidden for the LSP, e.g., *red* links cannot be used for a *green* LSP.

- 3) On the remaining graph, compute the optimal path for the LSP;
- 4) For those used links used by the LSP, deduct the resources (e.g., link bandwidth) used by the LSP;
- 5) Repeat steps 2-4 for the next LSP until all are done.

If backup LSPs need to be computed, repeat the above procedure on the remaining graph, that is, with resources used by the primary LSPs deducted. The only difference is in step 2. Now all links and routers used by

the corresponding primary LSP of this backup LSP are also pruned. This is to avoid any single point of failure for both the primary and the backup LSPs.

This algorithm may not find the globally optimal layout for LSPs. But it is simple. The problem of optimizing the *bandwidth-routing metric* product for all LSPs is *NP-complete*, because the BIN-PACKING [11] problem can be *reduced* to it. Therefore, finding the optimal solution is not practical except for small networks.

4. Quality of Service (QoS) in a Network with MPLS

With *Differentiated Services* (Diffserv), packets are classified at the edge of the network. The *Differentiated Services-fields* (DS-fields) [12] of the packets are set accordingly. In the middle of the network, packets are buffered and scheduled in accordance to their DS-fields by *Weighted Random Early Detection* (WRED) and *Weighted Round Robin* (WRR). Important traffic such as network control traffic and traffic from premium customers will be forwarded preferably [8].

With MPLS, QoS is provided in a slightly different way. Packets still have their DS-fields set at the edge of the network. In addition, the experimental fields in the MPLS headers are set at the ingress LSRs. In the middle of an LSP, packets are buffered and scheduled in accordance to the experimental fields. Whether MPLS is involved or not in providing QoS is transparent to end-users.

Sometimes it is desirable to use different LSPs for different classes of traffic. The effect is that the physical network is divided into multiple virtual networks, one per class. These virtual networks may have different topology and resources. The end effect is that premium traffic can use more resources than best effort traffic. Premium traffic will also have higher priority in getting the backup resources in the case of link or router failure. This process is largely automatic, driven by Constraint-based Routing.

It is envisioned that LSPs for VPN will be built between ARs, BRs, HRs across the network in the near future. Such LSPs will be built on the top of the current LSPs that are used for Traffic Engineering. In other words, a current LSP will be treated as a link in building the LSPs for VPN. Only the endpoints of current LSPs will be involved in the signaling process of building new LSPs for VPN. LSPs are therefore stacked. Such a scheme can significantly reduce state information in the core.

5. Conclusion

In this paper, we first discuss the general issues of designing an MPLS system for Traffic Engineering. We then present the design of GlobalCenter's MPLS system. Based on our experiences, a generic procedure for deploying MPLS systems is proposed.

Our experience shows that deploying an MPLS system for Traffic Engineering in a large ISP network is feasible. MPLS, Constraint-based Routing and an enhanced IGP have to be deployed in every router that par-

ticipates in the MPLS system. Because everything is done by software upgrade, such a wide-range deployment turns out not to be as difficult as many people might think.

Our experience shows that a hierarchical MPLS system consisting of a core LSP mesh and several regional LSP meshes is a good design. By first introducing LSPs to collect statistics, then deploying LSPs with bandwidth requirement, and updating the bandwidth requirement of the LSPs periodically, an MPLS system can be deployed in an evolutionary fashion in an ISP network.

An MPLS system is very useful for Traffic Engineering. It automatically achieves desired traffic behavior and significantly relieves network administrators from the tedious work of manipulating routing metrics. It is especially useful in regions of irregular network topology or tight capacity.

Many new tools are needed for managing an MPLS system. Among them are:

- simulators like WANDL for ensuring that the design of an MPLS system is technically sound before actual deployment;
- scripts for automatically configuring the routers and their interfaces;
- new SNMP *management information bases* (MIBs) and applications for collecting statistics, monitoring the condition of the MPLS system, accounting and billing;
- tools for auditing the system and reporting potential problems;

6. Acknowledgement

The authors would like to thank Y. Rekhter of Cisco, D. Gan of Juniper, J. Renwick of Ironbridge and J. Rexford of AT&T Research for their comments.

7. Reference

1. D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell and J. McManus, "Requirements for Traffic Engineering over MPLS", RFC 2702, Sept. 1999.
2. D. Awduche, L. Berger, D. Gan, T. Li, G. Swallow and V. Srinivasan, "Extension to RSVP for Traffic Engineering", Internet draft <draft-ietf-mpls-rsvp-lsp-tunnel-04.txt>, Sept. 1999.
3. D. Awduche, A. Hannan and X. Xiao, "Applicability Statement for Extensions to RSVP for LSP-LSPs," Internet Draft <draft-ietf-mpls-rsvp-tunnel-applicability-00.txt>, Sept. 1999.
4. E. Rosen, A. Viswanathan and R. Callon, "Multiprotocol Label Switching Architecture", Internet draft <draft-ietf-mpls-arch-06.txt>, Aug. 1999.
5. L. Andersson, P. Doolan, N. Feldman, A. Fredette and B. Thomas, "LDP Specification", Internet draft <draft-ietf-mpls-ldp-06.txt>, Oct. 1999.
6. E. Crawley, R. Nair, B. Jajagopalan and H. Sandick, "A Framework for QoS-based Routing in the Internet", RFC 2386, Aug. 1998
7. T. Li, G. Swallow and D. Awduche, "IGP requirements for Traffic Engineering with MPLS", Internet draft <draft-li-mpls-igp-te-00.txt>, Feb. 1999.

8. X. Xiao and L. Ni, "[Internet QoS: A Big Picture](#)", IEEE Network Magazine, March/April, pp. 8-18, 1999.
9. N. Shen and H. Smit, "Calculating IGP routes over Traffic Engineering LSPs", Internet draft <draft-hsmit-mpls-igp-spf-00.txt>, June 1999.
10. E. Rosen and Y. Rekhter, "BGP/MPLS VPN", RFC 2547, March 1999.
11. M. Garey and D. Johnson, "Computer and Intractability", W. H. Freeman and Company, 1979.
12. K. Nichols, S. Blake, F. Baker and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, Dec. 1998.
13. Wide Area Network Design Laboratory (WANDL), <http://www.wandl.com>

8. Author Information

Xipeng Xiao is a Senior Network Engineer and manager of the Advanced Technology group, GlobalCenter Inc, a Global Crossing company. He is the primary builder of GlobalCenter's MPLS system. Before joining GlobalCenter, he worked for the Multi-Gigabit Routing Division of Ascend Communication on QoS and MPLS. Xipeng is a Ph.D. candidate of Michigan State University. He won the First Prize in the ACM Graduate Research Contest '1998 with the project "The design of a scalable router".

Alan Hannan is the Senior Director of Network Architecture & Engineering, GlobalCenter. He is the principal designer of GlobalCenter's MPLS system.

Brook Bailey is a Senior Network Engineer of GlobalCenter.

Lionel M. Ni is a Fellow of IEEE and Distinguished Professor of Michigan State Univ. His research interests include parallel processing, computer architecture, IP telephony and networking.