Network Working Group
Internet Draft
Expiration Date: December 2001

K. Kompella (Juniper Networks)
Y. Rekhter (Juniper Networks)
A. Banerjee (Calient Networks)
J. Drake (Calient Networks)
G. Bernstein (Ciena)
D. Fedyk (Nortel Networks)
E. Mannie (GTS Network)
D. Saha (Tellium)
V. Sharma (Tellabs)
D. Basak (AcceLight Networks)

Routing Extensions in Support of Generalized MPLS

draft-many-ccamp-gmpls-routing-00.txt

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

2. Abstract

This document specifies routing extensions in support of Generalized Multi-Protocol Label Switching (GMPLS).

3. Summary for Sub-IP Area

3.1. Summary

This document specifies routing extensions in support of Generalized Multi-Protocol Label Switching (GMPLS).

3.2. Where does it fit in the Picture of the Sub-IP Work

This work fits squarely in the CCAMP box.

3.3. Why is it Targeted at this WG

This draft is targeted at the CCAMP WG, because this draft specifies the extensions to the link state routing protocols in support of GMPLS, and because GMPLS is within the scope of CCAMP WG.

3.4. Justification

The WG should consider this document as it specifies the extensions to the link state routing protocols in support of GMPLS.

4. Introduction

This document specifies routing extensions in support of carrying link state information for Generalized Multi-Protocol Label Switching (GMPLS). This document enhances the routing extensions [3] required to support MPLS Traffic Engineering.

5. GMPLS TE Links

Traditionally, a TE link is advertised as an adjunct to a "regular" link, i.e., a routing adjacency is brought up on the link, and when the link is up, both the regular SPF properties of the link (basically, the SPF metric) and the TE properties of the link are then advertised.

GMPLS challenges this notion in three ways. First, links that are not capable of sending and receiving on a packet-by-packet basis may yet have TE properties; however, a routing adjacency cannot be brought up on such links. Second, a Label Switched Path can be advertised as a point-to-point TE link (see [LSP-HIER]); thus, an advertised TE link may be between a pair of nodes that don't have a routing adjacency with each other. Finally, a number of links may be advertised as a single TE link (perhaps for improved scalability), so again, there is no longer a one-to-one association of a regular routing adjacency and a TE link.

Thus we have a more general notion of a TE link. A TE link is a "logical" link that has TE properties. The link is logical in a sense that it represents a way to map the information about certain physical resources (and their properties) into the information that is used by Constrained SPF for the purpose of path computation. Some of the properties of a TE link may be configured on the advertising Label Switching Router (LSR), others which may be obtained from other LSRs by means of some protocol, and yet others which may be deduced from the component(s) of the TE link.

A TE link between a pair of LSRs doesn't imply the existence of a routing adjacency between these LSRs.

A TE link must have some means by which the advertising LSR can know of its liveness (this means may be routing hellos, but is not limited to routing hellos). When an LSR knows that a TE link is up, and can determine the TE link's TE properties, the LSR may then advertise that link to its (regular) neighbors.

In this document, we call the interfaces over which regular routing adjacencies are established "control channels".

[ISIS-TE] and [OSFP-TE] defines the canonical TE properties, and says how to associate TE properties to regular (packet-switched) links. This document extends the set of TE properties, and also says how to associate TE properties with non-packet-switched links such as links between Optical Cross-Connects (OXCs). [LSP-HIER] says how to associate TE properties with links formed by Label Switched Paths; [LINK-BUNLDE] says how to associate TE properties with a "bundle" of

links.

5.1. Excluding data traffic from control channels

The control channels between nodes in a GMPLS network, such as OXCs, SONET cross-connects and/or routers, are generally meant for control and administrative traffic. These control channels are advertised into routing as normal links as mentioned in the previous section; this allows the routing of (for example) RSVP messages and telnet sessions. However, if routers on the edge of the optical domain attempt to forward data traffic over these channels, the channel capacity will quickly be exhausted.

In order to keep these control channels from being advertised into the user data plane a variety of techniques can be used.

If one assumes that data traffic is sent to BGP destinations, and control traffic to IGP destinations, then one can exclude data traffic from the control plane by restricting BGP nexthop resolution. (It is assumed that OXCs are not BGP speakers.) Suppose that a router R is attempting to install a route to a BGP destination D. R looks up the BGP nexthop for D in its IGP's routing table. Say R finds that the path to the nexthop is over interface I. R then checks if it has an entry in its Link State database associated with the interface I. If it does, and the link is not packet-switch capable (see [LSP_HIER]), R installs a discard route for destination D. Otherwise, R installs (as usual) a route for destination D with nexthop I. Note that R need only do this check if it has packet-switch incapable links; if all of its links are packet-switch capable, then clearly this check is redundant.

In other instances it may be desirable to keep the whole address space of a GMPLS routing plane disjoint from the endpoint addresses in another portion of the GMPLS network. For example, the addresses of a carrier network where the carrier uses GMPLS but does not wish to expose the internals of the addressing or topology. In such a network the control channels are never advertised into the end data network. In this instance, independent mechanisms are used to advertise the data addresses over the carrier network. The Optical VPNs architecture [OVPN] discusses a mechanism for automating the distribution of independent addresses.

Other techniques for excluding data traffic from control channels may also be needed.

6. GMPLS Routing Enhancements

In this section we define the enhancements to the TE properties of GMPLS TE links. Encoding of this information in IS-IS is specified in [ISIS-GMPLS]. Encoding of this information in OSPF is specified in [OSPF-GMPLS].

6.1. Support for unnumbered interfaces

Supporting unnumbered interfaces includes carrying the information about the identity of the interfaces.

6.1.1. Outgoing Interface Identifier

A link from LSR A to LSR B may be assigned an "outgoing interface identifier". This identifier is a non-zero 32-bit number that is assigned by LSR A. This identifier must be unique within the scope of A.

6.1.2. Incoming Interface Identifier

Suppose there is a link L from A to B. Suppose further that the link L' from B to A that corresponds to the same interface as L has been assigned an outgoing interface identifier by B. The "incoming interface identifier" for L (from A's point of view) is defined as the outgoing interface identifier for L' (from B's point of view).

If no such L' exists (e.g., the interface is unidirectional), A MUST NOT advertise an Incoming Interface Identifier. If A knows that such an L' exists, but does not know the outgoing interface identifier assigned to L' by B, A MAY include the Incoming Interface Identifier with a value of 0.

6.2. Link Protection Type

The Link Protection Type represents the protection capability that exists for a link. It is desirable to carry this information so that it may be used by the path computation algorithm to set up LSPs with appropriate protection characteristics. This information is organized in a hierarchy where typically the minimum acceptable protection is specified at path instantiation and a path selection technique is used to find a path that satisfies at least the minimum acceptable protection. Protection schemes are presented in order from lowest to highest protection.

This document defines the following protection capabilities:

Extra Traffic
If the link is of type Extra Traffic, it means that the link is protecting another link or links. The LSPs on a link of this type will be lost if any of the links it is protecting fail.

Unprotected
If the link is of type Unprotected, it means that there is no other link protecting this link. The LSPs on a link of this type will be lost if the link fails.

Shared
If the link is of type Shared, it means that there are one or more disjoint links of type Extra Traffic that are protecting this link. These Extra Traffic links are shared between one or more links of type Shared.

Dedicated 1:1
If the link is of type Dedicated 1:1, it means that there is one dedicated disjoint link of type Extra Traffic that is protecting this link.

Dedicated 1+1
If the link is of type Dedicated 1+1, it means that a dedicated disjoint link is protecting this link. However, the protecting link is not advertised in the link state database and is therefore not available for the routing of LSPs.

Enhanced
If the link is of type Enhanced, it means that a protection scheme that is more reliable than Dedicated 1+1, e.g., 4 fiber BLSR/MS-SPRING, is being used to protect this link.

The Link Protection Type is optional, and if an LSA doesn't carry this information, then the Link Protection Type is unknown.

6.3. Shared Risk Link Group Information

A set of links may constitute a 'shared risk link group' (SRLG) if they share a resource whose failure may affect all links in the set. For example, two fibers in the same conduit would be in the same SRLG. A link may belong to multiple SRLGs. Thus the SRLG Information describes a list of SRLGs that the link belongs to. An SRLG is identified by a 32 bit number that is unique within an IGP domain. The SRLG Information is an unordered list of SRLGs that the link belongs to.

The SRLG of a LSP is the union of the SRLGs of the links in the LSP. The SRLG of a bundled link is the union of the SRLGs of all the component links.

If an LSR is required to have multiple diversely routed LSPs to another LSR, the path computation should attempt to route the paths so that they do not have any links in common, and such that the path SRLGs are disjoint.

The SRLG Information starts with a configured value and does not change over time, unless manually reconfigured. The SRLG Information is optional and if an LSA doesn't carry the SRLG Information, then it means that SRLG of that link is unknown.

6.4. Interface Switching Capability Descriptor

In the context of this document we say that a link is connected to a node by an interface. In the context of GMPLS interfaces may have different switching capabilities. For example an interface that connects a given link to a node may not be able to switch individual packets, but it may be able to switch channels within a SONET payload. Interfaces at each end of a link need not have the same switching capabilities. Interfaces on the same node need not have the same switching capabilities.

The Interface Switching Capability Descriptor describes switching capability of an interface. The switching capabilities of an interface are defined to be the same in either direction. I.e., for data entering the node through that interface and for data leaving the node through that interface.

For a bidirectional link, each LSA carries just the Interface Switching Capability Descriptor of the interface that connects the link to the LSR that originates the LSA. For a unidirectional link each LSA should carry the Interface Switching Capability Descriptor for interfaces at both end of the link.

This document defines the following Interface Switching Capabilities:

```
Packet-Switch Capable-1 (PSC-1)
Packet-Switch Capable-2 (PSC-2)
Packet-Switch Capable-3 (PSC-3)
Packet-Switch Capable-4 (PSC-4)
Layer-2 Switch Capable  (L2SC)
Time-Division-Multiplex Capable (TDM)
Lambda-Switch Capable   (LSC)
Fiber-Switch Capable    (FSC)
```

If there is no Interface Switching Capability Descriptor for an interface, the interface is assumed to be packet-switch capable (PSC-1).

Interface Switching Capability Descriptors present a new constraint for LSP path computation.

Irrespective of a particular Interface Switching Capability, the Interface Switching Capability Descriptor always includes information about the encoding supported by an interface. The defined encodings are the same as LSP Encoding as defined in [GMPLS-SIG].

Depending on a particular Interface Switching Capability, the Interface Switching Capability Descriptor may include additional information, as specified below.

6.4.1. Layer-2 Switch Capable

If an interface is of type L2SC, it means that the node receiving data over this interface can switch the received frames based on the layer 2 address. For example, an interface associated with a link terminating on an ATM switch would be considered L2SC.

6.4.2. Packet-Switch Capable

If an interface is of type PSC-1 through PSC-4, it means that the node receiving data over this interface can switch the received data on a packet-by-packet basis. The various levels of PSC establish a hierarchy of LSPs tunneled within LSPs.

For Packet-Switch Capable interfaces the additional information includes Maximum LSP Bandwidth.

For a simple (unbundled) link its Maximum LSP Bandwidth at priority p is defined to be the smaller of its unreserved bandwidth at priority p and its Maximum Reservable Bandwidth.

The Maximum LSP Bandwidth of a bundled link at priority p is defined to be the maximum of the Maximum LSP Bandwidth at priority p of each component link.

The Maximum LSP Bandwidth takes the place of the Maximum Bandwidth ([ISIS-TE], [OSPF-TE]). However, while Maximum Bandwidth is a single fixed value (usually simply the link capacity), Maximum LSP Bandwidth is carried per priority, and may vary as LSPs are set up and torn down.

Although Maximum Bandwidth is to be deprecated, for backward compatibility, one MAY set the Maximum Bandwidth to the Maximum LSP Bandwidth at priority 7.

6.4.3. Time-Division Multiplex Capable

If an interface is of type TDM, it means that the node receiving data over this interface can multiplex or demultiplex channels within a SONET/SDH payload.

For Time-Division Multiplex Capable interfaces the additional information includes Maximum LSP Bandwidth, the information on whether the interface supports Standard or Arbitrary SONET/SDH, and Minimum LSP Bandwidth.

For a simple (unbundled) link the Maximum LSP Bandwidth at priority p is defined as the maximum bandwidth an LSP at priority p could reserve.

The Maximum LSP Bandwidth of a bundled link at priority p is defined to be the maximum of the Maximum LSP Bandwidth at priority p of each component link.

The Minimum LSP Bandwidth specifies the minimum bandwith an LSP could reserve.

Typical values for the Minimum LSP Bandwidth and for the Maximum LSP Bandwidth are enumerated in [GMPLS-SIG].

On an interface having Standard SONET (or Standard SDH) multiplexing, an LSP at priority p could reserve any bandwidth allowed by the branch of the SONET/SDH hierarchy, with the leaf and the root of the branch being defined by the Minimum LSP Bandwidth and the Maximum LSP Bandwidth at priority p.

On an interface having Arbitrary SONET (or Arbitrary SDH) multiplexing, an LSP at priority p could reserve any bandwidth between the Minimum LSP Bandwidth and the Maximum LSP Bandwidth at priority p, provided that the bandwidth reserved by the LSP is a multiple of the Minimum LSP Bandwidth.

To handle the case where an interface supports multiple branches of the SONET (or SDH) multiplexing hierarchy, multiple Interface Switching Capability Descriptors would be advertised, one per branch. For example, if an interface supports VT-1.5 and VT-2 (which are not part of same branch of SONET multiplexing tree), then it could advertise two descriptors, one for each one.

6.4.4. Lambda-Switch Capable

If an interface is of type LSC, it means that the node receiving data over this interface can recognize and switch individual lambdas within the interface. An interface that allows only one lambda per interface, and switches just that lambda is of type LSC.

The additional information includes Reservable Bandwidth and priority, which specifies the bandwidth of an LSP that could be supported by the interface, and the (numerically) largest priority number at which the bandwidth could be reserved. Note that the priority needs to be present only when an interface has more than one Interface Switching Capability Descriptor with LSC as the Interface Switching Capability.

To handle the case of multiple data rates or multiple encodings within a single TE Link, multiple Interface Switching Capability Descriptors would be advertised, one per supported data rate and encoding combination. For example, an LSC interface could support the establishment of LSC LSPs at both OC-48c and OC-192c data rates.

6.4.5. Fiber-Switch Capable

If an interface is of type FSC, it means that the node receiving data over this interface can switch the entire contents to another interface (without distinguishing lambdas, channels or packets). I.e., an interface of type FSC switches at the granularity of an entire interface, and can not extract individual lambdas within the interface. An interface of type FSC can not restrict itself to just one lambda.

6.4.6. Multiple Switching Capabilities per interface

An interface that connects a link to an LSR may support not one, but several Interface Switching Capabilities. For example, consider a fiber link carrying a set of lambdas that terminates on an LSR interface that could either cross-connect one of these lambdas to some other outgoing optical channel, or could terminate the lamdba, and extract (demultiplex) data from that lambda using TDM, and then cross-connect these TDM channels to some outgoing TDM channels. To support this an LSA may carry a list of Interface Switching Capabilities Descriptors.

6.4.7. Examples of Interface Switching Capability Descriptor

6.4.7.1. STS-48 POS Interface on a LSR

```
Interface Switching Capability Descriptor:
   Interface Switching Capability = PSC-1
   Encoding = SONET ANSI T1.105-1995
   Max LSP Bandwidth[p] = 2.5 Gbps, for all p
```

If multiple links with such interfaces were to be advertised as one TE link, link bundling techniques should be used.

6.4.7.2. GigE Packet Interface on a LSR

```
Interface Switching Capability Descriptor:
   Interface Switching Capability = PSC-1
   Encoding = Ethernet 802.3
   Max LSP Bandwidth[p] = 1.0 Gbps, for all p
```

If multiple links with such interfaces were to be advertised as one TE link, link bundling techniques should be used.

6.4.7.3. OC-192 SONET Interface on a Digital Cross Connect with Standard SONET

Consider a branch of SONET multiplexing tree : VT-1.5, STS-1, STS-3c, STS-12c, STS-48c, STS-192c. If it is possible to establish all these connections on a OC-192 interface, the Interface Multiplexing Capability Descriptor of that interface can be advertised as follows:

```
Interface Switching Capability Descriptor:
   Interface Switching Capability = TDM [Standard SONET]
   Encoding = SONET ANSI T1.105-1995
   Min LSP Bandwidth = VT1.5
   Max LSP Bandwidth[p] = STS192, for all p
```

If multiple links with such interfaces were to be advertised as one TE link, link bundling techniques should be used.

6.4.7.4. OC-192 SONET Interface on a Digital Cross Connect with two types of SONET multiplexing hierarchy supported

```
Interface Switching Capability Descriptor 1:
   Interface Switching Capability = TDM [Standard SONET]
   Encoding = SONET ANSI T1.105-1995
   Min LSP Bandwidth = VT1.5
   Max LSP Bandwidth[p] = STS192, for all p

Interface Switching Capability Descriptor 2:
   Interface Switching Capability = TDM [Arbitrary SONET]
   Encoding = SONET ANSI T1.105-1995
   Min LSP Bandwidth = VT2
   Max LSP Bandwidth[p] = STS192, for all p
```

If multiple links with such interfaces were to be advertised as one TE link, link bundling techniques should be used.

6.4.7.5. Interface on a transparent OXC (PXC) with external DWDM, that understands SONET framing

From a GMPLS perspective a combination of PXC and external DWMD is treated as a single unit.

```
                     _______
                    |       |
              /|___|       |
             | |___|  PXC  |
      =======X| |___|       |
             | |___|       |
              \|   |_______|
            DWDM
        (SONET framed)
```

The interface at X is advertised as:

```
Interface Switching Capability Descriptor:
   Interface Switching Capability = LSC
   Encoding = SONET ANSI T1.105-1995 (comes from DWDM)
   Reservable Bandwidth = Determined by DWDM (say OC192)
```

If multiple links with such interfaces were to be advertised as one TE link, one way to do this is to use link bundling. Another way to advertise multiple links with such interfaces as one TE link is just to require that all ports on the PXCs have identifiers unique to the PXC (as each interface identifier would act as a label).

6.4.7.6. Interface on an opaque OXC (SONET framed)

An "opaque OXC" is considered operationally an OXC, as the whole lambda (carrying the SONET line) is switched transparently, that is either none of the SONET overhead bytes are changed or at least the important ones are not changed.

An interface on an opaque OXC handles a single wavelength on the fiber. It fits the definition of LSC and not FSC as it cannot switch an interface with multiple wavelengths as a whole. Thus, an interface on an opaque OXC is always LSC, irrespective of whether there is DWDM external to it. Note, if there is external DWDM then the framing understood by the DWDM must be same as that understood by the OXC.

The following is an example of a interface switching capability decriptor on a SONET framed opaque OXC:

```
Interface Switching Capability Descriptor:
   Interface Switching Capability = LSC
   Encoding = SONET ANSI T1.105-1995
   Reservable Bandwidth = Determined by SONET Framer (say OC192)
```

If multiple links with such interfaces were to be advertised as one TE link, one way to do this is to use link bundling. Another way to advertise multiple links with such interfaces as one TE link is just to require that all interfaces on the OXCs have identifiers unique to the OXC (as each interface identifier would act as a label).

6.4.7.7. Interface on a PXC with no external DWDM

The absence of DWDM in between two PXCs, implies that an interface is not limited to one wavelength. Thus, the interface is advertised as FSC.

```
Interface Switching Capability Descriptor
   Interface Switching Capability = FSC
   Encoding = Photonic
   Reservable Bandwidth = Determined by optical technology limits
```

Note that this example assumes that the PXC does not restrict each port to carry only one wavelength.

6.4.8. Example of interfaces that support multiple switching capabilities

There can be many combinations possible, some are described below.

6.4.8.1. Interface on a PXC+TDM device with external DWDM

As discussed earlier, the presence of the external DWDM limits that only one wavelength be on a port of the PXC. On such a port, the attached PXC+TDM device can do one of the following. The wavelength may be cross-connected by the PXC element to other out-bound optical channel, or the wavelength may be terminated as a SONET interface and SONET channels switched.

From a GMPLS perspective DWDM and the PXC+TDM functionality is treated as a single unit. The interface is described using two Interface descriptors, one for the LSC and another for the TDM, with appropriate parameters. For example,

```
Interface Switching Capability Descriptor:
   Interface Switching Capability = LSC
   Encoding = SONET ANSI T1.105-1995 (comes from WDM)
   Reservable Bandwidth = OC192

   and

Interface Switching Capability Descriptor:
   Interface Switching Capability = TDM [Standard SONET]
   Encoding = SONET ANSI T1.105-1995
   Min LSP Bandwidth = VT1.5
   Max LSP Bandwidth[p] = STS192, for all p
```

6.4.8.2. Interface on an opaque OXC+LSR device with external DWDM

An interface on an "opaque OXC+TDM" device would also be advertised as LSC+TDM much the same way as the previous case.

6.4.8.3. Interface on a PXC+LSR device with external DWDM

As discussed earlier, the presence of the external DWDM limits that only one wavelength be on a port of the PXC. On such a port, the attached PXC+LSR device can do one of the following. The wavelength may be cross-connected by the PXC element to other out-bound optical channel, or the wavelength may be terminated as a Packet interface and packets switched.

From a GMPLS perspective DWDM and the PXC+LSR functionality is treated as a single unit. The interface is described using two Interface descriptors, one for the LSC and another for the PSC, with appropriate parameters. For example,

```
       Interface Switching Capability Descriptor:
          Interface Switching Capability = LSC
          Encoding = SONET ANSI T1.105-1995 (comes from WDM)
          Reservable Bandwidth = OC192

         and

       Interface Switching Capability Descriptor:
          Interface Switching Capability = PSC-1
          Encoding = SONET ANSI T1.105-1995
          Max LSP Bandwidth[p] = 10 Gbps, for all p
```

6.4.8.4. Interface on a TDM+LSR device

On a TDM+LSR device that offers a channelized SONET/SDH interface the following may be possible:

- A subset of the SONET/SDH channels may be uncommitted. That is, they are not currently in use and hence are available for allocation.

- A second subset of channels may already be committed for transit purposes. That is, they are already cross-connected by the SONET/SDH cross connect function to other out-bound channels and thus are not immediately available for allocation.

- Another subset of channels could be in use as terminal channels. That is, they are already allocated by terminate on a packet interface and packets switched.

The interface is advertised as TDM+PSC with example descriptors as:

```
       Interface Switching Capability Descriptor:
          Interface Switching Capability = TDM [Standard SONET]
          Encoding = SONET ANSI T1.105-1995
          Min LSP Bandwidth = VT1.5
          Max LSP Bandwidth[p] = STS192, for all p

         and

       Interface Switching Capability Descriptor:
```

Interface Switching Capability = PSC-1
Encoding = SONET ANSI T1.105-1995
Max LSP Bandwidth[p] = 10 Gbps, for all p

6.4.9. Other issues

It is possible that Interface Switching Capability Descriptor will change over time, reflecting the allocation/deallocation of LSPs. In general, creation/deletion of an LSP on a link doesn't necessarily result in changing the Interface Switching Capability Descriptor of that interface. For example, assume that STS-1, STS-3c, STS-12c, STS-48c and STS-192c LSPs can be established on a OC-192 interface whose Encoding Type is SONET (or to be more precise, SONET ANSI T1.105-1995). Thus, initially in the Interface Multiplexing Capability Descriptor the Minimum LSP Bandwidth is set to STS-1, and Maximum LSP Bandwidth is set to STS-192 for all priorities. As soon as an LSP of STS-1 size at priority 1 is established on the interface, it is no longer capable of STS-192c for all but LSPs at priority 0. Therefore, the node advertises a modified Interface Switching Capability Descriptor indicating that the Maximum LSP Bandwidth is no longer STS-192, but STS-48 for all but priority 0 (at priority 0 the Maximum LSP Bandwidth is still STS-192). If subsequently there is another STS-1 LSP, there is no change in the Interface Switching Capability Descriptor. The Descriptor remains the same until the node can no longer establish a STS-48c LSP over the interface (which means that at this point more than 144 time slots are taken by LSPs on the interface). Once this happened, the Descriptor is modified again, and the modified Descriptor is advertised to other nodes.

7. Security Considerations

The routing extensions proposed in this document do not raise any new security concerns.

8. Acknowledgements

The authors would like to thank Suresh Katukam, Jonathan Lang and Quaizar Vohra for their comments on the draft.

9. References

[1] Awduche, D., Rekhter, Y., Drake, J., Coltun, R., "Multi- Protocol Lambda Switching: Combining MPLS Traffic Engineering Control With Optical Crossconnects", draft-awduche-mpls-te-optical-02.txt (work in progress)

[2] Basak, D., Awduche, D., Drake, J., Rekhter, Y., "Multi- protocol Lambda Switching: Issues in Combining MPLS Traffic Engineering Control With Optical Crossconnects", draft-basak-mpls-oxc-issues-01.txt (work in progress)

[ISIS-TE] Smit, H., Li, T., "IS-IS Extensions for Traffic Engineering", draft-ietf-isis-traffic-02.txt (work in progress)

[4] Kompella, K., Rekhter, Y., Berger, L., "Link Bundling in MPLS Traffic Engineering", draft-kompella-mpls-bundle-05.txt (work in progress)

[5] Kompella, K., Rekhter, Y., "LSP Hierarchy with MPLS TE", draft-ietf-mpls-lsp-hierarchy-01.txt (work in progress)

[6] Generalized MPLS Group, "Generalized MPLS - Signaling Functional Description", draft-ietf-mpls-generalized-signaling-02.txt (work in progress)

[7] Lang J., Mitra K., Drake J., Kompella K., Rekhter Y., Berger L., Saha, D., Sandick, H., and Basak D., "Link Management Protocol", draft-ietf-mpls-lmp-02.txt (work in progress)

[OSPF-TE] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF", draft-katz-yeung-ospf-traffic-05.txt

[ISIS-GMPLS]

[OSPF-GMPLS]

[OVPN] Ould-Brahim, H., Rekhter, Y., Fedyk, D., Ashwood-Smith, P., Rosen, E., Mannie, E., Fang, L., Drake, J., "BGP/GMPLS Optical VPNs", draft-ouldbrahim-bgpgmpls-ovpn-01.txt (work in progress)

10. Authors' Information

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave
Sunnyvale, CA 94089
Email: kireeti@juniper.net

Yakov Rekhter
Juniper Networks, Inc.
1194 N. Mathilda Ave
Sunnyvale, CA 94089
Email: yakov@juniper.net

Ayan Banerjee
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138
Phone: +1.408.972.3645
Email: abanerjee@calient.net

John Drake
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138
Phone: (408) 972-3720
Email: jdrake@calient.net

Greg Bernstein
Ciena Corporation
10480 Ridgeview Court
Cupertino, CA 94014
Phone: (408) 366-4713
Email: greg@ciena.com

Don Fedyk
Nortel Networks Corp.
600 Technology Park Drive
Billerica, MA 01821
Phone: +1-978-288-4506
Email: dwfedyk@nortelnetworks.com

Eric Mannie
GTS Network Services
RDI Department, Core Network Technology Group
Terhulpsesteenweg, 6A
1560 Hoeilaart, Belgium
Phone: +32-2-658.56.52
Email: eric.mannie@ebone.com

Debanjan Saha
Tellium Optical Systems
2 Crescent Place
P.O. Box 901
Ocean Port, NJ 07757
Phone: (732) 923-4264
Email: dsaha@tellium.com

Vishal Sharma
Jasmine Networks, Inc.
3061 Zanker Rd, Suite B
San Jose, CA 95134
Phone: (408) 895-5000
Email: vsharma@jasminenetworks.com

Debashis Basak
AcceLight Networks,
70 Abele Rd, Bldg 1200
Bridgeville PA 15017
Email: dbasak@accelight.com