

Internet Engineering Task Force
INTERNET-DRAFT

Expiration Date: January 2001

Daniel O. Awduche
UUNET (Worldcom)

Yakov Rekhter
Cisco Systems

John Drake
Calient Networks

Rob Coltun
Redback Networks
July 2000

**Multi-Protocol Lambda Switching:
Combining MPLS Traffic Engineering Control With Optical Crossconnects**

[draft-awduche-mpls-te-optical-02.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) except that the right to produce derivative works is not granted.

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>.

Abstract

This document describes an approach to the design of control planes for optical crossconnects (OXCs), which leverages existing control plane techniques developed for MPLS Traffic Engineering. The proposed approach combines recent advances in MPLS traffic engineering control plane constructs with OXC technology to: (1) provide a framework for real-time provisioning of optical channels in automatically switched optical networks, (2) foster the expedited development and deployment of a new class of versatile OXCs, and (3) allow the use of uniform semantics for network management and operations control in hybrid networks consisting of OXCs and label switching routers (LSRs). The proposed approach is particularly advantageous for OXCs intended for data-centric optical internetworking systems. In such environments, it will help to simplify network administration. This approach also paves the way for the eventual incorporation of DWDM multiplexing capabilities in IP routers.

1. Introduction

This document describes an approach to the design of control planes for optical crossconnects (OXCs), which is based on the Multiprotocol Label Switching (MPLS) traffic engineering control plane model. In this approach, the main idea is to leverage recent advances in control plane technology developed for MPLS traffic engineering (see [[1](#),[2](#),[3](#),[4](#),[8](#),[9](#),[10](#)]). This approach is driven by pragmatic considerations, as it exploits an existing technology base to foster rapid development and deployment of a new class of versatile OXCs that address the optical transport needs of the rapidly evolving Internet. This approach will assist in optical channel layer bandwidth management, dynamic provisioning of optical channels, and network survivability through enhanced protection and restoration capabilities in the optical domain.

As used in this document, an OXC is a path switching element in an optical transport network that establishes routed paths for optical channels by locally connecting an optical channel from an input port (fiber) to an output port (fiber) on the switch element. Additional characteristics of OXCs, as used in this document, are discussed in [Section 4.1](#).

The proposed OXC control plane uses the IGP extensions for MPLS traffic engineering (with additional enhancements) to distribute relevant optical transport network state information, including topology state information. This state information is subsequently

used by a constraint-based routing system to compute paths for point-to-point optical channels through the optical transport network. The proposed OXC control plane also uses an MPLS signaling protocol (see [3,4]) to instantiate point-to-point optical channels between access points in the optical transport network.

This document does not specify the details of the extensions and domain specific adaptations required to map the MPLS traffic engineering control plane model onto the optical domain. These aspects will be covered in a number of supplementary documents that will follow. However, in [Section 7](#) of this memo, we provide a high level overview of the architectural issues involved in making such adaptations.

2. Advantages

The advantages of the proposed approach are numerous.

- It offers a framework for optical bandwidth management and the real-time provisioning of optical channels in automatically switched optical networks.
- It exploits recent advances in MPLS control plane technology and also leverages accumulated operational experience with IP distributed routing control.
- It obviates the need to reinvent a new class of control protocols for optical transport networks and allows reuse of software artifacts originally developed for the MPLS traffic engineering application. Consequently, it fosters the rapid development and deployment of a new class of versatile OXCs.
- It facilitates the introduction of control coordination concepts between data network elements and optical network elements.
- It simplifies network administration in facilities based service provider networks by providing uniform semantics for network management and control in both the data and optical domains.
- It paves the way for the eventual introduction of DWDM multiplexing capabilities on IP routers
- Lastly, it establishes a preliminary framework for the notion of an optical Internet.

3. Background

The growth, performance, and survivability requirements of the Internet (which is also becoming very mission critical) are mandating IP-centric networks to be cost effective, survivable, scalable, and to provide control capabilities that facilitate network performance optimization. Some of these requirements are being addressed by the Multiprotocol Label Switching (MPLS) traffic engineering capabilities under development by the IETF [[1](#),[2](#),[3](#),[4](#)].

The underlying optical transport network also needs to be versatile, reconfigurable, cost effective, and to support a variety of protection and restoration capabilities due to the rapidly changing requirements of the Internet.

A result of these trends, therefore, is the evolution of optical transport networks from simple linear and ring topologies to mesh topologies. By a mesh (not necessarily fully meshed) topology, we mean a connected (not necessarily fully connected) network of arbitrary topology in which the node degree is typically more than two. In mesh optical networks, optical crossconnects engender versatility and reconfigurability by performing switching and rearrangement functions.

Underscoring the importance of versatile networking capabilities in the optical domain, a number of standardization organizations and interoperability forums have initiated work items to study the requirements and architectures for reconfigurable optical networks. Refer, for example, to ITU-T recommendation G.872 [[5](#)]. This document defines a functional architecture for an optical transport network (OTN) that supports the transport of digital client signals. ITU-T G.872 speaks of an OTN as "a transport network bounded by optical channel access points"[[5](#)]. The ITU-T G.872 OTN architecture is based on a layered structure, which includes:

- (a) an optical channel (OCh) layer network,
- (b) an optical multiplex section (OMS) layer network, and
- (c) an optical transmission section (OTS) layer network.

The optical channel layer is the most relevant to the discussions in this document. The optical channel layer network supports end-to-end networking of optical channel trails between access points. The OCh layer network provides the following functions: routing, monitoring, grooming, and protection and restoration of optical channels. In this situation, programmable Optical crossconnects, with rearrangeable switch fabrics and relatively smart control planes, will be critical to the realization of the OCh layer functions, especially in mesh optical networks. In the data network domain,

routing, monitoring, and survivability are crucial functions performed by the MPLS traffic engineering control plane (see [[1](#),[2](#),[3](#),[4](#),[8](#),[9](#),[10](#)]).

Note: Although we mention the ITU-T recommendation G.872, the OXC control plane design approach described here is not restricted to G.872 compliant networks. Instead, it can be applied to any mesh optical network that uses programmable and reconfigurable OXCs.

Other standards organizations and interoperability forums that are actively pursuing projects related to dynamically reconfigurable optical networks include the ANSI T1X1.5 committee, the Optical Internetworking Forum (OIF), and now by virtue of this memo the IETF.

Recent contributions to ANSI T1X1.5 emphasize the need for automation of the OCh layer network by using appropriate signaling protocols to establish optical channels in real time (see [[12](#)] and [[13](#)]).

The Optical Internetworking Forum (<http://www.oiforum.com>), an international organization engaged in the development and promotion of interoperability agreements for optical internetworking systems, is also evaluating architectural and signaling options related to the internetworking of data network elements with reconfigurable optical networks -- to facilitate rapid provisioning, efficient protection/restoration, and other services in optical internetworking systems.

In all these cases, the successful realization of the vision of versatile optical networking depends very much on the synthesis of appropriate control plane technologies with programmable and reconfigurable OXCs.

[4.](#) OXCs, LSRs, Optical Trails, and Explicit LSPs

Consider a hybrid, IP-centric optical internetworking environment consisting of both label switching routers (LSRs) and OXCs, where the OXCs are programmable and support wavelength conversion/translation.

At a level of abstraction, an LSR and an OXC exhibit a number of isomorphic relations. It is important to enumerate these relations because they help to expose the reusable software artifacts from the MPLS traffic engineering control plane model. Architecturally, both LSRs and OXCs emphasize problem decomposition by decoupling the control plane from the data plane.

The data plane of an LSR uses the label swapping paradigm to transfer a labeled packet from an input port to an output port (see e.g.,

[6,7]). The data plane of an OXC uses a switching matrix to connect an optical channel trail from an input port to an output port.

An LSR performs label switching by first establishing a relation between an <input port, input label> tuple and an <output port, output label> tuple. Likewise, an OXC provisions optical channel trails by first establishing a relation between an <input port, input optical channel> tuple and an <output port, output optical channel> tuple. These relations are determined by the control plane of the respective network elements, and are locally instantiated on the device through a switch controller. In the LSR, the next hop label forwarding entry (NHLFE) maintains the input-output relations. In the OXC, the switch controller reconfigures the internal interconnection fabric to establish the relations. These relations cannot be altered by the payload of the data plane.

The functions of the control plane (for both LSRs and OXCs) include resource discovery, distributed routing control, and connection management. In particular, the control plane of the LSR is used to discover, distribute, and maintain relevant state information associated with the MPLS network, and to instantiate and maintain label switched paths (LSPs) under various MPLS traffic engineering rules and policies. An LSP is the path through one or more LSRs followed by a specific forwarding equivalence class (FEC) (see [7]). An explicit LSP is one whose route is defined at its origination node.

The control plane of the OXC, on the other hand, is used to discover, distribute, and maintain relevant state information associated with the OTN, and to establish and maintain optical channel trails under various optical internetworking traffic engineering rules and policies. An optical channel trail provides a point-to-point optical connection between two access points. An optical channel trail may consist of just one wavelength or a concatenation of multiple wavelengths. If an optical trail consists of just one wavelength, then it is said to satisfy the "wavelength continuity property." At each intermediate OXC along the route of an optical channel trail, OXC switch fabric connects the trail from an input port to an output port.

A distinction between the current generation of OXCs and LSRs is that the former do not perform packet level processing in the data plane, while the latter are datagram devices which may perform certain packet level operations in the data plane. A significant conceptual difference is that with LSRs the forwarding information is carried explicitly as part of the labels appended to data packets, while with OXCs the switching information is implied from the wavelength or optical channel.

4.1 Review of Relevant OXC Characteristics

The following section contains a brief overview of relevant OXC characteristics, focusing on the switching functions and underlying technologies.

As used in this document, the switching function of an OXC may be electrical or optical. If the switching fabric is purely electrical, then the crossconnect is typically referred to as a digital crossconnect (DXC), or a broadband digital cross-connect (BBDXC) -- if the capacity and port density are sufficiently high. Optical-Electrical-Optical (OEO) conversion is required in BBDXCs. A BBDXC may or may not have WDM multiplexing capabilities. If a BBDXC has WDM multiplexing capabilities, then it may be connected directly to other compatible WDM devices through optical fiber links that carry multiple wavelengths per fiber. If a BBDXC does not have WDM multiplexing capabilities, then it may be connected to an external DWDM multiplexer through a set of discrete fibers, where each fiber carries only one wavelength. A BBDXC may also perform regeneration, reshaping, and re-timing functions.

If the switching fabric of an OXC is completely photonic, then we refer to the cross-connect as a pure OXC. If the granularity of channel switching is the wavelength, then the OXC is called a wavelength routing switch (WRS), or simply a wavelength router. The problem of establishing optical channel trails using WRS is called the "Routing and Wavelength Assignment problem" (RWA) [[11](#)].

An OXC may also be equipped with both electrical and optical switching capabilities. In this situation, some channels may be switched in the electrical domain and others in the optical domain.

Other terms commonly used within the context of optical transport network switching elements include: wavelength selective crossconnects (WSXC) and wavelength interchanging crossconnects (WIXC).

In this document, we use the generic term OXC to refer to all categories of programmable and reconfigurable crossconnects for optical transport networks, irrespective of the technologies that underlie them.

The OXC control plane design approach described in this document is independent of the underlying OXC switch technologies. It is also independent of specific OXC implementation details. Local adaptation mechanisms can be used to tailor the control plane onto various OXC implementations with different hardware capabilities. As an example, a local adaption function can map a channel/port input/output

relation into specialized low level instructions to actuate a rearrangement of the crossconnect switch fabric such that the required input/output relation is realized.

A number of forthcoming supplementary documents will describe in some detail the extensions needed to adapt the control plane approach described in this memo to various OXC technologies and optical transport network contexts.

4.2 Explicit LSPs and Optical Channel Trails

At a conceptual level, explicit LSPs and optical channel trails exhibit certain commonalities. Essentially, they are both fundamentally unidirectional, point-to-point virtual path connection abstractions. An explicit LSP provides a parameterized packet forwarding path (traffic-trunk) between an ingress LSR and an egress LSR. Correspondingly, an optical channel trail provides a (possibly parameterized) optical channel between two endpoints for the transport of client digital signals.

The payload carried by both LSPs and optical trails are transparent to intermediate nodes along their respective paths. Both LSPs and optical trails can be parameterized to stipulate their performance, behavioral, and survivability requirements from the network. A set of LSPs induce a virtual graph on a data network topology, while a set of optical trails induce a virtual graph on the topology of a fiber plant.

A constraint-based routing scheme can be used to select appropriate paths for both LSPs and optical trails. Generally such paths may satisfy some demands and policy requirements subject to some constraints imposed by the operational environment.

There are also commonalities in the allocation of labels to LSPs and in the allocation of wavelengths to optical trails. Two different LSPs that traverse through a given LSR port or interface cannot be allocated the same label. The exception is for LSP aggregation using label merge or label stacking. Similarly, two different optical trails that traverse through a given OXC port cannot be allocated the same wavelength.

5. Generic Requirements for the OXC Control Plane

The following section contains the requirements for the OXC control plane, with emphasis on the routing components of these requirements. There are three key aspects to these requirements:

- (1) the capability to establish optical channel trails expeditiously, (in seconds or even milliseconds rather than days or months).
- (2) the capability to support traffic engineering functions, (see note below)
- (3) the capability to support various protection and restoration schemes.

Note: the introduction of DWDM and automatically switched optical networks is unlikely to eliminate the need for traffic engineering. Instead, it will simply mandate OXCs to also support some traffic engineering capabilities.

Historically, the "control plane" of optical transport networks has been implemented via network management. This approach has the following drawbacks:

- (1) It leads to relatively slow convergence following failure events (typical restoration times are measured in minutes, or even days and weeks especially in systems that require explicit manual intervention). The only way to expedite service recovery in such environments is to pre-provision dedicated protection channels.
- (2) It complicates the task of interworking equipment from different manufacturers, especially at the management level (generally, a custom "umbrella network management system -NMS- or operations support system -OSS-" is required to integrate otherwise incompatible Element Management Systems from different vendors)
- (3) It precludes the use of distributed dynamic routing control capabilities in such environments.
- (4) It complicates the task of inter-network provisioning (due to the lack of EDI between operator NMSs).

Thus, another important motivation for the approach described in this document is to improve the responsiveness of the optical transport network, and to increase the level of interoperability within and between service provider networks.

6. MPLS Traffic Engineering as a Generic Control Plane for OXCs

The requirements for the OXC control plane described in the previous section have already been addressed by the MPLS Traffic Engineering control plane, under development by the IETF (see e.g., [[1](#),[2](#),[3](#),[4](#),[8](#),[9](#)]).

6.1 Overview of the MPLS Traffic Engineering Control Plane

Let us now discuss the components of the MPLS traffic engineering control plane model. The MPLS traffic engineering control plane is a synthesis of new concepts in IP traffic engineering (enabled by MPLS) and the conventional IP network layer control plane. The high level requirements for traffic engineering over MPLS were articulated in [RFC-2702](#) [[1](#)]. It is the combination of the notions defined in [RFC-2702](#) (including relevant extensions) along with the conventional IP control plane constructs that effectively establishes a framework for the MPLS traffic engineering control plane model [[1](#)] (see also [[2](#)]).

The components of the MPLS traffic engineering control plane model include the following modules:

- Resource discovery
- State information dissemination, which is used to distribute relevant information concerning the state of the network, including topology and resource availability information. In the MPLS context, this is accomplished by extending conventional IP link state interior gateway protocols to carry additional information in their link state advertisements (see [[8](#),[9](#)]).
- Path selection, which is used to select an appropriate route through the MPLS network for explicit routing. It is implemented by introducing the concept of constraint-based routing which is used to compute paths that satisfy certain specifications subject to certain constraints, including constraints imposed by the operational environment (see [[1](#)]).
- Path management, which includes label distribution, path placement, path maintenance, and path revocation. These are used to establish, maintain, and tear down LSPs in the MPLS context. The label distribution, path placement, and path revocation functions are implemented through a signaling protocol, such as the RSVP extensions [[3](#)] or through CR-LDP [[4](#)].

These components of the MPLS traffic engineering control plane are separable, and independent of each other. This is a very attractive feature because it allows an MPLS control plane to be implemented using a composition or synthesis of best of breed modules.

In [RFC-2702](#) [1], several new MPLS control plane capabilities were proposed that allow various traffic engineering policies to be actualized in MPLS networks. Many of these capabilities are also relevant and applicable to automatically switched optical transport networks with reconfigurable OXCs.

We summarize some of these capabilities below, focusing on the set of attributes that can be associated with traffic-trunks. A traffic-trunk is an aggregation of traffic belonging to the same class which are forwarded through a common path. In general, the traffic-trunk concept is a technology independent abstraction. In [1], it was used within the context of MPLS and allowed certain attributes of the traffic transported through LSPs to be parameterized. The traffic-trunk concept can also be extended, in an obvious manner, to the optical transport network.

As stipulated in [RFC-2702](#) [1], the attributes that can be associated with traffic-trunks include: (1) traffic parameters which indicate the bandwidth requirements of the traffic-trunk, (2) adaptivity attributes which specify the sensitivity of the traffic-trunk to changes in the state of the network and in particular indicates whether the traffic-trunk can be re-routed when "better" paths become available, (3) priority attributes which impose a partial order on the set of traffic-trunks and allow path selection and path placement operations to be prioritized, (4) preemption attributes which indicate whether a traffic-trunk can preempt an existing traffic-trunk from its path, (5) resilience attributes which stipulate the survivability requirements of the traffic-trunk and in particular the response of the system to faults that impact the path of the traffic-trunk, and (6) resource class affinity attributes which further restrict route selection to specific subsets of resources and in particular allow generalized inclusion and exclusion policies to be implemented. Other policy attributes and options are also defined by Awduche et al in [RFC-2702](#) [1] for traffic-trunks, including policing attributes [1] (policing is irrelevant in the OXC context). Concepts of subscription (booking) factors are also supported to either bound the utilization of network resources through under-subscription or to exploit statistical multiplexing gain through over-subscription (this aspect is also not very relevant in the OXC context).

It should be clear that a subset of these capabilities can be mapped onto an optical transport network by substituting the term "traffic-

trunk" with the term "optical channel trail."

The MPLS control plane also supports the notion of abstract nodes. An abstract node is essentially a set of nodes (e.g., a subnet, an autonomous system, etc) whose internal topology is opaque to the origination node of an explicit LSP. So, in the most general manner, the route of an explicit LSP (or traffic-trunk) can be specified as a sequence of single hops and/or as a sequence of abstract nodes.

The MPLS control plane is very general and is also oblivious of the specifics of the data plane technology. In this regard, the MPLS control plane can be used in conjunction with a data plane that (a) does not necessarily process IP packet headers and (b) does not know about IP packet boundaries. For an existence proof, note that the MPLS control plane has been implemented on IP-LSRs, ATM-LSRs, and Frame Relay-LSRs.

The MPLS control plane may also be implemented on OXC's as discussed in this document.

[6.2](#) Synthesizing the MPLS Traffic Engineering Control Plane with OXC's

Given that both OXC's and LSRs require control planes, one option would be to have two separate, independent, and incompatible control planes - one for OXC's, and another for LSRs. To understand the drawbacks of this approach, especially in IP-centric optical internetworking systems, one need to look no further than the experience with IP over ATM, where IP has its own control plane (BGP, IS-IS, OSPF), and ATM its own control plane (PNNI) [[12](#)]. For some of the drawbacks see [[1,2](#)].

Given that the control planes for both OXC's and LSRs have relatively similar requirements, an alternative approach is to develop a coherent control plane technology that can be used for LSRs and for OXC's. Such a uniform control plane will eliminate the administrative complexity of managing hybrid optical internetworking systems with separate, dissimilar control and operational semantics. Specializations may be introduced in the control plane, as necessary, to account for inherent peculiarities of the underlying technologies and networking contexts.

All of the above observations suggest, therefore, that the MPLS Traffic Engineering control plane (with some minor extensions) would be very suitable as the control plane for OXC's. An OXC that uses the MPLS traffic engineering control plane would effectively become an IP addressable device. Thus, this proposition also solves the problem of addressing for OXC's. The distribution of topology state information,

establishment of optical channel trails, OTN traffic engineering functions, and protection and restoration capabilities would be facilitated by the MPLS Traffic Engineering control plane.

An out-of-band IP communications system can be used to carry and distribute control traffic between the control planes of OXCs, perhaps through dedicated supervisory channels (using e.g., dedicated wavelengths or channels, or an independent out-of-band IP network). In this environment, SNMP, or some other network management technology, could be used for element management. From the perspective of control semantics, an OXC with an MPLS Traffic Engineering control plane would resemble a Label Switching Router.

If the OXC is a wavelength routing switch, then the physical fiber between a pair of OXCs would represent a single link in the OTN network topology. Individual wavelengths or channels would be analogous to labels. If there are multiple fibers between a pair of OXCs, then as an option, these multiple fibers could be logically grouped together through a process called bundling and represented as a single link in the OTN network topology.

If a fiber terminates on a device that functions as both an OXC and an IP router, then the following situation may be possible:

- A subset of optical channels within the fiber 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 OXC element to other out-bound optical channels and thus are not immediately available for allocation.
- Another subset of optical channels (within the same fiber) could be in use as terminal channels. That is, they are already allocated but terminate on the local OXC/router device, for example, as SONET interfaces.

In the above scenario one way to represent the fiber in the OTN network topology is to depict it as several links, where one of these links would represent the set of uncommitted channels which constitute the residual capacity of the fiber; while each terminal channel that terminates on the OXC/router could be represented as an individual link.

In the control plane model described here, IS-IS or OSPF, with

extensions for traffic engineering ([8] or [9]) and possibly additional optical network specific extensions would be used to distribute information about the optical transport network topology and information about available bandwidth and available channels per fiber, as well as other OTN network topology state information. This information will then be used to compute explicit routes for optical channel trails. An MPLS signaling protocol, such as RSVP extensions (see [RSVP]), will be used to instantiate the optical channel trails. Using the RSVP extensions, for example, the wavelength information or optical channel information (as the case may be) will be carried in the LABEL object, which will be used to control and reconfigure the OXCs.

The use of a uniform control plane technology for both LSRs and OXCs introduces a number of interesting (and potentially advantageous) architectural possibilities. One such possibility is that a single control plane (MPLS Traffic Engineering) may be able to span both routers and OXCs. In such an environment a Label Switched Path could traverse an intermix of routers and OXCs, or could span just routers, or just OXCs. This offers the potential for real bandwidth-on-demand networking, in which an IP router may dynamically request bandwidth services from the optical transport network. Another possibility is that OXCs and LSRs may run different instances of the control plane which are decoupled with little or no interaction between the control plane instances.

To bootstrap the system, OXCs must be able to exchange control information. One way to support this is to pre-configure a dedicated control wavelength between each pair of adjacent OXCs, or between an OXC and a router, and to use this wavelength as a supervisory channel for exchange of control traffic. Another possibility, which has already been mentioned, is to construct a dedicated out of band IP network for the distribution of control traffic.

Even though an OXC equipped with an MPLS traffic engineering control plane would (from a control perspective) resemble a Label Switching Router, there are some important distinctions and limitations. One distinction concerns the fact that there are no analogs of label merging in the optical domain. This implies that an OXC cannot merge several wavelengths into one wavelength. Another distinction is that an OXC cannot perform the equivalent of label push and pop operations in the optical domain. This is because the analog of a label in the OXC is a wavelength or an optical channel, and the concept of pushing and popping wavelengths is infeasible with contemporary commercial optical technologies.

In the proposed control plane approach, an OXC will maintain a WFIB (Wavelength Forwarding Information Base) per interface (or per

fiber). This is because lambdas and/or channels (labels) are specific to a particular interface (fiber), and the same lambda and/or channel (label) could be used concurrently on multiple interfaces (fibers).

The MPLS traffic engineering control plane is already being implemented on data plane technologies that exhibit some of the aforementioned distinctions. For example, an ATM-LSR supports only a subset of the MPLS functionality. In particular, most ATM-LSRs are incapable of merging Label Switching Paths, and may not be able to perform label push/pop operations as well. Also, similar to the approach proposed here for OXCs, ATM-LSRs have per interface LFIB (Label Forwarding Information Base).

Yet another important distinction concerns the granularity of resource allocation. An MPLS Label Switching Router which operates in the electrical domain can potentially support an arbitrary number of LSPs with arbitrary bandwidth reservation granularities (bounded by the maximum reservable bandwidth per interface, the label space, and the amount of required control overhead). In sharp contrast, an OXC can only support a relatively small number of optical channel trails (this may change as the technology evolves), each of which will have coarse discrete bandwidth granularities (e.g., OC-12, OC-48, OC-192, and OC-768). A special degenerate case occurs when the control plane is used to establish optical channel trails which all have a fixed bandwidth (e.g., OC-48).

If the bandwidth associated with an LSP is small relative to the capacity of an optical channel trail, then very inefficient utilization of network resources could result if only one LSP is mapped onto a given optical channel trail. To improve utilization of resources, therefore, it is necessary to be able to map several low bandwidth LSPs onto a relatively high capacity optical channel trail. For this purpose, a generalized notion of "nested LSPs" may be used. Note that since an OXC cannot perform label push/pop operations, the start/end of a nested LSP has to be on a router (as nesting requires label push/pop). Also note that in this nesting situation, it is the wavelength of the "container" optical channel trail itself that effectively constitutes the outermost label.

The transparency and multiprotocol properties of the MPLS Control Plane approach would allow an OXC to route optical channel trails carrying various types of digital payloads (including IP, ATM, Sonet, etc) in a coherent and uniform way.

7. Control Adaptation

This section provides a high level overview of the architectural considerations involved in tailoring the MPLS traffic engineering control plane model to the optical domain. More detailed discussions of these issues will be provided in a number of supplementary documents to follow.

In adapting the MPLS traffic engineering control plane model to OXCs, a number of critical issues should be considered. One critical issue concerns the development of OTN specific domain models which abstracts and captures relevant characteristics of the OTN. The domain models help to delineate the design space for the control plane problem in OXCs, and to construct domain specific software reference architectures.


A domain model includes functional and structural aspects. For the purpose of the present discussions, however, we have grouped the considerations pertaining to OTN domain models into two broad categories: (1) a horizontal dimension and (2) a vertical dimension. The horizontal dimension pertains to the specialized networking requirements of the OTN environment. It indicates the enhancements needed to the MPLS TE control plane model to address the peculiar OTN networking requirements. The vertical dimension pertains to localized hardware and software characteristics of the OXCs, which helps to determine the device specific adaptations and support mechanisms needed to port and reuse the MPLS TE control plane software artifacts on an OXC.

Horizontal dimension considerations include the following aspects:

- What type of OTN state information should be discovered and disseminated to support path selection for optical channel trails? Such state information may include domain specific characteristics of the OTN (encoded as metrics), such as attenuation, dispersion (chromatic, PMD), etc. This aspect will determine the type of additional extensions that are required to IGP link state advertisements to allow distribute such information.
- What infrastructure will be used to propagate the control information?
- How are constrained paths computed for optical channel trails which fulfill a set of performance and policy requirements subject to a set of system constraints?
- What are the domain specific requirements for setting up optical

channel trails and what are the enhancements needed to existing MPLS signaling protocols to address these requirements?

Vertical dimension considerations include the aspects required to practically port MPLS control plane software onto an OXC. In terms of vertical dimensions, a candidate system architecture for an OXC equipped with an MPLS control plane model is shown in Figure 1 below.



Essentially, there are two extremal architectural options for deployment of the proposed control plane in an operational context consisting of LSRs and OXCs.

- Overlay Option: One option is to use different instances of the control plane in the OTN (OXC) and IP (LSR) domains. In this situation, each instance of the control plane will operate independent of the other. Interworking (including control coordination) between the two domains can be established through static configuration or through some other procedures that are outside the scope of this document. This partitioned and explicitly decoupled deployment option allows maximal control isolation between the OTN and IP domains. This scheme is conceptually similar to the model in use today, whereby the OTN simply provides point-to-point channels between IP network elements with very minimal control interaction between the two domains.
- Peer Option: Another option is to use a single instance of the control plane that subsumes and spans LSRs and OXCs.

Other architectural options are also possible which allow various degrees of control isolation and control integration between the OXCs and LSRs.

To improve scalability the control plane may use routing hierarchy (e.g., routing areas). Hierarchy may be applied in either of the situations mentioned above. Furthermore, in the overlay option with different control plane instances for OXCs and LSRs, hierarchy could be enabled for each control plane instance independent of the other.

In the deployment option with a single instance of the control plane, each routing area may maintain a link state database that contains: (1) physical LSPs (fiber links), (2) optical LSPs (optical channel trails), and (3) logical LSPs (conventional label switched paths). As a general rule, all of these path-oriented connection entities could simply be considered as LSPs with different characteristics. The origination LSR (the head-end) of each LSP entity may locally decide whether to advertise the LSP (with appropriate attributes), so that other LSRs could use it as a link for subsequent path computations.

There are significant tradeoffs to the above deployment options, including aspects related to scalability and fault isolation. Additional documents to follow may elaborate on some of these aspects.

One of the advantages of the control plane design approach described in this memo is that it potentially allows network administrators the

leeway to make these deployment architectural decisions based on their specific objectives, network contexts, and service models.

9. Summary

This document outlined how the MPLS Traffic Engineering control plane could be adapted and reused as the control plane for optical crossconnects. Such a control plane would be used to distribute optical transport network topology state information and to setup optical channel trails. Such a control plane would support various traffic engineering functions in the optical domain, and enable a variety of protection and restoration capabilities. Furthermore, such a control plane technology would expedite the development and deployment of a new class of versatile data-centric OXCs. Additionally, the proposed control plane approach would simplify integration of OXCs and label switching routers. Finally, the proposed control plane approach would provide coherent semantics for network management and operations control in hybrid optical internetworking systems consisting of LSRs and OXCs.

References:

- [1] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus, "Requirements for Traffic Engineering Over MPLS," [RFC-2702](#), September 1999.
- [2] D. Awduche, "MPLS and Traffic Engineering in IP Networks," IEEE Communications Magazine, December, 1999.
- [3] D. Awduche, L. Berger, D. Gan, T. Li, G. Swallow, and V. Srinivasan, "Extensions to RSVP for LSP Tunnels," Internet Draft, Work in Progress, 1999
- [4] B. Jamoussi et al, "Constraint-Based LSP Setup using LDP," Internet Draft, Work in Progress, 1999
- [5] ITU-T G.872, "Architecture for Optical Transport Networks," 1999.
- [6] R. Callon, P. Doolan, N. Feldman, A. Fredette, G. Swallow, A. Viswanathan, "A Framework for Multiprotocol Label Switching," Internet Draft, Work in Progress, 1999
- [7] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label

Switching Architecture," Internet Draft, Work in Progress, 1999

[8] H. Smit and T. Li, "IS-IS extensions for Traffic Engineering," Internet Draft, Work in Progress, 1999

[9] D. Katz, D. Yeung, "Traffic Engineering Extensions to OSPF," Internet Draft, Work in Progress, 1999

[10] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," [RFC-1171](#), March 1995

[11] B. Mukherjee, "Optical Communications Networks," McGraw-Hill, 1997. [[14](#)]

[12] G. Newsome and P. Bonenfant, "The Automatic Switched Optical Network," Contribution to T1 STANDARDS PROJECT - T1X1.5, 1999.

[13] P. Bonenfant and X. Liu, "A Proposal for Automatically Switched Optical Channel Networks (ASON)," Contribution to T1 STANDARDS PROJECT - T1X1.5, 1999.

[14] "T. Worster, "General Switch Management Protocol," Internet Draft, Work in Progress, 1999.

[15] E. Crawley, R. Nair, B. Rajagopalan, and H. Sandick, "A Framework for QoS-Based Routing in the Internet," [RFC-2386](#), August, 1998.

Security Considerations

It is imperative to guarantee the integrity and confidentiality of control information used by the proposed OXC control plane. This can be accomplished by using existing security mechanisms for the various components of the MPLS traffic engineering control plane model.

Author's Addresses

Daniel O. Awduche
UUNET (Worldcom)
22001 Loudoun County Parkway
Ashburn, VA 20147
Phone: 703-886-5277
Email: awduche@uu.net

Yakov Rehkter
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134
Email: yakov@cisco.com

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

Rob Coltun
Redback Networks
300 Ferguson Drive
Mountain View, CA 94043
Phone: (650) 390-9030
Email: redback.com

