

Contribution Number: oif2000.096

Working Group: Signaling

TITLE: Deployment Considerations for the User to Network Interface (UNI)

SOURCE: D. Pendarakis, B. Rajagopalan and D. Saha

Tellium, Inc.

Email: {dpendarakis, braja, dsaha} @ tellium.com

DATE: May 2, 2000

ABSTRACT:

The objective of this contribution is to describe the UNI service definition and its relation to the optical internetworking model. It discusses the UNI design objectives and the information exchanged between an electrical device and the optical network as required to support different internetworking models. The signaling messages exchanged in order to perform address registration and discovery are described. In addition, deployment considerations related to policy control and security are discussed.

This contribution discusses the optical internetworking model and the associated User to Network Interface service definition.

1. Definition

The User to Network Interface (UNI) defines the interface between the optical network and a client to the optical network. Clients include electrical devices such as IP routers, ATM switches and SONET add/drop multiplexers. The interface will allow dynamic bandwidth provisioning with a desired set of attributes between two endpoints residing on clients, across an optical network.

2. Internetworking Model

2.1 Design Objectives

The objective of UNI specification is twofold. First, it defines the exact service offered by the optical network, as expressed by the message types, formats, attributes and signaling mechanisms used across the interface. Second, UNI specification implicitly defines the internetworking model between an optical network and networks of electrical devices. Decisions related to the internetworking model include:

- Naming and addressing scheme used to identify endpoints on optical devices.
- Level of detail of topology and routing information advertised by the optical network to electrical networks.

Optical switches (OXC) differ fundamentally from IP routers and ATM switches in the sense that they switch lightpaths, in an operation analogous to circuit switching, and do not perform packet by packet forwarding. Given the high amount of bandwidth offered, it is expected that optical switches (OXC) will be deployed at the network core, with electrical devices performing grooming and traffic aggregation at the periphery. In an optical network a lightpath has to be set-up prior to any data transfer. These lightpaths are composed of point-to-point connections between adjacent OXCs. The internetworking model is responsible for defining the following:

- What is the target topology of lightpaths interconnecting a set of N electrical devices? Options include:
 - Full mesh topology where every router is connected to every other router through a dedicated lightpath, resulting in an N^2 count of lightpaths.
 - A topology with less than N^2 lightpaths, where some pairs of routers are connected through other routers. In this case, the average number of lightpaths that have to be traversed to reach a router from another router is larger than 1.
- What is the time scale of setting up and tearing-down the above lightpaths?
 - Lightpaths could be set up for days, weeks or months, in which case they are called *permanent*.
 - Lightpaths could be set up and torn down on demand and for considerably smaller periods of time (seconds or minutes), in which case they are called *dynamic* or *switched*?
- How is the path of lightpaths selected across the optical network?
 - This can be done entirely by the optical network, either in a distributed or centralized manner.
 - Electrical devices such as routers could maintain complete topology and routing information to attempt circuit set up by specifying the complete path.
 - An alternative approach could include elements of both of the above solutions; path selection is, in general, the responsibility of the optical domain, while there is also an option available to electrical devices to specify the path or portions thereof.

2.2 Peering Model

The relation of the optical networking layer to existing protocol layers, such as IP, yields different internetworking models. The *peer model* uses IP addressing schemes within optical networks. OXC endpoints are identified by existing network layer (IP) addresses which are carried in signaling requests. Existing network layer routing protocols such as OSPF can be used within the optical network to route lightpath set-up requests. In the peer model the optical layer is a peer of IP and therefore can take advantage of routing and signaling protocols defined for IP.

The alternate model, known as *overlay model*, seeks to decouple the optical layer from existing protocols. In the case of ATM the overlay model required the definition of both a new addressing structure, and an associated routing protocol that could take advantage of ATM's unique QoS capabilities. The ATM overlay model introduced substantial complexity by requiring ATM address resolution in order to map IP addresses to the corresponding ATM addresses, which proved to be impractical over wide area networks.

The optical internetworking model should try to combine the *best of both overlay and peer* models. The internetworking model will resemble the peer model in that optical network addressing, signaling and routing protocols will be modeled closely to existing schemes defined in the IP domain. Modifications should be introduced only to the extent they are necessary to reflect unique characteristics and requirements of optical networks, such as physical line attributes, redundancy and restoration parameters, etc. In this way, development of protocols and interoperability can be sped up.

At the same time, a network of optical switches is fundamentally different than a one consisting of IP routers. Therefore, in order to expedite development, it is desirable to include a level of decoupling, as provided in the overlay model, between optical and router networks. The level of decoupling and the amount of information exchanged between the two networks is defined by the *User Network Interface*.

The optical internetworking model allows decoupling between the optical and IP router domains by resembling hierarchical routing employed in the Internet. Examples of hierarchical routing include the two-level hierarchical routing scheme used by *OSPF areas* [1], for intra-domain routing, and BGP4 [2] which allows two different *autonomous systems (AS)* to exchange routing information so that data can be forwarded across the AS border, for inter-domain routing.

Support for hierarchical IP routing protocols in the internetworking model requires the exchange of BGP messages or OSPF summary LSAs across the UNI. An alternative is to define custom registration and reachability advertisement messages, which are exchanged between routers and border OXCs across the UNI. This allows a client of the optical network to retrieve the set of remote addresses reachable from a given port. This approach will allow dynamic set-up of signaled lightpaths with reduced routing complexity.

3. UNI Signaling Protocols

UNI signaling requires the definition of protocols used and the messages exchanged. UNI signaling provides service discovery, address registration and reachability and provisioning capabilities. The specific signaling protocol can be based on Multi-Protocol Label Switching (MPLS) signaling protocols such as RSVP and CR-LDP. Signaling messages are sent over an IP control channel, which is established either in-band or out-of-band between a border router and an OXC. It is recommended that the control channel is established in-band via use of the SONET overhead bytes. The format of signaling messages follows the RSVP or CR-LDP specifications.

3.1 Service Discovery

Service discovery is the process whereby an IP router or other electrical device communicates with an OXC across the UNI to determine the service features and supported parameters that are available to it. Service definition across the UNI is likely to have optional and negotiable components. Service discovery allows the IP router to establish the service environment.

3.2 Address Registration and Reachability Propagation

In order to enable dynamic provisioning, border routers must register their own IP address with the optical network. A router does this by sending an "Address Register" message with parameters indicating IP addresses to the border OXC using UNI signaling. The OXC notes the address parameters and returns an acknowledgement to the router. A router may also send an "Address De-Register" message, indicating IP addresses that must be removed. The border OXC is assumed to update its local state and propagate appropriate external reachability information within the optical network based on the registration and de-register messages received.

The reachability information is sent from the OXC to a router using "Reachability Update" messages. These messages indicate a set of IP addresses that are reachable across the optical network. "Address Withdraw" messages are sent from border OXCs to routers to indicate that a set of IP addresses are no longer reachable across the optical network. Using these mechanisms, border routers are aware at all times of the addresses of other border routers reachable across the optical network.

3.3 Provisioning

A router may request that an optical path be provisioned to a remote destination by specifying the IP address of the destination and the parameters of the path including bandwidth, QoS parameters, priority, protection scheme desired, etc. The provisioning request is sent as a "Set-Up Path" message across the UNI. In response, the border OXC will initiate the provisioning of the end-to-end path. If provisioning is successful, the border OXC will send a "Path Established" message to the router, indicating a path identifier. A router may tear down an existing path by specifying a "Release Path" message across the UNI, indicating the path identifier. The border OXC, after tearing down the path, sends a "Path Released" message to the router. Finally, a router may modify some of the path parameters by sending the "Modify Path" message, indicating the path identifier and new values for the parameters being modified. The border OXC will attempt to modify the path parameters. If it succeeds in this attempt, it sends a "Path Modified" message. Otherwise, the path is left intact and a "Modify Failed" message is sent back to the router.

4. Service Definition

4.1 Message Functional Specification

The following types of messages can be carried across the interface.

1. Set-up Lightpath: Allows a router or other electrical device to request that an optical path be provisioned to a remote destination.
2. Lightpath Established: If provisioning is successful, the border OXC sends this message to the router, including a path identifier.
3. Release Lightpath: Used by a router to tear-down a path, given by the path identifier.
4. Modify Lightpath: Allows a router to modify some of the path characteristics.
5. Lightpath Modified: Acknowledges successful modification of path parameters.
6. Modify Failed: Indicates failure in attempt to modify path parameters.

7. Address Register. This message allows a router connected to an OXC to register one or more local addresses that the OXC should associate with this router interface.
8. Address De-Register. Indicates IP addresses that must be removed.
9. Reachability Update. These messages are sent from an OXC to a router to indicate a set of IP addresses that are reachable across the optical network.
10. Address Withdraw messages are sent from border OXCs to routers to indicate that a set of IP addresses are no longer reachable across the optical network.
11. Address Query. Allows a router to query a border OXC for remote addresses reachable from a particular OXC port.

4.2 Naming and Addressing

The naming and addressing scheme allows connection endpoint specification. Endpoints could be identified using a combination of:

1. IP addresses of router interfaces or Optical Cross Connects
2. Port Index
3. Sub-rate channel; for example a Virtual Tributary mapped in the SONET frame.

4.3 Lightpath Attributes

- Bandwidth characteristics of the optical connection.
 1. **Channel Type:** the specific encoding and rate used for the optical channel. Currently the following usage types are defined:
 0. Reserved
 1. Transparent
 2. GE
 3. 10 GE
 4. OC-3
 5. OC-3c
 6. OC-12
 7. OC-12c
 8. OC-48
 9. OC-48c
 10. OC-192
 11. OC-192c
 12. OC-768
 13. OC-768c
 2. **Set-up Priority:** Indicates the level of priority at connection set-up. Currently defined values are 0–7.
 3. **Holding Priority:** Used to decide whether a connection in progress can be preempted when a call set-up request with higher set-up priority is requested. Currently defined values are 0-7.
 4. Optionally, Quality of Service specifications such as propagation delay and error rate.
- **Protection mode:** type of protection against failures that is requested. Currently the following protection modes are defined.
 1. Unprotected
 2. 1+1
 3. 1:8
 4. 1:16

- Additional parameters used by request authentication and verification, accounting and billing purposes. These might include cryptographic authentication of the request messages, account information, digital signatures, etc.

5. Deployment Considerations

5.1 Policy Control

Policy control refers to the set of administrative criteria used to decide whether a service request should be granted to a router [3]. These are beyond the regular resource considerations that have to be taken into account to route an optical path within the optical network. Policy control decisions can be based on rules that define conditions on parameters such as source and destination addresses, priorities, bilateral agreements among service providers, time-of-day constraints, cost constraints, etc. It is expected that, at least initially, policy control will rely on very simple rules, like, for example, approve all requests received from a given adjacent domain.

A policy decision module can be implemented locally on an OXC or remotely on a policy server. The decision will depend on the complexity of policies employed as well as the capabilities of the OXC. In the case of an external policy server, a standard protocol like COPS [4] should be used to pass requests from the OXC to the policy server.

5.2 Security Mechanisms

The most important security property from the internetworking perspective is *authentication* of control messages and service requests exchanged across the UNI. Authentication may initially be implicit, based on the interface from which a message is received. However, since an optical path consumes considerable network resources and has substantial cost, stronger security mechanisms should be put in place to safeguard an optical network against attacks on the control plane and unauthorized use of network resources.

One first step in enhancing security is to authenticate call requests received across the UNI. Authentication provides origin verification and non-repudiation, which is desirable for accounting and billing purposes; it is generally achieved using either symmetric authentication algorithms or public-key cryptography. Some of the commonly used IP signaling protocols, such as RSVP, have already introduced symmetric authentication. RSVP cryptographic authentication [5], requires pair-wise secret key configuration across adjacent networking devices, which increases management complexity and, since it is based on MD5, does not provide non-repudiation. Authentication via digital signatures, in the form of signing UNI messages, provides non-repudiation, but relies on the existence of a public key infrastructure. Public keys could also be negotiated in a pair-wise manner between the two parties across a UNI.

6. Summary and Conclusions

This contribution considered the UNI service definition and its relation to the optical internetworking model. It discussed the UNI design objectives and the information exchanged between an electrical device and a border OXC required to support different internetworking models. The signaling messages exchanged in order to perform address registration and discovery are described. In addition, deployment considerations related to policy control and security are discussed.

This contribution should be considered supplementary to document #2000.061.

References

1. J. Moy, "OSPF Version 2," RFC 1247, July, 1991.
2. Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, March, 1995.
3. R. Yavatkar, D. Pendarakis and R. Guerin, "A Framework for Policy-based Admission Control", RFC 2753, January 2000.
4. D. Durham, Ed., J. Boyle, R. Cohen, S. Herzog, R. Rajan and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.
5. F. Baker, B. Lindell and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.