INTERNET-DRAFT
Document: draft-ietf-ipo-carrier-requirements-05.txt
Category: Informational
Expiration Date: June, 2003
Yong Xue
(Editor)
WorldCom, Inc

December 2002

Optical Network Service Requirements

Status of This Memo

Abstract

This Internet Draft describes the major carrier's optical service
requirements for the Automatically Switched Optical Networks
(ASON)
from both an end-user's as well as an operator's perspectives. Its
focus is on the description of the service building blocks and
service-related control plane functional requirements. The
management
functions for the optical services and their underlying networks
are
beyond the scope of this document.

Table of Contents

1. Introduction

Optical transport networks are evolving from the current TDM-based SONET/SDH optical networks as defined by ANSI T1.105 and ITU Rec.

G.803 [ansi-sonet, itu-sdh] to emerging WDM-based optical transport networks (OTN) as defined by ITU Rec. G.872 in [itu-otn]. Therefore in

the near future, carrier optical transport networks are expected
to
consist of a mixture of the SONET/SDH-based sub-networks and the
WDM-
based wavelength or fiber switched OTN sub-networks. The OTN
networks
can be either transparent or opaque depending upon if O-E-O
functions
are utilized within the optical networks. Optical networking
encompasses the functionalities for the establishment,
transmission,
multiplexing and switching of optical connections carrying a wide
range of user signals of varying formats and bit rate. The optical
connections in this document include switched optical path using
TDM
channel, WADM wavelength or fiber links.

Some of the challenges for the carriers are efficient bandwidth
management and fast service provisioning in a multi-technology and
possibly multi-vendor networking environment. The emerging and
rapidly
evolving Automatically Switched Optical Network (ASON) technology
[itu-astn, itu-ason] is aimed at providing optical networks with
intelligent networking functions and capabilities in its control
plane
to enable rapid optical connection provisioning, dynamic rerouting
as
well as multiplexing and switching at different granularity
levels,
including
fiber, wavelength and TDM channel. The ASON control plane should
not
only enable the new networking functions and capabilities for the
emerging OTN networks, but significantly enhance the service
provisioning capabilities for the existing SONET/SDH networks as
well.

The ultimate goals should be to allow the carriers to automate
network
resource and topology discovery, to quickly and dynamically
provision
network resources and circuits, and to support assorted network
survivability using ring and mesh-based protection and restoration
techniques. The carriers see that this new networking platform
will
create tremendous business opportunities for the network operators
and
service providers to offer new services to the market, and in the
long

run to reduce their network operation cost (OpEx saving), and to improve their network utilization efficiency (CapEx saving).

1.1 Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT","SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL"
in
this document are to be interpreted as described in RFC 2119.

1.2 Value Statement

By deploying ASON technology, a carrier expects to achieve the following benefits from both technical and business perspectives:
Automated Discovery: ASON technology will enable automatic network inventory management, topology and resource discovery which
eliminates
the manual or semi-manual process for maintaining the network information database that exist in most carrier environment.

Rapid Circuit Provisioning: ASON technology will enable the dynamic end-to-end provisioning of the optical connections across the optical network by using standard routing and signaling protocols.

Enhanced Protection and Restoration: ASON technology will enable the network to dynamically reroute an optical connection in case of failure using mesh-based network protection and restoration techniques, which greatly improves the cost-effectiveness compared to the current line and ring protection schemes in the SONET/SDH network.

- Service Flexibility: ASON technology will support provisioning of an assortment of existing and new services such as protocol and bit-rate independent transparent network services, and bandwidth-on-demand services.

- Enhanced Interoperability: ASON technology will use a control plane utilizing industry and international standards-based architecture and protocols, which facilitate the interoperability of the optical network equipment from different vendors.

In addition, the ASON control plane may offer the following potential value-added benefits:

- Reactive traffic engineering at optical layer that allows network resources to be dynamically allocated to traffic flow.

- Reduce the need for service providers to develop new operational support systems (OSS) software for the network control and new service provisioning on the optical network, thus speeding up the deployment of the optical network technology and reducing the software development and maintenance cost.

- Potential development of a unified control plane that can be used for different transport technologies including OTN, SONET/SDH, ATM

and
PDH.

1.3. Scope of this document

This document is intended to provide, from the carriers perspective, a service framework and some associated requirements in relation to the optical transport services to be offered in the next generation optical transport networking environment and their service control and management functions. As such, this document concentrates on the requirements driving the work towards realization of automatically switched optical networks. This document is intended to be protocol-neutral, but the specific goals include providing the requirements to guide the control protocol development and enhancement within IETF in terms of reuse of IP-centric control protocols in the optical transport network.

Every carrier's needs are different. The objective of this document is NOT to define some specific service models. Instead, some major service building blocks are identified that will enable the carriers to use them in order to create the best service platform most suitable to their business model. These building blocks include generic service types, service enabling control mechanisms and service control and management functions.

The Optical Internetworking Forum (OIF) carrier group has developed a comprehensive set of control plane requirements for both UNI and NNI [oif-carrier, oif-nnireq] and they have been used as the base line input to this document.

The fundamental principles and basic set of requirements for the control plane of the automatic switched optical networks have been provided in a series of ITU Recommendations under the umbrella of ITU ASTN/ASON architectural and functional requirements as listed below:

Architecture:
- ITU-T Rec. G.8070/Y.1301 (2001), Requirements for the Automatic Switched Transport Network (ASTN)[itu-astn]

- ITU-T Rec. G.8080/Y.1304 (2001), Architecture of the Automatic Switched Optical Network (ASON)[itu-ason]

Signaling:
- ITU-T Rec. G.7713/Y.1704 (2001), Distributed Call and Connection Management (DCM)[itu-dcm]

Routing:
- ITU-T Draft Rec. G.7715/Y.1706 (2002), Architecture and Requirements for Routing in the Automatically Switched Optical Network [itu-rtg]

Discovery:
- ITU-T Rec. G.7714/Y.1705 (2001), Generalized Automatic Discovery [itu-disc]

Signaling Communication Network:

- ITU-T Rec. G.7712/Y.1703 (2001), Architecture and Specification
of
Data Communication Network [itu-dcn]

This document provides further detailed requirements based on the
ASTN/ASON framework. In addition, even though for IP over Optical
we
consider IP as a major client to the optical network in this
document,
the same requirements and principles should be equally applicable
to
non-IP clients such as SONET/SDH, ATM, ITU G.709, Ethernet, etc.
The
general architecture for IP over Optical is described in the IP
over
Optical framework document [ipo-frame]

2. Contributing Authors

This document was the combined effort of the editors and the following authors who contributed to this document:

Monica Lazer
Jennifer Yates
Dongmei Wang
Ananth Nagarajan
Hirokazu Ishimatsu
Olga Aparicio
Steven Wright

3. Acronyms

APON	ATM Passive Optical Network
ASON	Automatic Switched Optical Networking
ASTN	Automatic Switched Transport Network
CAC	Connection Admission Control
EPON	Ethernet Passive Optical Network
ESCON	Enterprise Storage Connectivity
FC	Fiber Channel
FICON	Fiber Connectivity
NNI	Node-to-Node Interface
UNI	User-to-Network Interface
I-NNI	Internal NNI
E-NNI	External NNI
NE	Network Element
OTN	Optical Transport Network
CNE	Customer/Client Network Element
ONE	Optical Network Element
OLS	Optical Line System
PI	Physical Interface
PDH	Plesiosynchronous Digital Hierarchy
CI	Control Interface
SLA	Service Level Agreement
SCN	Signaling Communication Network
SONET	Synchronous Digital Hierarchy
SDH	Synchronous Optical Network

4. General Requirements

In order to provide the carriers with flexibility and control of the optical networks, the following set of architectural requirements are essential.

4.1. Separation of Networking Functions

A fundamental architectural principle of the ASON network is to segregate the networking functions within each layer network into three logical functional planes: control plane, data plane and management plane. They are responsible for providing network control functions, data transmission functions and network management functions respectively. The crux of the ASON network is the networking intelligence that contains automatic routing, signaling and discovery functions to automate the network control functions.

Control Plane: includes the functions related to networking control capabilities such as routing, signaling, and policy control, as well as resource and service discovery. These functions are automated.

Data Plane (Transport Plane): includes the functions related to bearer channels and signal transmission.

Management Plane: includes the functions related to the management functions of network element, networks and network resources and services. These functions are less automated as compared to control plane functions.

Each plane consists of a set of interconnected functional or control entities, physical or logical, responsible for providing the networking or control functions defined for that network layer.

Each plane has clearly defined functional responsibilities. However, the management plane is responsible for the management of both control and data planes, thus playing an authoritative role in overall control and management functions as discussed in Section 9.

The separation of the control plane from both the data and management plane is beneficial to the carriers in that it:

- Allows equipment vendors to have a modular system design that will be more reliable and maintainable.

- Allows carriers to have the flexibility to choose a third party vendor control plane software systems as the control plane solution for its switched optical network.

- Allows carriers to deploy a unified control plane and OSS/ management systems to manage and control different types of transport networks it owns.

- Allows carriers to use a separate control network specially
designed

and engineered for the control plane communications.

The separation of control, management and transport function is required and it shall accommodate both logical and physical level separation. The logical separation refers to functional separation while physical separation refers to the case where the control, management and transport functions physically reside in different equipment or locations.

Note that it is in contrast to the IP network where the control messages and user traffic are routed and switched based on the
same

network topology due to the associated in-band signaling nature of the IP network.

When the physical separation is allowed between the control and data plane, a standardized interface and control protocol (e.g. GSMP [ietf-gsmp]) should be supported.

4.2. Separation of call and connection control

To support many enhanced optical services, such as scheduled bandwidth on demand, diverse circuit provisioning and bundled connections, a call model based on the separation of call control and connection control is essential.

The call control is responsible for the end-to-end session negotiation, call admission control and call state maintenance while connection control is responsible for setting up the connections associated with a call across the network. A call can correspond to zero, one or more connections depending upon the number of connections needed to support the call.

The existence of the connection depends upon the existence of its associated call session and connection can be deleted and re-established while still keeping the call session up.

The call control shall be provided at an ingress port or gateway port to the network such as UNI and E-NNI [see Section 6 for definition].
The connection control is provided at the originating node of the circuit as well as on each link along the path.

The control plane shall support the separation of the call control from the connection control.

The control plane shall support call admission control on call setup and connection admission control on connection setup.

4.3. Network and Service Scalability

Although some specific applications or networks may be on a small scale, the control plane protocol and functional capabilities

shall
support large-scale networks.

In terms of the scale and complexity of the future optical network,
the following assumption can be made when considering the scalability
and performance that are required of the optical control and management functions.

- There may be up to thousands of OXC nodes and the same or higher order of magnitude of OADMs per carrier network.

- There may be up to thousands of terminating ports/wavelength per OXC
node.

- There may be up to hundreds of parallel fibers between a pair of
OXC
nodes.

- There may be up to hundreds of wavelength channels transmitted
on
each fiber.

As for the frequency and duration of the optical connections:

- The expected end-to-end connection setup/teardown time should be
in
the order of seconds, preferably less.

- The expected connection holding times should be in the order of minutes or greater.

- There may be up to millions of simultaneous optical connections switched across a single carrier network.

4.4. Transport Network Technology

Optical services can be offered over different types of underlying optical transport technologies including both TDM-based SONET/SDH network and WDM-based OTN networks.

Standards-based transport technologies SONET/SDH as defined in the
ITU
Rec. G.803 and OTN implementation framing as defined in ITU Rec.
G.709
[itu-g709] shall be supported.

Note that the service characteristics such as bandwidth
granularity
and signaling framing hierarchy to a large degree will be
determined
by the capabilities and constraints of the server layer network.

4.5. Service Building Blocks

One of the goals of this document is to identify a set of basic service building blocks the carriers can use to create the best suitable service models that serve their business needs.

The service building blocks are comprised of a well-defined set of capabilities and a basic set of control and management functions. These capabilities and functions should support a basic set of services and enable a carrier to build enhanced services through extensions and customizations. Examples of the building blocks

include
the connection types, provisioning methods, control interfaces,
policy
control functions, and domain internetworking mechanisms, etc.

5. Service Model and Applications

A carrier's optical network supports multiple types of service
models.
Each service model may have its own service operations, target
markets, and service management requirements.

5.1. Service and Connection Types

The optical network is primarily offering optical paths that are fixed bandwidth connections between two client network elements, such as IP routers or ATM switches, established across the optical network. A connection is also defined by its demarcation from ingress access point, across the optical network, to egress access point of the optical network.

The following connection capability topologies must be supported:

- Bi-directional point-to-point connection

- Uni-directional point-to-point connection

- Uni-directional point-to-multipoint connection

The point-to-point connections are the primary concerns of the carriers. In this case, the following three types of network connections based on different connection set-up control methods shall be supported:

- Permanent connection (PC): Established hop-by-hop directly on each ONE along a specified path without relying on the network routing and signaling capability. The connection has two fixed end-points and fixed cross-connect configuration along the path and stays up until it is deleted. This is similar to the concept of PVC in ATM and there is no automatic re-routing capability.

- Switched connection (SC): Established through UNI signaling interface and the connection is dynamically established by network using the network routing and signaling functions. This is similar to the concept of SVC in ATM.

- Soft permanent connection (SPC): Established by specifying two PC at end-points and let the network dynamically establishes a SC connection in between. This is similar to the SPVC concept in ATM.

The PC and SPC connections should be provisioned via management plane

to control interface and the SC connection should be provisioned via
signaled UNI interface.

Note that even though automated rapid optical connection provisioning
is required, the carriers expect the majority of provisioned circuits,
at least in short term, to have a long lifespan ranging from months to
years.

In terms of service provisioning, some carriers may choose to perform
testing prior to turning over to the customer.

5.2. Examples of Common Service Models

Each carrier may define its own service model based on it business strategy and environment. The following are example service models that carriers may use.

5.2.1. Provisioned Bandwidth Service (PBS)

The PBS model provides enhanced leased/private line services provisioned via service management interface (MI) using either PC or

SPC type of connection. The provisioning can be real-time or near real-time. It has the following characteristics:

- Connection request goes through a well-defined management interface

- Client/Server relationship between clients and optical network.

- Clients have no optical network visibility and depend on network intelligence or operator for optical connection setup.

5.2.2. Bandwidth on Demand Service (BDS)

The BDS model provides bandwidth-on-demand dynamic connection services

via signaled user-network interface (UNI). The provisioning is real-

time and is using SC type of optical connection. It has the following

characteristics:

- Signaled connection request via UNI directly from the user or its

proxy.

- Customer has no or limited network visibility depending upon the control interconnection model used and network administrative policy.

- Relies on network or client intelligence for connection set-up depending upon the control plane interconnection model used.

5.2.3. Optical Virtual Private Network (OVPN)

The OVPN model provides virtual private network at the optical layer

between a specified set of user sites. It has the following characteristics:

- Customers contract for specific set of network resources such as optical connection ports, wavelengths, etc.

- Closed User Group (CUG) concept is supported as in normal VPN.

- Optical connection can be of PC, SPC or SC type depending upon the
provisioning method used.

- An OVPN site can request dynamic reconfiguration of the connections
between sites within the same CUG.

- A customer may have visibility and control of network resources up to the extent allowed by the customer service contract.

At a minimum, the PBS, BDS and OVPN service models described above shall be supported by the control functions.

6. Network Reference Model

This section discusses major architectural and functional components of a generic carrier optical network, which will provide a reference model for describing the requirements for the control and management of carrier optical services.

6.1. Optical Networks and Sub-networks

As mentioned before, there are two main types of optical networks that are currently under consideration: SDH/SONET network as defined in ITU Rec. G.803, and OTN as defined in ITU Rec. G.872.

In the current SONET/SDH-based optical network, digital cross-connects (DXC) and add-drop multiplexer (ADM) and line multiplexer terminal (LMT) are connected in ring or linear topology. Similarly, we assume an OTN is composed of a set of optical cross-connects (OXC) and optical add-drop multiplexer (OADM) which is interconnected in a general mesh topology using DWDM optical line systems (OLS).

It is often convenient for easy discussion and description to treat an optical network as an sub-network cloud, in which the details of the network become less important, instead focus is on the function and the interfaces the optical network provides. In general, a subnetwork can be defined as a set of access points on the network boundary and a set of point-to-point optical connections between those access points.

6.2. Control Domains and Interfaces

A generic carrier network reference model describes a multi-

carrier
network environment. Each individual carrier network can be
further
partitioned into sub-networks or administrative domains based on
administrative, technological or architectural reasons. This
partition
can be recursive. Similarly, a network can be partitioned into
control
domains that match the administrative domains and are controlled
by a
single administrative policy. The control domains can be
recursively
divided into sub-domains to form control hierarchy for
scalability.
The control domain concept can be applied to routing, signaling
and
protection & restoration to form an autonomous control function
domain.

The demarcation between domains can be either logical or physical
and
consists of a set of reference points identifiable in the optical
network. From the control plane perspective, these reference
points
define a set of control interfaces in terms of optical control and
management functionality as illustrated in Figure 1.

```
                    +----------------------------------------+
                    |    Single carrier network              |
 +------------+     |                                        |
 |Customer    |     |  +------------+        +------------+  |
 |IP          |     |  |            |        |            |  |
 |Network     +-UNI|-  +  Optical   +--UNI--+CarrierÆs IP|  |
 |            |     |  | Subnetwork |        |  network   |  |
 +------------+     |  | (Domain A) +--+     |            |  |
                    |  +------+-----+  |     +------+-----+  |
                    |         |        |            |        |
                    |       I-NNI    E-NNI         UNI       |
 +------------+     |         |        |            |        |
 |Customer    |     |  +------+-----+  |     +------+-----+  |
 |IP          +-UNI|-  +            |  +-----+            |  |
 |Network     |     |  | Optical    |        |  Optical   |  |
 |            |     |  | Subnetwork +-E-NNI-+ Subnetwork |  |
 +------------+     |  | (Domain A) |        | (Domain B) |  |
                    |  +------+-----+        +------+-----+  |
                    |         |                     |        |
                    +----------------------------------------+
                             UNI                  E-NNI
                              |                     |
                       +------+-------+    +-------+--------+
                       |              |    |                |
                       | Other Client |    |  Other Carrier |
                       |Network       |    | Network        |
                       | (ATM/SONET)  |    |                |
                       +--------------+    +----------------+
```

Figure 1 Generic Carrier Network Reference Model

The network interfaces encompass two aspects of the networking functions: user data plane interface and control plane interface. The

former concerns about user data transmission across the physical network interface and the latter concerns about the control message

exchange across the network interface such as signaling, routing, etc.

We call the former physical interface (PI) and the latter control interface (CI). Unless otherwise stated, the CI is assumed in the remaining of this document.

6.2.1. Control Plane Interfaces

The Control Interface defines the relationship between two connected

network entities on both sides of the interface. For each control

interface, we need to define the architectural function that each
side
plays and a controlled set of information that can be exchanged
across
the interface. The information flowing over this logical interface
may
include, but not limited to:
- Interface endpoint name and address

- Reachability/summarized network address information

- Topology/routing information

- Authentication and connection admission control information

- Connection management signaling messages

- Network resource control information

Different types of interfaces can be defined for network control and architectural purposes and can be used as the network reference points in the control plane. In this document, the following set of interfaces is defined as shown in Figure 1.

User-Network Interface (UNI): is a bi-directional control interface between service requester and service provider control entities. The service request control entity resides outside the carrier network control domain.

Network-Network/Node-Node Interface (NNI): is a bi-directional signaling interface between two optical network elements or sub-networks.

We differentiate between internal NNI (I-NNI) and external NNI (E-NNI) as follows:

- E-NNI: A NNI between two control plane entities belonging to different control domains.

- I-NNI: A NNI between two control plane entities within the same control domain in the carrier network.

Different types of interface, internal vs. external, have different implied trust relationship for security and access control purposes. The trust relationship is not binary. Instead a policy-based control mechanism need to be in place to restrict the type and amount of information that can flow cross each type of interfaces depending on the carrier's service and business requirements.

Generally, two networks have a fully trusted relationship if they belong to the same administrative domain. In this case, the control information exchanged across the control interface between them should be unlimited. Otherwise, the type and amount of the control information that can go across the information should be constrained by the administrative policy.

An example of fully trusted interface is an I-NNI between two optical network elements in a single control domain. Non-trusted interface examples include an E-NNI between two different carriers or a UNI interface between a carrier optical network and its customers. The trust level can be different for the non-trusted UNI or E-NNI

interface depending upon if it within the carrier or not. In general,
intra-carrier E-NNI has higher trust level than inter-carrier E-NNI.

The control plane shall support the UNI and NNI interface described
above and the interfaces shall be configurable in terms of the type
and amount of control information exchange and their behavior shall be
consistent with the configuration (i.e., external versus internal interfaces).

6.3. Intra-Carrier Network Model

Intra-carrier network model concerns the network service control and
management issues within networks owned by a single carrier.

6.3.1. Multiple Sub-networks

Without loss of generality, the optical network owned by a carrier service operator can be depicted as consisting of one or more optical
sub-networks interconnected by direct optical links. There may be many
different reasons for more than one optical sub-network. It may be the
result of using hierarchical layering, different technologies across
access, metro and long haul (as discussed below), or a result of business mergers and acquisitions or incremental optical network technology deployment by the carrier using different vendors or technologies.

A sub-network may be a single vendor and single technology network.
But in general, the carrier's optical network is heterogeneous in terms of equipment vendor and the technology utilized in each sub-network.

6.3.2. Access, Metro and Long-haul networks

Few carriers have end-to-end ownership of the optical networks. Even
if they do, access, metro and long-haul networks often belong to different administrative divisions as separate optical sub-networks.

Therefore Inter-(sub)-networks interconnection is essential in terms of supporting the end-to-end optical service provisioning and management. The access, metro and long-haul networks may use different technologies and architectures, and as such may have different network properties.

In general, end-to-end optical connectivity may easily cross multiple sub-networks with the following possible scenarios:

Access -- Metro -- Access
Access - Metro -- Long Haul -- Metro - Access

6.4. Inter-Carrier Network Model

The inter-carrier model focuses on the service and control aspects between different carrier networks and describes the internetworking relationship between them. The inter-carrier connection is often not only constrained technical and business requirements, but by the government regulations as well,

Inter-carrier interconnection provides for connectivity between optical network operators. To provide globally reachable end-to-end optical services, optical service control and management between different carrier networks becomes essential. For example, it is possible to support distributed peering within the IP client layer network where the connectivity between two distant IP routers can be achieved via an inter-carrier optical transport connection.

6.5. Implied Control Constraints

The intra-carrier and inter-carrier models have different implied control constraints. For example, in the intra-carrier model, the address for routing and signaling only need to be unique with the carrier while the inter-carrier model requires the address to be globally unique.

In the intra-carrier network model, the network itself forms the largest control domain within the carrier network. This domain is usually partitioned into multiple sub-domains, either flat or hierarchical. The UNI and E-NNI interfaces are internal to the carrier network, therefore higher trust level is assumed. Because of this, direct signaling between domains and summarized topology and resource information exchanged can be allowed across the internal UNI or intra-carrier E-NNI interfaces.

In the inter-carrier network model, each carrier's optical network is a separate administrative domain. Both the UNI interface between the user and the carrier network and the NNI interface between two carrier's networks are crossing the carrier's administrative boundary and therefore are external interfaces by definition.

In terms of control information exchange, the topology information shall not be allowed to cross both E-NNI and UNI interfaces.

7. Optical Service User Requirements

This section describes the user requirements for optical services, which in turn impose the requirements on service control and management for the network operators. The user requirements reflect the perception of the optical service from a user's point of view.

7.1. Common Optical Services

The basic unit of an optical transport service is fixed-bandwidth optical connectivity between applications. However different services are created based on its supported signal characteristics (format, bit rate, etc), the service invocation methods and possibly the associated Service Level Agreement (SLA) provided by the service provider.

At present, the following are the major optical services provided in the industry:

- SONET/SDH, with different degrees of transparency
- Optical wavelength services, transparent or opaque
- Ethernet at 10Mbps, 100Mbps, 1 Gbps and 10 Gbps
- Storage Area Networks (SANs) based on Fiber Connectivity (FICON), Enterprise Storage Connectivity (ESCON) and Fiber Channel (FC).

Optical Wavelength Service refers to transport services where signal framing is negotiated between the client and the network operator (framing and bit-rate dependent), and only the payload is carried transparently. SONET/SDH transport is most widely used for network-wide transport. Different levels of transparency can be achieved in the SONET/SDH transmission.

Ethernet Services, specifically 1Gb/s and 10Gbs Ethernet services, are gaining more popularity due to the lower costs of the customers' premises equipment and its simplified management requirements (compared to SONET or SDH).

Ethernet services may be carried over either SONET/SDH (GFP mapping) or WDM networks. The Ethernet service requests will require some service specific parameters: priority class, VLAN ID/Tag, traffic aggregation parameters.

ESCON and FICON are proprietary versions of the SAN service, while Fiber Channel is the standard alternative. As is the case with Ethernet services, SAN services may be carried over either SONET/SDH (using GFP mapping) or WDM networks.

The control plane shall provide the carrier with the capability functionality to provision, control and manage all the services listed above.

7.2. Bearer Interface Types

All the bearer interfaces implemented in the ONE shall be supported by the control plane and associated signaling protocols.

The signaling shall support the following interface types protocol:

- SDH/SONET
- Ethernet
- FC-N for Fiber Channel services
- OTN (G.709)
- PDH (Plesiosynchronous Digital Hierarchy)

- Passive Optical Network (PON) based on ATM (APON) and Ethernet (EPON)
- ESCON and FICON

7.3. Optical Service Invocation
As mentioned earlier, the methods of service invocation play an important role in defining different services.

7.3.1. Provider-Initiated Service Provisioning

In this scenario, users forward their service request to the provider via a well-defined service management interface. All connection management operations, including set-up, release, query, or modification shall be invoked from the management plane. This provisioning method is for PC and SPC connections.

7.3.2. User-Initiated Service Provisioning

In this scenario, users forward their service request to the provider via a well-defined UNI interface in the control plane (including proxy signaling). All connection management operation requests, including set-up, release, query, or modification shall be invoked from directly connected user devices, or its signaling proxy. This provisioning method is for SC connection.

7.3.3. Call set-up requirements
In summary the following requirements for the control plane have been identified:
- The control plane shall support action result codes as responses to any requests over the control interfaces.

- The control plane shall support requests for call set-up, subject to policies in effect between the user and the network.

- The control plane shall support the destination client device's decision to accept or reject call set-up requests from the source client's device.

- The control plane shall support requests for call set-up and deletion across multiple (sub)networks.

- NNI signaling shall support requests for call set-up, subject to policies in effect between the (sub)networks.

- Call set-up shall be supported for both uni-directional and bi-directional connections.

- Upon call request initiation, the control plane shall generate a network unique Call-ID associated with the connection, to be used for information retrieval or other activities related to that connection.

- CAC shall be provided as part of the call control functionality. It
is the role of the CAC function to determine if the call can be
allowed to proceed based on resource availability and authentication.

- Negotiation for call set-up for multiple service level options shall
be supported.

- The policy management system must determine what kinds of call setup
requests can be authorized.

- The control plane elements need the ability to rate limit (or pace)
call setup attempts into the network.

- The control plane shall report to the management plane, the
success/failures of a call request.

- Upon a connection request failure, the control plane shall report to
the management plane a cause code identifying the reason for the
failure and all allocated resources shall be released. A negative
acknowledgment shall be returned to the source.

- Upon a connection request success a positive acknowledgment shall be
returned to the source when a connection has been successfully
established.

- The control plane shall support requests for call release by Call-
ID.

- The control plane shall allow any end point or any intermediate node
to initiate call release procedures.

- Upon call release completion all resources associated with the call
shall become available for access for new requests.

- The management plane shall be able to release calls or connections
established by the control plane both gracefully and forcibly on
demand.

- Partially deleted calls or connections shall not remain within

the
network.

- End-to-end acknowledgments shall be used for connection deletion requests.

- Connection deletion shall not result in either restoration or protection being initiated.

- The control plane shall support management plane and neighboring device requests for status query.

- The UNI shall support initial registration and updates of the client with the network via the control plane.

7.4. Optical Connection granularity

The service granularity is determined by the specific technology, framing and bit rate of the physical interface between the ONE and the client at the edge and by the capabilities of the ONE. The control plane needs to support signaling and routing for all the services supported by the ONE. In general, there should not be a one-to-one correspondence imposed between the granularity of the service provided and the maximum capacity of the interface to the user.

The control plane shall support the ITU Rec. G.709 connection granularity for the OTN network.

The control plane shall support the SDH/SONET connection granularity.

The optical control plane shall support sub-rate interfaces such as VT /TU granularity (as low as 1.5 Mb/s).

The following fiber channel interfaces shall be supported by the control plane if the given interfaces are available on the equipment:

- FC-12
- FC-50
- FC-100
- FC-200

Encoding of service types in the protocols used shall be such that new service types can be added by adding new code point values or objects.

7.5. Other Service Parameters and Requirements

7.5.1 Classes of Service

We use "service level" to describe priority related characteristics of connections, such as holding priority, set-up priority, or restoration priority. The intent currently is to allow each carrier to define the actual service level in terms of priority, protection, and restoration options. Therefore, individual carriers will determine mapping of individual service levels to a specific set of quality features.

The control plane shall be capable of mapping individual service classes into specific priority or protection and restoration options.

7.5.2. Diverse Routing Attributes

Diversity refers to the fact that a disjoint set of network resources (links and nodes) is utilized to provision multiple parallel optical connections terminated between a pair of ingress and egress ports.

There are different levels of diversity based on link, node or administrative policy as described below. In the simple node and link

diversity case:
- Two optical connections are said to be node-disjoint diverse, if the

two connections do not share any node along the path except the ingress and egress nodes.
- Two optical connections are said to be link-disjoint diverse, if the

two connections do not share any link along the path.

A more general concept of diversity is the Shared Risk Group (SRG) that is based on a risk-sharing model and allows the definition of administrative policy-based diversity. A SRG is defined as a group of

links or nodes that share a common risk component, whose failure can

potentially cause the failure of all the links or nodes in the group.

When the SRG is applied to the link resource, it is referred to as shared risk link group (SRLG). For example, all fiber links that go

through a common conduit under the ground belong to the same SRLG group, because the conduit is a shared risk component whose failure,

such as a cut, may cause all fibers in the conduit to break. Note that

SRLG is a relation defined within a group of links based upon a specific risk factor that can be defined based on various technical or

administrative grounds such as ôsharing a conduitö, ôwithin 10 miles

of distance proximityö etc. Please see ITU-T G.7715 for more discussion [itu-rtg].

Therefore, two optical connections are said to be SRG-disjoint diverse

if the two connections do not have any links or nodes that belong to

the same SRG along the path.

The ability to route service paths diversely is a required control feature. Diverse routing is one of the connection parameters and is

specified at the time of the connection creation.

The control plane routing algorithms shall be able to route an optical

connection diversely from a previously routed connection in terms
of
link disjoint path, node disjoint path and SRG disjoint path.

8. Optical Service Provider Requirements

This section discusses specific service control and management requirements from the service provider's point of view.

8.1. Service Access Methods to Optical Networks

In order to have access to the optical network service, a customer needs to be physically connected to the service provider network
on
the transport plane. The control plane connection may or may not
be
required depending upon the service invocation model provided to
the
customer: provisioned vs. signaled. For the signaled, either
direct or
indirect signaling methods can be used depending upon if the UNI
proxy

is utilized on the client side. The detailed discussion on the UNI signaling methods is in [oif-uni].

Multiple access methods blow shall be supported:

- Cross-office access (CNE co-located with ONE)

- Direct remote access (Dedicated links to the user)

- Remote access via access sub-network (via a multiplexing/distribution sub-network)

8.2. Dual Homing and Network Interconnections

Dual homing is a special case of the access network. Client devices can be dual homed to the same or different hub, the same or different access network, the same or different core networks, the same or different carriers. The different levels of dual homing connectivity result in many different combinations of configurations. The main objective for dual homing is for enhanced survivability.

Dual homing must be supported. Dual homing shall not require the use of multiple addresses for the same client device.

8.3. Inter-domain connectivity

A domain is a portion of a network, or an entire network that is controlled by a single control plane entity. This section discusses the various requirements for connecting domains.

8.3.1. Multi-Level Hierarchy

Traditionally current transport networks are divided into core inter-city long haul networks, regional intra-city metro networks and access networks. Due to the differences in transmission technologies, service, and multiplexing needs, the three types of networks are served by different types of network elements and often have different capabilities. The network hierarchy is usually implemented through the control domain hierarchy.

When control domains exists for routing and signaling purpose,

there
will be intra-domain routing/signaling and inter-domain
routing/signaling. In general, domain-based routing/signaling
autonomy
is desired and the intra-domain routing/signaling and the inter-
domain
routing/signaling should be agnostic to each other.

Routing and signaling for multi-level hierarchies shall be
supported
to allow carriers to configure their networks as needed.

8.3.2. Network Interconnections

Sub-networks may have multiple points of inter-connections. All relevant NNI functions, such as routing, reachability information exchanges, and inter-connection topology discovery must recognize and support multiple points of inter-connections between subnetworks. Dual inter-connection is often used as a survivable architecture.

The control plane shall provide support for routing and signaling for subnetworks having multiple points of interconnection.

8.4. Names and Address Management

8.4.1. Address Space Separation

To ensure the scalability of and smooth migration toward to the optical switched network, the separation of three address spaces are required as discussed in [oif-addr]:

- Internal transport network addresses: This is used for routing control plane messages within the transport network. For example, if GMPLS is used then IP address should be used.

- Transport Network Assigned (TNA) address: This is a routable address in the optical transport network and is assigned by the network.

- Client addresses: This address has significance in the client layer. For example, if the clients are ATM switches, the NSAP address can be used. If the clients are IP router, then IP address should be used.

8.4.2. Directory Services

Directory Services shall support address resolution and translation between various user/client device names or address and the corresponding TNA addresses. UNI shall use the user naming schemes for connection request. The directory service is essential for the implementation of overlay model.

8.4.3. Network element Identification

Each control domain and each network element within a carrier
network

shall be uniquely identifiable. Similarly all the service access
points shall be uniquely identifiable.

8.5. Policy-Based Service Management Framework

The optical service must be supported by a robust policy-based
management system to be able to make important decisions.

Examples of policy decisions include:
- What types of connections can be set up for a given UNI?

- What information can be shared and what information must be restricted in automatic discovery functions?

- What are the security policies over signaling interfaces?

- What routing policies should be applied in the path selection?
E.g

The definition of the link diversity.

Requirements:
- Service and network policies related to configuration and provisioning, admission control, and support of Service Level Agreements (SLAs) must be flexible, and at the same time simple
and

scalable.

- The policy-based management framework must be based on
standards-

based policy systems (e.g., IETF COPS [rfc2784]).

- In addition, the IPO service management system must support and
be

backwards compatible with legacy service management systems.

9. Control Plane Functional Requirements for Optical Services
This section addresses the requirements for the optical control
plane

in support of service provisioning.

The scope of the control plane includes the control of the
interfaces

and network resources within an optical network and the interfaces between the optical network and its client networks. In other
words,

it should include both NNI and UNI aspects.

9.1. Control Plane Capabilities and Functions

The control capabilities are supported by the underlying control functions and protocols built in the control plane.

9.1.1. Network Control Capabilities

The following capabilities are required in the network control
plane

to successfully deliver automated provisioning for optical
services:

- Network resource discovery

- Address assignment and resolution

- Routing information propagation and dissemination

- Path calculation and selection

- Connection management

These capabilities may be supported by a combination of functions across the control and the management planes.

9.1.2. Control Plane Functions for Network Control

The following are essential functions needed to support network control capabilities:
- Signaling
- Routing
- Automatic resource, service and neighbor discovery

Specific requirements for signaling, routing and discovery are addressed in Section 10.

The general requirements for the control plane functions to support optical networking and service functions include:

- The control plane must have the capability to establish, teardown and maintain the end-to-end connection, and the hop-by-hop connection segments between any two end-points.

- The control plane must have the capability to support optical traffic-engineering (e.g. wavelength management) requirements including resource discovery and dissemination, constraint-based routing and path computation.

- The control plane shall support network status or action result code responses to any requests over the control interfaces.

- The control plane shall support call admission control on UNI and connection-admission control on NNI.

- The control plane shall support graceful release of network resources associated with the connection after a successful connection teardown or failed connection.

- The control plane shall support management plane request for connection attributes/status query.

- The control plane must have the capability to support various protection and restoration schemes.

- Control plane failures shall not affect active connections and shall not adversely impact the transport and data planes.

- The control plane should support separation of control function

entities including routing, signaling and discovery and should allow
different control distributions of those functionalities, including
centralized, distributed or hybrid.

- The control plane should support physical separation of the control
plane from the transport plane to support either tightly coupled or
loosely coupled control plane solutions.

- The control plane should support the routing and signaling proxy to participate in the normal routing and signaling message exchange and processing.

- Resilience and security are crucial issues for the control plane and will be addressed in Section 11 and 12 of this document respectively.

9.2. Signaling Communication Network (SCN)

The signaling communication network is a transport network for control plane messages and it consists of a set of control channels that interconnects the nodes within the control plane. Therefore, the signaling communication network must be accessible by each of the communicating nodes (e.g., OXCs). If an out-of-band IP-based control message transport network is an overlay network built on top of the IP data network using some tunneling technologies, these tunnels must be standards-based such as IPSec, GRE, etc.

- The signaling communication network must terminate at each of the nodes in the transport plane.

- The signaling communication network shall not be assumed to have the same topology as the data plane, nor shall the data plane and control plane traffic be assumed to be congruently routed.

A control channel is the communication path for transporting control messages between network nodes, and over the UNI (i.e., between the UNI entity on the user side and the UNI entity on the network side).
The control messages include signaling messages, routing information messages, and other control maintenance protocol messages such as neighbor and service discovery.

The following three types of signaling in the control channel shall be supported:

- In-band signaling: The signaling messages are carried over a
logical
communication channel embedded in the data-carrying optical link
or
channel. For example, using the overhead bytes in SONET data
framing
as a logical communication channel falls into the in-band
signaling
methods.

- In fiber, Out-of-band signaling: The signaling messages are
carried
over a dedicated communication channel separate from the optical
data-
bearing channels, but within the same fiber. For example, a
dedicated
wavelength or TDM channel may be used within the same fiber as the
data channels.

- Out-of-fiber signaling: The signaling messages are carried over
a
dedicated communication channel or path within different fibers to
those used by the optical data-bearing channels. For example,
dedicated optical fiber links or communication path via separate
and

independent IP-based network infrastructure are both classified as out-of-fiber signaling.

The UNI control channel and proxy signaling defined in the OIF UNI 1.0 [oif-uni] shall be supported.

The signaling communication network provides communication mechanisms between entities in the control plane.

- The signaling communication network shall support reliable message transfer.

- The signaling communication network shall have its own OAM mechanisms.

- The signaling communication network shall use protocols that support congestion control mechanisms.

In addition, the signaling communication network should support message priorities. Message prioritization allows time critical messages, such as those used for restoration, to have priority over other messages, such as other connection signaling messages and topology and resource discovery messages.

The signaling communication network shall be highly reliable and implement failure recovery.

9.3 Control Plane Interface to Data Plane

In the situation where the control plane and data plane are decoupled, this interface needs to be standardized. Requirements for a standard control-data plane interface are under study. The specification of a control plane interface to the data plane is outside the scope of this document.

Control plane should support a standards based interface to configure switching fabrics and port functions via the management plane.

Data plane shall monitor and detect the failure (LOL, LOS, etc.) and quality degradation (high BER, etc.) of the signals and be able to provide signal-failure and signal-degrade alarms to the control plane accordingly to trigger proper mitigation actions in the control plane.

9.4. Management Plane Interface to Data Plane

The management plane shall be responsible for the network resource management in the data plane. It should be able to partition the network resources and control the allocation and the deallocation of the resource for use by the control plane.

Data plane shall monitor and detect the failure and quality
degradation of the signals and be able to provide signal-failure
and

signal-degrade alarms plus associated detailed fault information
to

the management plane to trigger and enable the management for
fault

location and repair.

Management plane failures shall not affect the normal operation of
a

configured and operational control plane or data plane.

9.5. Control Plane Interface to Management Plane

The control plane is considered a managed entity within a network.
Therefore, it is subject to management requirements just as other
managed entities in the network are subject to such requirements.

The control plane should be able to service the requests from the
management plane for end-to-end connection provisioning (e.g. SPC
connection) and control plane database information query (e.g.
topology database)

The control plane shall report all control plane faults to the
management plane with detailed fault information

The control, management and transport plane each has its well-
defined

network functions. Those functions are orthogonal to each other.
However, this does not imply total independency. Since the
management

plane is responsible for the management of both control plane and
transport plane, the management plane plays an authoritative role

In general, the management plane shall have authority over the
control

plane. Management plane should be able to configure the routing,
signaling and discovery control parameters such as hold-down
timers,

hello-interval, etc. to affect the behavior of the control plane.

In the case of network failure, both the management plane and the
control plane need fault information at the same priority. The
control

plane shall be responsible for providing necessary statistic data
such

as call counts and traffic stats to the management plane. They
should

be available upon query from the management plane. The management plane shall be able to tear down connections established by the control plane both gracefully and forcibly on demand.

9.6. IP and Optical Control Plane Interconnection

The control plane interconnection model defines how two control networks can be interconnected in terms of controlling relationship and control information flow allowed between them. There are three basic types of control plane network interconnection models: overlay, peer and hybrid, which are defined in the IETF IPO WG document [ipo-frame]. See Appendix A for more discussion.

Choosing the level of coupling depends upon a number of different factors, some of which are:

- Variety of clients using the optical network

- Relationship between the client and optical network

- Operating model of the carrier

Overlay model (UNI like model) shall be supported for client to optical control plane interconnection.

Other models are optional for client to optical control plane interconnection.

For optical to optical control plane interconnection all three models shall be supported. In general, the priority for support of interconnection models should be overlay, hybrid and peer, in decreasing order.

10. Requirements for Signaling, Routing and Discovery

10.1. Requirements for information sharing over UNI, I-NNI and E-NNI

Different types of interfaces shall impose different requirements and functionality due to their different trust relationships. Specifically:

- Topology information shall not be exchanged across inter-carrier E-NNI and UNI.

- The control plane shall allow the carrier to configure the type and extent of control information exchange across various interfaces.

- Address resolution exchange over UNI is needed if an addressing directory service is not available.

10.2. Signaling Functions

Call and connection control and management signaling messages are used for the establishment, modification, status query and release of an end-to-end optical connection. Unless otherwise specified, the word "signaling" refers to both inter-domain and intra-domain signaling.

- The inter-domain signaling protocol shall be agnostic to the intra-domain signaling protocol for all the domains within the network.

- Signaling shall support both strict and loose routing.

- Signaling shall support individual as well as groups of connection requests.

- Signaling shall support fault notifications.

- Inter-domain signaling shall support per connection, globally unique identifiers for all connection management primitives based on a well-defined naming scheme.

- Inter-domain signaling shall support crank-back and rerouting.

10.3. Routing Functions

Routing includes reachability information propagation, network topology/resource information dissemination and path computation. Network topology/resource information dissemination is to provide each node in the network with information about the carrier network such that a single node is able to support constraint-based path selection.

A mixture of hop-by-hop routing, explicit/source routing and hierarchical routing will likely be used within future transport networks.

All three mechanisms (Hop-by-hop routing, explicit / source-based routing and hierarchical routing) must be supported. Messages crossing untrusted boundaries must not contain information regarding the details of an internal network topology.

Requirements for routing information dissemination:
- The inter-domain routing protocol shall be agnostic to the intra-domain routing protocol within any of the domains within the network.

- The exchange of the following types of information shall be supported by inter-domain routing protocols:
- Inter-domain topology
- Per-domain topology abstraction
- Per domain reachability summarization

Major concerns for routing protocol performance are scalability and stability, which impose the following requirement on the routing protocols:
- The routing protocol shall scale with the size of the network

The routing protocols shall support following requirements:

- Routing protocol shall support hierarchical routing information dissemination, including topology information aggregation and summarization.

- The routing protocol(s) shall minimize global information and keep information locally significant as much as possible. Over external interfaces only reachability information, next routing hop and service capability information should be exchanged. Any other network related information shall not leak out to other networks.

- The routing protocol shall be able to minimize global information and keep information locally significant as much as possible (e.g., information local to a node, a sub-network, a domain, etc). For example, a single optical node may have thousands of ports. The ports with common characteristics need not to be advertised individually.
- The routing protocol shall distinguish static routing information and dynamic routing information. The routing protocol operation shall update dynamic and static routing information differently. Only dynamic routing information shall be updated in real time.

- Routing protocol shall be able to control the dynamic information updating frequency through different types of thresholds. Two types of thresholds could be defined: absolute threshold and relative threshold.

- The routing protocol shall support trigger-based and timeout-based information update.

- Inter-domain routing protocol shall support policy-based routing information exchange.

- The routing protocol shall be able to support different levels of protection/restoration and other resiliency requirements. These are discussed in Section 11.

All the scalability techniques will impact the network resource representation accuracy. The tradeoff between accuracy of the routing information and the routing protocol scalability is an important consideration to be made by network operators.

10.4. Requirements for path selection

The following are functional requirements for path selection:
- Path selection shall support shortest path routing.

- Path selection shall also support constraint-based routing. At least

the following constraints shall be supported:
- Cost
- Link utilization
- Diversity
- Service Class

- Path selection shall be able to include/exclude some specific network resources, based on policy.

- Path selection shall be able to support different levels of diversity, including node, link, SRLG and SRG.

- Path selection algorithms shall provide carriers the ability to support a wide range of services and multiple levels of service classes. Parameters such as service type, transparency, bandwidth, latency, bit error rate, etc. may be relevant.

Constraint-based routing in the optical network in significantly complex Compared to the IP network. There are many optical layer constraints to consider such as wavelength, diversity, optical layer impairments, etc. A detailed discussion on the routing constraints at the optical layer is in [ietf-olr].

10.5. Discovery Functions

The discovery functions include neighbor, resource and service discovery. The control plane shall support both manual configuration and automatic discovery

10.5.1. Neighbor discovery

Neighbor Discovery can be described as an instance of auto-discovery that is used for associating two network entities within a layer network based on a specified adjacency relation.

The control plane shall support the following neighbor discovery capabilities as described in [itu-disc]:

- Physical media adjacency that detects and verifies the physical layer network connectivity between two connected network element ports.

- Logical network adjacency that detects and verifies the logical network layer connection above the physical layer between network layer specific ports.

- Control adjacency that detects and verifies the logical neighboring relation between two control entities associated with data plane network elements that form either physical or logical adjacency.

The control plane shall support manual neighbor adjacency configuration to either overwrite or supplement the automatic neighbor discovery function.

10.5.2. Resource Discovery

Resource discovery is concerned with the ability to verify physical connectivity between two ports on adjacent network elements, improve inventory management of network resources, detect configuration mismatches between adjacent ports, associating port characteristics of adjacent network elements, etc. Resource discovery shall be supported.

Resource discovery can be achieved through either manual provisioning or automated procedures. The procedures are generic while the specific mechanisms and control information can be technology dependent.

After neighbor discovery, resource verification and monitoring must be performed periodically to verify physical attributes to ensure compatibility.

10.5.3. Service Discovery

Service Discovery can be described as an instance of auto-discovery that is used for verifying and exchanging service capabilities of a network. Service discovery can only happen after neighbor discovery.

Since service capabilities of a network can dynamically change, service discovery may need to be repeated. Service discovery is required for all the optical services supported.

11. Requirements for service and control plane resiliency

Resiliency is a network capability to continue its operations under the condition of failures within the network. The automatic switched optical network assumes the separation of control plane and data plane. Therefore the failures in the network can be divided into those affecting the data plane and those affecting the control plane. To provide enhanced optical services, resiliency measures in both data plane and control plane should be implemented. The following failure-handling principles shall be supported.

The control plane shall provide optical service failure detection and recovery functions such that the failures in the data plane within the control plane coverage can be quickly mitigated.

The failure of control plane shall not in any way adversely affect the normal functioning of existing optical connections in the data plane.

In general, there shall be no single point of failure for all major control plane functions, including signaling, routing etc. The control plane shall provide reliable transfer of signaling messages and

flow
control mechanisms for easing any congestion within the control
plane.

11.1. Service resiliency

In circuit-switched transport networks, the quality and
reliability of
the established optical connections in the transport plane can be
enhanced by the protection and restoration mechanisms provided by
the
control plane functions. Rapid recovery is required by transport
network providers to protect service and also to support stringent
Service Level Agreements (SLAs) that dictate high reliability and
availability for customer connectivity.

The protection and restoration actions are usually in reaction to
the
failure in the networks. However, during the network maintenance
affecting the protected connections, a network operator needs to
proactively force the traffic on the protected connections to
switch
to its protection connection. Therefore in order to support easy

network maintenance, it is required that management initiated protection and restoration be supported.

The failure and signal degradation in the transport plane is usually technology specific and therefore shall be monitored and detected by the transport plane.

The transport plane shall report both physical level failure and signal degradation to the control plane in the form of the signal failure alarm and signal degrade alarm.

The control plane shall support both alarm-triggered and hold-down timers based protection switching and dynamic restoration for failure recovery.

Clients will have different requirements for connection availability. These requirements can be expressed in terms of the "service level", which can be mapped to different restoration and protection options and priority related connection characteristics, such as holding priority (e.g. pre-emptable or not), set-up priority, or restoration priority. However, how the mapping of individual service levels to a specific set of protection/restoration options and individual carriers will determine connection priorities.

In order for the network to support multiple grades of service, the control plane must support differing protection and restoration options on a per connection basis.

In order for the network to support multiple grades of service, the control plane must support setup priority, restoration priority and holding priority on a per connection basis.

In general, the following protection schemes shall be considered for all protection cases within the network:
- Dedicated protection: 1+1 and 1:1
- Shared protection: 1:N and M:N.

- Unprotected

The control plane shall support "extra-traffic" capability, which allows unprotected traffic to be transmitted on the protection circuit.

The control plane shall support both trunk-side and drop-side protection switching.

The following restoration schemes should be supported:
- Restorable
- Un-restorable

Protection and restoration shall be supported on both an end-to-end

basis and a link-by-link basis.

Protection and restoration configuration should be based on
software

only.

The control plane shall allow the modification of protection and restoration attributes on a per-connection basis.
The control plane shall support mechanisms for reserving bandwidth resources for restoration.

The control plane shall support mechanisms for normalizing
connection

routing (reversion) after failure repair.

Normal connection management operations (e.g., connection
deletion)

shall not result in protection/restoration being initiated.

11.2. Control plane resiliency

The control plane may be affected by failures in signaling network connectivity and by software failures (e.g., signaling, topology
and

resource discovery modules).

The control plane should implement signaling message priorities to ensure that restoration messages receive preferential treatment, resulting in faster restoration.

The optical control plane signaling network shall support
protection

and restoration options to enable it to be self-healing in case of failures within the control plane.

Control network failure detection mechanisms shall distinguish
between

control channel and software process failures.

The control plane failure shall only impact the capability to provision new services.

Fault localization techniques for the isolation of failed control resources shall be supported.

Recovery from control plane failures shall result in complete
recovery

and re-synchronization of the network.

There shall not be a single point of failure in the control plane systems design.

Partial or total failure of the control plane shall not affect the existing established connections. It should only lose the capability to accept the new connection requests.

12. Security Considerations

In this section, security considerations and requirements for optical services and associated control plane requirements are described.

12.1. Optical Network Security Concerns

Since optical service is directly related to the physical network, which is fundamental to a telecommunications infrastructure, stringent security assurance mechanism should be implemented in optical networks.

In terms of security, an optical connection consists of two aspects. One is security of the data plane where an optical connection itself belongs, and the other is security of the control plane.

12.1.1. Data Plane Security

- Misconnection shall be avoided in order to keep the user's data confidential. For enhancing integrity and confidentiality of data, it may be helpful to support scrambling of data at layer 2 or encryption of data at a higher layer.

12.1.2. Control Plane Security

It is desirable to decouple the control plane from the data plane physically.

Restoration shall not result in miss-connections (connections established to a destination other than that intended), even for short periods of time (e.g., during contention resolution). For example, signaling messages, used to restore connectivity after failure, should not be forwarded by a node before contention has been resolved.

Additional security mechanisms should be provided to guard against intrusions on the signaling network. Some of these may be done with the help of the management plane.

- Network information shall not be advertised across external interfaces (UNI or E-NNI). The advertisement of network information across the E-NNI shall be controlled and limited in a configurable

policy based fashion. The advertisement of network information
shall
be isolated and managed separately by each administration.

- The signaling network itself shall be secure, blocking all unauthorized access. The signaling network topology and addresses shall not be advertised outside a carrier's domain of trust.

- Identification, authentication and access control shall be rigorously used by network operators for providing access to the control plane.

- Discovery information, including neighbor discovery, service discovery, resource discovery and reachability information should be

exchanged in a secure way.

- Information on security-relevant events occurring in the control plane or security-relevant operations performed or attempted in the

control plane shall be logged in the management plane.

- The management plane shall be able to analyze and exploit logged data in order to check if they violate or threat security of the control plane.

- The control plane shall be able to generate alarm notifications about security related events to the management plane in an adjustable

and selectable fashion.

- The control plane shall support recovery from successful and attempted intrusion attacks.

12.2. Service Access Control

From a security perspective, network resources should be protected from unauthorized accesses and should not be used by unauthorized entities. Service access control is the mechanism that limits and controls entities trying to access network resources. Especially on

the UNI and E-NNI, Connection Admission Control (CAC) functions should

also support the following security features:
- CAC should be applied to any entity that tries to access network resources through the UNI (or E-NNI). CAC should include an authentication function of an entity in order to prevent masquerade

(spoofing). Masquerade is fraudulent use of network resources by pretending to be a different entity. An authenticated entity should be

given a service access level on a configurable policy basis.

- The UNI and NNI should provide optional mechanisms to ensure origin

authentication and message integrity for connection management requests such as set-up, tear-down and modify and connection signaling

messages. This is important in order to prevent Denial of Service attacks. The UNI and E-NNI should also include mechanisms, such as usage-based billing based on CAC, to ensure non-repudiation of connection management messages.

- Each entity should be authorized to use network resources
according

to the administrative policy set by the operator.

13. Acknowledgements

The authors of this document would like to extend our special
appreciation to John Strand for his initial contributions to the
carrier requirements. We also want to acknowledge the valuable
inputs

from, Yangguang Xu, Zhiwei Lin, Eve Verma, Daniel Awduche, James

Luciani, Deborah Brunhard, Lynn Neir, Wesam Alanqar, Tammy Ferris,
and
Mark Jones.

14. References

14.1 Normative References

[rfc2026] S. Bradner, "The Internet Standards Process -- Revision
3,"
BCP 9, RFC 2026, IETF October 1996.

[rfc2119] S. Bradner, ôKey words for use in RFC to indicate
requirement levelsö, BCP 14, RFC 2119, 1997
[itu-astn] ITU-T Rec. G.8070/Y.1301 (2001), ôRequirements for the
Automatic Switched Transport Network (ASTN)ö.

[itu-ason] ITU-T Rec. G.8080/Y.1304 (2001), ôArchitecture of the
Automatic Switched Optical Network (ASON)ö.

[itu-dcm] ITU-T Rec. G.7713/Y.1704 (2001), ôDistributed Call and
Connection Management (DCM)ö.

[itu-rtg] ITU-T Rec. G.7715/Y.1706 (2002), ôArchitecture and
Requirements for Routing in the Automatic Switched Optical
Networksö.

[itu-disc] ITU-T Rec. G.7714/Y.1705 (2001), ôGeneralized Automatic
Discovery Techniques.

14.2 Informative References

[itu-otn] ITU-T G.872 (2000) û ôArchitecture of Optical Transport
Networksö.

[itu-g709] ITU-T G.709 (2001)û ôNetwork Node Interface for the
Optical
Transport Networkö.

[itu-sdh] ITU-T Rec. G.803 (2000), ôArchitecture of Transport
Networks
based on the Synchronous Digital Hierarchyö

[ipo-frw] B. Rajagopalan, et. al ôIP over Optical Networks: A
Frameworkö, work in progress, IETF

[oif-addr] M. Lazer, "High Level Requirements on Optical Network
Addressing", oif2001.196, 2001

[oif-carrier] Y. Xue and M. Lazer, et al, ôCarrier Optical Service
Framework and Associated Requirements for UNIö, OIF2000.155, 2000

[oif-nnireq] M. Lazer et al, ôCarrier NNI Requirementsö,
OIF2002.229,
2002

[ipo-olr] A Chiu and J. Strand et al., "Impairments and Other
Constraints on Optical Layer Routing", work in progress, IETF

[ietf-gsmp] A. Doria, et al ôGeneral Switch Management Protocol
V3ö,
work in progress, IETF, 2002

[rfc2748] D. Durham, et al, ôThe COPS (Common Open Policy
Services)
Protocolö, RFC 2748, Jan. 2000

[oif-uni] Optical Internetworking Forum (OIF), "UNI 1.0 Signaling
Specification," December, 2001.

[ansi-sonet] ANSI T1.105-2001, ôSynchronous Optical Network
(SONET) -
Basic Description including Multiplex Structure, Rates and
Formatsö,
2001

[itu-dcn]ITU-T Rec. G.7712/Y.1703 (2001), ôArchitecture and
Specification of Data Communication Networkö.

14 Author's Addresses

Yong Xue
UUNET/WorldCom
22001 Loudoun County Parkway
Ashburn, VA 20147
Email: yxue@cox.net

Monica Lazer
AT&T
900 ROUTE 202/206N PO BX 752
BEDMINSTER, NJ 07921-0000
mlazer@att.com

Jennifer Yates
AT&T Labs
180 PARK AVE, P.O. BOX 971
FLORHAM PARK, NJ 07932-0000
jyates@research.att.com

Dongmei Wang
AT&T Labs
Room B180, Building 103
180 Park Avenue
Florham Park, NJ 07932
mei@research.att.com

Ananth Nagarajan

Sprint
6220 Sprint Parkway
Overland Park, KS 66251, USA
ananth.nagarajan@mail.sprint.com

Hirokazu Ishimatsu
Japan Telecom Co., LTD
2-9-1 Hatchobori, Chuo-ku,
Tokyo 104-0032 Japan
Phone: +81 3 5540 8493
Fax: +81 3 5540 8485
hirokazu.ishimatsu@japan-telecom.co.jp

Olga Aparicio
Cable & Wireless Global
11700 Plaza America Drive
Reston, VA 20191
Phone: 703-292-2022
Email: olga.aparicio@cwusa.com

Steven Wright
Science & Technology
BellSouth Telecommunications
41G70 BSC
675 West Peachtree St. NE.
Atlanta, GA 30375
Phone +1 (404) 332-2194
Email: steven.wright@snt.bellsouth.com

Appendix A: Interconnection of Control Planes

The interconnection of the IP router (client) and optical control planes can be realized in a number of ways depending on the required level of coupling. The control planes can be loosely or tightly coupled. Loose coupling is generally referred to as the overlay model and tight coupling is referred to as the peer model. Additionally there is the augmented model that is somewhat in between the other two models but more akin to the peer model. The model selected determines the following:

- The details of the topology, resource and reachability information advertised between the client and optical networks

- The level of control IP routers can exercise in selecting paths across the optical network

The next three sections discuss these models in more details and the last section describes the coupling requirements from a carrier's perspective.

Peer Model (I-NNI like model)

Under the peer model, the IP router clients act as peers of the optical transport network, such that single routing protocol instance runs over both the IP and optical domains. In this regard the optical network elements are treated just like any other router as far as the control plane is concerned. The peer model, although not strictly an internal NNI, behaves like an I-NNI in the sense that there is sharing of resource and topology information.

Presumably a common IGP such as OSPF or IS-IS, with appropriate extensions, will be used to distribute topology information. One tacit assumption here is that a common addressing scheme will also be used for the optical and IP networks. A common address space can be trivially realized by using IP addresses in both IP and optical domains. Thus, the optical networks elements become IP addressable entities.

The obvious advantage of the peer model is the seamless interconnection between the client and optical transport networks. The tradeoff is that the tight integration and the optical specific routing information that must be known to the IP clients.

The discussion above has focused on the client to optical control plane inter-connection. The discussion applies equally well to inter-connecting two optical control planes.

Overlay (UNI-like model)

Under the overlay model, the IP client routing, topology distribution, and signaling protocols are independent of the routing, topology distribution, and signaling protocols at the optical layer. This model is conceptually similar to the classical IP over ATM model, but applied to an optical sub-network directly.

Though the overlay model dictates that the client and optical network are independent this still allows the optical network to re-use IP layer protocols to perform the routing and signaling functions.

In addition to the protocols being independent the addressing scheme used between the client and optical network must be independent in the overlay model. That is, the use of IP layer addressing in the clients must not place any specific requirement upon the addressing used within the optical control plane.

The overlay model would provide a UNI to the client networks through which the clients could request to add, delete or modify optical connections. The optical network would additionally provide reachability information to the clients but no topology information would be provided across the UNI.

Augmented model (E-NNI like model)

Under the augmented model, there are actually separate routing instances in the IP and optical domains, but information from one routing instance is passed through the other routing instance. For example, external IP addresses could be carried within the optical

routing protocols to allow reachability information to be passed
to IP
clients. A typical implementation would use BGP between the IP
client
and optical network.

The augmented model, although not strictly an external NNI,
behaves
like an E-NNI in that there is limited sharing of information.

Generally in a carrier environment there will be more than just IP routers connected to the optical network. Some other examples of clients could be ATM switches or SONET ADM equipment. This may
drive
the decision towards loose coupling to prevent undue burdens upon
non-
IP router clients. Also, loose coupling would ensure that future clients are not hampered by legacy technologies.

Additionally, a carrier may for business reasons want a separation between the client and optical networks. For example, the ISP business unit may not want to be tightly coupled with the optical network business unit. Another reason for separation might be
just
pure politics that play out in a large carrier. That is, it would seem unlikely to force the optical transport network to run that
same
set of protocols as the IP router networks. Also, by forcing the
same
set of protocols in both networks the evolution of the networks is directly tied together. That is, it would seem you could not
upgrade
the optical transport network protocols without taking into consideration the impact on the IP router network (and vice
versa).

Operating models also play a role in deciding the level of
coupling.

Four main operating models envisioned for an optical transport network:
Category 1: ISP owning all of its own infrastructure (i.e.
including
fiber and duct to the customer premises)

Category 2: ISP leasing some or all of its capacity from a third party

Category 3: Carriers carrier providing layer 1 services

Category 4: Service provider offering multiple layer 1, 2, and 3 services over a common infrastructure

Although relatively few, if any, ISPs fall into category 1 it would seem the mostly likely of the four to use the peer model. The other operating models would lend themselves more likely to choose an overlay model. Most carriers would fall into category 4 and thus would most likely choose an overlay model architecture.

Full Copyright Statement