Towards Automated Interactions between the Internet and the Carrier-Grade Management Ecosystems

Small or medium scale focused research project (STREP)
Co-funded by the European Commission within the Seventh Framework Programme
Grant Agreement no. 258300

Strategic objective: The Network of the Future (ICT-2009.1.1)

Start date of project: September 1st, 2010 (36 months duration)

Deliverable D2.1
Definition of requirements and use cases

Due date: 03/16/2011
Submission date: 04/5/2011
Deliverable leader: TID
Author list: Marcelo Yannuzzi (UPC), Anny Martínez (UPC), Xavi Masip-Bruin (UPC), René Serral-Gracià (UPC), Mohit Chamania (TUBS), Admela Jukan (TUBS), Fernando Muñoz del Nuevo (TID), Óscar González de Dios (TID), Carlos García Argos (TID), Javier Jiménez Chico (TID), Jörn Altmann (SNU), Mohammad Hassan (SNU), Maciek Maciejewski (ADVA), Christine Brunn (ADVA), Gabriela Aronovici (MySoft).

Dissemination Level

☒	**PU:**	Public
☐	**PP:**	Restricted to other programme participants (including the Commission Services)
☐	**RE:**	Restricted to a group specified by the consortium (including the Commission Services)
☐	**CO:**	Confidential, only for members of the consortium (including the Commission Services)

SEVENTH FRAMEWORK PROGRAMME

Table of Contents

Figure Summary

Table Summary

1 Executive Summary

In order to cope with the ever-growing Internet traffic, telecom carriers have deployed flexible IP/MPLS network infrastructures, supported over high capacity optical transport networks. Despite the efforts in convergence, currently there are two logically and administratively separate ecosystems within each telecom carrier, one being exclusively based on networks of IP and IP/MPLS routers whereas the other on deploying circuit switches (Ethernet virtual circuit switches, WDM switches, SDH/OTN nodes, etc.). As a result, telecom carriers have no but manual means to coordinate the provisioning between the routers and the Layer2/Layer1 switches, and much less to communicate and coordinate policy rules, or failures between the IP networks and Ethernet/optical networks.

In light of this, ONE proposes an easy-to-deploy solution aimed at addressing the current isolation between the IP and transport Network Management Systems (NMSs). The key innovation of this project is the design and implementation of an adapter, which can enable communication and coordinated management between the IP and transport management planes (see Figure 1). Note that ONE does not have as a goal to integrate various NMSs, but to achieve, facilitate, and foster the evolution of automated interactions between the Internet and a broad category of emerging carrier-grade NMSs, including those designed for carrier grade Ethernet standards, optical switching and transmission systems based on WDM, as well as interactions with third party systems, such as the Path Computation Element (PCE). More specifically, ONE will design and prototype an ontology-based communication adapter, selectively enabling automated management functions and operations between the Internet and the transport NMSs, without the need for large scale system integration or drastic changes to the current telecom management teams practice.

ONE will demonstrate that the interaction of the network management systems of the future Internet and powerful optical and carrier-grade Ethernet technologies is not only possible, but will also allow carriers who run both types of networks to make use of coordinated actions for the following use cases: i) IP service provisioning; ii) IP/MPLS offloading; and iii) co-ordinated self-healing. These use cases can be summarized as follows:

- The **IP service provisioning** use case focuses on reducing the manual interactions between the IP and the transport layer departments, until achieving fully automated service provisioning. The role of the ONE adapter is to interpret the operator's IP requests and convert them into comprehensible and unambiguous transport layer resource requests, which are required to provision the desired IP service in an automated fashion. In this use case, the ONE adapter can be used either to provision new IP services which demand resources from the transport layer (e.g., a new VPN), or to enable network re-planning to cope with long-term changes in the traffic profile.

- The **IP/MPLS offloading** use case presents a novel multi-layer operation targeted at reducing the IP traffic across unnecessary intermediate routers. More precisely, when the traffic of specific services (in the IP offloading case) or when a set of traffic flows (in the MPLS offloading case) become higher than a desired throughput, the ONE adapter will drive network resource requests in order to bypass the intermediate routers affected by the increase in traffic load. The main novelty in this use case is that the offloading is enabled through an adapter, which can be used for automating and coordinating the overall operation (including the configuration processes in both layers), thereby avoiding the introduction of significant changes in current management practices and their corresponding management systems.

- The **coordinated self-healing** use case aims at performing coordinated actions between IP/MPLS and transport networks in the event of a network failure, so as to recover services in an efficient and rapid way. In the short term, the ONE adapter will orchestrate coordinated recovery actions first in the transport and then in the IP/MPLS layer. Increased coordination and automation via the ONE adapter will also be used to demonstrate fast recovery from severe failures (such as multiple failures or catastrophic events). Coordinated self-healing processes can help to reduce the Capital Expenditures (CAPEX) as well as the operational costs, while improving the network availability.

This document describes in detail the use cases outlined above, and identifies the requirements imposed on the design of the ONE adapter for demonstration of these three use cases. It is worth anticipating that, the ONE adapter will be designed and implemented with an open approach, facilitating in this way its utilization for other possible use cases that may require coordinated interactions between the IP and transport management layers in the future.

2 Introduction

During the last decades, the Internet traffic has grown exponentially, which has led Internet Service Providers (ISPs) and carriers to heavily invest in IP-based infrastructures to handle this traffic. These IP infrastructures are typically designed in a segmented and hierarchical way, resulting in different networks, e.g.: a) the aggregation network, which has the role to collect traffic from end users; b) core networks in order to interconnect the users; and c) the Internet backbone to transport traffic among countries and carriers. In parallel, a new optical transport infrastructure was built to connect the IP routers and carry the ever-increasing level of data. The transport network has evolved both in transmission capacity (10G, 40, 100G and more) and in flexibility (OTN, MPLS-TP, Wavelength Switching, etc.), from pure point to point transmission, to complex meshed structures with dynamically reconfigurable switching features.

In light of such a heterogeneous environment, the convergence of the Internet data services and the optical transport network services has been at the heart of carriers' investments and business strategies. In today's telecom world, however, significant challenges remain towards convergence. Despite the seemingly converged evolution ("everything is data"), the operational complexity as well as organizational and technological separation of the different networks, i.e., the Internet and the optical transport network, remain as large as ever. Indeed, the segmentation of the IP routing and transport networking has not only created profound differences in the way how the Internet and optical transport systems have evolved, but has also led to organizational separation and fragmentation of technical competencies and solutions inside the network carrier. These are still two logically and administratively separate ecosystems within each telecom carrier today, one being exclusively based on networks of IP and IP/MPLS routers whereas the other on deploying circuit switches (Ethernet virtual circuit switches, WDM switches, SDH/OTN nodes, etc.).

With ever increasing cost in operating multiple networks in isolation, the key question to be addressed is what the optimal level of **coordination** and **integration** between the packet and circuit switched technologies and services should be, and which innovative network architectures will offer better service quality at a lower price. From the point of view of network management, which is the main focus of this project, each of the two types of networks has its own requirements and characteristics to be considered. For instance, proprietary Element Management Systems (EMSs) based on Transaction Language One (TL-1) are commonly used for transport Network Management System (NMS) in the U.S. [Sub10], whereas the Simple Network Management Protocol (SNMP) is generally used for IP networks as well as for transport NMSs in Europe [Har02].

Though it is debatable whether the segmentation of the management has produced separate business processes within carriers, or it was due to the much pronounced differences in technologies, the technical competencies within a telecom provider are mainly split between "packet" and "circuit"

networks. In fact, despite several past approaches to integrate the two NMSs (and even ongoing initiatives such as [Cya11]), carriers remain reluctant to adopt any integrated solution. One of the main reasons is the complexity; for instance, the functions of multi-layer routing, signaling, and recovery are significantly more complex than if they are handled in each layer separately. In addition, the ever decreasing price of bandwidth could hardly make any new technology solution attractive, due to the cost of new deployment. Even though telecom carriers are facing duplications of network management functions, the premium on stability and simplicity prevailed over any integrated solution.

In this context, ONE proposes an **easy-to-deploy solution** aimed at enabling **communication** and **coordinated management** of the two networks. More specifically, the key innovation of this project is the **design and implementation of a communication adapter between the IP and carrier-grade management planes** (see Figure 1), which can enable **coordinated and automated provisioning of IP services over transport circuits**, either in response to a specific request or based on pre-established policies to dynamically adapt to traffic changes, including the **automated configuration** of Network Elements (NEs) in both layers as well as self-healing functionality. It is worth highlighting that the ONE adapter does not intend to replace the functions of current IP and carrier-grade NMSs, but rather to offer an interface capable of interpreting and coordinating functions between these. ONE will use an **Ontology Mapper** as a basic functional block for interpreting the requests coming from the Internet management interface, since ontologies offer an explicit and formal specification of a set of concepts and their relations in a way that can be understand unambiguously by computer systems. This approach provides a solid basis where a set of semantics can be used for automated configuration of NEs, while each layer can keep running its own NMS and functionality without the need to adopt a unified management system.

Figure 1 illustrates the application scenario in which this project will be developed. We assume that a network operator has both an IP and a transport network infrastructure, the first composed of IP and IP/MPLS routers, and the second composed of optical switches and possibly carrier-grade Ethernet switches. The transport network is assumed to be managed by one or more Transport Network Management System (TNMS), each of which is typically provided by the vendor of the carrier equipment, such as, ADVA's Optical Networking FSP Service Manager [Adv11] or the Nokia Siemens Networks' NetAct manager [Nsn11]. The IP network devices, on the other hand, can be managed either by one or more IP NMS, such as OpenNMS [Ope11] or Nagios [Nag11], or directly by the network administrators using the Command Line Interface (CLI) [Cis11] or NetConf [Enn06]. In its basic configuration, the ONE adapter shall be able to communicate both with the TNMSs and the IP NMSs, e.g., using Web services over a standardized interface, such as the Multi-Technology Operations System Interface (MTOSI), which is an XML-based Operations System (OS)-to-OS interface suite [Mto11].

In case that the operator lacks a proper IP NMS solution and relies on a set of proprietary tools, the administrators of the IP network shall be able to interact directly with the ONE adapter, and issue high-level requests that will be interpreted by the latter and adapted to the appropriate semantics and syntax of the corresponding TNMS. As shown in Figure 1, in a more advanced configuration, the ONE adapter shall also be able to interact with other management sub-systems, such as a Path Computation Element (PCE) [Far06], a measurement system, or a policy-based Traffic Engineering (TE) management system.

Under these assumptions, which we contend are realistic and representative of most operational networks, this document proposes three well-defined use cases describing the potential application of the ONE adapter, and examines the requirements that such use cases impose on the latter. The proposed use cases entail configurations in both management layers and encompass:

i) Automated provisioning of IP links and services, enabling enhanced network planning;

ii) Optical bypass triggered by pre-configured policies to accommodate an increase of traffic;

iii) Coordinated self-healing actions.

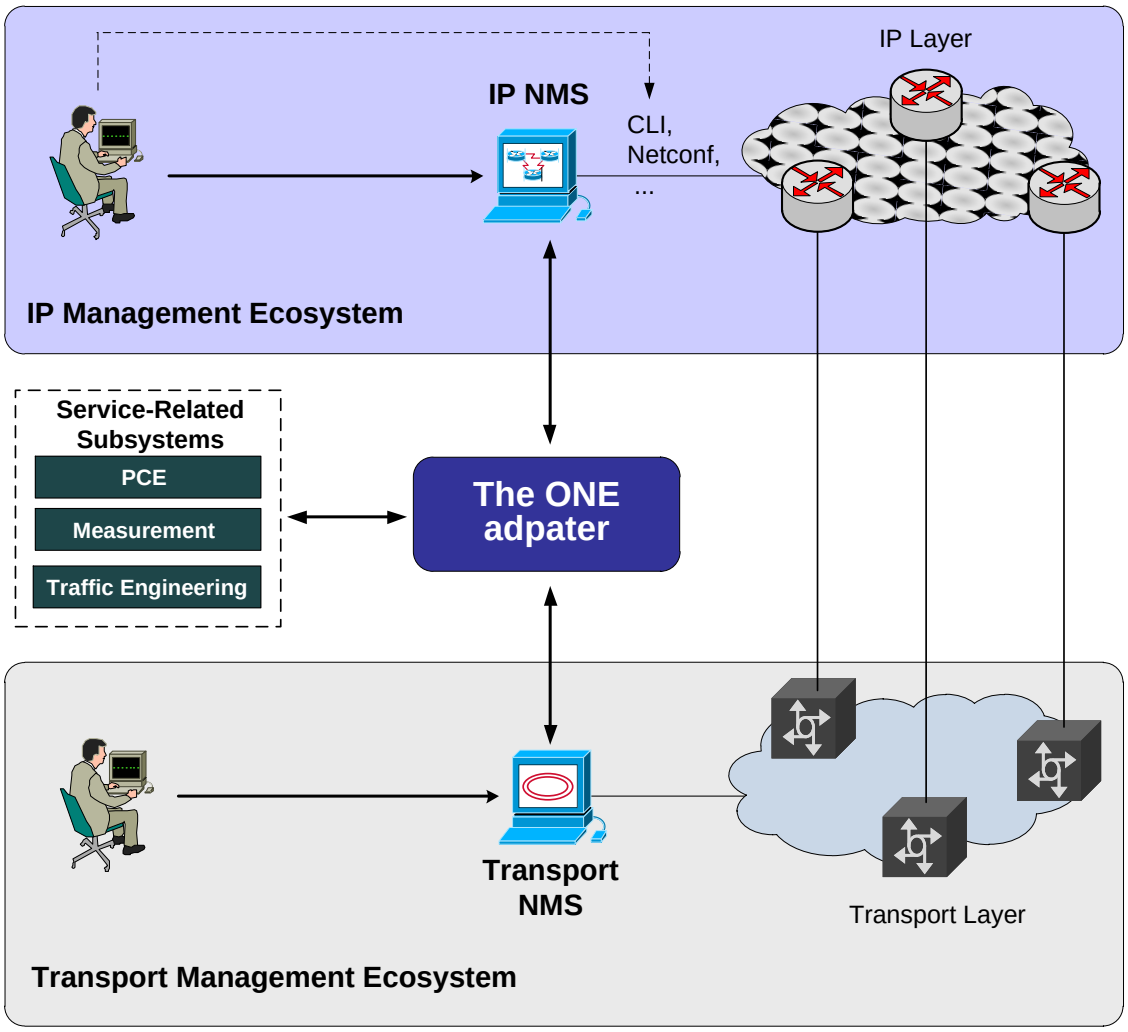

Figure 1: Application scenario for the ONE adapter.

We shall describe the motivations behind these use cases and revisit the question of integration versus coordination of the Internet and transport NMSs. In particular, we shall provide our vision of the importance of pushing towards the convergence of the IP and transport NMSs without requiring the integration of these systems in a unified network management solution. We contend that due to its simplicity and scalability, the ONE adapter is a true enabler of controlled convergence of "packet" and "circuit" switching networks and their corresponding management ecosystems, since it not only eliminates the need for large scale system integration or drastic changes to the current telecom management teams practice, but also bridges the communication gap between two management ecosystems that are currently isolated.

In addition to this, the ONE adapter eliminates the need for 1+1 protection in IP and Carrier layer which means the ONE adapter can increases the network utilization significantly and therefore, it reduces significantly the service cost, since more data can be served with the same infrastructure.

To quantify this, the ONE consortium aims to provide a techno-economic analysis of the ONE adapter during the course of this project. Particularly, we aim to assess the overall network performance increase which is caused by using the ONE adapter as a facilitator of coordination and interaction. We hope it will help network operators to understand the cost efficiency of a mediator tool like the ONE adapter.

The rest of this document is structured as follows. We first analyze the current structure of telecom operator's networks, and examine some of the most common network management tools and practices for both the Internet and the transport network. In particular, we outline the state of a set of relevant standards, and analyze the reasons for the current separation between the IP and transport NMSs. After that, we expose the three use cases considered in this project, and discuss the evolution of these use cases and the application of the ONE adapter in the context of short-, mid-, and long-term scenarios. We then analyze and list the requirements for the adapter, and finally, conclude this document providing relevant insights from an industrial perspective.

3 Scenario Analysis

3.1 Current Network Scenario

In this section, we describe the typical structure of a telecom operator's IP and transport networks, their operation, and the management systems employed for monitoring and configuring the network elements in the two layers. We describe the main features of some of the most representative Network Management Systems (NMSs) both for the IP and the transport layer, and discuss in detail the pros of coordination versus integration of NMSs, with the aim to bridge the communication gap between these two separate management ecosystems. We also outline the state of standardized frameworks, such as GMPLS and ASON, and discuss their limitations not only in terms of flexibility for deploying services in packet networks, but also in overcoming the existing isolation between the IP and transport network management layers.

The typical structure of a European nation-wide operator is based on a stratified hierarchy. In the bottom level of the hierarchy, the access and mobile networks collect the traffic from the users, having the infrastructure to reach the homes in the case of fixed access, and the spectrum to reach the mobile users. In the next level, the aggregation networks (which can be both metropolitan and regional networks) collect the traffic from the access nodes (DSLAMs, FTTH headers), and forward it to the core IP/MPLS network. The aggregation networks also collect traffic from business customers, who usually have dedicated equipment at their premises (note that typically, the aggregation networks for residential and business customers are separated). Then, the IP/MPLS core provides the next level of the hierarchy, interconnecting all the aggregation networks with the Internet access (interconnection exchange points, both international and national) and the services. On the other hand, a transport network is used to connect the IP/MPLS routers this hierarchy.

We now proceed to describe in more detail the main aspects of the core, aggregation, and transport networks and their management, since they are essential for the purpose of this project.

3.1.1 Core IP/MPLS Networks

The core network aggregates and transports IP traffic among IP access centers, and provides connectivity among users, service centers, datacenters, and the Internet. Core networks are in turn designed in a hierarchical way, in order to perform an efficient aggregation and leverage from statistical multiplexing of IP packets. As an example, Figure 2 shows the typical composition of an IP/MPLS backbone network structured in, access, transit, and interconnection nodes in a hierarchical

topology, which efficiently aggregates the national and international traffic. The figure depicts the typical IP/MPLS core network of a telecom operator providing Internet services to end customers. As it can be seen, this IP backbone is usually based on a hierarchy of routers interconnected through high speed point-to-point links.

An IP/MPLS backbone network typically runs the IS-IS protocol for topology discovery, the Label Distribution Protocol (LDP) for MPLS label distribution, and BGP for the distribution of reachability information in the Internet. In these networks, switching is based on the IP (layer 3) header at the border routers (i.e. access and interconnection nodes), and on the MPLS label (layer 2.5 header) at the transit routers. In general, operators configure their IP/MPLS backbones so as to provide best effort IP services form residential Internet customers, i.e., neither bandwidth nor QoS guarantees are provided for these customers. Traffic from/to business and VPN customers, on the other hand, is typically marked with a "high priority" tag (e.g., by using the 3 EXP bits), so that the traffic will be sent over lower delay queues across the network and would not be discarded in case of congestion (as long as it remains a low amount of traffic when compared to best effort services). An acceptable quality can be achieved based on this strategy, since in practice this approach is usually complemented by the operator with significant network over-provisioning.

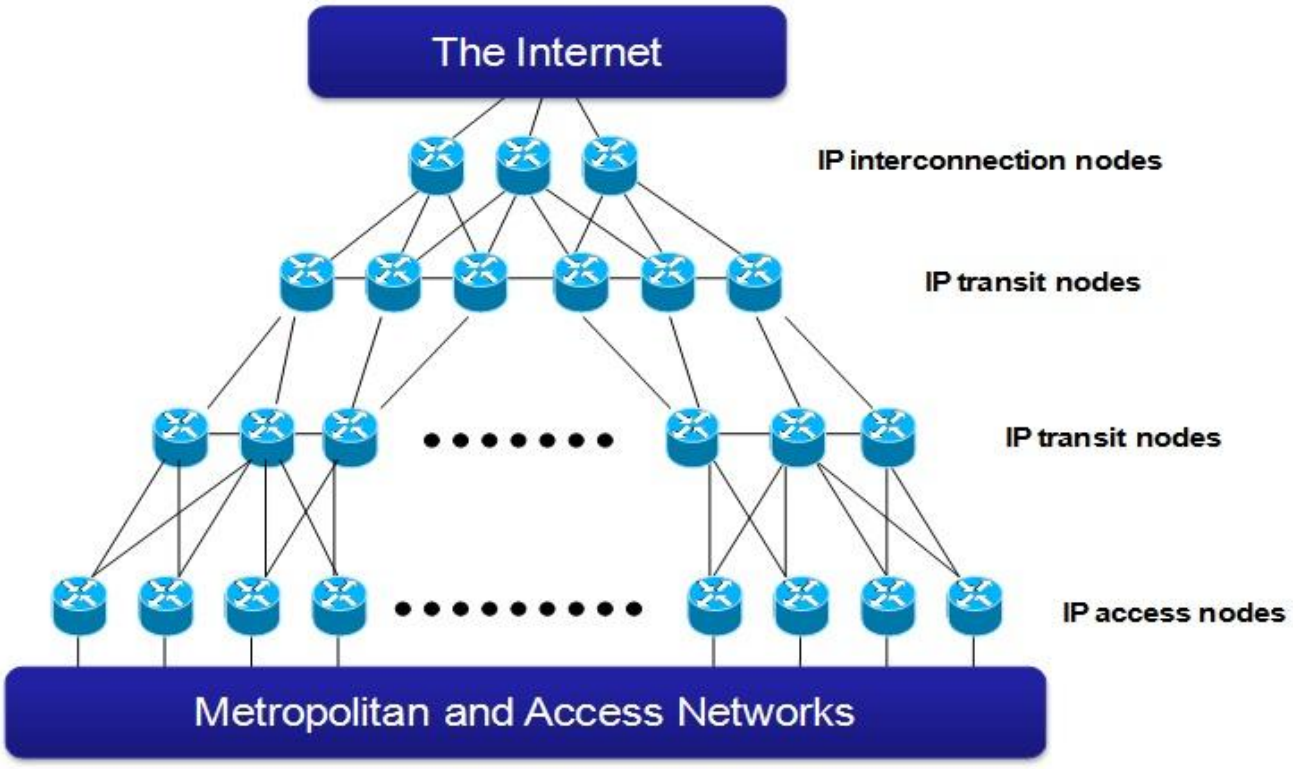

Figure 2: IP/MPLS Backbone structure

For resilience purposes, IP/MPLS network operators usually design partially meshed networks, so that connectivity failures can be often solved by the IGP routing protocols (e.g. IS-IS or OSPF). In addition, and to provide robust protection for Internet services, a dual network design is frequently

implemented by large operators. With this approach, every router in the network is actually duplicated, which means that the network is composed of two overlapping and interconnected topologies: *the even network and the odd one*. This architecture offers a 1+1 protection scheme in case of router failures, as there is always a backup element for every single router in the IP network, which, in combination with the IGP recovery, provides IP/MPLS services with a very high degree of availability.

3.1.2 Aggregation Networks

The aggregation networks are in charge of collecting the traffic form the access networks and of business customers. The Customer Premises Equipment (CPE) is located at the customer's premises (either residential or business customer), and is typically connected to the network through the access segment by means of:

- xDSL or GPON technologies for residential customers and also SMEs.
- A direct fiber link at configurable line rates (e.g. 100M, 1G, 10G, 40G, or even 100G) for larger business and corporation customers.

The access link for residential customers is terminated at the so called "access node" in the operator side (namely, a DSLAM for xDSL clients or an OLT for GPON clients). This node aggregates the traffic from many customers and is typically connected to the metropolitan network by means of 1GbE interfaces. The access link for corporations is typically terminated directly at the metropolitan network border node, generally called Multi-Tenant Unit (MTU).

The aggregation network (also called metropolitan, metro/aggregation or regional network) aggregates traffic from many users towards the IP access centers. This network also provides layer 2 VPN services to business customers within a region. The metro/aggregation region, though implemented by means of VPLS technologies, can be perceived as an Ethernet network (with typical support for multicast, VLANs, etc.) from the customer's perspective.

The residential traffic from a whole region is aggregated at the BRAS (Broadband Remote Access Server) located at the IP access center, which also performs AAA functionalities and applies per-user policies. The BRAS is subsequently connected to an IP/MPLS access node (see Fig. 2), which provides IP connectivity with the rest of users, with service datacenters and with the Internet. On the other side, the business traffic does not traverse the BRAS, but arrives directly to the IP/MPLS access nodes.

3.1.3 Transport Networks

As mentioned above, the IP/MPLS core network provides layer 3 connectivity among IP access centers, interconnecting end users, service centers, datacenters, and the Internet. The IP/MPLS routers are typically interconnected by high speed optical links (or circuits) that are provided by the transport network. The latter is in charge of transporting high capacity traffic and providing point to point connections between IP/MPLS network elements. A new connection or circuit is typically established

in the transport network when the traffic demand grows above a certain threshold, in order to satisfy the established Service Level Agreements (SLAs) and ensure that the network can comply with the planned reliability.

With the advent of optical networking, reconfigurable optical switches (ROADMs) are now available in the transport network, bringing the possibility to establish on demand direct connections among switches (usually called lightpaths). At present, most network operators rely on lightpath provisioning via a Transport Network Management System (TNMS), sometimes also referred to as Carrier-Grade NMS. When a TNMS is used, the overall process for the establishment of a lightpath can be fully automated, drastically simplifying the operation and management of the transport network.

GMPLS [Man04], a generalization of MPLS to support packet switching, time division and wavelength multiplexing, is the current state of the art in terms of control plane for transport networks. The GMPLS protocol stack allows for the dynamic establishment, teardown and modification of transport connections. For this reason, GMPLS stands as a strong candidate to provide on-demand transport connections. In addition, and due to its multi-layer nature, a GMPLS integrated control plane can theoretically allow the implementation of multi-layer path restoration solutions, which could be used to reconsider the 1+1 protection strategy. However, the implementation of integrated GMPLS solutions is uncommon in operational networks. On one side, most of already deployed IP/MPLS networks are based on IP/MPLS protocols (not GMPLS), and therefore are not interoperable with GMPLS. On the other side, a pure multi-layer GMPLS approach represents a significant challenge for multi-vendor scenarios (different vendors in each layer). Indeed, such scheme would force the different GMPLS implementations to be fully compatible, something that is not available by default. Moreover, even though GMPLS might hypothetically bring such functionality, GMPLS does not address the existing isolation between the IP and transport network management systems, so GMPLS is not expected to actually bridge the communication gap between the two management layers.

3.2 Management Ecosystems

In this context, the operator's network paradigm is moving towards a two-layer scheme, composed of high capacity IP/MPLS routers connected through reconfigurable optical switches (ROADMs). Despite its apparent simplicity, many issues remain unsolved, particularly, in the subject of network management when considering both layers jointly. The IP/MPLS and transport networks have been traditionally designed and operated by separate departments within network operators. Likewise, both layers have always represented different business areas for network providers, maintaining different product lines in each of them. In light of the potential cost efficiency that an integrated network could bring about, this historic separation is now being questioned, especially, in a context where a constant traffic increase is combined with stalling revenue growth. In this sense, the role of network management, and particularly, the development of realistic and easy-to-deploy solutions facilitating the convergence of the IP and transport NMSs are essential, and thus this is the main focus of the ONE project.

In general terms, a Network Management System (NMS) is a tool capable of remotely configuring, controlling, and monitoring Network Elements (NE) or devices. Current NMSs define the service parameters of the network and automate the configuration of vendor specific equipment using proprietary Element Management Systems (EMSs). Modern transport NMSs support many service-oriented functions while using vendor specific platforms, which dramatically reduces the operational

overheads of telecom carriers and the complexity of the management tasks involved. One of the main features that an NMS platform should provide is awareness, in order to alert operators of performance issues over network devices.

We now proceed to describe the main features of some of the most important NMSs currently available both in the IP and transport layers.

3.2.1 IP Network Management Systems

Many IP Network Management Systems and solutions have been developed during the past years. Current tools can be categorized in two major groups (see Figure 3) *proprietary* and *open source* solutions. Commercial deployments encompass powerful management tools with high licensing fees. HP OpenView, IBM Tivoli, CA Unicenter and BMC Business Service Management are four recognized proprietary alternatives currently available for network management. Other commercial tools such as AdRem NetCrunch, AccelOps, CiscoWorks LMS, Dhyan NMS, dopplerVUE, ExtraHop, NetMRI are also deployed for network management. Despite the existence of easy use and highly supported commercial tools, free open source solutions have successfully arose in the past years.

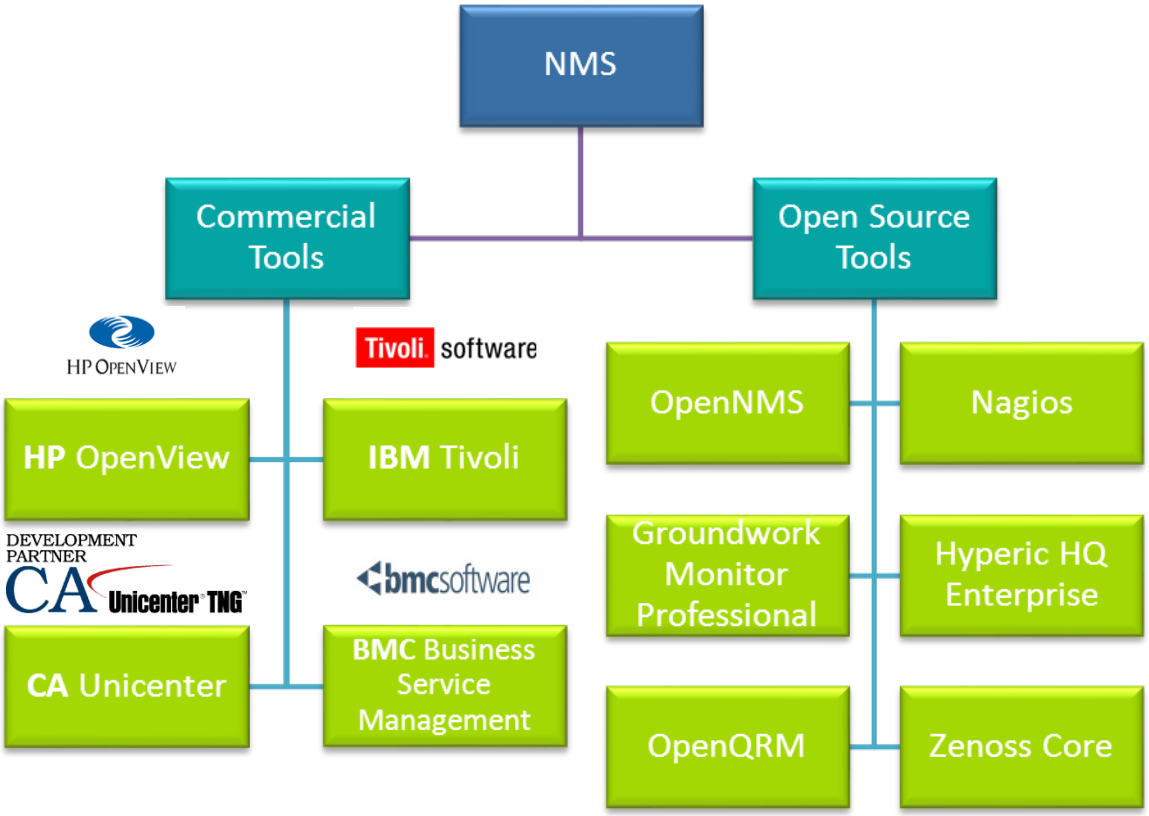

Figure 3: Tree Diagram of NMS Solutions

Open source systems have become an attractive and popular alternative for enterprise network management, not only because they represent cost free tools, but also because they provide a flexible and extensible solution for satisfying customized demands and requirements, while counting with a strong and up to date support community. A large number of network management open source tools have been developed, some of the most popular are: OpenNMS, Nagios, Zenoss Core, OpenQRM, NETDISCO, Hyperic HQ Enterprise, Groundwork Monitor Professional, among many others. These tools differ on their features and operating system support.

In the following lines we will present a brief description of two IP NMS tools: OpenNMS [Ope11] and Nagios [Nag11].

OpenNMS is a free open source solution for networking management, released under the GNU Generic Public License in 2000. This management platform focuses not only on network elements but also on services that network resources provide. OpenNMS works over three main functional areas: Service Polling, Data Collection and Event Management and Notifications. This java-based tool provides Windows support and a variety of open operating systems, including Linux, Mandrake and Solaris, as well as Mac OS X. It has been designed to manage large enterprise networks, with proven scalability. It works over XML based configurations and is extremely configurable and customizable. OpenNMS provides auto-discovery functionality, network map generation, and customizable event management, even allowing the execution of external programs given the occurrence of an event.

Nagios, on the other hand, was originally released under the name *NetSaint* and was created in 1999. It's a free open source solution licensed under the terms of the GNU General Public License. Nagios stands as a powerful tool for network, server, and application monitoring. It has in fact become one of the most commonly used server monitoring platforms. Nagios supports active and passive checks and stands out as a flexible, configuration file-based alternative. One of the main features that this tool provides is the development of language independent plug-ins that allows extending Nagios functionalities in a customized way. Companies such as Amazon, AT&T, Domino's Pizza, ebay, Google, Twitter, Symantec, among many others, are some well-known organizations that use Nagios as part of their IT management toolset.

The following table summarizes the main features of the two open source NMS described.

Name	OpenNMS	Nagios
IP SLA Reports	Yes	Via plugin

Auto Discovery	Yes	Via plugin
Agent	Supported	Supported
SNMP	Yes	Via plugin
Syslog	Yes	Via Plugin
Plugins	Yes	Yes
Triggers / Alerts	Yes	Yes
WebApp	Full Control	Full Control
Distributed Monitoring	Yes	Yes
Inventory	Limited	Via plugin
Data Storage Method	JRobin, PostgreSQL	Flat File, SQL
License	GPL	GPL
Maps	Yes	Yes
Access Control	Yes	Yes
IPv6	Limited	Yes

Table 1: Comparison of open source IP NMS.

3.2.2 Carrier-Grade Network Management Systems

Most of the optical transport networks of large telecom operators are managed by NMSs provided by vendors of the carrier equipment, like Nokia Siemens Networks NetAct[1], ADVA Optical Networking FSP Service Manager[2], Ciena One Software Suite[3]. Although there are third party solutions, such as NOCVue[4] and WebNMS Framework[5], they are not commonly used in practice, due to significant differences in the way they handle vendor specific equipment.

NMSs for the transport layer are required to be extremely stable and show high performance in a large network environment. The presentation of the network data is in most of the cases clearly organized in accordance to a well-defined standard like ITU-T G.709. Network management systems typically communicate with the hardware through SNMP or TL1 protocols [Sub10] [Har02]. Management itself is offered through a Graphical User Interface (GUI) and with optional interfaces to support custom operator's OSSs based on various technologies like Corba or MTOSI.

A typical structure of a carrier-grade NMS is presented on Figure 4. The Service Management Layer (SML) is responsible for the presentation and configuration of customer services in the entire

[1] http://www.nokiasiemensnetworks.com/portfolio/products/network-management-oss/netact

[2] http://www.advaoptical.com/en/products/automated-network-management.aspx

[3] http://www.ciena.com/products/ciena-one

[4] http://www.nocvue.com/

[5] http://www.develcon.com/nms/webnms.html

network. Each service is presented with relevant fault and performance data. The Network Management Layer (NML) is responsible for managing the equipment in the entire network to satisfy the needs of the SML. This includes provisioning of requests, and communication of network changes and events. The Element Management Layer (EML) is the mediator between a specific device and the NMS, and it is responsible for responding to NML requests, and informing the NML of any events related to the device EML and its handling.

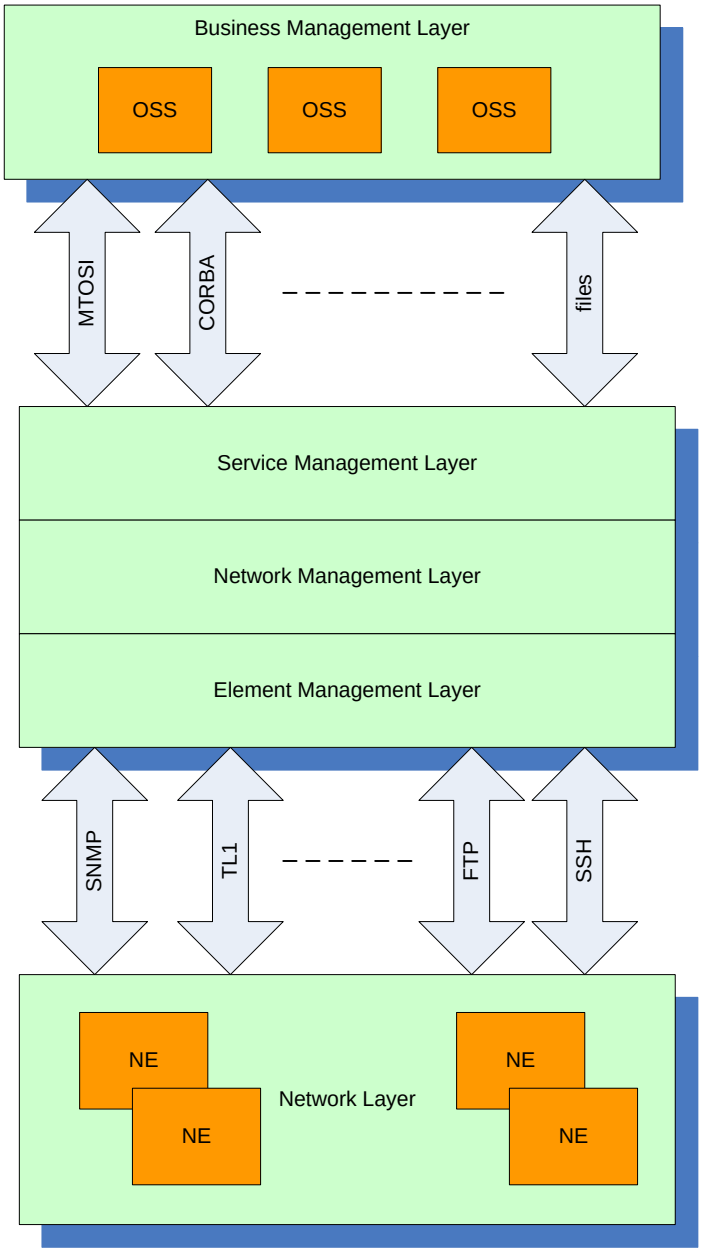

Figure 4: Carrier grade NMS structure.

3.3 Towards coordinated and automated Network Management

The rapid expansion of IP-based technologies has significantly increased the complexity and diversified the tasks and specialization required for managing telecom networks. Despite this, operators and service providers have been systematically forced to reduce their operational costs in order to retain revenue while remaining competitive. A considerable part of the savings has come from the simplification of transport network operations, and the convergence of traditional and new telecom services in IP over a common transport infrastructure. Indeed, the speed of innovation at the IP layer has strongly conditioned the strategies adopted by operators for operating and maintaining their Internet and transport networks, and explains in part the marked differences in the evolution of management tools in support of these operations.

The reasons behind the development of separate administrative and technical ecosystems of the IP and the transport network ecosystems are embedded in the service demands from both technologies. The transport network was designed to deliver a small number of services with fairly static demands on network operation. However, the ever increasing size of the network and the constant demand for more bandwidth have led the industry to heavily invest in R&D, in order to cope with the increased transmission capacity and network scale, while simplifying the operation and maintenance of transport networks as much as possible. In practice, transport networks are operated via NMSs, which define the service parameters for the network, and the configuration of vendor-specific equipment is facilitated using proprietary EMSs (see, e.g., [Adv11] [Nsn11], [Nsn11b], [Alc11], [Fuj11], [Hua11]). Modern transport NMSs today support many service-oriented functions while using vendor-specific platforms, which dramatically reduces the operational overheads of telecom carriers and the complexity of the management tasks involved.

On the other hand, the changing dynamics of the Internet drove the deployment of a wide spectrum of IP enabled equipment by telecom carriers. The IP network was expected to support a large number of services and quickly adopt new upcoming services to reduce time to market. However, in order to provide the desired flexibility, IP network configuration became increasingly complex and vendor-specific. Currently, monitoring of IP devices is mostly managed by the SNMP protocol [Har02], whereas their configuration is typically performed through direct access to the command line of the specific device. The configuration process can be either manual or assisted by means of custom tools that are tailored to automate the interactions through device specific interfaces, which are generally based on the Command Line Interface (CLI) [Cis11] or the NETCONF interface [Emn06].

As a result, telecom carriers have been forced to find a reasonable balance between the complexity and associated cost of the operations required at the IP layer and the simplicity and cost savings of operating and configuring the equipment at the transport layer. The fragmentation of technical competencies has also led to isolated management islands, where even simple tasks involving operations on both the IP and transport layers require multiple human-assisted interactions, which, paradoxically, end up cutting back part of the savings obtained by the operators. The fragmentation of management islands also means that common services such as network monitoring, AAA, and PCE are often duplicated in both layers leading to additional costs.

The left-hand side of Figure 5 illustrates some of the main consequences of the isolation between the management systems. It is worth noting that even the provisioning of a new IP link (A) requires the intervention of human operators from two different departments, each of which is responsible for the configurations in the corresponding layer. These operations not only lead to long times for service provisioning and to potential configuration inconsistencies, but also impede the instrumentation of more advanced mechanisms, such as policy-based resource provisioning (e.g., in response to traffic churn (B)), or any type of coordinated self-healing action (C). The lack of automated coordination

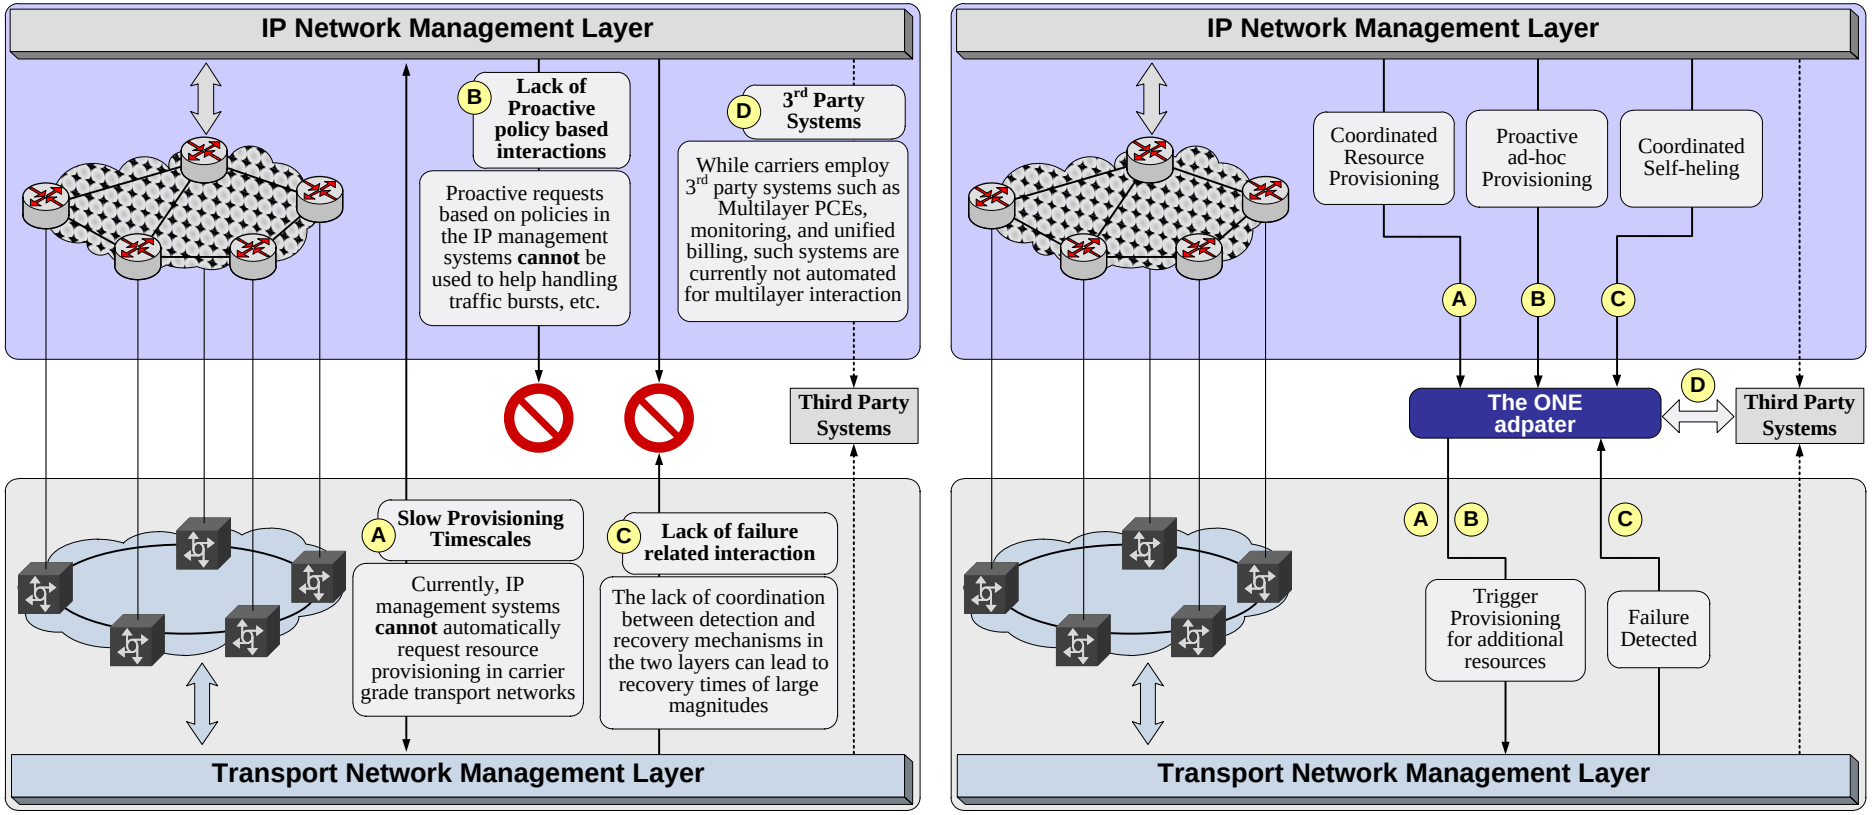

Figure 5: Main limitations and targeted goals in ONE.

upon a failure is particularly critical, since adequately planning the recovery of certain failures may demand a significant amount of interactions between the two network management teams before or even during the occurrence of the failure.

Considering that several management processes are still far from being automated at the IP layer itself, and also that the position adopted by telecom carriers has led to the isolation of management technologies, it seems hard to conceive that the efforts towards a unified and integrated multilayer NMS could eventually succeed. Moreover, operators have thus far been reluctant to introduce automated mechanisms for requesting connections without explicit human control. In fact, this is probably one of the central issues that the research in automated planning has failed to address to date. In this context, it seems reasonable to push for solutions towards the convergence of the Internet and transport NMSs that do not require the integration of these systems, and include the possibility of having human control during the automation of the management tasks. A starting point in this direction is to overcome their current isolation by means of an adapter (a "middle-box") that can provide a simple, reliable, and automated communication channel between the two management layers (see the right-hand side of Figure 5). The goal is to enable the coordination of a set of selected and rather basic management tasks, such as service provisioning (A) (B), and coordinated self-healing (C). In addition, carriers may desire an automated communication with external control and management sub-systems, such as with the PCE, a measurement system, or a TE management system (see case (D) in Figure 5 and also Figure 1). This approach provides a reasonable basis for operators, facilitating their operations and making more cost effective the interactions between the Internet and the transport layer NMSs.

It should be noted that attempts have been made towards developing a unified control framework to support both "packet" and "circuit" services. One of the main efforts in this direction is led by the Internet Engineering Task Force (IETF), in the form of the GMPLS-based standard solution for inter-technology and inter-layer control plane interactions to support both packet and circuit switched networks in the same control plane [Man04]. Another important effort is the ASON framework developed by the ITU-T [ITU11], which aims at reducing the human intervention in the process of service provisioning as well as providing means to automatically switch circuits between different networks. In particular, the ASON approach enables network interactions in a standardized way, keeping the network-internal operation protocols independent.

In the rest of this section, we shall outline these standardized frameworks and identify some of their main shortcomings. It is important highlighting that the ONE adapter does not intend to compete with GMPLS or ASON, but rather to offer an easy-to-deploy solution enabling automated management interactions between the IP and transport NMSs, in a way that can perfectly coexist with, and leverage on, the GMPLS and ASON control plane frameworks.

3.4 Standardized Frameworks: the roles of GMPLS and ASON

The rapid development of the Internet put strong pressure on operators' data communication infrastructures, especially, in terms of the dynamicity and the requirements for provisioning transport circuits to support Internet services. The static nature of the SONET's control-plane model could no longer offer the flexibility and dynamicity required by IP traffic demands, so a new switching technology was introduced by the IETF, namely the Multiprotocol Label Switching (MPLS), which is based on the assignment of labels for efficient packet switching between the ingress and egress nodes. The wide acceptance and popularity of MPLS encouraged the IETF to extend the MPLS technology to GMPLS [Man04], in order to provide support for spatial switching, optical wavelength switching, and TDM switching.

The GMPLS framework is basically the result of the attempt of the IETF to bridge the gap between electronic and optical switching through a standardized control-plane. In order to address the need for IP routing and signaling in GMPLS, a set of protocols were extended, such as OSPF-TE, IS-IS-TE, and RSVP-TE. The introduction of the Link Management Protocol (LMP) in GMPLS has also facilitated the way of handling link connectivity as well as the management of the routing and signaling processes. The following table summarizes the GMPLS protocol suite.

Protocols		Description
Routing	OSPF-TE, IS-IS - TE	Routing protocols for the auto-discovery of network topology, advertise resource availability.
Signaling	RSVP - TE, CR - LDP	Signaling protocols for the establishment of traffic-engineered Label Switched Paths (LSPs).
Link Management	LMP	• Control-Channel Management • Link-Connectivity Verification • Link-Property Correlation • Fault Isolation

Table 2: GMPLS protocol suite overview.

Although the common control plane promises to simplify network operation and management by automating end-to-end provisioning of connections as well as managing network resources and providing the levels of QoS that are expected for the new applications, GMPLS suffers from a number of shortcomings which have so far hindered its massive adoption by network operators. On the one hand, GMPLS introduces significant changes in the software and hardware of the network, which requires considerable capital expenditures, and thus is costly to deploy.

The lack of involvement of ITU-T in the development of the MPLS and GMPLS standards, together with the need to overcome the limitations of manual provisioning in transport networks, led the ITU-T to the development of a complete definition of an Automatic Switched Optical Network [G8080]. ASON is an optical transport network capable of dynamically provision and release optical connections. A network conforming to the ASON recommendation is composed of domains which interact in a standardized way, but whose internal operation is protocol-independent and is not subject to standardization. In the ASON framework, different networks interact through standardized interfaces, namely, the User Network Interface (UNI), the Internal Network-Network Interface (I-NNI), and the External Network-Network Interface (E-NNI). Figure 6 shows the basic elements in the ASON control plane and the interfaces with the transport network.

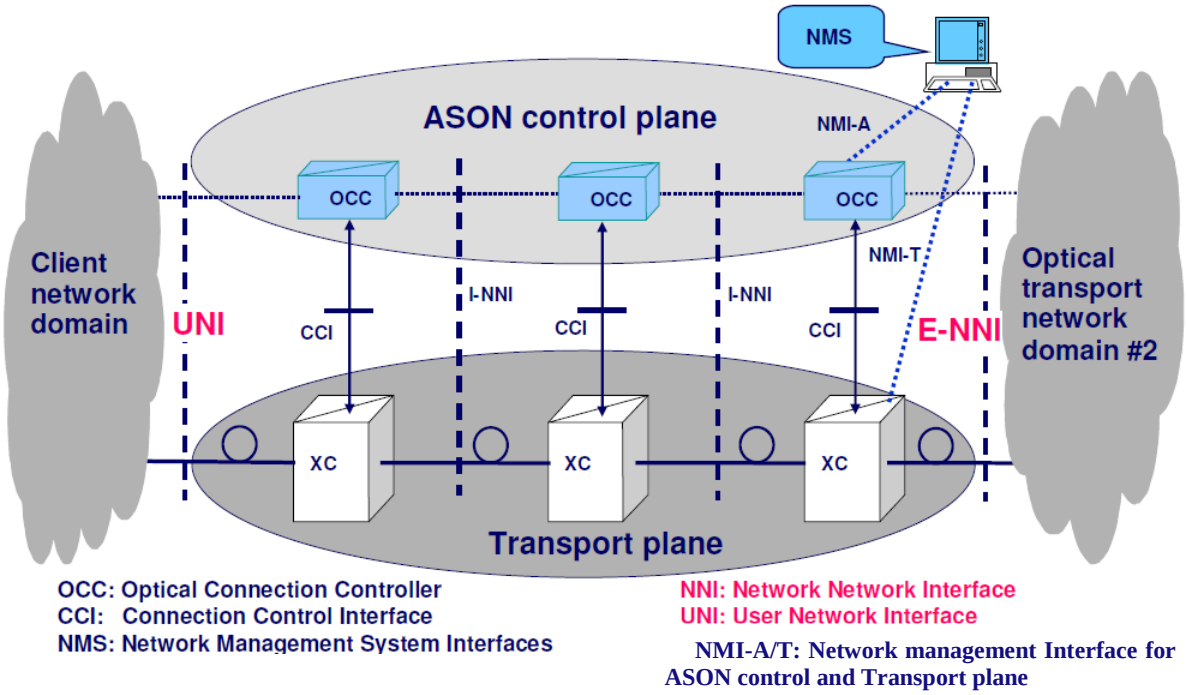

Figure 6: The ASON network architecture (source OIF).

Although the ASON architecture provides a solid basis and interfaces for an automatically switched optical network, ASON per se does also not address the current isolation between the IP and transport management layers. In an attempt to bring together the architecture and requirements defined by ITU-T in ASON and the protocols defined by IETF in GMPLS, the Optical Internetworking Forum (OIF) worked towards a unified view in order to make possible the interoperability between the two standards, but unfortunately the OIF standards also do not solve the isolation between the IP and transport management layers.

Indeed, many carriers are still reluctant to deploy both framework in a way which would include automatically initiated setup and release of IP links involving carrier-grade dynamic circuits. This is because carriers would like to exercise a controlled interaction between the two layers, rather than their ad-hoc operations, due to the implications on the stability of both networks and their corresponding management systems. The development of coordinated interactions between the two NMSs, which can be supervised by network operators, might be a true enabler for automated interactions in a controlled and acceptable way.

4 Use Cases for the ONE Adapter Module

4.1 Use case 1: IP Service Provisioning

As mentioned above, the management of IP and transport networks is carried out by different administrative departments even inside the same provider network. The existence of different management planes as well as administrative entities for the IP and the carrier-grade transport networks makes even basic multi-layer interactions such as the provisioning of a new IP link a tedious task requiring manual intervention of different operational teams, each of which is responsible for configurations in different networks. This practice not only leads to long times for provisioning and solving potential configuration inconsistencies, but also hinders the implementation of advanced inter-layer interactions required for policy-driven resource provisioning and co-ordinated self-healing actions. This use case is designed to demonstrate the ability of the ONE adapter to facilitate IP provisioning using automated interactions between the IP and the carrier-grade transport networks. This use case presents three scenarios where the network operator can use the ONE adapter to facilitate multi-layer co-ordination.

4.1.1 Automated IP Link Provisioning

In this scenario, we use the IP link provisioning process to demonstrate the ability of the ONE adapter to replace manual interventions/communications between the IP and the carrier-grade management entities with automated interactions facilitated over the ONE adapter. Here, as shown in Figure 7, the operator in the IP network will request a new IP link between a pair of routers in the IP network, and the ONE adapter will facilitate coordination and inter-layer communication required for the same. The ONE adapter will first check if IP interfaces (virtual/physical) are available at the IP routers in question and will identify corresponding transport network switches (1). If IP interfaces are available at the requested end-points, the ONE adapter will request the carrier-grade transport network to reserve a circuit between the corresponding transport network end-points (2). If reservation is successful, the ONE adapter will configure the data plane parameters for the circuit end-points and the inter-layer interconnects and will then configure the IP interfaces including IP interface addresses and routing rules as defined by the operator to initialize the newly formed IP link (3).

4.1.2 Automated Multi-layer IP service Provisioning

While a new link may be provisioned to support network engineering activities, more frequent multi-layer operations may be required for provisioning IP services (e.g. VPNs or IPTV) in an existing network. In a more advanced setting, the ONE adapter will be used to facilitate IP *service* provisioning. In this scenario, as shown in Figure 8, if an IP service cannot be provisioned due to lack of capacity in the IP network, the network operator will request the ONE adapter to initiate inter-layer interactions required to provision the service request. Here, the ONE adapter will communicate with a path computation entity like the PCE [Far06], [PCE09], to determine the location and the required capacity for new IP links that would be required to facilitate service provisioning (1). Service

templates are mapped to different requirements in terms of network parameters (e.g. bandwidth, delay, jitter, hop-count) by different network operators, and the path computation entity used by the ONE adapter will compute paths based on the network requirements as defined by the operator. If the path computation is successful, The ONE adapter will then initiate inter-layer communication and coordination functions as described in 4.1.1 to introduce the computed IP links in the network and will then instruct the IP NMS to initiate service provisioning (2).

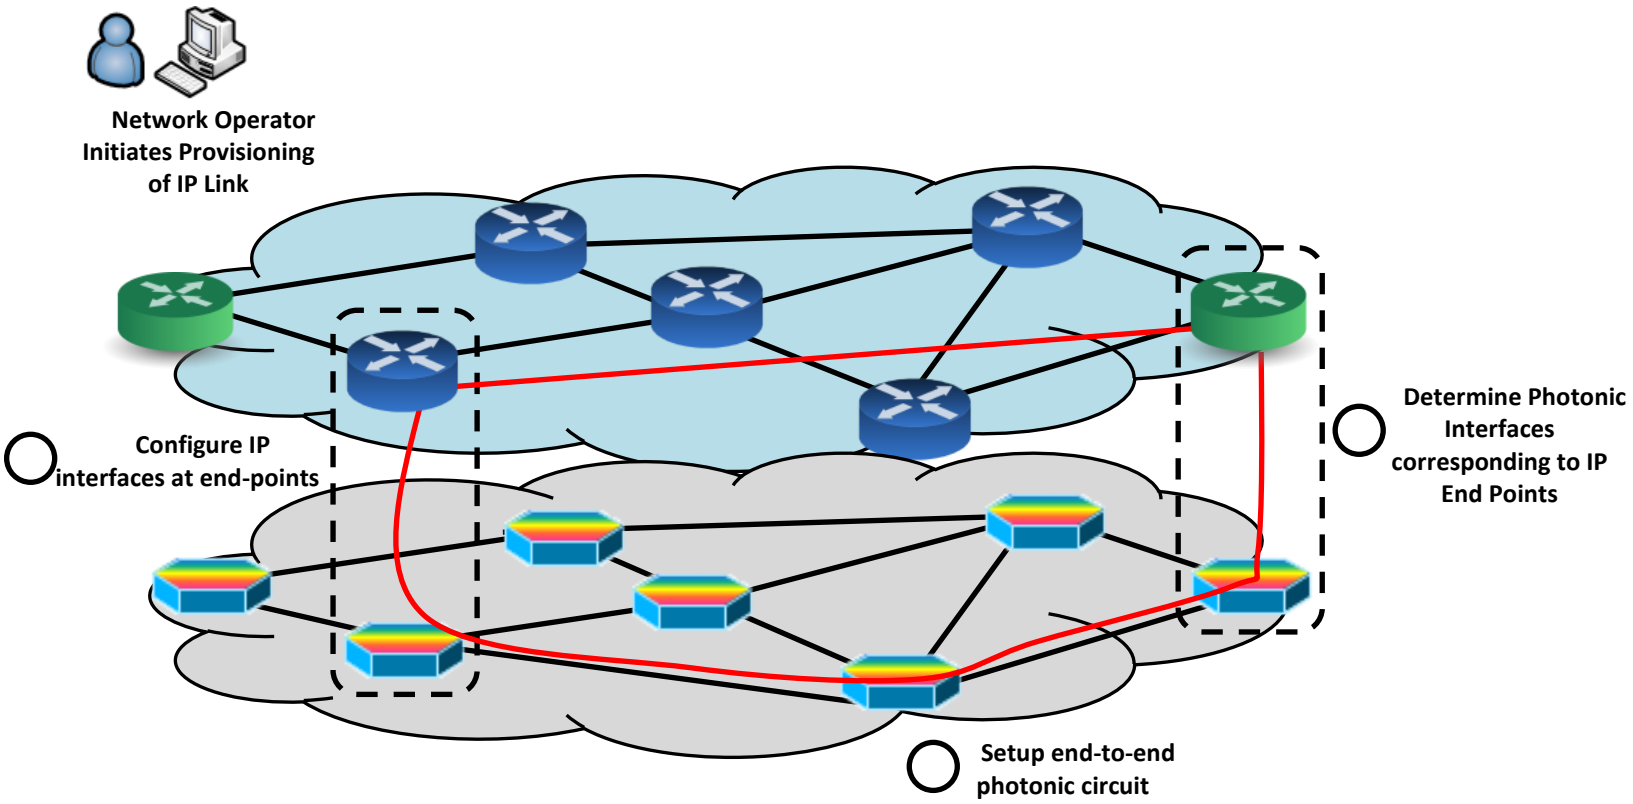

Figure 7: Automated IP Link Provisioning.

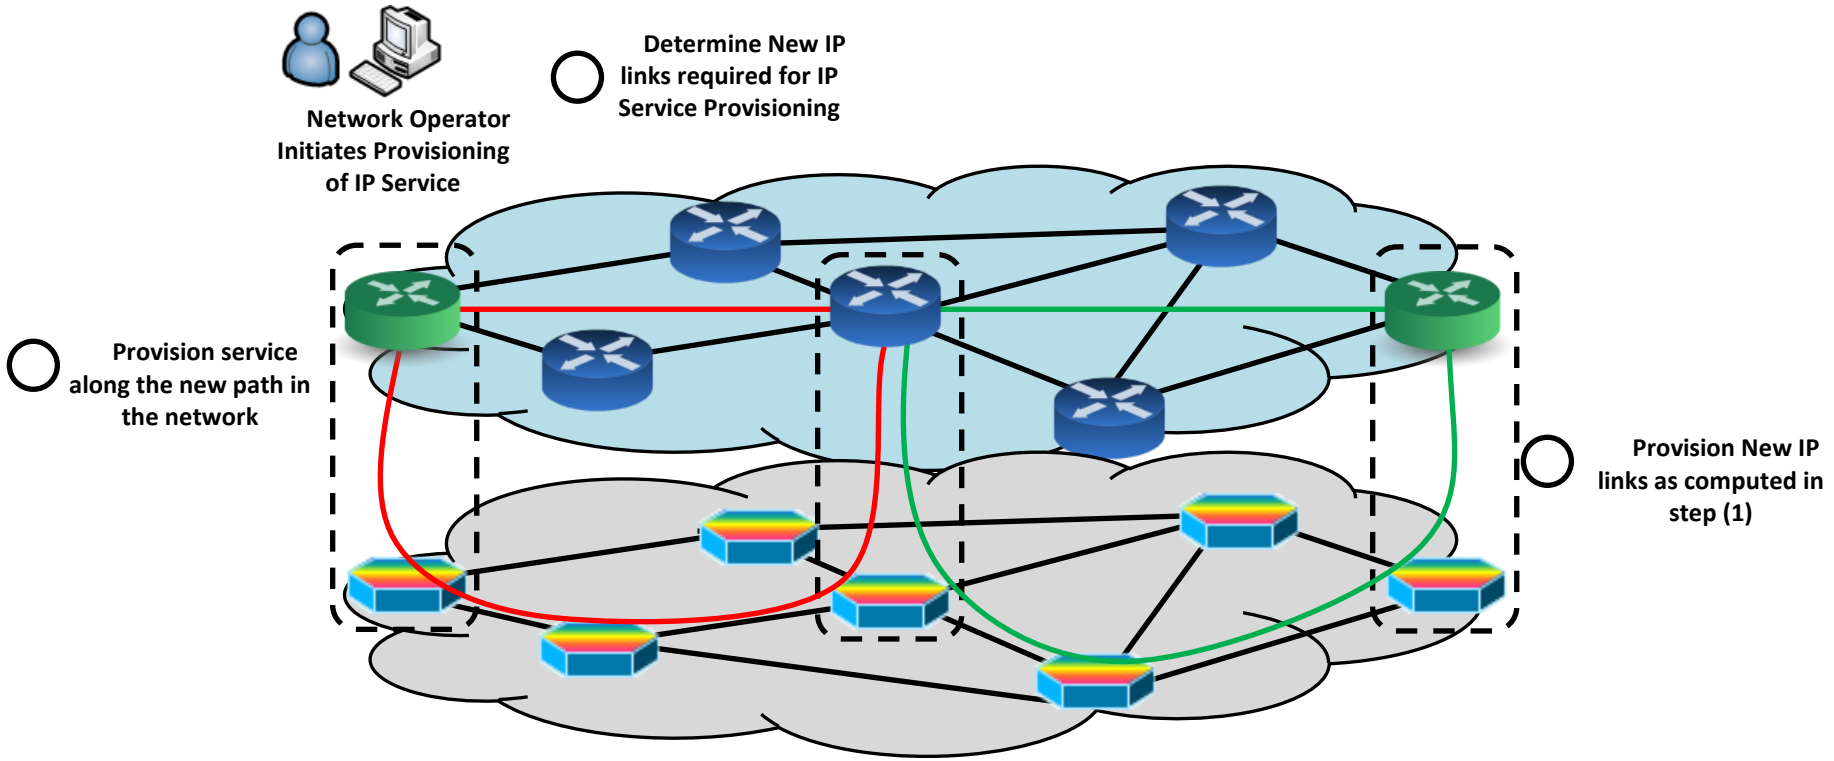

Figure 8: Automated multi-layer IP service provisioning.

4.1.3 Coordinated IP/MPLS Network Re-planning

The ability to facilitate coordinated multi-layer operations can also be exploited by network operators to facilitate network re-planning operations. Taking into account the current trend in which the traffic profile in operators' networks changes moderately every year, both in terms of traffic distribution and by overall traffic growth, and that from time to time new applications significantly change the traffic distribution, the deployed IP/MPLS network topology that was *planned* according to the company estimates, may become suboptimal and a periodic re-planning application may be considered in order to address this issue. While the issue of designing and re-planning the network has been extensively addressed in the literature, and may be performed using an external NMS or custom optimization tools and algorithms, the migration from the existing IP/MPLS network topology to the new network topology currently involves a set of manual interventions, prone to errors, which may lead to network operation breakdown. Thus, network operators tend to avoid re-planning as much as possible and prefer to cope with sub-optimality. Using the mechanisms as described in 4.1.1 and 4.1.2, the ONE adapter can help to execute a network re-planning operation, and, in a coordinated and automated way facilitate the set-up and release of the necessary transport circuits, as well as orchestrate the necessary configuration actions in the IP/MPLS layer. Here, the ONE adapter must also take care of the appropriate coordination when initiating different configuration operations in the IP and transport NMS so that topology and service reconfigurations are performed in a specific (e.g. make-before-break) order to ensure that no traffic is lost during this operation.

In this use case, the main role of the ONE adapter is to interpret high-level requests from an operator's IP unit and convert them into a set requests that can be understood by a path computation entity and the underlying carrier-grade NMSs to facilitate IP provisioning. As stated before, the ONE adapter does not replace or duplicate functions of the existing IP and carrier-grade NMSs, but rather offers an interface to *translate* requests for link and or service provisioning into operations and interactions between the IP and the transport NMSs, thereby bridging the *communication* gap between two management ecosystems. ONE targets an overall response time in the order of tens up to few hundreds of seconds to facilitate IP service provisioning in carrier networks, depending on underlying technologies used.

4.2 Use case 2: IP/MPLS Offloading

In the use case scenarios presented in Section 4.1, the driver to initiate multi-layer operations was a manual (operator-generated) request for provisioning IP links and services. In this use case, we plan to demonstrate the ability to automatically trigger multi-layer operations based on policies defined by the network operator using the ONE adapter. A typical application of policy-driven activity in multi-layer (IP/MPLS over optical) networks is the so-called "IP Offloading" paradigm [IPo01], wherein the network responds to the increase in traffic of a particular service or network segment in the IP network by offloading IP traffic onto *optical circuits* to reduce the load on intermediate IP routers and links. The increase in traffic is typically caused by an elevated number of high traffic flows/services traversing an IP network segment which requires the intermediate routers to forward large IP traffic volumes. As the intermediate routers can be eliminated in the path, offloading excess traffic onto circuits helps reduce unnecessary traffic load on intermediate routers and can significantly reduce the headroom requirements in static IP links.

The decision process driving IP offloading is complex and is constrained by a set of rules to reduce the cost of running the network while ensuring that the network remains stable. The basic rules governing IP offloading are:

- A circuit may be established to *offload* applications across the core network (end-to-end) or can be established to offload aggregate service traffic across overloaded network segments.
- Routing rules must be configured at the ingress and egress of the established circuit to ensure that only specific traffic (as determined by the offloading logic) are offloaded onto the circuit, and that the routing of other flows remains unaffected by the offloading operation.
- The creation of new connections and enabling new interfaces should not pose a threat to the adjacency scalability in terms of possible flooding upon elements failure.
- The offloading logic used to calculate bypasses should take into account costs incurred both in the IP/MPLS and transport network layers in order to optimize overall cost. For example, a single 1 Gbps flow should not be offloaded over a single 10 Gbps circuit, as it would be underused and would prove to be cost inefficient.

Automation of the policy-based IP offloading process poses a number of challenges. First, the policy configuration should initiate the offloading operation automatically when a critical traffic load is achieved in the network. Thereafter, using available network information such as IP link loads and the multi-layer topology information, the offloading logic must determine (1) the application/service traffic to be offloaded, (2) the routers/links which should be *bypassed* and (3) the required bandwidth for the circuit(s) to support the offloaded traffic. This information is used in conjunction with the operator's defined policies on: (a) the assignment of IP interface addresses, and (b) the configuration of routing rules for offloading to determine the set of operations required for IP offloading. Note that the routing rules must be configured in a manner in which only selected application/service traffic is offloaded and routing of other traffic in the network is not affected.

As can be seen, the computation and coordination of operations required for IP offloading is complex, and misconfigurations can significantly affect the regular operation of the network. If not configured properly, the introduction of multiple interfaces in the network can significantly affect the performance of the latter. For example, the possibility of constructing a dense mesh among a large number of routers may cause adjacency explosions, meaning that, in case of a link or a node failure, a large number of routers would be flooded with updates leading to routing instability. Also if routing rules are configured incorrectly, they can lead to re-routing of other traffic in the network, causing temporary network instability. For instance, if static routing rules are used to offload traffic onto circuits, in case of a failure of the circuit, the traffic will not be re-routed automatically via the routing protocols and the offloading paradigm must intervene and re-configure the static routing rules to restore the offloaded traffic.

In this use case, we will demonstrate the capability of the ONE adapter to proactively react to network situations based on pre-configured policies. We present two scenarios that demonstrate the ability of the ONE adapter to offload both: i) application-based traffic and ii) aggregated IP traffic.

4.2.1 End-to-end Application Traffic Offloading

In this scenario, we consider an application-based offloading scenario for commercial Video-on-demand applications. In this use case, the network attempts to offload high-bandwidth application-specific traffic end-to-end across the IP core network using optical circuits, in an attempt to drive

down cost. This approach has been successfully applied to scientific applications requiring high-bandwidth long-lived application flows in different projects [Lam01, I2Ph01], and in this use case, we will apply the same principles to a commercial Video-on-Demand (VoD) application.

Figure 9 depicts the traditional VoD application scenario, where on-demand video is served by the VoD server across the core network to the access network. The characteristics of VoD traffic make it a good candidate for offloading over the transport network across the network core.

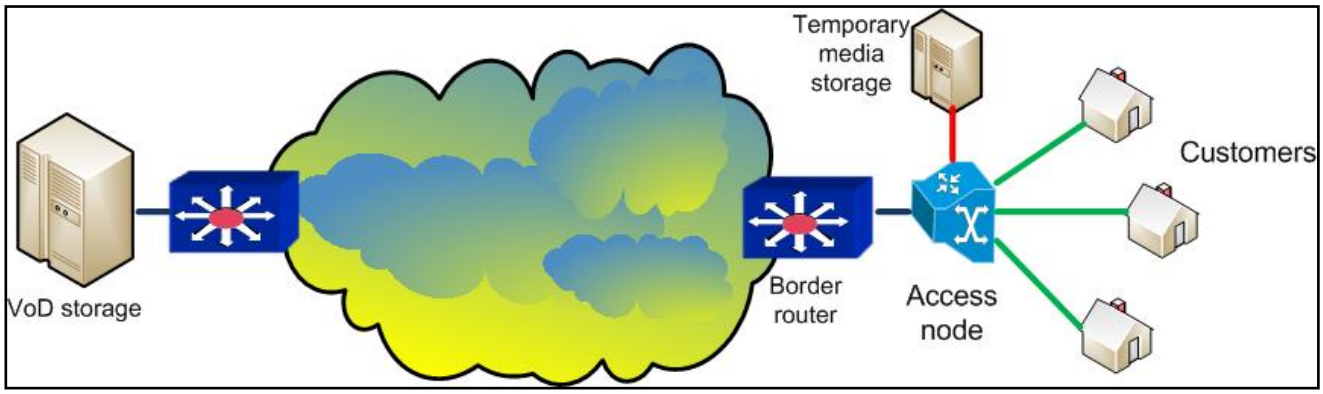

Figure 9: IP offloading use case: VoD traffic.

At first, the network is configured so that the VoD application traffic traverses the core IP routers as shown in Figure 9, because the amount of VoD traffic is not sufficient to justify the cost of creating a dedicated circuit across the core. As a result, the VoD traffic traverses multiple IP hops in the core network from the VoD storage site to the border router. An increase in the demand for the VoD application can lead to overloading of network segments inside the core and loss in QoS for the services offered over these segments. Typically, network operators upgrade the link capacities and IP router interfaces to address these situations, which lead to negative economic margins. However, as the photonic equipment is cheaper and the cost-per-bit is much lower than that of the IP equipment, large application-specific flows could be offloaded (in this case VoD traffic) across the IP network using an optical circuit.

In this scenario, the ONE adapter will initiate operations when the VoD application traffic increases beyond a pre-defined threshold and will attempt to offload this traffic across the core network using an optical circuit. When the application traffic increases beyond this threshold, the ONE adapter will be notified by an external entity (e.g., a monitoring system assessing the application load). The ONE adapter will first identify the border routers associated with the application end-points, and will check with the IP NMS if these routers have available interfaces to support a new optical circuit. The ONE adapter will then request the PCE to compute an optical circuit between these routers in the transport network. If the circuit is available, the ONE adapter will instruct the transport network to setup the computed circuit, and will then provide instructions to the IP NMS to configure the IP interfaces. The IP interface configurations will consist of IP addresses which are outside the OSPF area of the core network, and will also consist of forwarding policies, which will instruct the VoD application traffic to be re-directed onto the optical circuit. In this way, the establishment of the new circuit will not affect the traffic inside the IP network core, and only the application-specific traffic will be offloaded effectively across the network (see Figure 10).

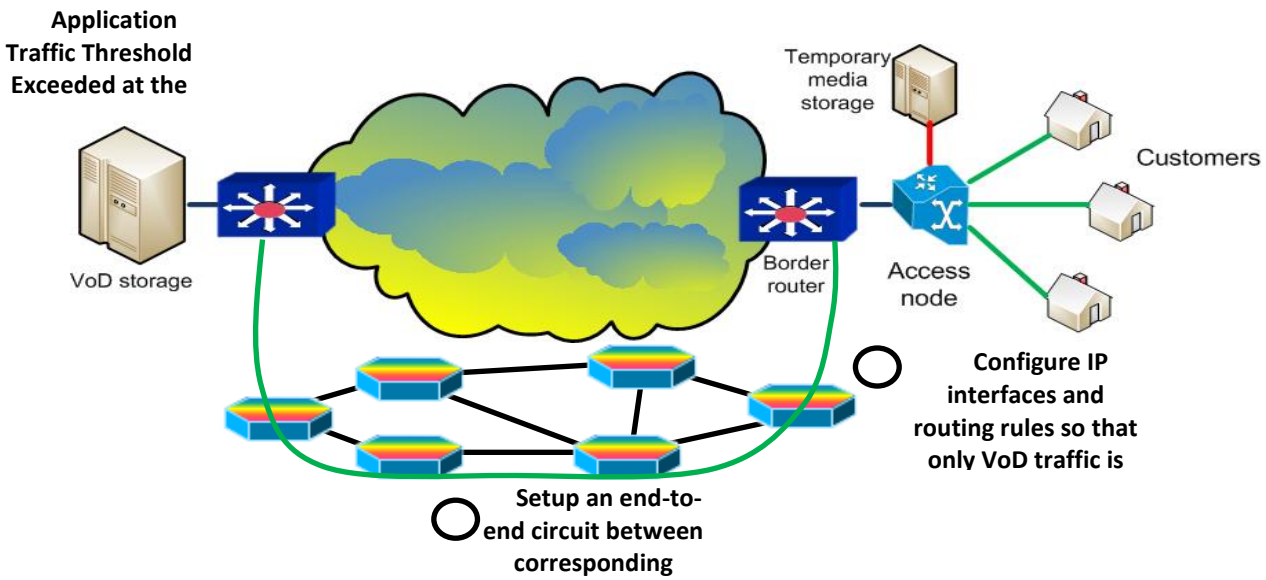

Figure 10: Application offloading use case: VoD traffic bypassed over the photonic mesh.

4.2.2 Aggregated IP Traffic Offloading

In a more advanced avatar, IP offloading can be used to offload traffic across overloaded links/routers/network segments inside the core of the network. In this scenario, no single application/service may justify the cost for being offloaded over an optical circuit (unlike 4.2.1), but traffic aggregated over multiple services/applications can better justify cost for deploying a new optical circuit, while at the same time alleviating overloading in the IP network.

The typical network scenario assumed here is shown in Figure 11, where different applications/services are served over the IP core. Each IP core router is co-located with a photonic switch which can be used to set-up optical circuits. In this scenario, the measurement infrastructure in the IP network monitors the network load on each IP link, which is facilitated by either using dedicated infrastructure to measure IP link loads or via special SNMP traps on IP equipment supporting link-load measurements. The IP offloading process is initiated when the utilization of IP link(s) in the network increases beyond a pre-defined threshold as seen in Figure 12. At this point, the ONE adapter will communicate with the IP monitoring and measurement infrastructure to retrieve network information including information about the IP links loads and IP routing information, and will use this information in conjunction with the multi-layer topology information to compute:

1) IP router pairs between which an optical circuit will be established to offload traffic.
2) Services/application traffic which must be offloaded onto these optical circuits.
3) Required bandwidth of the optical circuits to support offloaded traffic.

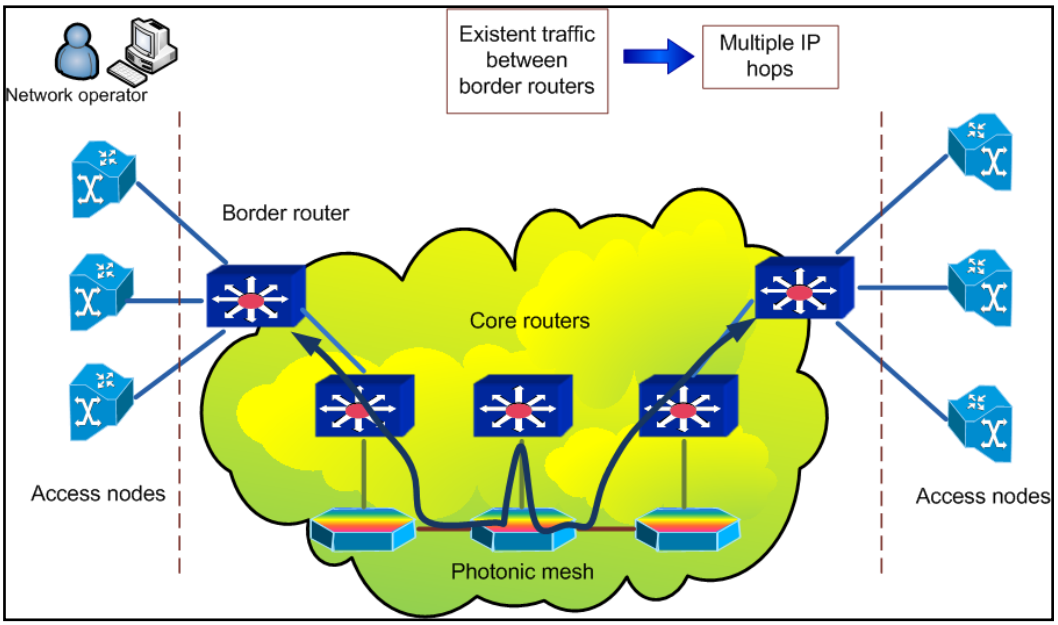

Figure 11: IP offloading use case: initial network configuration.

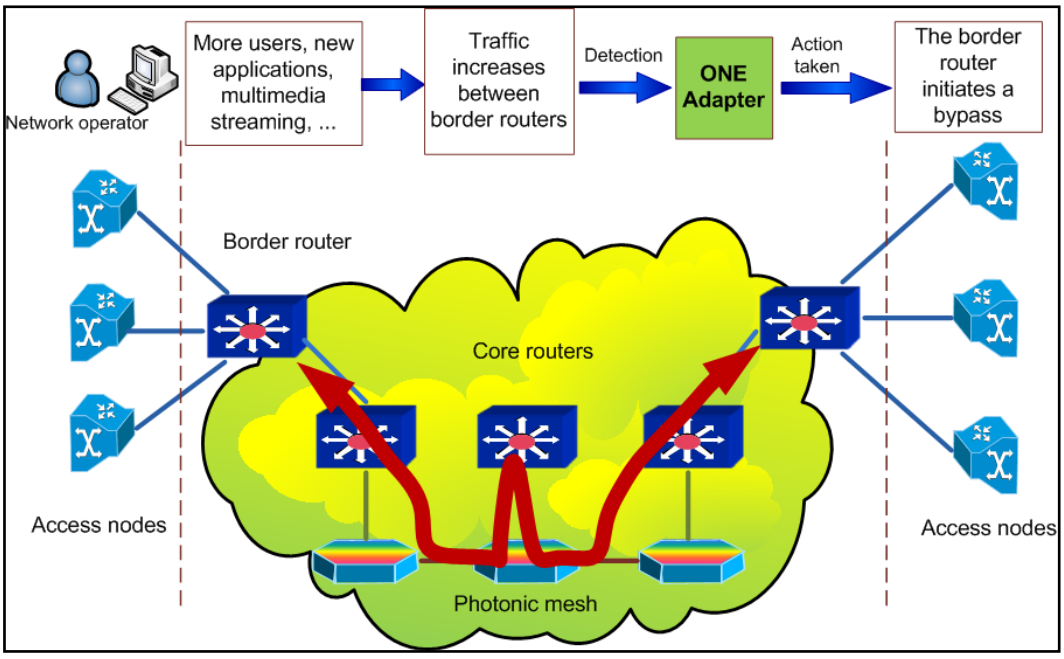

Figure 12: IP offloading use case: traffic increase.

Using these parameters, the transit traffic can be re-routed across the congested IP network segment over the optical circuit, as seen in Figure 13. To facilitate the same, the ONE adapter will first communicate with the transport NMS to initiate the setup of the required optical circuits. If all circuits are established in the network, the ONE adapter will then instruct the IP NMS to configure the corresponding IP interfaces at the ingress and egress routers. The interfaces will be assigned IP

addresses from a private IP address pool which is outside the OSPF routing area, to ensure that these optical circuits are not advertised as new links by the OSPF routing protocol. After driving the configuration of the IP addresses, the ONE adapter will use the IP NMS to configure routing rules to offload traffic onto the created optical circuits. Note that in core IP networks (given the high traffic volumes) it is difficult to apply forwarding rules to individual application flows, so typically the offloading operation will offload aggregate MPLS tunnels. This can be facilitated by modifying the MPLS forwarding entries at the ingress and the egress routers of the established optical circuits.

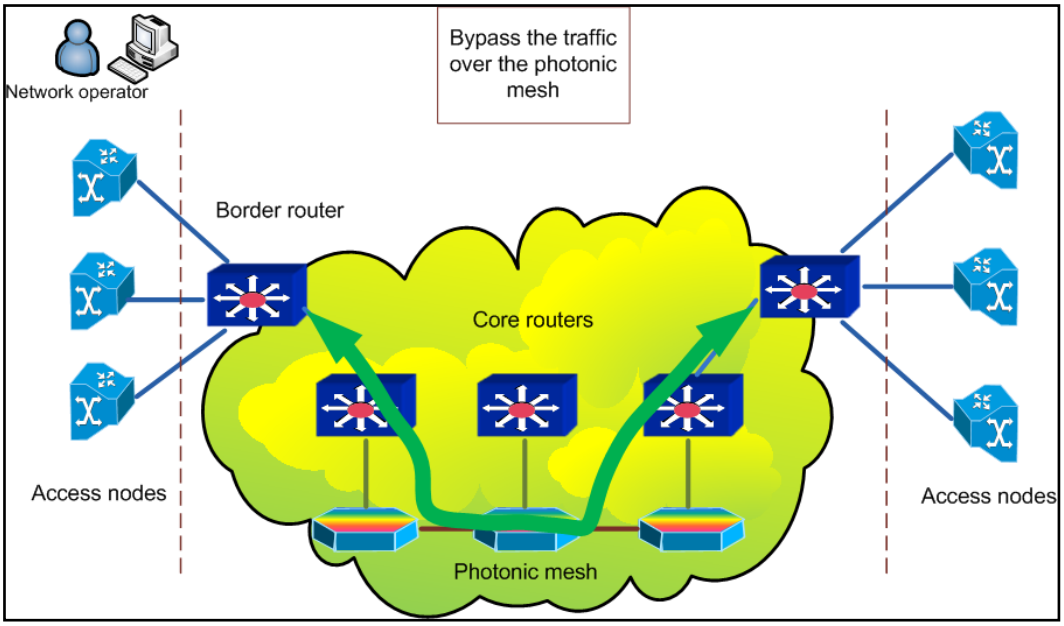

Figure 13: IP offloading use case: bypass over the photonic mesh.

This scenario will also demonstrate the operation of decommissioning IP offloading or *IP uploading*. This is a rarely addressed and complex scenario [Inf02], and addresses the issue of when to decommission an existing IP or bypass. The decommissioning of a link is not only dependent on the traffic on the given link but is also dependent on the traffic on the alternate path in the IP network, and the ONE adapter will perform periodic computations to check if a provisioned IP link/bypass should be decommissioned to save resources, while not affecting the network adversely.

These use case scenarios are designed to demonstrate the ability of the ONE adapter to integrate existing policy and alarm features in the IP network to automatically initiate multi-layer operations. As these operations are potentially intrusive, the ONE adapter may also be configured to suggest a possible solution and wait for approval from the network operator before initiating configuration and provisioning tasks.

4.3 Use case 3: Coordinated Self-healing

In the current network scenario, the separation of the IP and optical management layers has led to highly redundant and un-coordinated protection schemes, where each layer is equipped with its own

protection capacity and mechanisms. In current carrier networks, each circuit used to provision an IP link in the transport network is protected using a dedicated 1+1 protection scheme, and each IP interface is duplicated to protect from interface failures. Moreover, each IP link is designed with large headroom (peak load link utilization ~ 30-40%) so as to ensure recovery in the IP network in case recovery in the transport network fails. In order to assure that the recovery mechanisms of the IP and optical transport layers do not interfere, resilience mechanisms in different networks operate in different timescales.

Currently, in order to initiate protection first in the transport network and then in the IP network, network operators configure timers for responding to failures in a static manner, with response times that differ in several orders of magnitude. Therefore, while response times of transport networks are kept below 50 ms, response times in IP networks are configured to respond in the order to tens of seconds, thereby ensuring that any activity to restore the circuit in the transport layer is finished before a recovery in the IP network is initiated. This procedure leads to very high costs for network survivability and does not provide additional recovery guarantees.

Despite this redundancy in the protection schemes, there is a need for coordination, to assure the reliability of the recovery actions. Also, automated recovery from unplanned failures is not feasible at present. Recovering from a complex unplanned failure requires manual interactions between the IP and the optical network management teams, and yet network recovery cannot be entirely guaranteed. Indeed, configuration inconsistencies as well as failures that can hardly be protected by planning (e.g., multiple link failures) are practical examples that can lead to multiple human interactions to recover a failure in spite of the protection redundancy.

By using coordinated protection through the ONE adapter, carriers could:

1. Compose mixed protection schemes, where some services or traffic flows may be protected in the IP layer while others may be protected in the transport layer, thereby reducing the protection redundancy in a controlled way.
2. Make use of their NMSs to trigger coordinated restoration actions in addition to coordinated protection, facilitating recovery from unplanned failures and thereby reducing the overall restoration time.
3. Reduce CAPEX by decreasing the level of duplicated protection schemes.
4. Increase network availability by quickly recovering from catastrophic failures (e.g. double failures, natural disasters, etc.).

We now present two scenarios which are targeted to demonstrate the advantages of multi-layer co-ordination in recovering from simple as well as complex network failures using the ONE adapter.

4.3.1 Coordinated Recovery from Traditional Single-link Failures

Note that coordinated multi-layer protection of traditional single link failures can significantly reduce the CapEx by reducing both the number of redundant protection paths in the transport layer and the number of redundant IP network elements, as well as the excess protection capacity maintained in the IP links. Moreover, the time to recover a failure can be reduced as well: for example, in case that transport network cannot recover a failed link, the ONE adapter can automatically trigger restoration in the IP layer instead of waiting for the timers in the IP layer to expire.

In this scenario and as an example of coordinated protection, we will present the capability of the ONE adapter to coordinate recovery in a multi-layer environment. In our example, we consider a network scenario where circuits do not use a dedicated 1+1 link disjoint path protection, but instead employ a shared path/segment protection variant. Moreover, in the IP network, links are provisioned with smaller headroom (peak load link utilization ~ 60 %). In case of a failure in the transport network, the ONE adapter is notified by the transport NMS of the same, and waits for the recovery mechanism in the transport layer to finish the recovery operation. In case the path cannot be recovered in the transport layer, the ONE adapter then instructs the IP NMS to initiate recovery. Here, we will attempt to demonstrate the time-saving that can be achieved by facilitating direct coordination between the IP and the transport layers. Furthermore, in case that all services cannot be recovered in the IP layer due to lack of excess capacity, the ONE adapter will then use a mechanism similar to the IP service provisioning scenario presented in 4.1.2 to compute an alternate route to a different router in the IP network to recover IP traffic. Using this mechanism, we intend to demonstrate the capability of the ONE adapter to save on CapEx incurred due to over-provisioning in both the IP as well as the transport network, while providing the same level of protection as current networks.

4.3.2 Coordinated restoration from unplanned dual network failures

In this advanced scenario, we will present the ability of the ONE adapter to facilitate recovery from an unplanned failure. An example of such a failure event is the occurrence of multiple failures that may leave a router or even a complete network segment isolated. While multiple failures are typically considered very rare events and are not considered in typical network planning exercises, recent studies show that these events may not be as rare. Given the diversity of potential factors for failures in a network, varying from hardware failures, power source failures, software mis-configurations, and fibre cuts to name a few, failures can be very common events in a large network. For example, as published in a Federal Communications Commission report in the U.S., it was found that a typical long-haul network experienced 3 fibre-cuts annually per 1500 km of fibre, which translates to approximately a fibre cut every 4 days in a typical long-haul network with 45,000 km of fibre (see [Gro03]). In such a scenario, it is possible that a new error may occur before the network can recover from an existing error. Moreover, with IP networks now supporting numerous mission-critical applications, the economic implications of long network outages can be catastrophic. An analysis of the financial impact of network outages presented in [Fis03] (see Table 3) shows the economic impact of these failures, which has now led network operators to ensure that all their transport switches supporting core IP routers are at least 3-connected, in order to recover from even double fibre cuts.

Financial Impact Analysis: Average Hourly Loss of Downtime	
Package Shipping Service	$ 28,250
900 Number Services	$ 54,000
Airline Reservation Centres	$ 89,500
Catalogue Sales Centres	$90,000
Pay Per View Events	$ 150,000
Home Shopping Channels	$ 139,000
Credit Card Authorization	$ 2,600,000
Brokerage Operations	$ 6,450,000

Table 3: Average loss during downtime.

As traditional recovery mechanisms only consider a single failure, recovery from a double failure is currently a manual task which requires significant time. In such a case, we will demonstrate how the ONE adapter can request the computation of restoration paths to an external subsystem (e.g., the PCE). Upon receiving the computed paths, the ONE adapter can then use the already existing workflows for circuit setup and network configuration to automate the necessary changes in the network.

Note that while self-healing in its strict meaning implies that the network can recover any failure, we here refer to self-healing as the capability of the network to recover most important customer or business services (whenever possible) in different failure environments.

The reference network for this scenario is shown in Figure 14. The network is initially stable, with MPLS tunnels coupled with load balancing driving traffic towards the Internet. We assume that every MPLS tunnel has a backup tunnel through the duplicated transit node (T1a and T1b in the figure). The transport network has its own topology which may differ from the IP topology. This will be useful in order to get recovery capabilities through, at first, inexistent IP adjacencies.

In this scenario, we will consider IP equipment failures, and differences in case of other failures will be explained later to offer the complete picture of what we can expect from the ONE adapter in every failure scenario.

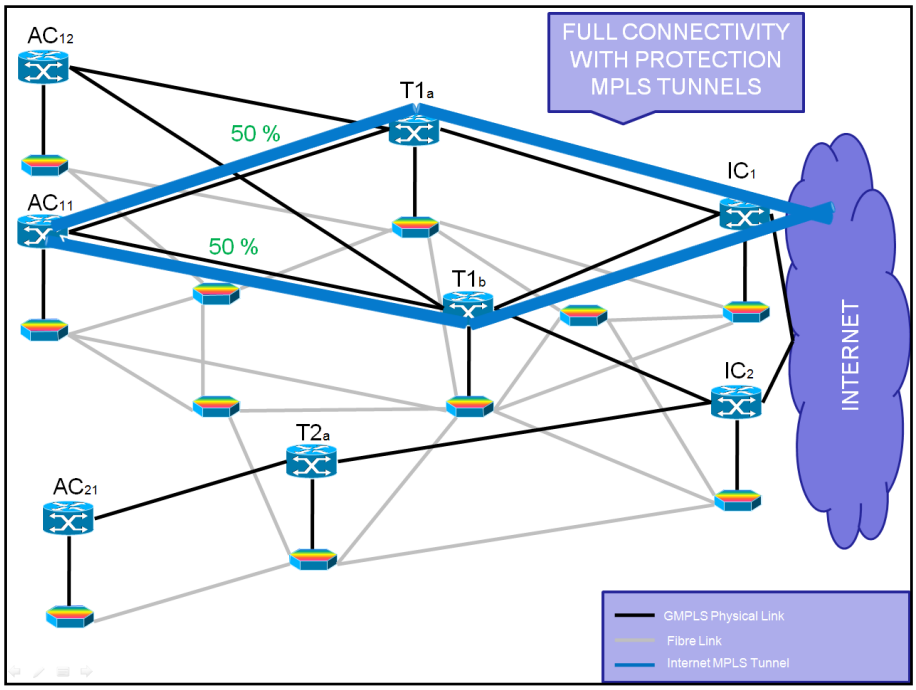

Figure 14: Self-healing use case reference network.

In the first failure event, as indicated in Figure 15, the transit router T1a fails. As mentioned above, this can be caused due to a multiple reasons such as old equipment, natural catastrophes, lack of maintenance or simply bad luck. In this scenario, the protection path via T1b now absorbs the full traffic load. Note that the network continues working correctly but it is near to lost connectivity in

case of another failure before router T1a is repaired. Currently, without the self-healing mechanism, it will be mandatory to repair in a short period of time the node T1a because if a double failure occurs, it will be hard to get connectivity recovered quickly with all the economic implications that this entails. Getting repaired in a short time period has high OpEx implications because a quick fixing service is in fact expensive.

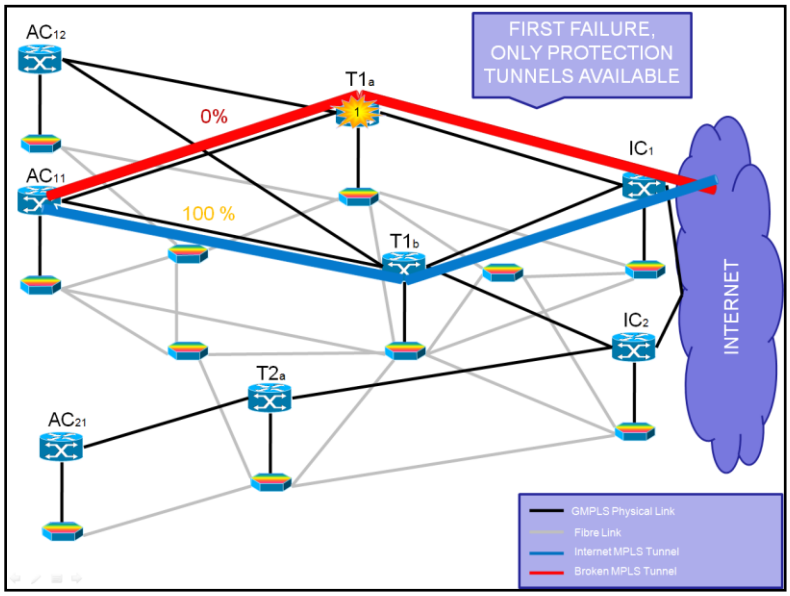

Figure 15: Self-healing use case: first failure.

In the event of a second failure (Figure 16), both routers AC_{11} and AC_{12} now loose Internet connectivity, which means loss of Internet connectivity for a large number of users. As indicated previously, such a situation can be very damaging financially, and also derives in considerable damage to the Internet service provider's reputation.

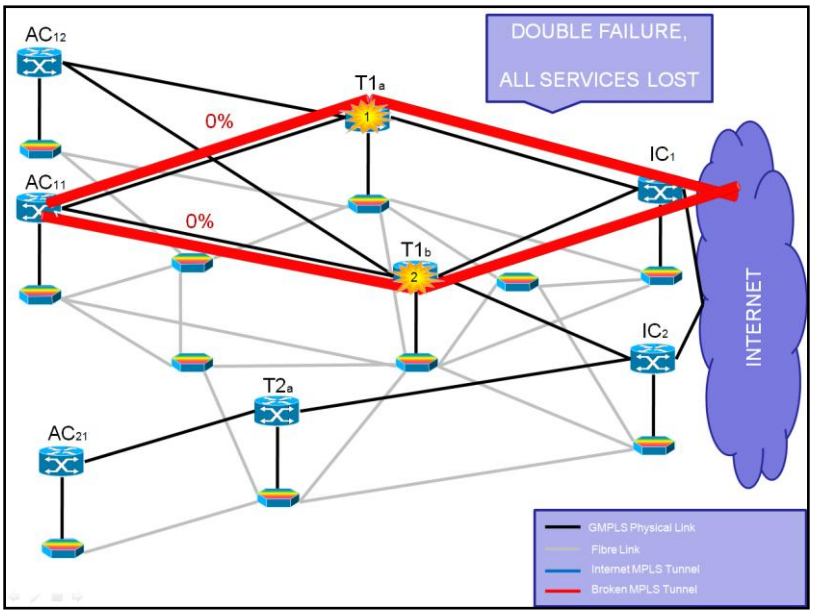

Figure 16: Self-healing use case: double failure.

In such a scenario, the ONE adapter can be used in order to create new IP adjacencies in order to restore connectivity and re-provision (at least) the premium services from the disconnected routers (see Figure 17). In this manner, connectivity will be recovered in a relatively short amount of time while the operator attempts to fix/replace the damaged routers in the network.

Note that while we present this scenario with two catastrophic router failures occurring simultaneously, such a failure situation can also be realized by a combination of other failures such as multiple card or port failures, and even multiple fibre cuts. The ONE adapter in this scenario, must identify the loss of Internet connectivity of the router via alarms from the IP NMS indicating loss of connectivity to both routers T1a and T1b, and use an external path computation entity to compute new IP adjacencies to alternate transit routers (in this case T2a in Figure 17). The capability to automatically recover connectivity will not only improve service availability and reduce user complaints but will also significantly reduce the OpEx, as existing failures may now be repaired under more relaxed time constraints.

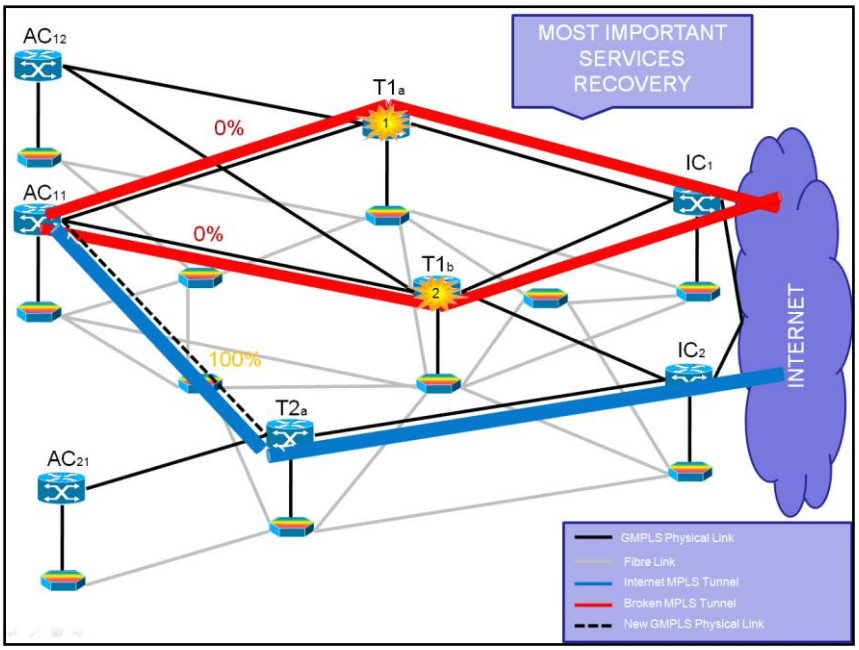

Figure 17: Self-healing use case: connectivity recovery.

Other self-healing consequences should be letting an operator consider reducing elements in the transit and core networks that are nowadays redundant. In case this reduction is not necessary or convenient, the element would allow automatic restoration against double failures, a feature that nowadays is mostly done manually, and that would come up with OpEx reduction.

4.4 Overview of the Use cases for the ONE adapter

In this section, we present a brief summary of the progression we plan to exhibit during the course of the project and also present directions for future applications using multi-layer co-ordination. For each use case, we present three scenarios, namely: short-term, mid-term, and long-term. Short-term capabilities can be deployed by facilitating communication of the ONE adapter with basic management functions in the IP and the transport NMS, and can model basic management capabilities such as process co-ordination inside the ONE adapter. Mid-term scenarios will demonstrate more complex multi-layer coordination functions via the ONE adapter, where logic processes will be required to compute multi-layer operations through the ONE adapter. This scenario will also require a tighter integration of the ONE adapter with certain network services, such as network and routing monitoring to facilitate multi-layer operations.

Finally long-term scenarios present our vision for the future of the ONE adapter, which can be realized by combining network centric services with service awareness. Note that while the short-term and mid-term scenarios mimic the use case scenarios presented in Sections 4.1-4.3, long term scenarios only indicate our vision and will not be implemented during the course of this project. A summary of these operations is presented in Table **4**.

Challenges	Short Term Scenario	Mid Term Scenario	Long Term Scenario
Coordinated IP Provisioning	Provisioning of an IP link by coordinated actions between IP and transport networks	Multi-layer path computation using PCE, and automatic provisioning of IP links to facilitate IP service provisioning	Service Template awareness, QoS monitoring and load balancing
Policy-driven IP Offloading	Threshold-triggered application offloading onto end-to-end photonic circuits	Threshold triggered aggregate traffic offloading onto photonic circuits inside the IP core; policy driven traffic *uploading*	Service and energy-aware traffic offloading; providing re-planning advice
Coordinated Self Healing	Automated coordination between recovery operations in IP and transport networks	Recovery from unplanned complex failure scenarios; maintaining connectivity and delivering services over newly computed paths	Proactive detection to prevent catastrophic failures, use of QoR differentiators.
	Management Centric Process Integration	**Network Centric Automated Operation**	**Service Aware Proactive Action**

Table 4: A roadmap towards coordinated actions between the Internet and the transport management systems.

4.4.1 Short Term

The Short Term scenario will encompass integration of basic network operations available in the IP and transport NMSs and will focus on multi-layer process integration while incorporating features of multi-layer topology awareness. In the scenarios presented here, the ONE adapter will demonstrate coordination of both IP and Transport layer mechanisms to get first simple integration of actions for resource provisioning, policy-based operations and coordinated recovery.

In the provisioning scenario, as described in Section 4.1.1, we expect the ONE adapter to use knowledge of the multi-layer topology and inter-layer interconnectivity to automatically determine interfaces in the IP layer and corresponding circuits (and their configurations) in the transport layer, and provision the same using basic interactions with the IP and the transport NMS.

For policy based operations, as presented in Section 4.2.1, we will integrate automated triggers based on application traffic thresholds into the ONE adapter, which will then initiate the process of offloading application traffic end-to-end across photonic circuits.

Coordinated Self-healing and recovery, as presented in Section 4.3.1, will require from the ONE adapter to coordinate recovery actions in the transport and the IP network without using static timers, while avoiding simultaneous recovery in both layers. Here, we will develop mechanisms in the ONE adapter to communicate with the IP and transport NMS to receive failure alarms, and to initiate actions of recovery so as to recover from a failure in one layer only.

4.4.2 Mid-Term

The mid-term scenario will require the ONE adapter to interact with advanced network services such as traffic monitoring, IP routing monitoring, multi-layer path computation functions, etc., in order to perform complex coordinated multi-layer operations.

For the resource provisioning case (as presented in 4.1.2 and 4.1.3), the ONE adapter will use external path computation entities to compute optimal multi-layer paths for incoming service requests, and will provision and configure new IP links as required to provide these services. The logic for service provisioning will also integrate business logic which will consider the cost of setting up a multi-layer path against the service request before initiating service provisioning.

In the policy-based operation use case (presented in 4.2.2), the ONE adapter will integrate offloading logic which will allow to respond to overloading of network segments, including components such as routing processors, IP links etc., and will allow offloading of aggregate traffic inside the core of the network. The ONE adapter will also be expected to provide advice to decommission existing IP links (IP uploading) based on utilization information in order to save network resources when possible.

Coordinated self-healing and recovery challenges in the mid-term scenario will target enabling multi-layer recovery using the ONE adapter to recover from unplanned multiple failure scenarios, such as isolated nodes as presented in 4.3.2.

4.4.3 Long-Term

While the short and mid-term scenarios are designed to show that the ONE adapter can facilitate integration and coordination of complex network processes, we envision that the ONE adapter could be extended after the end of this project to incorporate features of service awareness in order to provide network operators with insights into improving the customers' experience.

In case of provisioning, the ONE adapter when requested, may perform functions of service aware load balancing, and facilitate complex computations such as multi-layer availability and QoR computations, thus helping the operator in delivering the required QoS.

In case of policy based management, the ONE adapter can, in conjunction with facilitating offloading, also track service arrival and traffic trends, and use these to provide re-planning advice to the operator in order to reduce the frequency of offloading operations.

Finally, in the case of coordinated self-healing and recovery, the ONE adapter is envisioned to include dynamic recovery computation mechanisms which also include application service aware parameters and Quality of Resilience (QoR) differentiators.

5 Requirements for the ONE Adapter

5.1 High Level Functional Requirements

In order to facilitate multi-layer interactions for IP provisioning, automated IP offloading, and coordinated self healing and recovery as described in the use cases above, the ONE adapter requires functions to communicate and correlate information from different measurement and service management systems. The functional blocks are needed by the ONE adapter to facilitate the aforementioned inter-layer operations. They can be classified into six categories of functional blocks.

5.1.1 Functional Block: Topology Information

Topological information functions are required to leverage the existing mechanisms used by the network operators to discover and identify network elements in both the IP and the Transport NMS. Leveraging existing functions of topology discovery will facilitate simpler correlation with existing systems. The topology functions should not only describe the connectivity but have to also describe the physical layer and the data layer supported by interfaces in the IP as well as the transport network interfaces, facilitating seamless provisioning. The topology functions should provide mechanisms to specify used/unused inter-layer inter-connections such as port/VLAN mappings, which are essential when commissioning/decommissioning IP links in both the IP network as well as the transport network.

5.1.2 Functional Block: Routing Information

Routing information functions are required to provide necessary details about the routing in IP as well as transport networks. Routing information (e.g., the current routing and link weights in the IP network) is essential to evaluate the effect of adding/removing a link in the network. It is also necessary to have routing information for existing services provisioned via explicit configuration to ensure that any operation in the IP network does not disrupt service delivery. In the transport network, routing information is necessary to facilitate computation of circuits for provisioning IP links.

5.1.3 Functional Block: Network Measurement

Network measurement functions use existing mechanisms to measure network parameters. Network measurements such as available capacity, QoS parameters (delay, jitter), link utilization, and traffic matrix information are required in conjunction with the topology and routing information to evaluate the effect of changing the IP network topology and, hence, are essential to the operation of the ONE adapter. Alternately, alarms such as neighbour loss in the IP layer and link/interface failures in the transport network layer can automatically be exchanged between the layers. The ONE adapter can, in such cases leverage, the alarm correlation functions of existing monitoring systems. The ONE adapter can also use other performance metrics such as packet loss, bit-error rate, to identify impending failures and take preventive measures to ensure service delivery is not affected.

5.1.4 Functional Block: Service Description

In order to facilitate service provisioning, the ONE adapter requires information to map abstract service requirements of the IP network operator to their corresponding network performance requirements, as defined by the network operators. The description can contain mapping of service to parameters such as availability, constraints on delay and jitter, technology used.

5.1.5 Functional Block: Configuration Logic and Policies

The ONE adapter needs information in the form of algorithmic logic or policies to compute the required operation set when responding to network or user-initiated multi-layer events. For example, the ONE adapter needs an algorithm to compute new IP links when facilitating multi-layer failure recovery which needs to be provided by the network operator in order to conform to their specific requirements.

The ONE adapter also needs information about the policies used by the network operator to configure new equipment, connections as well as services in the network. In conformation with the ONE adapter approach, the way to ensure uniformity in configuration would be to interact with the existing infrastructure such as T-NMS and IP-NMS (if present) which are currently used to configure devices. However, for operations such as automated IP link provisioning which are not common in current IP networks and are largely governed via manual intervention in current networks, the ONE adapter needs additional information. These configuration details comprise routing policy rules, IP addresses for new interfaces, QoS constraints on transport layer circuits etc.

5.1.6 Functional Block: Operator Interaction

Functions to facilitate operator interaction with the ONE adapter are necessary for three primary reasons. Operators require a mechanism to initiate multi-layer operations by interacting with the ONE adapter, and may be required to confirm/decline automated operations as computed by the ONE adapter for policy driven functions such as IP offloading. The operators might also be required to provide the configuration logic and policies used to compute the ONE adapter multi-layer operations.

To facilitate the same, it is necessary to have an interface using which the network operator can interact with the ONE adapter via a secure interface after authentication and authorization to perform the abovementioned operations.

5.1.7 Functional Block: Authentication, Authorization and Security Functions

The ONE adapter must support functions required to communicate with external systems and services such as the T-NMS, IP-NMS, PCE, and the Monitoring and Measurement systems to facilitate reliable and safe interaction. The ONE adapter must also support functions to authenticate and authorize users (typically IP or transport network operators), who can approve or trigger operations in the ONE adapter.

5.2 Business Requirements

Last decade was evidence of several standardized and non-standard frameworks and solutions to address the issues of multi-layer and multi-domain networks, covering also the IP network and the carrier network issues. Although network operators look for these frameworks to address coordinated traffic engineering, fault management, and protection, none of the frameworks found wide deployment in practice. The unwillingness of network operators toward the deployment of these frameworks is largely related to business requirements.

Considering the above mentioned points, any new solution, which aims to address the issues of multi-layer, should have an advantage over the existing practice and proposed frameworks. That means, the operational and capital expenditure should be justified by the technological and economical benefits. It should fulfill the real network operators' demand and address current shortcomings, enabling the operators to run their business in a cost effective way.

As approaches to cover the issues of full system automation have been found to be undesirable solutions to the above-mentioned problems for now, there is the need to address the issues in an effective and integrative way. An effective and integrative management system must focus on the successful transmission of the packets, utilization of network, service quality, network availability, and network maintenance [Rol01], while making possible the realization of new services, quality of service improvement, customer satisfaction improvement, and reduction of human intervention, and taking into account the network operators current needs.

The current isolation and lack of coordination in the management of IP networks and Carrier networks is the main cause of duplication of functions, long provisioning times, and large amount of human resources which leads to high OpEx and CapEx. The ONE consortium aims to address these issues in an easy to deploy and cost effective way by allowing interaction and coordination between the two layers. We believe that the interaction and coordination between the IP and the Carrier layer will significantly reduce the total cost of ownership (TCO) and the operational expenditure in both IP and transport layers.

The following four subsections address four business requirements for the design of the ONE adapter in detail.

5.2.1 Attractiveness through low costs, low deployment time, and low network changes

Attractiveness has been found to be an important driver of success for any framework. As stated in [Mad02], attractiveness drivers can be classified into economical drivers, resource-based drivers, and social-interpersonal drivers. One of the most important reasons for a low deployment rate of a new framework in practice is the low attractiveness of the solution, caused by the longer deployment time, high operational cost (OPEX), high capital cost (CAPEX), and huge change in the structure and operation. Most of the proposed network management frameworks, however, require huge changes in the network structure and need long deployment times. Changing the network structure, adding or removing new hardware and software may require long non-operating times of the network.

Additionally, deployment of new technologies and frameworks require familiarization and knowledge dissemination about the operational processes, which is also time consuming and costly And current practice show that operators are not willing to implement little known processes and operations, in order to minimize the risk of catastrophic problems in the network.

Similarly, the capital expenditure for the new hardware and the software is also an important factor for the attractiveness of the solutions. Network operators prefer cost-efficient solutions in order to keep their networks competitive, especially when considering the fact that the price of bandwidth decreases and network traffic increases over time.

Therefore, no framework will find a wide deployment, if it is not possible to deploy it in the shortest time with no or little change to the network structure and, of course, with low operational and capital expenditure. In order to address the IP-Carrier layer issues according to today's IP and the Carrier networks real needs , the ONE consortium works closely with world-wide IP and carrier network operators to find and address the real need and shortcomings in both layer. The design of the ONE adapter will also ensure that multi-layer coordination is achieved while existing network and management systems in the IP and the transport network layer remains separate.

5.2.2 Improved performance without loss of autonomy

Existing control and management plans of IP and Carrier layers lack coordination in terms of multi-layer resilience, and route trace optimization. The lack of coordination leads to capacity waste and low quality of service. This isolation of carrier and IP network management prevents dynamicity in the heterogeneous network environment.

Research shows that network autonomy and neutrality is a sensitive issue in today's network environment, as no operator is willing to lose control over its network [Tre03]. The research efforts toward the integration of IP and Carrier layer management are not of interest to operators, since the reason for that is the unwillingness of operators toward use converged management systems. The ONE model aims to offer a solution which allows network efficiency improvement and dynamicity while assure the autonomy of each layer.

5.2.3 Customer satisfaction

Practice shows that no business can be successful without satisfying the consumers with an acceptable quality of service. Today's carrier networks lack customer satisfaction, especially in terms

of low service provisioning times, required quality of service for different applications at a reasonable price, and low recovery times. These issues must be considered when designing a new solution, as the multi-layer and multi-domain frameworks proposed by standardization organizations already suggest solutions.

Previous research in the area shows that full system automation is not a desirable solution for the network operators at this point in time, as most of the network operators prefer price negotiation rather than fixed prices to capture consumer surplus. On the other hand, the existing long provisioning times with large human intervention is vulnerable to inserting errors into the network, which is also not acceptable from the carrier and IP network operators' point of view. Therefore, a balance has to be found between the two requirements.

Although it is difficult to measure the QoS as agreed in SLAs and provided by a carrier provider in today's multi-layer and multi-domain environments because of the 'multi-domain dependency', IP network operators look for ways to measure QoS they receive and accordingly deal with their transport providers. However, traditional methods of control and management in both layers do not support free exchange of information between both layers. There is a need for a reasonable change. We expect that the exchange of information between the two layers leads to an increase in the quality of service not only for IP network providers but also for end users.

5.2.4 Complexity reduction

Providing support for new services derived from new protocols and applications is essential to run the business. But the current complex network environment does not allow achieving these goals easily. Today, network enterprises run complex processes for managing their networks, advertising the products and conducting business with upstream and downstream suppliers. Their business is dispersed world-wide, which creates a complex mesh consisting of many point-to-point connections and requiring a complex network management system and large amount of human resources to run the network. Consequently, network operators seek new ways of operation to reduce the network complexity and to increase the flexibility and dynamicity of the network.

A desirable solution of today's multi-layer environment issues can easily be accepted and deployed, if it offers a simple route of changes toward system automation while providing flexible management functions. The network operator must have the opportunity to adjust their network capabilities (flexibility, full automation) based on the real market demand. Such triggers in the future could be the demand for services such as bandwidth on demand or virtual IP network provisioning.

6 Conclusions

Current network management is complex and slow due to the partition of the operator's networks, which is translated even into different departments which only handle particular network parts. The project proposes the ONE adapter solution that aims at reducing this complexity by increasing network coordination.

The network scenario considered in the deliverable is composed mainly by two layers. The IP/MPLS layer is based on a hierarchical IP/MPLS network with three different levels. The bottom level of the hierarchy is composed of the access and mobile networks. The second level is the aggregation networks the top level of the hierarchy consist in interconnecting all the aggregation networks with the Internet access. On the other hand, the transport layer used to connect the IP/MPLS routers, focusing mainly on the long distance transport of date and relying on optical based networks.

This document has also presented an identification of the state of the art of the current network management solutions for the mentioned network scenario, paying special attention for the IP/MPLS network management, from open source tools to commercial versions. A first analysis has shown the lack of coordination with multiple layers of current NMS solution, and has been able to understand the benefits of adding an increase of network coordination.

The main contribution of the document is the definition of three use cases which benefit from the proposed network coordination and are well suited to be solved by the ONE adapter:

- The **IP link/service provisioning** use case focused on reducing the manual interactions between the IP and the transport layer departments to create the desired link/service.

- The **IP/MPLS offloading** use case is targeted at reducing the IP traffic across unnecessary intermediate routers. The main novelty in this use case is that the offloading is enabled through the ONE adapter which will coordinate all necessary interactions between layers.

- The **coordinated self-healing** use case aims at performing coordinated actions between IP/MPLS and transport networks in the event of a network failure.

The above mentioned use cases show the usefulness of the ONE adapter for network operators. The presented use cases allow the reduction of operation expenditures in network operators, thanks to the easiness in multilayer coordination, automation and ability to recover from failures. In this sense, future work in the project will analyse in depth the potential OPEX and CAPEX savings by using the proposed ONE approach.

Finally, the next steps of the project consist in defining the architecture of the ONE adapter that allows performing the use cases and satisfies the requirements identified in current deliverable. The architecture will be designed and applied to every use case, and will be finally assessed in the economic studies.

7 References

[Sub10] M. Subramanian, T. A. Gonsalves, and N. U. Rani, "Network Management: Principles and Practice," Pearson, 2010.

[Har02] D. Harrington, R. Presuhn, and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks," IETF RFC 3411, December 2002.

[Cya11] CYAN CyMS Multi-Layer Management System: http://cyaninc.com/cyms/cyan-cyms (URL last checked on Feb 2011).

[Adv11] ADVA Optical Networking FSP Service Manager, http://www.advaoptical.com /en/products/automated-network-management/fsp-service-manager.aspx. (URL last checked on Feb 2011).

[Nsn11] Nokia Siemens Networks NetAct, http://www.nokiasiemensnetworks.com /portfolio/products/network-management-oss/netact. (URL last checked on Feb 2011).

[Ope11] OpenNMS, http://www.opennms.org/. (URL last checked on Feb 2011)

[Nag11] NAGIOS, http://www.nagios.org/.

[Cis11] Cisco Systems, Inc. "Using the Cisco IOS Command-Line Interface," April 2010.

[Enn06] R. Enns, "NETCONF Configuration Protocol," IETF RFC 4741, December 2006.

[Mto11] Multi-Technology Operations System Interface (MTOSI), release 2.0, Tele-Management Forum, http://www.tmforum.org.

[Far06] A. Farrel, J. P. Vasseur, and J. Ash, "A Path Computation Element (PCE)-Based Architecture," IETF RFC 4655, August 2006.

[Man04] E. Mannie, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," IETF RFC 3945, October 2004.

[Nsn11b] Nokia Siemens Networks, Telecommunication Network Management System (TNMS), http://www.nokiasiemensnetworks.com/.

[Alc11] Alcatel-Lucent 5529 OSS Alarm Dispatcher (OAD) and 5620 Network Manager, http://www. alcatel-lucent.com/.

[Fuj11] Fujitsu, NETSMART 1500 Element Management System, http://www.fujitsu.com/.

[Hua11] Huawei, the U2000 system, http://www.huawei.com.

[ITU11] OTN ITU-T Recommendations on ASTN/ASON Control Plane, http://www.itu.int/ITU-T/.

[G8080] ITU-T Recommendation G.8080/Y.1304, Architecture for the Automatically Switched Optical Networks, Nov.2001, and Amendment 1, March 2003.

[PCE09] J. P. Vasseur, and J. L. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)," IETF RFC 5440, March 2009.

[IPo01] M. Chamania, A. Jukan, O. Gonzales de Dios, J. Jimenez Chico, "Offloading Excess IP Traffic with Optical Bypass – A Simple Capacity Upgrade, or More?" ONTC PRISM Newsletter, vol. 1, no. 3, Aug. 2010

[Lam01] LambdaStation, http://www.lambdastation.org/

[I2Ph01] Phoebus, http://e2epi.internet2.edu/phoebus.html

[Inf02] M. Chamania, M. Caria, A. Jukan, "A Comparative Performance Analysis of IP Traffic Offloading Schemes over Dynamic Circuits," IEEE INFOCOM, Shanghai, April 2011.

[Gro03] W. D. Grover "Mesh-based Survivable Transport Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking", Prentice Hall 2003.

[Fis03] Business Recovery Over Wide Area Networks: Are You Ready? By Randolph A. Fisher – CBCP: http://www.wancom.net/business_continuity.htm.

[Rol01] Roland T. Rust, 1993, Customer Satisfaction, Customer Relation and Market share, 1993, Journal of Retailing.

[Mad02] M. H. Mortensen, "The influence of attractiveness in business relationship development," Nov 2010, International IPSERA workshop on Customer attractiveness.

[Tre03] Trevor R. Roycoroft, "Economic analyses and network neutrality," June 2006, consumer Federation of America, Consumer Union and Free Press.

8 Acronyms

[AAA]	Authentication, Authorization, Accounting
[ASON]	Automatic Switched Optical Network
[BGP]	Border Gateway Protocol
[BRAS]	Broadband Remote Access Server
[CPE]	Customer Premises Equipment
[DSLAM]	Digital Subscriber Line Access Multiplexer
[EML]	Element Management Layer
[GMPLS]	Generalized Multiprotocol Label Switching
[GPL]	GNU General Public License (free software license)
[GPON]	Gigabit Passive Optical Network
[IGP]	Interior Gateway Protocol
[IP NMS]	Internet Protocol Network Management System
[IS-IS]	Intermediate System to Intermediate System
[ISO]	International Organization of Standardization
[L2]	Layer 2
[LDP]	Label Distribution Protocol
[LMP]	Link Management Protocol
[MPLS]	Multiprotocol Label Switching
[MTOSI]	Multi-Technology Operations System Interface
[MTU]	Multi-Tenant Unit
[NML]	Network Management Layer
[NMS]	Network Management System
[NNI]	Network Network Interface
[OLT]	Optical Line Termination
[OIF]	Optical Internetworking Forum
[OSI]	Open System Interconnection
[OSPF]	Open Shortest Path First
[PCE]	Path Computation Element
[QoR]	Quality of Resilience
[QoS]	Quality of Service
[ROADM]	Reconfigurable Optical Add-Drop Multiplexer
[RRD]	Round Robin Database

[RSVP]	Resource Reservation Protocol
[RSVP-TE]	RSVP Traffic Engineering (enhancement with object classes for TE)
[SLA]	Service Level Agreement
[SME]	Small or Medium Enterprises
[SML]	Service Management Layer
[SNMP]	Simple Network Management Protocol
[TE]	Traffic Engineering
[TNMS]	Transport Network Management System
[UNI]	User Network Interface
[VLAN]	Virtual Local Area Network
[VoD]	Video on-Demand
[VPLS]	Virtual Private LAN Services
[VPN]	Virtual Private Network
[xDSL]	Digital Subscriber Line (x is a placeholder for the specific technique)