An Experience on Implementing Network Management for a GMPLS Network

Bela Berde   Alcatel Research & Innovation (France)

Carolina Pinart   CTTC (Spain)

Javier Gonzales Ordas, Javier Jimenez   Telefnica I+D (Spain)

Piet Demeester, Koen Casier   University of Ghent - IMEC (Belgium)

Remi Theillaud, Veronique Piperaud   Atos Origin (France)

Dimitri Papadimitriou   Alcatel Bell (Belgium)

Abstract

The management of optical networks and their components is usually performed through proprietary systems. This paper describes an experience on implementing a management system for multilayer Generalized Multi-Protocol Label Switching (GMPLS) [1] networks and the associated integrated network provisioning strategy. The main characteristics of this solution are the interface connecting the control plane and the management plane, the management application framework itself, and further its joint use with a Policy-based Management (PBM) system for accurate provisioning. The experience is based on the definition, design, and implementation of the overall management system in the ITEA Project TBONES for the support of both switched and soft-permanent connections.

1 Introduction

In current optical networks, management systems are typically focused on centralized data plane management, with minimal interactions between the client (typically IP) and transport layers. As optical networking architecture is shifting from a static point-to-point architecture towards more dynamic and re-configurable, the increased flexibility and agility in networking equipment has espoused two important trends: 1) a shift from static planned towards dynamic on-demand resource allocation and service provisioning, and 2) a shift from centralized management and off-line optimization strategies towards distributed control and on-line incremental heuristics in network and traffic management. In other words, there is a clear trend to endow more intelligence into the network elements to produce a more autonomous network, thereby simplifying and reducing the cost of network operations.

Optical networking test-beds can be categorized into packet-switched, flow-switched, and circuit-switched. One of the first IP/WDM test-beds is the Winchester field trial [2], which deployed a management plane in a network scenario that did neither consider a control plane nor transparent nodes. The test-bed KomNet is a WDM ring network that provides a switching granularity at wavelength channel level [3]. This test-bed is focused on transparency issues, but it does not deploy a control plane.

Recently, the management of control plane-enabled optical networks have attracted considerable attention. The test-bed of the European IST project Layers Interworking

Optical Networks (LION) demonstrated an experimental feasibility of setup/teardown of soft-permanent connections (SPC) with GMPLS signaling and resilience experiments in a managed, multi-vendor, and multi-domain environment [4]. The LION management plane is built with distributed protocols (CORBA). Alternatively, in the so-called customer-owned and -managed optical networks, users are able to manage some resources in the wide area network, as in CANARIE's CA*Net4. In such a network, users will be able to perform their own restoration and protection, optical add-drop multiplexing, or cross-connecting to other users. The project Dynamic Resource Allocation via GMPLS Optical Networks (DRAGON) [5] will create dynamic, deterministic, and manageable end-to-end network transport services for high-end e-Science applications through deploying a GMPLS-capable optical core network.

The Transparent Backbone Optical Network Simulator (TBONES) test-bed [6] is formed of an optical network and a network simulator. In the test-bed, end-to-end Label Switched Paths (LSP) are set up and torn down dynamically by means of a GMPLS control plane. That is, TBONES provides not only SPCs, as done in the LION test-bed, but also switched connections (SC) that may be triggered from the simulator. TBONES supports the request of different granularities, ranging from lambda to packet switching, similarly as in DRAGON (packet, wavelength, and fiber cross-connection). TBONES allows different interconnection models, from overlay (such as LION) to integrated routing (such as DRAGON).

The integration of the control plane into the manageable functional components is of special interest to TBONES. It must be pointed out that many control plane features, for instance Traffic Engineering, may also use information imposed from the network management. Moreover, configuration, performance, and accounting management as well as fault management functions specific to control plane-enabled networks will be needed.

The paper is organized as follows. In Section 2, we describe the design background and present the requirements and constraints of the TBONES management plane as well as we outline the proposed framework. Finally in Section 3 we conclude and present future work.

2   Design of the overall management system

2.1   Capabilities for managing GMPLS-enabled networks

ITU-T defined functional areas of the generic management model [7] often referred to as FCAPS. These functions include path provisioning, which encompasses routing and QoS support, and automatic recovery of failures or performance degradations. The former can be included in configuration management, whereas the latter is part of fault and performance management. The remaining FCAPS functions are out of the scope of the TBONES project and therefore they will not be discussed any further. In [7], configuration management deals with the initialization and modification of network element state information. Unlike performance, fault, and accounting management, configuration management requires intrusive changes in managed devices. Configuration should be carried out in a concerted way for introducing network-wide state, e.g., a path setup in the network, and not on a per-device basis. The key issues in configuration are the

distributed transaction processing coupled with management protocols used to access managed devices. Traditionally, each network layer has been independently managed, since each layer had its own requirements and unique characteristics. It can be inferred that no universal management system exists that can deal with the available offerings of network technologies, but many proprietary management systems, resulting in a complex management infrastructure. The aim of simplifying the overall management of an optical network, and making it more cohesive, should be attained through opening communication between management systems and allowing common objectives to be met at all layers. Another important issue in simplifying optical management is the choice of the management approach, among full centralized and distributed including policies, etc. Here we identify some of the most relevant capabilities required by next-generation management planes for collaboration with the control plane, when focusing on service provisioning:

- Division of network resources between those visible to the control plane and those visible to the management plane. This is necessary to avoid contention when establishing and releasing services by the control (SC) and management (SPC) planes simultaneously.

- Assignment of certain capacity to a particular player in order to create virtual private networks (VPN), which is especially important for highly distributed services, such as Grid.

- Provision of configuration and policy information in the control plane.

- Setting and modification of the signaling system parameters, such as time-outs (e.g., call setup time-out), thresholds, congestion control mechanisms, maximum number of allowed connections, maximum signaling load (above which the signaling processor denies call establishment requests, for instance).

- Determination of the maximum number of connections that can be supported by a network element and to set, where appropriate, the maximum number to be supported.

- Distinguishing changes in the state of connections due to management or control plane actions and those resulting from network failures and suppress or generate alarms as appropriate.

- Setting or modifying survivability priority levels or QoS contract levels for all connections associated with a given 'performance class'.

- Assigning the maximum value of a connection identifier on a link, where appropriate, setting traffic management controls either manually resulting from a specific input or automatically in response to internal or external stimuli (in control plane enabled networks the management system sets the conditions under which control is applied and the strength of response).

- Management of the Data Communication Network (DCN) to ensure a consistent configuration of signaling resources, including determination of attributes of signaling links including their functional status, error indications, traffic data or maximum bandwidth will be as well required.

Note that in these capabilities we have not considered (centralized) routing occurring in the management system, but distributed in the control plane, and have not considered the management plane to calculate routes for SPCs or permanent connections, and use of management protocols for connection management. This is due to the fact that the distributed GMPLS control plane performs path computation (and routing, signaling), and in real time.

2.2   The TBONES approach

Spanning both the simulator and the real network areas Figure 1, control plane instances allow experimenting the behavior of a collection of (GMPLS) Label Switching Routers (LSR) by instantiating for each node a lower protocol stack and several control plane controllers. The OSPF-TE [8] as well as RSVP(-TE) [9] form the lower protocol stack. In the simulator areas, some minimal modeling of the data plane is required in the form of a data plane emulator for experimental support, while the control plane is emulated using the GMPLS protocol stack.
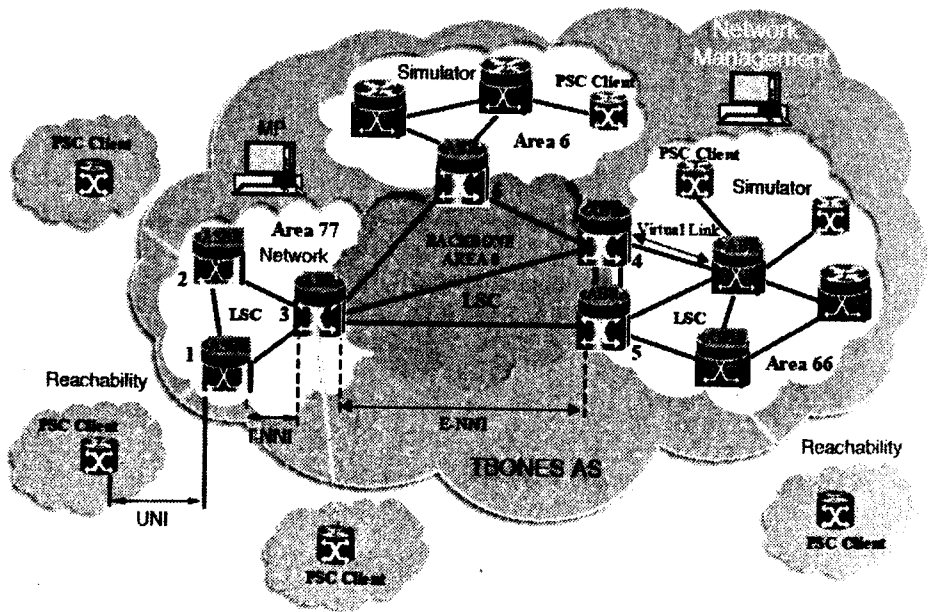

Figure 1: Interconnected (managed) areas in the TNONES AS.

There are mainly two control plane driven network management priorities in TBONES. Apart from the Operation, Administration, Maintenance, and Provisioning (OAM&P) tasks, the major innovations brought by TBONES are: Constraint-based, dynamic provisioning using control plane capabilities as well as LSP supervision including monitoring, setup, and deletion. The constraints include those given by users Service Level Agreements (SLA), current network state for Traffic Engineering, and business objectives of the Operator.

While the optical network has its own management plane (MP), the management system developed in the project for the simulator areas (identified in Figure 1 as Network Management) is composed of three main blocks. As shown in Figure 2, every LSR in a simulator area (e.g., Area 66) is interfaced with a Network Management System (NMS) and a Policy-based Management (PBM) system [10] through a management agent, located in the LSR hosting various control plane controllers. The NMS triggers SPC and performs supervision of LSPs. The PBM provides and enforces policies to the LSR, and specifically to controllers. Indeed, within the GMPLS module, the management agent communicates with the connection admission controller (CAC-C), the signaling controller (Sig- C), and the Traffic Engineering controller (TE-C), and the routing controller (R-C).

The control plane is organized around the TE-C, which interacts with the CAC-C, Sig-C, and the R-C. In particular, the R-C processes TE, topology and reachability information from multiple SC and SPC (TEDB without any specialization). A Traffic Engineering Database (TEDB) registers any controlled entity, i.e. Traffic Engineering link (TE link): LSP, Forwarding Adjacency-LSP (FA-LSP), un/bundle TE link. All the controllers are subordinated to the TE-C, which performs connection handling, retrieves notifications and policy reporting, and is feed, on its turn, by policies from the PBM. The management agent communicates with the NMS and the PBM through the Network Management Interface (MNI), which is carried in a control channel over the Data Communication Network (DCN).
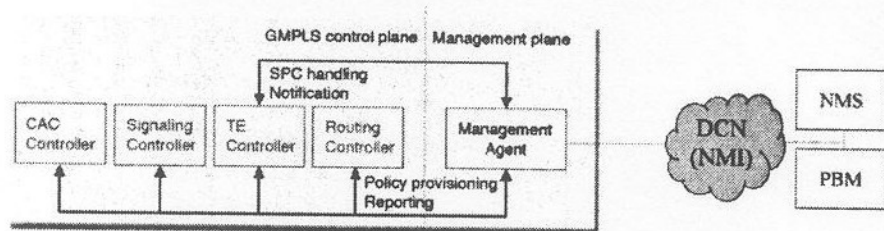

Figure 2: Logical interface between an LSR and the TBONES management plane.

A Dimensioning Tool (DT) provides to the overall simulator an adequate routing and dimensioning for the considered multilayer topology given a traffic-matrix on the top layer of this topology. The routing information calculated here can be used by the NMS as an SPC/SC request list. Additional settings for the DT allow splitting the dimensioning result for a layer into soft-Forwarding Adjacencies (FA) and hard-FAs. More precisely, a threshold defining the minimal usage in time of a circuit identifies the amount of circuits that need to be setup semi-permanently (corresponding to the hard-FAs) and for which it is not worth to tear them down in silent periods. The other circuits (the soft-FAs) need to be setup on demand as needed. It is also possible to calculate the dimensioning for a layer taking a specific resilience scheme into account.

2.3   Network model

TBONES NMS uses a simplified network model based on ITU-T recommendations [11] and on Telemanagement Forum Information Agreement [12] with proper extensions to cover TBONES specific use cases.
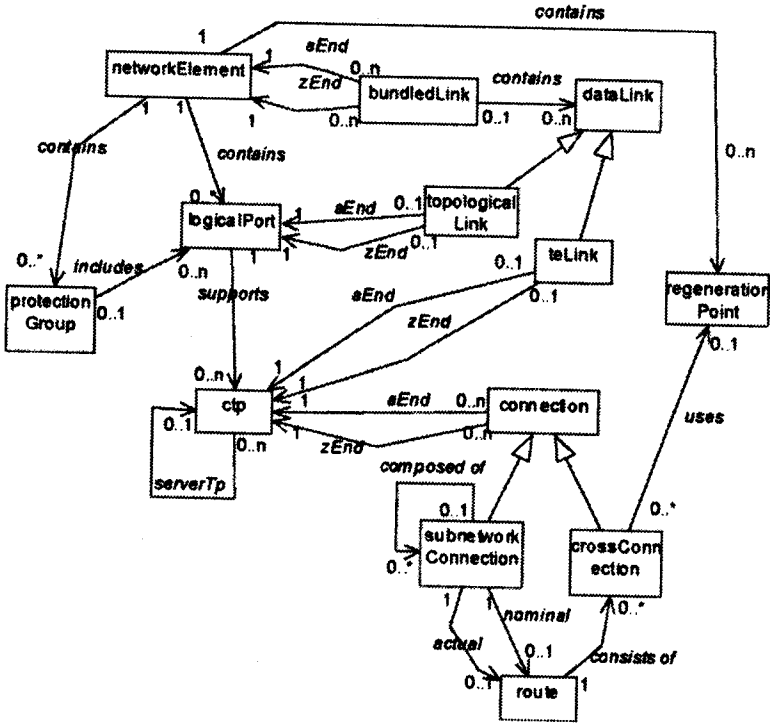

Figure 3: TBONES NMS network model.

Figure 3 provides the class diagram for this network model. The entities in the TBONES NMS network model include both physical and virtual topology components:

- Network Element (NE): A network element is a device that terminates or switches communication signals. For TBONES purposes, a network element represents an emulated LSR.

- Logical port (interface): Represents a transport plane bidirectional communication interface of a network element. Logical ports in TBONES may support one single wavelength or may be WDM capable.

- Data link: A data link at a given layer represents bidirectional capacity between sets of CTPs at different NEs. CTPs connected by a common link must be served by a single connection at the server layer, i.e. a data link at the client represents the capacity provided to such layer by a connection at the server layer. Two types of data links can be identified: topological links and FA links.

- TE link: Topological links connect adjacent NEs, as viewed by the management system; as such, they represent fixed connectivity between logical ports

- Topological link: FA links are established as a result of setting up LSPs at the server layers, and therefore they can be dynamically instantiated and removed.

- Bundled link: A bundled link is a collection of data links, which are grouped together in the control plane routing protocol

- Connection termination point (CTP): A connection termination point represents, in GMPLS terms, a specific bidirectional label at an specific interface. In TBONES context, it represents a wavelength or a waveband supported by a logical port. Packet labels are not modeled

- Connection: A connection represents the common characteristics of subnetwork connections and crossconnections.

- Subnetwork connection (SNC): A subnetwork connection represents a bidirectional transport entity that transfers information between CTPs. In TBONES context, a SNC represents a wavelength, waveband, or packet LSP. Connections at a server layer create FA links at its client layer.

- Cross-connection: A cross connection represents a subnetwork connection set up within a single network element, i.e. a connection set up through a switching matrix.

- Route: A route is a partially ordered list of LSRs defining the resources that a subnetwork connection traverses. TBONES NMS keeps track of the nominal and actual routes of all subnetwork connections.

- Regeneration point: A regeneration point represents a device inside a network element with the capability to regenerate in the electrical domain an optical signal. A network element may contain multiple regenerators (and thus multiple regeneration points).

- Protection group: A protection group models link protection configurations at the optical layer.

2.4 NMS architectural aspects

Figure 4 illustrates the architecture of the TBONES NMS, including the relationships with two other elements of the TBONES solution: control plane simulator/emulator and the DT. The system relies on a relational database for data persistency and on a CORBA bus for distributed computation, on top of which a set of management modules performs the required functionality in the configuration, fault, and performance management areas.

The Provisioning manager module is in charge of setting-up, tearing-down, re-routing, and reverting to nominal path SPCs. It receives from the control plane notifications on the establishment and release of automatically setup SCs.

The SLA manager provides additional information on the status and quality of connections associated to particular service agreements. The Alarm manager performs all

fault management related functionality, including reception of failure notifications, identification of affected entities, and handling of alarm logs.
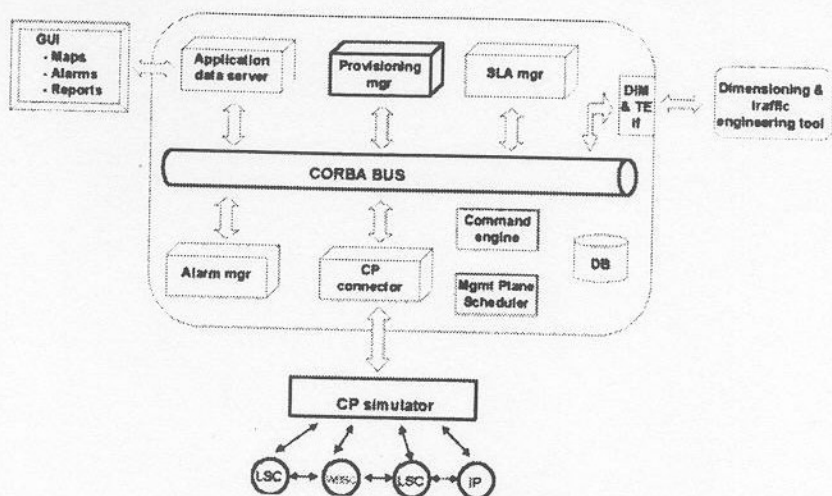

Figure 4: TBONES NMS architecture.

The Management Plane Scheduler and the Management Command Engine provide capabilities for performing automated experiments. The presentation layer of the system is realized through a WEB based GUI client complemented by the so-called Application Data Server, which processes requests generated at the GUI. Two specific interconnection modules support the interfacing with the control plane emulator and the DT.

2.5   Control plane-Management plane interface for the simulator

specified and implemented. This interface is used to exchange provisioning requests and reports as well as alarm information between both functional planes. The requirements that guided the interface specification are:

- Easy extensibility to accommodate a wide range of test scenarios and new control or management capabilities

- Simplicity, both to speed up development and facilitate writing automated test tools at any of the sides of the interface

- Flexibility to support multiple provisioning approaches, the comparison of which is one of the objectives of the project

- Similarity to real world implementations, following standards if possible

Management specific protocols, such as CMIP and SNMP, were considered but the final choice was to exchange XML messages over TCP sockets in an asynchronous manner. This approach avoids overloading the control plane implementation with an additional protocol stack, allowing leveraging the large number of available XML tools and

- Link or connection alarms

The NMS uses these messages to keep a consistent and complete view of the status of the network. An example of a SC setup notification, including its end-to-end recorded route and per link labels, is:

```
<MP-CP-Msg>
    <notification message-id="001">
        <SC-setup>
            <ID Source="10.1.1.1" Destination="10.2.2.2" TunnelID="50"/>
            <ActualPath>
                <Hop type="strict">
                    <If IPv4="10.1.2.2" Id="1"/>
                    <Label type="LSC">00010101</Label>
                </Hop>
                <Hop type="strict">
                    <If IPv4="10.1.2.3" Id="1"/>
                    <Label type="LSC">00010102</Label>
                </Hop>
            </ActualPath>
        </SC-setup>
    </notification>
</MP-CP-Msg>
```
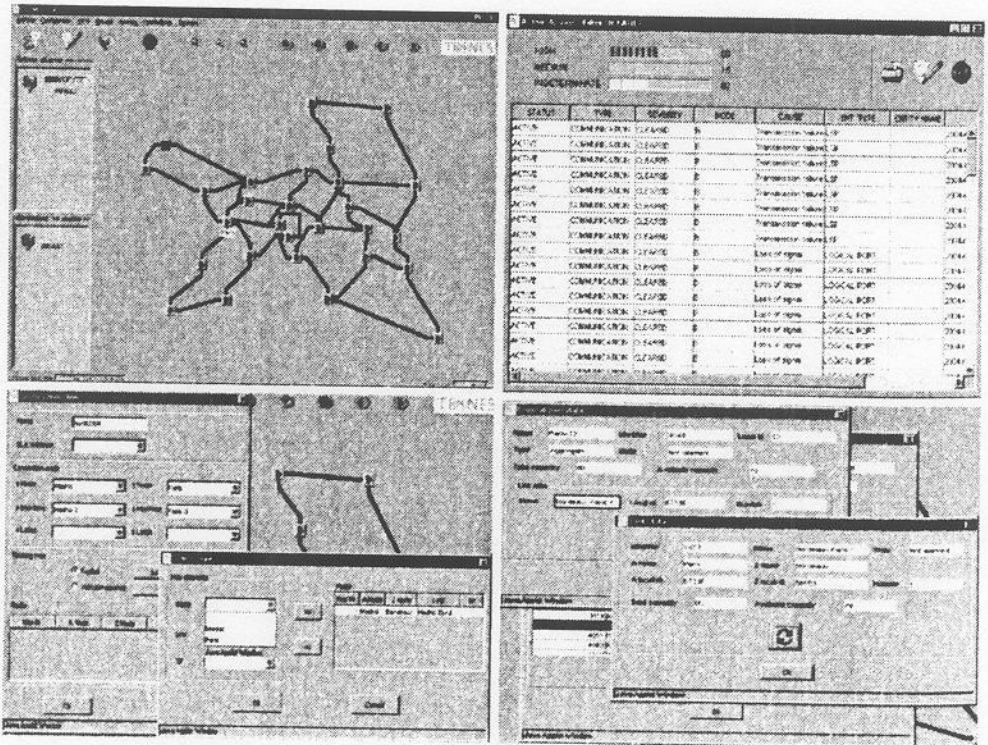

Figure 5: NMS GUI snapshots. a) Network map (top-left), b) Alarm table (top-right), c) SPC setup form (botton-left) d) Link view (botton-right).

testing and debugging, since interface primitives are human readable. It must be noted that existing standards do not meet all the TBONES functional requirements and thus all options would need proprietary extensions. Two communication channels are setup between the NMS and each emulated network node (LSR).

One channel carries connection provisioning requests from the NMS and the associated responses, while the second channel is only used for upstream spontaneous notifications from the control plane to the NMS. Provisioning messages support multiple options:

- The connection route may be absent or explicitly defined in an strict or loose manner

- Exclusion constraints are supported

- For resilience, pre-planned rerouting or dynamic rerouting may be selected

- Crankback mechanisms may be enabled on a per-connection basis

- Diversity with respect to shared risk ling groups may be requested

Any of these options may be missing in the message. If so, a default value, or a policy defined one, would be used instead. Simulation scenarios will be run for different combinations of options as a way of benchmarking the different approaches. An example message for the setup of a SPC with a node inclusion constraint follows:

```
<MP-CP-Msg>
    <rpc message-id="1">
        <SPC-setup>
            <ID Destination="10.1.0.8" Source="10.1.0.2"/>
            <SwitchingGranularity>
                <EncType>8</EncType>
                <SwgType>150</SwgType>
                <GPID>31</GPID>
                <Bandwidth>9953.28</Bandwidth>
            </SwitchingGranularity>
            <SrcLink>
                <If IPv4="10.1.0.2" Id="3"/>
            </SrcLink>
            <DestLink>
                <If IPv4="10.1.0.8" Id="4"/>
            </DestLink>
            <Path>
                <Hop type="loose">
                    <If IPv4="172.16.0.11"/>
            </Path>
            <MaxNbOfHops>3</MaxNbOfHops>
            <MaxNbOfHopsReg>8</MaxNbOfHopsReg>
            <SetupPriority>3</SetupPriority>
            <HoldingPriority>3</HoldingPriority>
        </SPC-setup>
    </rpc>
</MP-CP-Msg>
```

Control plane nodes issue spontaneous notifications to report to the NMS any significant event that has occurred in the network. Three types of notifications are defined:

- SC setup and release messages

- Connection restoration notifications

2.6   User interface

There are four main elements in the TBONES NMS GUI:

1. The network map displays the dynamically changing topology of the data and optical network as a set of connected edges, as can be seen in Figure 5.a. Faulty nodes and links are highlighted in orange or red, where each code denotes a different level of severity.

2. The alarm table provides a summary of the failure state of nodes, ports and links in the network. Figure 5.b provides a snapshot of this table.

3. Connection and SLA creation forms. Figure 5,c is a snapshot of the SPC request form, where the GUI assists the user to select the explicit path of the connection

4. SLA, connection and topology reports, which provide access to the system database via queries and graphical browsing. Figure 5.d is an example of the link view, reached from the node and port views.

3   Conclusion

In this contribution, we described the design of a management plane for a test-bed that spans a real GMPLS network and an emulated GMPLS-enabled network. We presented the fundamental network management requirements, architectural features, and a pilot implementation. The combined PBM-NMS approach allows converging towards a management continuum with reference to efficient and accurate provisioning. The joint validation of the control/management collaboration is currently ongoing in the project. Future works will investigate several issues regarding connection management and usage. Firstly, since the provisioning relies on proactive policies, we intend to study its possible interoperability with reactive and hybrid control plane mechanisms. Secondly, since resource segmentation and incoherency problem can occur (TE issues), we aim to examine if TE policies can be used to tackle this phenomena. Finally, we envisage studying more possible use cases for the overall management system.

Acknowledgments

References

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

[2] S. Hubendick. L. Bystrm, Dynamic Connection Management System for Optical Networks: Applicability in a TINA environment, TINA Conference Paris, 2000.

[3] K-D. Langer, J. Vathke, G. Walf, KomNet - A modular platform to assess optical networking techniques and components, Proceedings 11th Tyrrhenian Workshop on Digital Communications, Optical Networking, Springer - Verlag, pp. 48 - 51, 1999.

[4] K. Shimano, A. Manzalini, C. Cavazzoni, Demonstrations of Layers Interworking Between IP and Optical Networks on the IST-LION Test-Bed, OECC2002 Conference, 2002.

[5] J. Sobieski, Dragon: Dynamic Resource Allocation via GMPLS Optical Networks, MCNC Optical Control Planes Workshop, Chicago, Illinois, 2004.

[6] D. Papadimitriou, B. Berde, R. Theillaud, K. Casier, TBONES: a GMPLS Unified Control Plane for Multi-Area Networks, Infocom 2005, March 2005.

[7] G.8080, Architecture for the Automatically Switched Optical Network (ASON), ITU-T Recommendation G8080/Y.1304, 2001.

[8] K. Kompella et al., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching", draft-ietf-ccamp-ospf-gmpls-extensions-12.txt (work in progress), October 2003.

[9] D. Awduche, et al., "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[10] B. Berde, et al., "Improving Network Performance through Policy-based Management applied to Generalized Multi-Protocol Label Switching", 10th IEEE ISCC 2005 Symposium, June 2005.

[11] ITU-T Recommendation, G.805, Generic Functional Architecture of Transport Network, 2000.

[12] Telemanagement Forum, MF608, Multi Technology Network Management Information, 2003.

IV Workshop
in G/MPLS
Networks

Edited by
Teodor JovÈ & Ramon Fabregat

Documenta
Universitaria