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

M. Drogon[1], M. Chamania[1], A. Jukan[1], M. Yanuzzi[2], X. Masip-Bruin[2], V. Lopez[3], O. González[3], M. Maciejewski[4], M. Roth[4], C. Brunn[4], J. Altmann[5]
[1]Technische Universität Braunschweig, [2]Technical University of Catalonia, [3]Telefónica I+D [4]ADVA Optical Networking, [5]Seoul National University

Abstract: For many years, the convergence of Internet data services and transport network services based on optical transmission has been at the heart of carriers' investments and business strategies. Over time, the inherent segmentation of IP and transport networks has created major technology differences, thus impacting their management and operation. In addition, the carrier's organizational separation and fragmentation of technical competencies between "Internet" and "transport" have been and continue to be profound. Consequently, the two networks, while increasingly deployed in tandem, lack automated coordination of management procedures and business practices. To address this challenge, we present a roadmap for convergence from the point of view of network management, and describe an easy-to-deploy solution that can enable coordination of the two networks. To the best of our knowledge, this approach is the first attempt to enable what we define as a "controlled" as opposed to "fully automated" evolution of converged packet and circuit switching networks, without the need for large scale management system integration or changes to the current telecom management practices, other than in improved speed and efficacy.

Keywords: Multi-layer, Network Management, IP-Optical

1. Introduction

Despite the seemingly converged evolution of Internet and optical transport networks, their operational and technological separation remains as large as ever. The carrier's organizational separation and fragmentation of technical competencies has resulted in two administratively separate networks, where management team expertise and operational practices significantly differ. These differences can be seen at the management, control and data plane layers.

From a network management perspective, the transport and IP Network Management Systems (NMSs) show profound differences and design philosophies. For example, in the transport layer, the NMSs specify the services and its associated functions in a standard format, such as QoS and protection, and proprietary Element Management Systems (EMSs) based on Transaction Language 1 (TL-1) are used to translate service requests into hardware configurations [1]. In contrast, IP NMSs rely on direct configuration of devices either by proprietary Command Line Interfaces (CLIs) or via SNMP [1], and configuration as well as service-specific decisions, such as the routing and protection, are left to the network administrator.These differences make any attempted integration of management functions significantly more complex that is commonly handled in each layer separately. Even though carriers are facing duplications of network management functions, such as routing and protection, the premium on stability and simplicity has prevailed over any integrated solution.

From a data plane perspective, two competing trends can be identified. On the one hand, the ever decreasing price of optical bandwidth creates a major barrier for entry to any new technology solutions that can combine Internet and optical transport. On the other hand, carriers are pressed to implement converged Internet and transport due to the large expected operational cost savings and faster service provisioning, thus creating increased revenue opportunities. Also the network equipment vendors are facing the challenge of potential physical as well as control and management process integration of packet and circuit switching elements, one being IP routers and the other circuit switches (Ethernet virtual circuit switches, WDM switches, etc). Despite the significant advances that have been made towards the development of a unified control plane framework to support both "packet" and "circuit" services, carriers remain reluctant to deploy any unified control plane solutions without a level of manual control and coordinated automation between these two networks.

In this paper, we present the issues of coordinated operations between the IP and the transport networks, and propose a roadmap towards automated multi-layer interactions via the management planes. To this end, we propose the ONE adapter, an ontology-based communication adapter, designed to facilitate coordination between the two management ecosystems, applicable to a generic class of business practices, such as multi-layer service provisioning and post-failure network management.

The key aspects of the ONE adapter are 1) Coordination-without-integration of NMSs for a multi-vendor, multi-layer network management system coordination, 2) Semantic adaptation between the NMSs over standardized, vendor-independent interfaces, 3) Smart analytics to create, manage and redefine network management procedures, 4) Programmability of legacy management workflows and systems involving multiple NMSs with semantic adaptation, 5) Ease of extension for standardized interfaces towards emerging third party network management systems, such as the Path Computation Element (PCE) [2].

In comparison to the above features, however, ONE has not been designed to be a) a new NMS (it is designed to be a light weight system independent of any NMS), b) a control plane client (but is control-plane friendly and can talk to any control plane implementation), c) yet-another virtualization or router programmability tool, nor it changes anything on legacy systems (it only adapts the semantics of whatever vendor is deployed in a carrier network.)

Due to its main design objectives presented here, we believe that our approach is a true enabler of what we define as *controlled convergence* of packet and circuit switching networks and their corresponding management ecosystems, as it eliminates the need for large scale system integration, or drastic changes to the current telecom management practices. The applicability of the ONE adapter is shown in a real testbed scenario located at the premises of Telefonica and ADVA, demonstrating working use cases.

2. Managing the Internet and Transport Networks: Two Separate Ecosystems

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. In addition, the speed of application innovation in the Internet has strongly conditioned the strategies adopted by operators for managing

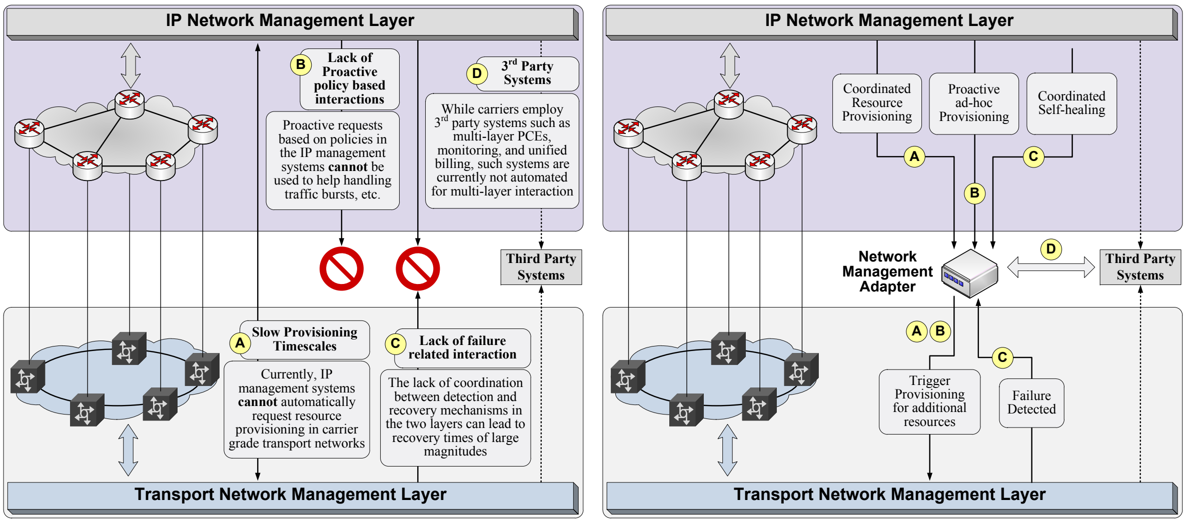

(a) Main limitations. (b) Targeted goals.

Figure 1: Research challenge: can these goals be achieved in a way that is sufficiently pragmatical to become adopted?

their Internet and transport networks, and explains in part the marked differences in the evolution of management tools in support of these operations.

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., [3]). 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.

The IP network configurations, on the other hand, have, over time, become increasingly complex and vendor-specific due to multiple factors. First, the changing dynamics of the Internet has driven the deployment of a wide spectrum of IP enabled equipment by telecom carriers. This equipment is usually bought from only one vendor and the whole network is aligned to the vendor specific constraints, making it infeasible to change equipment easily; Second, the IP network is expected to support a large number of services and quickly adopt new upcoming services to reduce time to market. Currently, monitoring of IP devices is mostly managed by the SNMP protocol [1], 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) or the NETCONF interface [4].

Figures 1(a) and (b) illustrate some of the consequences of the isolation between the management systems and the solution proposed. It is worth noting that even the provisioning of a new IP link (Ⓐ) requires multiple communications between human operators from two different departments, each responsible for the configurations in one layer. These operations not only lead to long service provisioning times and 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 Ⓑ), or any type of coordinated self-healing action (Ⓒ). The latter is especially critical, since the successful recovery from a failure often demands a significant amount of interactions and coordination between the two network management teams before the occurrence of the failure. Carriers also desire an automated communication with external control and management subsystems, such as the PCE (Ⓓ). The approach to develop an adapter that can meet these demands seems essential to the operators, facilitating the operations and cost-effective interactions between the Internet and transport NMSs.

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. To this end, significant advances have been made towards the development of a unified *control* framework to support both "packet" and "circuit" services. The ASON framework is an important effort from the ITU-T [5], aimed 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 using three standard interfaces: the User Network Interface (UNI), the Internal Network-Network Interface (I-NNI), and the External Network-Network Interface (E-NNI). A similar effort is also led by the IETF in the form of the GMPLS-based standard solution for inter-technology and inter-layer control plane interactions which support both packet and circuit switched networks in the same control plane [6].

However, it should be noted that the support for *packet* technologies in control planes extends to the MPLS and carrier-grade Ethernet technologies, and does not provide inherent support for specialized functions of traditional IP routing and forwarding. For instance, policy-based routing functions are not supported by the standardized control planes as their implementation differs significantly across different vendors. Also, control planes can address automation of specific management functions, but cannot automatically orchestrate actual management and business procedures. For instance, capacity upgrades during network planning are orchestrated by the operator as a series of individual provisioning and configuration actions based on the current state of the networks, which cannot be facilitated inherently by control planes. Finally, network operators increasingly use third party systems such as network planning tools, PCEs, network monitoring tools, etc., to make policy-based decisions on network operations. However, the control planes do not provide integration with these external subsystems, and need an external entity to integrate information from them and translate them into network operations.

In this context, it seems reasonable to seek solutions towards the convergence of the Internet and transport NMSs that do not require the integration of all different management systems (due to complexity), and include the possibility of having a level of manual control during the automation of the management tasks (due to business procedures). 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, vendor-agnostic and automated communication channel between the two management layers. The goal is to enable coordination to support a set of basic management tasks, such as provisioning, and coordinated post-failure management.

3. ONE Adapter

The ONE adapter is devised to be useful in future networks, which are expected to be composed of heterogeneous technologies. Our design attempts to facilitate multi-layer coordination while addressing the issues of multi-vendor support and user-defined multi-layer coordination. The proposed architecture for the ONE adapter addresses three major features in order to facilitate multilayer orchestration in commercial network settings, namely 1) programmable orchestration, 2) semantic multi-vendor adaptation and 3) integration of third-party systems.

Programmable orchestration provides the operator to flexibly define operational workflows for multilayer operations that mimic the existing business operations in the network. The ONE adapter architecture also takes into cognizance the fact that the process for defining a workflow within ONE should be simple in order for it to be usable by operators, who may not be specialized programmers. Semantic multi-vendor adaptation helps in the operation of multi-vendor networks and in supporting migration of underlying hardware while maintaining the business logic programmed within ONE.

Multi-vendor support is challenging in the different ecosystems, with multi-vendor support for IP being especially difficult. Configuration syntax and context can change significantly between vendors and even across the IOS versions of the same vendors which makes implementing multi-vendor support in IP especially challenging.

In the ONE adapter architecture, we do not integrate functionalities inside yet-another NMS. The ONE adapter is a central intelligence between layers that offers the functionality to coordinate third-party vendor equipment by the use of Ontological transformations to map operations from a given internal operation specification to the corresponding configuration operations on the different devices connected in the network. Ontological transformations are used during configuration operations on IP and transport networks. Ontological transformations are also applied on incoming information to ensure that information used by the different components within ONE (e.g. multi-layer topology information) is not affected with changes in the external systems used by the operator.

Third-party systems such as the Path Computation Element (PCE), Authentication Authorization and Accounting (AAA), and Service Level Agreement (SLA) management can significantly improve the performance of the network, either in terms of optimizing actual network operations or in reducing the OPEX involved in the management of the network. However, the integration of third-party systems, especially in multi-vendor settings is challenging as there always exist some minor discrepancies in the specifications and/or interpretation of standards, which in turn leads to different implementations that may not always match each other. The ONE adapter architecture uses the Service Oriented Architecture in order to ensure that third-party systems can be easily integrated within the ONE adapter as services and can interact with other management subsystems via the ONE adapter.

3.1 ONE use cases

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: 1) IP service provisioning; 2) IP Traffic offloading; and 3) Coordinated Self-healing.

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 Traffic 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.

3.2 ONE Architecture

The architecture of the ONE adapter consists of different modules, that handle different operations, reducing the overall complexity. In Figure 2, a typical process is shown, in how the different modules communicate. In this overview, we have the GUI, Trigger Module, Management Controller, AAA, Workflow Database and Workflow Processor modules.

The GUI is the graphical user interface of the ONE adapter, that can trigger events by an human operator or monitor the status of the network by receiving updates from the ONE adapter measurement functions. The GUI enforces user login by authenticating with the AAA (1) and receives an authorisation ticket in response (2). Once the user has been authenticated, the GUI initializes the link provisioning request. Sending out to the Trigger Module mandatory and optional configuration parameters required for performing such operation plus the granted authorization ticket. The Trigger Module analyzes the information retrieves the corresponding trigger by enforcing predefined

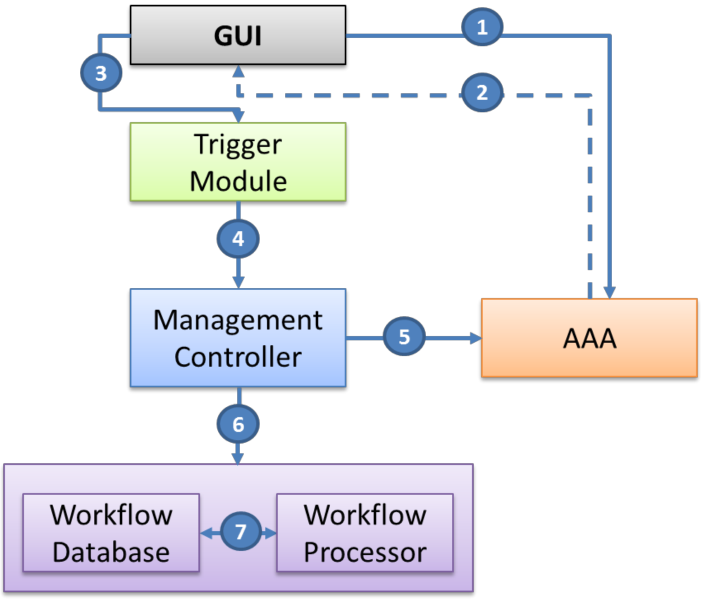

Figure 2: Link Provisioning Process in ONE

policies and sends it to the Management Controller (4). The main task of the Management Controller is to control the execution of workflows as instructed by the Trigger Module. To do so, it first need to check wether the workflow initiator (i.e. network administrator) is authorised or not to request the execution of a specific workflow (5). If the AAA module approves the execution, then the Management Controller passes the received information to the Workflow Database and Workflow Processor (6). The Workflow Processor looks up the workflow in the Database (7) and the Workflow Processor finally initiates the workflow.

4. Demonstration

The ONE adapter is able to provide IP Link Provisioning, IP Traffic Offloading and Multi-Layer Restoration in its current implementation status in a real testbed scenario. The focus of this demonstration is the ease of operation that the ONE adapter offers. The exact operation times are in the current stage not representative for the final product, but already show a significant improvement from several days down to roughly 1 minute in the case of path provisioning. Figure 3 shows the physical network topology. The network equipment in the IP layer consists of three MX-240 Juniper Router with 960 Gbps System capacity and 240 Gbps Throughput per slot. In the transport layer, four ADVA FSP 3000 Colorless, directionless and contentionless multi-degree reconfigurable optical add-drop multiplexer (CDC-ROADM) are used to set up circuits. The initial network does not have any circuits configured and thus the IP router are not connected. This initial network is used to describe the IP link provisioning in section 4.1, the IP traffic offloading in section 4.2 and the multi-layer restoration in section 4.3.

4.1 IP Link Provisioning

In the initial topology, a new IP link from *MX240-1* to *MX240-3* can be set up by the graphical user interface. The operator in the IP network requests a new IP link between

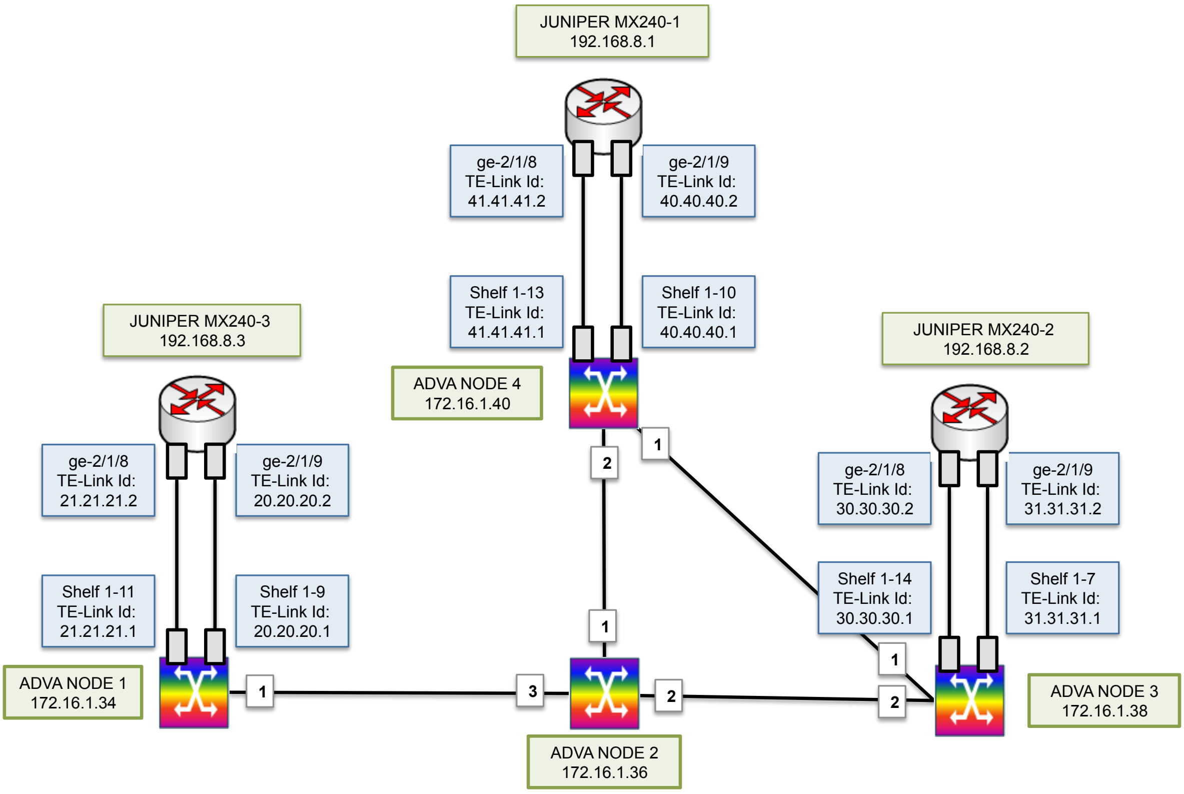

Figure 3: Physical Network Topology

both 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 and will identify corresponding transport network switches. 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 endpoints. 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 initialise the newly formed IP link.

In this topology setup a circuit is reserved between *ADVA Node 4, ADVA Node 2* and *ADVA Node 1* and the IP interfaces of *MX240-1* and *MX240-3* are configured to use this circuit as the new IP link.

4.2 IP Traffic Offloading

In the IP Traffic Offloading scenario, a path from *MX240-1* though *MX240-2* to *MX240-3* already exists and a traffic generator application is connected to *MX240-1* sending traffic to *MX240-3*. In the transport network, two circuits are reserved: *ADVA NODE 4, ADVA NODE 3* and *ADVA NODE 3, ADVA NODE 2, ADVA NODE 1* to establish the IP path.

In this scenario, the ONE adapter will initiate operations when the generated 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 endpoints, 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.

In this topology setup a new circuit is reserved between *ADVA Node 4, ADVA Node 2, ADVA Node 1*, resulting in a direct link between *MX240-1* and *MX240-3* for the traffic generator application. The load of the network is balanced out and once the traffic decreases on the newly created circuit below the predefined threshold, the new circuit is released and the previous configuration is used again to forward traffic.

4.3 Coordinated Self-healing

The Coordinated Self-healing scenario is similar to the IP Traffic Offloading scenario. We have again a path from *MX240-1* though *MX240-2* to *MX240-3* with the same circuit setup in the transport network. In this scenario, the connection from *MX240-2* to *ADVA NODE 3* was disabled to simulate a link failure.

The ONE adapter identifies the loss of Internet connectivity of the router via alarms from the IP NMS indicating loss of connectivity between *MX240-1* and *MX240-3*. The same procedure as in IP Traffic Offloading is triggered and a new circuit is reserved, redirecting the traffic. The broken link can be repaired without disrupting the network flows and once the link is repaired, the previous configuration is used again to forward traffic.

5. Conclusion

In this paper, we have outlined the problems derived from the isolation between the IP and the transport management ecosystems, and have highlighted the reasons why a disruptive all-encompassing solution will not be easily adopted by network carriers. At the same time, recognising the need for automated communications between these two ecosystems, we proposed the utilisation of the ONE adapter, an ontology-based communication adapter. This ONE adapter is envisioned as a non-disruptive technology, to facilitate communication between the two management ecosystems using existing standardised protocols, and also to integrate third party management subsystems in an efficient manner. The use of standard protocols, such as MTOSI, ensure that the adapter itself remains non-disruptive, while the integration of external subsystems can reduce the cost and duplication of roles in multiple layers, as well as enhance the inter-layer communications by incorporating service awareness. To this effect, we have highlighted the immediate advantages of using this approach in current networks and show in our demonstration working use-cases for IP Link Provisioning, IP Traffic Offloading and Coordinated Self-healing in a real testbed scenario.

Acknowledgements

This work was supported by the European Commission through the ONE project in the Seventh Framework Program (FP7), contract number INFSO-ICT-258300.

References

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

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

[3] ADVA Optical Networking, FSP Network Manager and FSP Service Manager, http://www.advaoptical.com/.

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

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

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

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

[8] A. K. Y. Wong, P. Ray, N. Parameswaran, and J. Strassner, "Ontology Mapping for the Interoperability Problem in Network Management," IEEE Journal on Selected Areas in Communications, Vol. 23, no. 10, pp. 2058–2068, October 2005.

[9] "Reducing complexity and minimizing mistakes in network configuration," IBM Tivoli Netcool Configuration Manager, White paper, September 2010.

[10] S. Kaza and H. Chen, "Evaluating ontology mapping techniques: An experiment in public safety information sharing," Decision Support Systems, Elsevier, Vol. 45, no. 4, pp. 714–728, 2008.

[11] FP7 Project ONE (INFSO-ICT-258300): "Towards Automated Interactions between the Internet and the Carrier-Grade Management Ecosystems," http://www.ict-one.eu/.