

The Internet and Transport Network Management Ecosystems: A Roadmap Toward Convergence

M. Yannuzzi¹, A. Jukan², X. Masip-Bruin¹, M. Chamania², R. Serral-Gracià¹, V. López-Álvarez³, O. González de Dios³, A. Azañón³, M. Maciejewski⁴, C. Brunn⁴, M. Roth⁴ and J. Altmann⁵

¹Technical University of Catalonia (UPC), ²Technische Universität Braunschweig (TUBS), ³Telefónica I+D

⁴ADVA Optical Networking, ⁵Seoul National University (SNU)

Abstract—For many years, the convergence of IP data services and transport network services based on optical transmission has been at the heart of carriers’ investments and business strategies. However, significant challenges still remain. Over time, the inherent technological differences between the IP and transport networks have deeply segmented their operation, leading to the carrier’s organizational fragmentation and segregation of management competencies. Although these two networks are typically deployed in tandem, carriers still lack of tools capable of providing automated coordination of management procedures and orchestration of business practices between them, especially, in multi-vendor deployments. Overall, the separation of management functions between the “IP” and “transport” networks continues to be profound, and remains one of the major impediments hindering the expected convergence of IP and transport network services. 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 interoperation and coordinated actions between the IP and transport management systems already in place in carrier networks. This solution is being prototyped with special attention on multi-vendor network management scenarios.

I. 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 fragmentation of technical competencies has resulted in two administratively separate networks, where management team expertise and operational practices significantly differ. 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 T-NMSs specify the services and associated functions in a standard format, and proprietary Element Management Systems (EMSs) mainly 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, 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 than is commonly handled in each layer separately, especially, in operational settings, which are typically composed of nodes provided by multiple vendors. 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 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 go-to-market, thus creating increased revenue opportunities. Also 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 toward the development of a unified control plane framework to support both “packet” and “circuit” services, carriers remain reluctant to deploy any unified control plane solution without a level of manual control and coordinated automation between these two networks.

In this paper, we revisit the issues of coordinated operations between the IP and the transport networks, and propose a roadmap toward coordinated multi-layer interactions via the management planes. We describe a new system, which is essentially a 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 proposed system can also be extended with standardized interfaces toward the emerging third party network management systems, such as the Path Computation Element (PCE) [2]. 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.

II. THE INTERNET AND TRANSPORT NETWORKS: TWO SEPARATE MANAGEMENT ECOSYSTEMS

Transport networks were designed to deliver a small number of services with fairly static demands on network operation. In practice, transport networks are operated via T-NMSs (see, e.g., [3]), which support many service-oriented functions

through vendor-specific platforms, dramatically reducing the operational overheads of telecom carriers as well as the complexity of the management tasks involved. In addition, the ever increasing demand for bandwidth has led the industry to heavily invest in R&D, in order to cope with the increased transmission capacity while simplifying the operation and maintenance of transport networks as much as possible.

The IP network configurations, on the other hand, became increasingly complex over time, and vendor-specific. First, the changing dynamics of the Internet has driven the deployment of a wide spectrum of IP enabled equipment by telecom carriers. 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].

As a result, telecom carriers have been forced to support the complexity and associated cost of the operations required at the IP layer, with the simplicity and cost savings of operating and configuring the equipment at the transport layer. It is worth highlighting that significant advances have been made toward the development of a unified *control* framework, with the aim of reducing the human intervention in the process of service provisioning as well as providing a standard solution for inter-layer control plane interactions supporting both packet and circuit switched networks [5], [6]. Even though the control planes can address automation of specific management functions, they cannot automatically orchestrate actual management and business procedures. Moreover, network op-

erators are increasingly using third party systems such as PCEs, network planning and 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 toward the convergence of the IP-NMSs and the T-NMSs that do not require the integration of these 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, and automated communication channel between the two management layers. The initial goal should be to enable coordination, so as to support a set of basic operations, such as provisioning, and coordinated post-failure management.

Figures 1(a) and (b) illustrate some of the consequences of the isolation between these management systems and the solution considered. Even the provisioning of a new IP link (A) 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 B), or any type of coordinated self-healing action (C). Carriers also desire an automated communication with external control and management subsystems, such as the PCE (D). The development of an adapter that can bridge the interoperability gap between these management seems essential to telecom carriers.

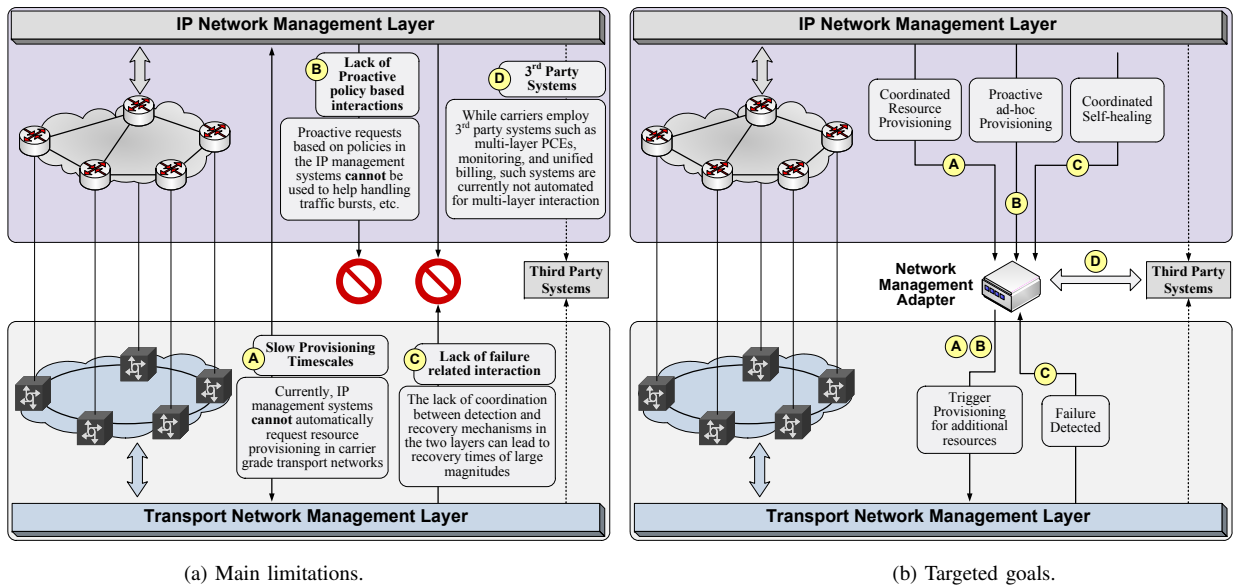


Fig. 1. Research challenges: Can the interactions between two management systems already deployed be enabled in a scalable and extensible fashion, and furthermore, could such “enabler” be introduced and adopted in the telecom management practice in a non-disruptive way?

III. COORDINATED MANAGEMENT INTERACTIONS: A PRACTICAL APPROACH

The goal is to make possible the interoperability between the IP-NMSs and T-NMSs already deployed in carrier networks, and thereby enable automated coordination of management procedures in a practically-relevant manner. To this end, a management adapter is proposed, which is being developed and prototyped within the FP7 project ONE [7]. In a nutshell, this adapter will be easy to integrate into the existing management ecosystems and will not depend on any specific vendor or NMS implementation. The adapter will also exploit existing standardized interfaces and protocols, including standardized interfaces to the external management subsystems, e.g., to the PCE and the AAA subsystems. In fact, the adapter that is being prototyped in [7] will act as a Path Computation Client (PCC) [2]. Moreover, the adapter will provide functions for *orchestration* of multi-layer interactions in an easy-to-implement manner, as well as include a controlled automation of coordinated management tasks.

Figure 2 illustrates the building blocks of the proposed Network Management (NM) adapter, which is composed of three primary modules: a *Front-end Management Module*, an *Ontology Mapper*, and the *Operation Workflow Database*.

Front-end Management Module: It is in charge of receiving the requests, and all the operations initiated through the Front-end Management Module are internally designed and processed as *workflows*. A workflow consists of the sequence of actions required to coordinate and automate a set of cross-layer operations. Each workflow is specified as a telecom practice and management procedure. One workflow can be used for multiple requests—the difference between two requests which are processed using the same workflow is in the input parameters. For instance, two requests for the provisioning of an IP link may differ in the end-points, the IP addresses, and the capacity required from the transport network, but the workflow used to orchestrate the provisioning of the link will be the same in both cases. It is worth noting that these workflows reflect current business processes within the carrier’s organization, which typically do not change with change in technology. In this regard, the role of the adapter is to facilitate the telecom business processes as they are, rather than modifying them.

The Front-end Management Module also provides a *programmable* framework through which carriers can build and orchestrate their own operations. As an example, consider the scenario shown in Fig. 2, which illustrates an operator-initiated action requesting the provisioning of a new IP link between a

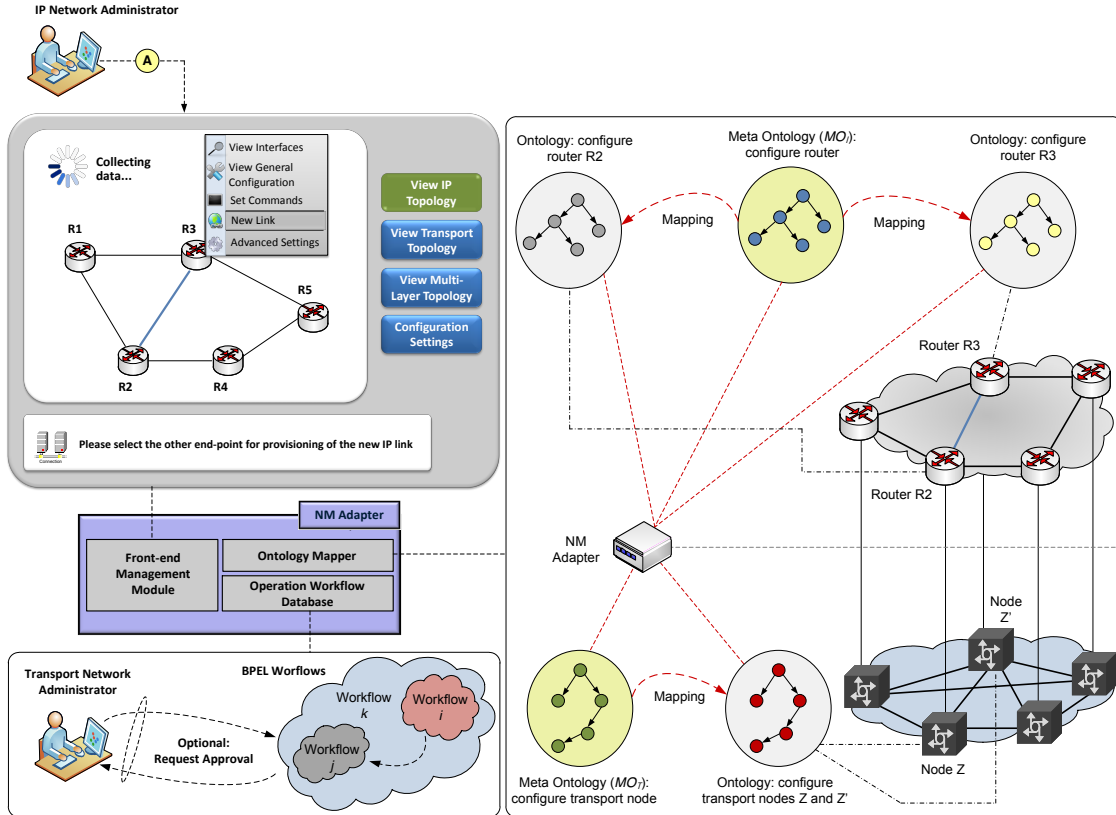


Fig. 2. The main blocks in the architecture of the management adapter. The orchestration in the example shows the case of the provisioning of a new IP link between routers R2 and R3. The orchestration is triggered as a request which is matched to a pre-configured workflow for internal processing and execution. The corresponding workflow is retrieved from the Operation Workflow Database, and the Ontology Mapper is the module in charge of the semantic interpretation of the configurations required and their corresponding mapping to the specific command set of the devices involved. After the mappings, the adapter can request the required configurations to the transport T-NMS, and use either the IP-NMS or the CLI for the IP layer. Note that the execution of a given workflow may optionally require the approval of an administrator from the transport layer.

pair of available interfaces at two routers in the IP network, namely, routers R2 and R3. This is facilitated by the Graphical User Interface (GUI) provided by the Front-end Management Module, which allows the operator in the IP layer to orchestrate a set of operations, part of which will require the configuration of resources in the transport network. Such orchestration could be launched as a one-time provisioning request, or may be recorded and stored in the Operation Workflow Database as a new workflow for its future reuse. The programmable nature of the adapter will allow carriers to customize their provisioning processes as well as to program the sequence of actions that need to be performed upon particular events.

The coordination of management tasks through the adapter is supported through web services, and in particular, via the Multi-Technology Operations System Interface (MTOSI) toward the T-NMS [8]. MTOSI is an increasingly important TeleManagement Forum standard that provides an open management interface between Element Management Systems (EMSs), NMSs, and/or a Service Management System (SMS). Due to its simplicity and platform independence, MTOSI is becoming one of the preferred web-service based standards in industry (see, e.g., [3]). In addition to the operations initiated through Web services, the adapter also provides support for coordinated actions triggered by SNMP traps, which may be originated by the NMSs, by the Network Elements (NEs) directly, or by any external management subsystem.

The Ontology Mapper: Due to the interoperability gap between the IP and transport management systems, the adapter needs to implement the corresponding semantic adaptations. In the proposed adapter, the formal representation of concepts is based on a set of *ontologies*, and the semantic interpretation of the configurations required and their corresponding mapping to the command set of the devices involved is solved by means of *mappings* between entities in these ontologies.

An ontology can be defined as a pair $O = (V, A)$, where V denotes a formal description of a vocabulary, and A represents a set of axioms that specify the interpretation of the vocabulary V in a certain domain of knowledge. The vocabulary is frequently modeled as an ordered set of concepts, each of which may have a collection of instances that are connected through a group of relations in the form of a hierarchical graph.

Any mapping process in this context requires of a mechanism that can reliably assess the semantic similarity between entities that belong to two different ontologies. The usual approach is to use a *similarity* function S , which basically quantifies how close the semantic meanings encoded by these entities are between each other [9]. More specifically, let e and e' be two entities in the vocabularies V and V' of two ontologies, O and O' , respectively. The similarity between e and e' can be captured by $S(e, e')$, with $0 \leq S(e, e') \leq 1$, where $S(e, e') = 1$ when e and e' are semantically identical and 0 when they have no semantic content in common. On this basis, the mapping from e to e' can be defined as a function $M : O \rightarrow O'$, such that $M(e) = e' \Leftrightarrow S(e, e') = \max S(e, u) \forall u \in V'$ and $S(e, e') > t$, with t being a threshold. The latter is usually adjusted to increase the precision of the mappings.

When a carrier switches from one router/switch vendor to another, this may considerably change the vocabulary used as well as the configurations required for carrying out a set of operations—although the operations per se often remain unchanged. To address this challenge, the workflows stored in the Operation Workflow Database of the network management adapter should be agnostic of any technology, and should use a standard data and process model to describe the operations and mappings required. In the proposed adapter, the workflows are based on the Business Process Execution Language (BPEL) [10], and the ontologies involved in the workflows are normalized to uniform representations, which we call *Meta Ontologies* (see the right-hand side of Fig. 2). A Meta Ontology offers a common and device-independent framework which normalizes the configuration tasks required at IP and transport layers.

In the adapter, the mappings are performed between entities in a Meta Ontology and entities in ontologies that conceptualize the configuration of proprietary systems. The market is visibly moving in this direction; for instance, IBM's Tivoli Netcool Configuration Manager [11] provides a similar approach, by hiding the complexity of proprietary configurations from the administrator. This tool uses XML schemas to provide mappings between vendor-specific vocabularies and XML. Despite this strength, this tool does not provide support for orchestrating operations involving configurations both in IP routers and transport layer EMSs.

Figure 2 illustrates the mappings required to provision a new IP link between routers R2 and R3. We assume that R2 and R3 are from different vendors (hence two mappings are required), while the transport nodes and their corresponding NMS are all by the same vendor (i.e., one mapping). A high-level description of these mapping processes is illustrated in Figure 3. The Meta Ontologies required at the Internet and transport layers are denoted as MO_I and MO_T , respectively. Once the workflow is loaded and processed, a set of ontological entities $\{e_I\}$ in MO_I and $\{e_T\}$ in MO_T are identified, which are those that need to be mapped. The destination ontologies for these mappings can be inferred from the input parameters that triggered the coordinated operation through the adapter. For instance, two destination ontologies are required at the IP layer (R2 and R3), while only one ontology is needed to manage the configuration of the transport nodes Z and Z' in Fig. 2. Observe that the mappings are based on maximizing the similarity function between the ontological entities (for simplicity, we omitted the potential constraint imposed by the threshold t). Also note that the vocabularies relative to the IP and transport layers may differ considerably. This means that the ontological representation of entities and their potential similarities could differ as well, so the similarity functions used for each layer may be different. The details of the similarity functions are out of the scope of this paper, for further reference please refer to the large body of literature on this subject. To the best of our knowledge, none of the previous work focused on coordinated interactions in multi-layer management scenarios.

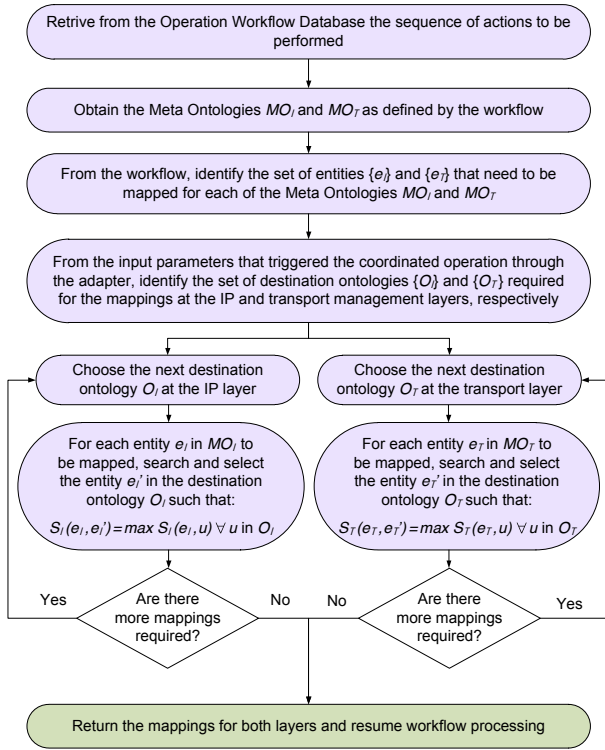


Fig. 3. The ontology mapping processes.

The Operation Workflow Database: It is used to store the workflows which contain the definition of a multi-layer management processes. These workflow definitions can then be executed using a standard workflow execution language such as BPEL [10], which is used extensively for business process automation. The workflow definitions can also be used as small function definitions for other more complex processes, thereby making complex process definition easier and more robust. As an example, the simple IP link provisioning workflow could be used as part of a failure recovery workflow to setup a new link. By providing the ability to store and retrieve a workflow, we ensure that *triggers* for events do not need to carry complex process definitions, and use the existing process definitions inside the proposed adapter to facilitate operation.

IV. APPLICATIONS

Today, the provision of an IP link requires multiple human interactions, and therefore may take *days* as opposed to *seconds*, as one would expect. Although the setup of an IP link is a basic operation, it offers an illustrative example on how the adapter can be used to facilitate coordinated operation across multiple layers. This coordinated action requires four basic steps: 1) determining link end-point available free IP interfaces at the corresponding routers; 2) determining the client transport interfaces connected to the selected free IP interfaces; 3) provisioning a circuit service between the corresponding client interfaces in the transport network with the capacity required and suitable framing adaptations; and finally, 4) configuring the IP interfaces at the two end-points to initialize the IP link.

In the example shown in Fig. 4, the *trigger* (①) for the operation is generated by the operator (e.g., through the GUI shown in Fig. 2) which includes information about the end-point routers in the IP network, the available IP interfaces, the required capacity, and the IP addresses to be used for the required link. Based on this, the IP link provisioning workflow is selected from the Operation workflow database (②, ③), so the adapter requests the PCE for a path between the two end-points in the transport network (④). The PCE uses the multi-layer TED to determine the correlation between the IP interfaces and the corresponding client transport interfaces in the transport network and computes a path between them in the transport network and returns this information to the adapter (⑤). Once the transport path is obtained, the Ontology Mapper module is invoked (⑥, ⑦), which provides the semantic adaptations for the configurations required at the IP and transport layers. Using this information, the adapter first requests the T-NMS via MTOSI to setup a circuit in the transport network (⑧, ⑨, ⑩), and if successful, it requests the IP-NMS to configure the IP interfaces (⑪, ⑫, ⑬).

While this basic coordination can be facilitated in a number of automated ways, such as by means of the control plane, the novelty of the proposed approach remains in its capability to flexibly develop and re-use workflows, and ensure that the workflow definitions themselves are not affected by changes in the external systems used (e.g., the workflows remain agnostic to the IP-NMSs and T-NMSs actually used). For instance, in case of a technology transition to a new IP-NMS, the operator would only be required to create ontology definitions based on the new interface, since the core operation of the adapter is immune to such changes. This ability to easily create, store, and re-use workflows facilitates the development of a rich set of complex multi-layer interaction scenarios, most of which cannot be entirely solved through control planes. Here-with, we illustrate a couple of areas of interest to providers for facilitating multi-layer coordination, which can be classified as: i) policy-driven optimization, and ii) failure recovery and post-failure optimizations.

Policy-driven actions deal with cases where operations are *triggered* automatically, e.g., in response to changes in the network state based on a pre-configured policy. The IP offloading paradigm is one such application which is designed to cope with a sudden increase in traffic in the IP network. In this case, the adapter may use incoming SNMP traps to determine the overloading level of an IP link, and use interfaces to the monitoring systems to determine the required set of bypasses and routing rules to be established. The adapter could then use the typical link provisioning workflows with IP addresses assigned from a private IP address pool, to ensure that bypasses do not disrupt regular IP routing, and then configure rules for re-routing specific traffic flows onto the established bypasses to alleviate overloading conditions.

In the failure recovery and post-recovery scenarios, we envision the possibility to use the adapter to facilitate coordinated healing and provide post recovery functions. The separation of the IP and optical management layers has led to highly re-

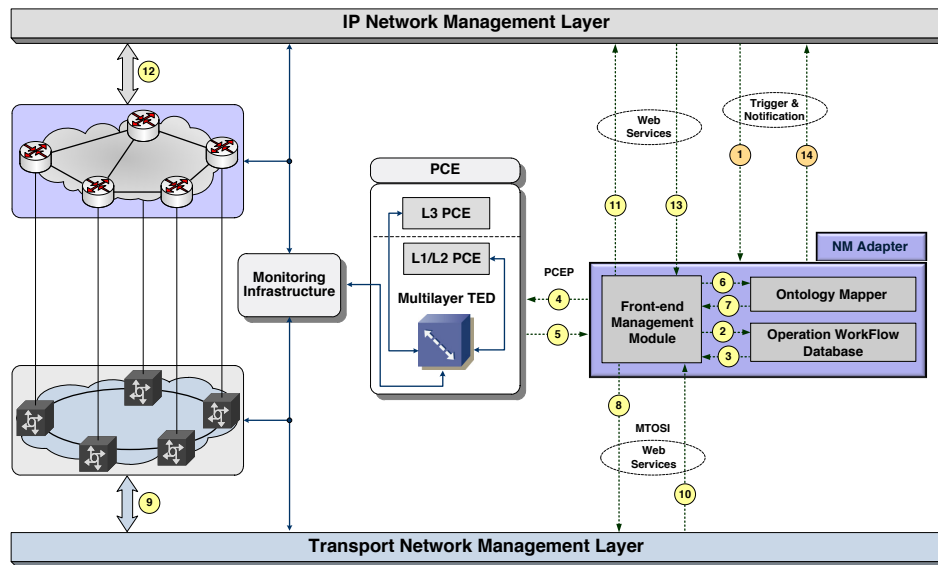


Fig. 4. An example showing coordinated service provisioning involving both the IP-NMS and the T-NMS.

dundant protection schemes, since each layer is equipped with its own protection capacity. Despite this redundancy, 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. To this end, the network management adapter can be used to facilitate:

- (1) Timing coordination, i.e. instead of using fixed timers to respond to failure operations, the adapter can monitor recovery operations in the transport network, and initiate recovery operations in the IP layer only when the recovery in the transport network fails.
- (2) Recovery of services from multiple failures by re-provisioning services over different multi-layer paths.
- (3) Post-recovery actions, such as setting up protection paths for services that were recently re-routed onto backup paths due to a failure.

Overall, the utilization of the proposed adapter is expected to allow faster recovery response times and a more efficient use of resources by reducing the redundancy in the protection equipment and operations. We expect lower CapEx with this technique, as networks can now be planned and dimensioned according to the availability of this self-healing ability. We also expect lower OpEx, as networks can now be operated more efficiently based on existing telecom practices.

V. CONCLUSIONS

In this paper, we presented a roadmap for convergence of Internet data services and transport network services from the point of view of network management, and described an easy-to-deploy solution that can enable their coordination via a network management adapter. We highlighted various application scenarios where the proposed network management adapter can facilitate service provisioning, coordinated policy-driven

operations as well as coordinated self and post-recovery actions. To the best of our knowledge, this approach is the first attempt to enable what we define as a controlled 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.

ACKNOWLEDGEMENTS

This work was supported by the European Commission through the ONE project (www.ict-one.eu) in the Seventh Framework Program (FP7), contract nr. INFOS-ICT-258300. UPC authors also acknowledge the support received by the Spanish Ministry of Science and Innovation under contract TEC2009-07041, and the Catalan Research Council (CIRIT) under contract 2009 SGR1508.

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] FP7 EU Project ONE, <http://www.ict-one.eu>.
- [8] Multi-Technology Operations System Interface (MTOSI), release 2.0, TeleManagement Forum, <http://www.tmfforum.org>.
- [9] 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.
- [10] Web Services Business Process Execution Language Version 2.0, OASIS Standard, 11 April 2007, <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.pdf>.
- [11] "Reducing complexity and minimizing mistakes in network configuration," IBM Tivoli Netcool Configuration Manager, White paper, September 2010.