TECHNOLOGY WHITE PAPER

Multi-Region Networks

Generalized Multi-Protocol Label Switching (GMPLS) as enabler for vertical integration

Multi-Region Networks (MRN) is solution for vertical integration based on the Generalized Multi-Protocol Label Switching (GMPLS). Considering the limitations of current interworking models and the requirements for integration, MRN offers the possibility of operating domains hosting several technologies as a single network. Using a single instance of the control plane, multi-layer operations are profiled and optimized thanks to a limited number of protocol extensions. Routing and signaling are extended to support internal node adaptation capability and deterministic switching capability indication along the explicit route, respectively.

ALCATEL

Introduction

Most vendors and carriers are showing a growing interest in control plane technology based on Generalized Multi-Protocol Label Switching (GMPLS). This offers a unique opportunity to develop a new set of advanced items that squeeze the most out of this unified approach in terms of (inter-)network models spanning diverse data plane technologies. In a background of multi-protocol (multi-service) interoperability constraints and inherited multi-switching layer architectures, carriers are making strategic decisions to continue to evolve their network towards an operationally integrated solution. A few concrete, simple issues illustrate the benefits of convergence. First, heterogeneity creates interoperability complexity, a lack of automation, and a lack of efficiency. It also implies dedicating - and replicating - developments, network infrastructures, control planes, and associated operations (management) and competencies. Finally, functional and operational model gaps should not mandate, nor do they justify, entirely rebuilding the network each time a new technology is introduced.

Building a protocol suite that bridges “data” and “telecom” philosophies in a unified way may seem utopic. Nevertheless, with a pragmatic re-use of dominant IP-centric technologies, but relating to transport world requirements, GMPLS has become the sole credible candidate for each existing sub-IP layer (MPLS being part of this superset). Since the protocol suite kernel is almost complete, the diversity of Switching Capabilities (SC), from packet to fiber, can benefit from the GMPLS control plane. Capitalizing on the intrinsic unified features of GMPLS, the time has come for integrated models where commonalities in protocols serve to overcome the traditional drawbacks mentioned above.

As networks are rarely composed of a single data plane technology, the applicability of a control plane being in charge of multiple switching layers (as well as their combinations) within the scope of a single network is of interest for a wide range of carriers. The architecture targeted here refers to vertical integration, where nodes hosting multiple and interworking SC are controlled by a single instance of the control plane. Therefore, it differs from horizontal integration, which is related to interworking between routing systems (such as areas and Autonomous Systems).

The issue is then to extend – if needed – the GMPLS protocol suite capabilities for such integrated (both in data and control plane) networks. The expected solution will benefit from a single set of tools and an almost complete perception of what is happening through the layers. As a consequence, the management plane will in turn converge into a single and consistent entity.

Multi-Region Networks

Most of the initial efforts on GMPLS have been dedicated to single instance controllers for single switching layer capable devices. However, the ability to operate various switching layers within networks using a unified control plane technology is a strong carrier requirement. Considering typical infrastructures, it becomes sensible to seek better coordination between multi-technology networks that previously merely coexisted. Therefore, the concept of a multi-switching layer capable network has been introduced. It is referred to as Multi Label Switched Path (LSP)-Region Network or simply Multi-Region Network (MRN) [1].

Multi-Region Networks have to be considered in the scope of vertical integration, that is environments in which at least two different switching layers, which may be hosted by the same device, are present. More precisely, vertical integration defines collaborative mechanisms allowing a single instance of the control plane to control multiple (at least two) data plane switching layers potentially integrated into a single system.

Vertical integration is complementary to the horizontal integration also under development [2]. Horizontal integration is defined when each entity constituting the network environment includes at least one common (data plane) switching layer, and the control plane topology extends over several partitions (i.e., routing systems), being either areas or Autonomous Systems (AS). In this case, integration is defined between nodes hosting the same SC. For instance, the control plane interconnection between SONET/SDH switching-capable routing areas defines a horizontal integration. Figure 1 depicts a network environment including three partitions. Partitions 1 and 2 host a single data plane switching layer. Partition 3 is vertically integrated since its hosted switching layers are controlled by a single control plane instance. One of these three layers (in the present case, Ethernet) is identical to the one hosted by partitions 1 and 2, so constituting a horizontally integrated network.

Fig. 1 Vertical and horizontal integration

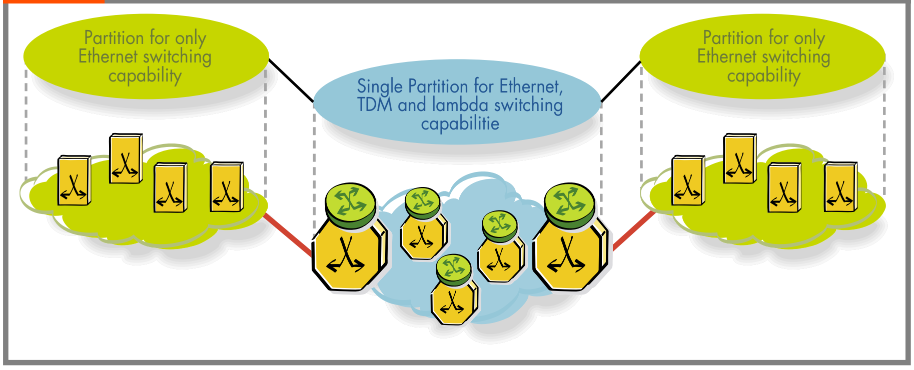

Realization of the MRN control plane is based on the GMPLS protocol architecture (as defined in RFC 3945). Even, if this concept is already positioned as an enabler for next-generation integrated network, the purpose is to identify issues and propose protocol extensions and specific profiles fitting new interworking models. While making intensive use of the GMPLS control plane capabilities, MRN is also expected to capitalize on existing building blocks with as little supplementary effort as possible, including protocol extensions.

Model applicability

Before going into the details of MRN applicability, a few key GMPLS concepts are briefly introduced (more details can be found in [3]):

- From the control plane viewpoint [4], a (set of) data plane switching layer is mapped to an LSP region. The MRN approach fully exploits the routing and signaling features of GMPLS. Particularly, extensions of Traffic Engineering (TE) links [5] allow control of network resources through the layers with the necessary characterization.
- A Switching Capability (SC) is associated to a TE link end-point (interface) that defines the capability of this interface to switch data traffic based on certain generic properties associated to the set of corresponding data plane layers. An example of such value is TDM that is associated to a SONET/SDH cross-connect interfaces that switch data traffic based on data's timeslot in a repeating cycle (e.g. from VC11 until VC4-256c). Thus, the link inherits its TE properties from its association to one (or more) SC, being either Packet (PSC), Layer 2 (L2SC), Time Division Multiplexing (TDM-SC) or Lambda (LSC). The SC information is part of the Interface Switching Capability Descriptor (ISCD) attribute associated to the TE link.
- A (bundled) TE link [6] is defined as a set of one or more component links; the TE capabilities associated to a bundle represent the aggregate of the TE information associated to its individual component links. Moreover, since this definition is recursive, component links can be defined themselves as TE links. Components of the TE link may follow different paths between the pair of nodes that they interconnect. As TE links can comprise a large number of component links, representation of un/allocated resource capacity is associated to these links and allows for efficient path computation and signaling. Note that this representation provides for resource allocation in discrete units (e.g., for time-division multiplexing) or continuous units (e.g., for statistical multiplexing).
  Since there is no longer a one-to-one association between a regular routing adjacency and a TE link, the number of routing adjacencies in the network can be kept proportional to the number of control plane adjacencies and not to the actual number of data plane links. Therefore, independently of the type of interface switching capability associated to the link, link bundling improves the routing scalability by reducing the amount of information to be processed by the link state routing protocol instance. Furthermore, TE links have been extended to non-adjacent devices by introducing the Forwarding Adjacency (FA) concept. Once TE link resources are allocated as part of an FA LSP, its actual capacity can be represented as an FA TE Link within another TE link. This, in turn, enables a further decrease in the number of control plane instances to control N transport layers. Finally, the capability to bundle FA links and TE links as part of the same TE link allows for additional flexibility in controlling large-scale backbone networks.

Once the GMPLS building blocks are introduced, understanding of the MRN functionality (and requirements) is strongly related to its positioning with respect to the existing control plane interworking models. It is assumed that efficient control plane interworking between different data plane switching layers is a practical requirement implied by network architecture evolution. Respective specifications and limitations of the existing interworking models are described below.

Overlay model

The overlay control plane interconnection model was designed for carriers or (bandwidth service) providers leasing their network facilities to Internet Service Providers (ISPs). It arose from carriers and providers owning an extensive installed base of SONET/SDH transmission equipment that today has to deal with explosive growth of IP traffic and a dramatic increase of demand for Virtual Private Networks (VPNs). This legacy model (layered circuit-switched flavor coming from the ITU-T Recommendation G.805) assumes a very low trust relationship between the involved parties, mandates a strict separation of the respective network control planes (including their addressing spaces), and strictly limits the exchange of signaling information. Under this limiting model, the routing and signaling protocols of each control plane layer act independently. The collaboration between control planes is reduced to interactions through a User Network Interface (UNI) defining a client/server relationship. Moreover, the overlay model requires, by definition, an address resolution mechanism between the client and the server layer addressing space. Such operations require manual configuration, since there is no automated mechanism available that allows client nodes to obtain this information. As a consequence, this model is the most opaque, the less featured, and the least flexible of the common interconnection models.

A dilemma inherently linked to the overlay model is referred to as the “unknown adjacency” problem. In link state routing protocols such as Open Shortest Path First (OSPF), each router maintains a link state database (describing the network topology) from which a routing table can be derived by means of shortest-path computation. However, it mandates a routing adjacency between each pair of client nodes (e.g., IP/MPLS LSRs) to exchange the link state information, which results in an undesirable full mesh of connections at the server layer. Consequently, this model does not allow for top-down triggering of connections (and therefore precludes any end-to-end dynamic re-routing), and it scales very badly to large networks. It is even less suited to control systems hosting multiple switching layers, since this problem is replicated between each sub-node component. This model, depicted in figure 2, is therefore inadequate for vertical integration of a MRN control plane and strictly limited in its applicability to horizontal integration.

Augmented model

The augmented model is intermediate between the overlay and the unified model. Although more featured than the former, the augmented model allows the exchange of a limited amount of routing information (mainly reachability) between the client and the server layer's network control plane. Moreover, the augmented model allows (but does not mandate) different addressing schemes and full or partial opacity of the server layer addressing space. When used with a common addressing scheme, the augmented model solves - de facto - the address resolution problem between the client and the server switching layer addressing spaces. However, as this model conveys primarily reachability information to the client layer, it is mainly appropriate for large single-carrier networks partitioned in multiple Autonomous Systems (AS). The implication is that the augmented model must complement an MRN control plane, if the target network environment comprises multiple switching-layer domains determined on the basis of separate administrative/topological/geographical units within the same operator towards distinct carriers.

Fig. 2 Overlay model

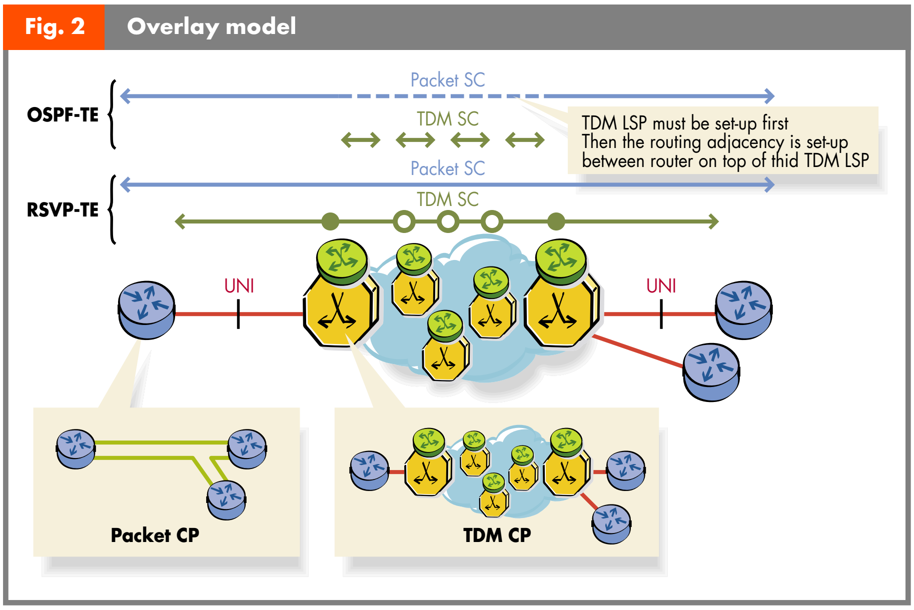

TDM LSP must be set-up first
Then the routing adjacency is set-up
between router on top of thid TDM LSP

Fig. 3 FA concept: LSP regions, FA-LSPs and TE links

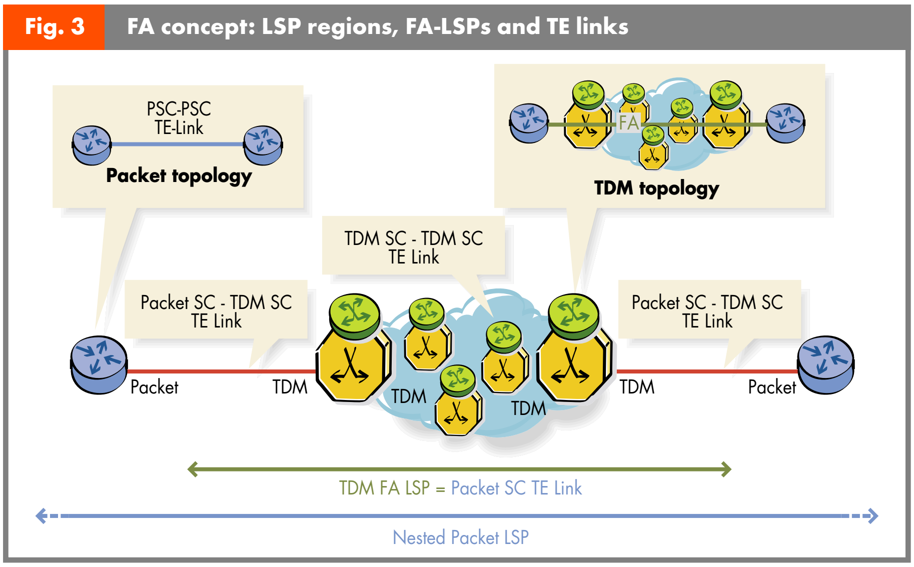

Unified model as starting point

Compared to the overlay model, the unified model assumes that the control plane applicability is ubiquitous, i.e., it applies independently of the data plane switching layer. Under this model, each control plane acts as a peer to the lower data plane switching layer. A complete exchange of routing information (involving all the interconnected layers) is possible, and a common addressing space is used.

One of the key concepts underlying the unified control plane interconnection model is the notion of Forwarding Adjacency (FA). Using this concept, a GMPLS-capable node may, under its local policy configuration, advertise a Label Switched Path (LSP) as a TE link into the same link state routing protocol instance (e.g., OSPF) as the one that determines the path taken for this LSP. Such a link is referred to as a “Forwarding Adjacency” and the corresponding LSP as a “Forwarding Adjacency LSP”. Afterwards, OSPF floods the link-state information about FAs, allowing other nodes to use FAs as any other TE link for path computation purposes.

Figure 3 illustrates the FA concept. The LSPs referred to as TDM FA-LSPs established through the lower LSP region (i.e., TDM) appear as TE links at the higher LSP region (i.e., PSC). When an FA-LSP is triggered, the TE attributes of the corresponding TE link are inherited from the incoming LSP request that induced its creation. Therefore, once set up, the TDM FA-LSP appears at the higher region as a TE link with an SC value corresponding to the triggering LSP, i.e., a PSC TE link.

The use of FAs provides an efficient mechanism for improving both 1) routing scalability, as the number of control plane adjacencies becomes independent of the number of data plane adjacencies (and expected to be much lower) and 2) signaling scalability, as the state of a nested LSP is now only maintained at the terminating nodes of the nesting FA-LSP.

The FA concept can be considered as the peer/unified model enabler, by overcoming the full mesh scalability problem encountered in any overlay-based environment such as IP-over-ATM or its modern version, Pseudo-Wire (PW)-over-Packet Switched Networks (PSN). As a consequence, the unified model entails the tightest coupling between layers by considering all nodes as Label Switching Routers (LSR). In turn, the scalability of such a network is equivalent to that of a "regular" IP/MPLS network with the same total number of nodes. In addition, as the "unknown adjacency" problem does not apply, this model supports pre-provisioned and dynamically triggered connections. Therefore, this model is essential for operators who control both the IP/MPLS and the optical transport infrastructure, as it allows them to optimize their network design and end-to-end operations.

Fig. 4a **Multiple control plane instances**

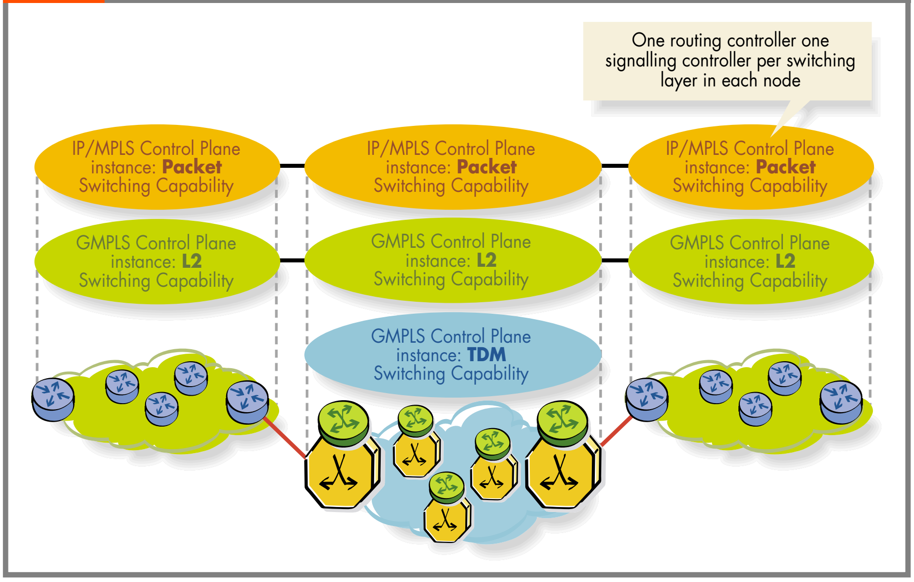

Fig. 4b **versus a single control plane instance**

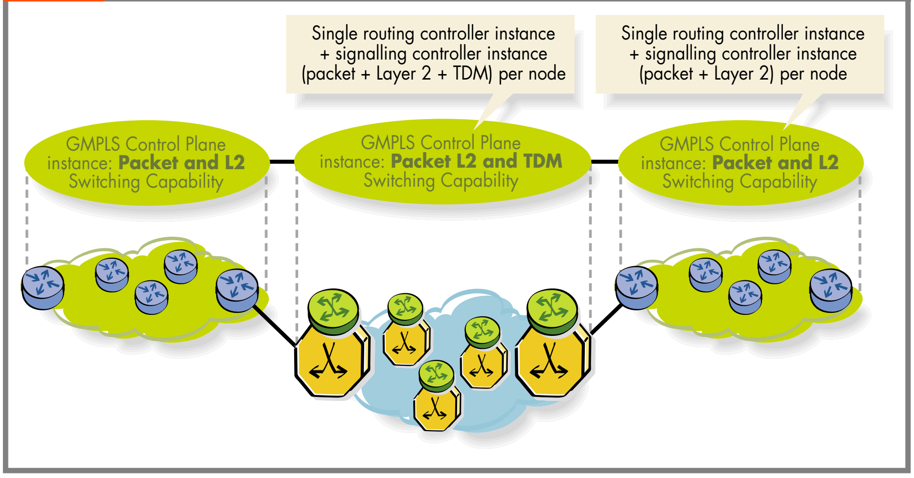

MRN application

Environments that include nodes hosting more than one (data plane) switching layer are already part of an operator's day-to-day life. For such environments, control plane integration is a key enabler for network resource optimization and more notably operation simplification. A good example of such a node would include IP/MPLS, Ethernet, and SDH switching capabilities under the supervision of a single GMPLS controller. For instance, such a system allows MPLS Packet LSPs (P-LSP) to be set up on top of Layer 2 LSPs (L2-LSP), themselves nested into TDM LSPs. Note that Ethernet L2-LSPs [7] rely on VLAN swapping capability, where the VLAN label is defined as part of the regular Ethernet frame header.

In the unified model context, the GMPLS protocol suite currently assumes that each of these LSPs can be established using a common instance of the control plane. However, it does not specify how these LSPs can interact with each other. MRN enables a single controller (i.e., a single GMPLS control plane instance) to handle multi-layer capable nodes. This single control plane instance advertises, in addition to the canonical single layer routing information, the information that represents the cross-region TE constraints to be processed in a multi-layer provisioning framework. Moreover, introduction of a few signaling extensions allows the indication of the deterministic location of cross-region points when requesting resource provisioning across the various data plane switching layers.

Consistent with these new capabilities, the inheritance of respective region characteristics (TE attributes, protection and restoration information) should be profiled, allowing actual multi-layer networking. It is expected that this additional information becomes an enabler for better coordination in terms of FA-LSP usage and their triggering depending not only on the TE link resource usage but also on the actual node capabilities. It should be noted that such constraint consideration is achieved without adding node state information to the link state routing protocol instance.

Being agnostic to the increasing number of data plane switching layers, the MRN concept allows a single Data Communication Network (DCN) and a single addressing space to be maintained, together with the optimized and automated operational mode of one single integrated network. Figure 4 depicts the result of applying the MRN concept (4b) to an existing multi-layer network with independent, non-collaborating and diverse control plane instances (4a).

The MRN solution

Behind the GMPLS-based protocol extensions and profiles, the MRN solution is positioned as a set of architectural principles (system and network) and a specific theory of operations.

- MRN is the solution for vertical control plane integration where carriers want to overcome traditional multi-layer limitations and complexity. As such, it applies first in scenarios of a single domain/single carrier facing multiple technologies in the infrastructure for which there is no fundamental reasons for control plane duplication, since an intrinsic trust relationship is assumed.
- MRN is particularly applicable to the control of equipment hosting several SCs. Many combinations of those systems exist or are foreseen, while the MRN solution matches vendors' and carriers' interests in terms of supported features, development, and competency capitalization with a view to unified and simplified networking. Based on the GMPLS unified properties, only one single controller per node is required (and that runs a single control plane instance). This single instance deals locally with the various SCs and processes the domain-wide control information.
- MRN avoids classical "blind" operations performed on a per-layer basis and consequently the inherent complexity induced to reach consistent synchronized states and efficient multi-layer networking. As such, MRN opens the door for automated operations aiming at engineering traffic and resources across the various data plane switching layers.
- MRN introduces the features needed for dynamically and deterministically controlling the provisioning of network resources. By enabling a "multi-layer TE vision" of the network infrastructure, MRN allows for control-driven resource provisioning without any pre-provisioning. However, during the provisioning phase, as each node may integrate multi-switching capability, the explicit route (result of the path computation) carried as part of the signaling request is not sufficient to determine, on a per link basis, the exact SC to be used. The fundamental reason lies in the fact that explicit routes (as currently defined) do not include any TE information on the selected links but only topological information (e.g., interface identification). Therefore, a new sub-object of the eXclude Route Object (XRO) is introduced to indicate the decision regarding SC selection. In the case of nodes hosting several SCs that do not exhibit the same capability for each of their links (and that are not advertised as part of the TE routing information), blocking situations may occur. The Interface Adaptation Capability Descriptor (IACD) has been introduced to solve this issue. Its role consists in describing the node's link adaptation capabilities without introducing a "node state" within the link-state routing protocol instance. Finally, inheritance of respective region characteristics (TE attributes, protection and restoration information) should be profiled, allowing actual multi-layer networking.

Integrated network and system architectures

Integration is a paradigm that applies to both data and control planes. The current state-of-the-art results from the experience of combining multiple technologies. Since each of them associates data and control, their combination generates a wide range of interworking relationships. End-to-end networking is therefore dependent on multiple architectural and operational issues. And it becomes important to enable step-by-step convergence to build networks with simplified data and control plane architectures. While the need for value-added points of service (typically layer 3 and above) is not essential the full length of the path between end-users, transport networks may regarded as a single function supporting provisioning and survivability. In view of this commonly accepted separation, there remain various constraints to be further integrated before the days of a single, homogeneous infrastructure.

First of all, it seems unreasonable to freeze network evolution, as requirements depend on continuously evolving applications. On the contrary, integration is expected to bring enough simplicity and flexibility to accommodate such an evolution while keeping a common control plane basis. Another major constraint is the traffic itself, which has by definition non-uniform properties in time and space domains. Together with the current cost constraints, this leads to a multi-switching capability infrastructure. Moreover, as IP technologies apply to numerous networks and devices, they also constitute the basis of an IP-centric control plane. Consequently, GMPLS, as unified control plane technology, capitalizes also on a common set of widely applicable IP protocols and building blocks.

For transport networks in particular, which will have to deal with several data plane switching layers, it is of crucial importance to design an integrated architecture using a single instance of the control plane. From packet to lambda, the integrated architecture should be a substitute for the sum of separate but interdependent networks. MRN is the GMPLS application for controlling such integrated architectures. This vertical integration is of relevance in every situation where the various switching layers constitute the resources to be controlled by a single instance of the control plane. In particular, mono-carrier, multi-technology networks are the primary target, since no security issue should restrict full collaboration between the resources.

So applying the same integration paradigm at the node level becomes natural. This situation refers to systems hosting several SCs, for instance, transport elements providing both packet and circuit switching (as depicted in Figure 5). In principal, there is no reason to introduce separate instances of the control plane (one per SC) within a single device while GMPLS offers the capability to control all the SCs using a single instance (and thus a single controller). A direct benefit of this approach is that the number of network elements to be controlled is significantly lower. In turn, this prevents multiple identification of the same object and so allows the scalability of the control plane to be optimized. Another noticeable benefit is that it prevents any additional entity from controlling interactions (and so ensures collaboration) between these instances. Moreover, centralized intelligence for all the SCs hosted by a node paves the way for unprecedented optimization through local but multi-layer policies.

Fig. 5 Example of a node hosting multiple switching capabilities: Optical Transport Hierarchy (G.709), High Order SDH and Layer 2 SC (e.g., Ethernet)

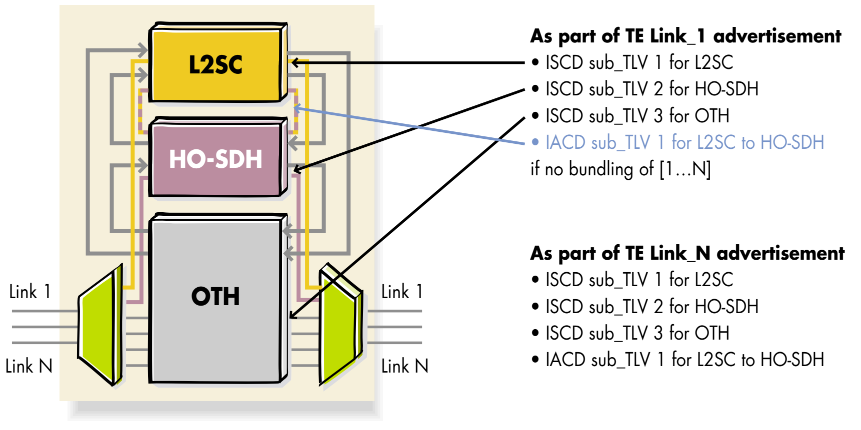

Enhanced theory of operations

Forwarding Adjacency usage

As the MRN approach is applicable within a given routing area, or a sub-domain within this area, with or without TE Link State Advertisement (LSA) filtering at its edges, the key issue is to efficiently attract traffic over existing FA-LSPs. But also, making sure that, when an incoming Resource reSerVation Protocol - Traffic Engineering (RSVP-TE) path message arrives at one of its edges, the right provisioning decision is taken (i.e., trigger new FA-LSP(s) - and at which SC - or re-use existing FA-LSPs). Note that in the present context, the situation is different from the one described in [5], since LSP region boundaries are defined between each node and not only at the edges of the network that defines the LSP region. Therefore, using MRN, the decision to trigger an FA-LSP should be deterministic (but not mandatory). Enhanced usage of FAs enables optimization of pre-provisioned or pre-computed virtual topologies but also dynamic triggering of resources at various SCs.

TE and FA attributes inheritance

The (OSPF-)TE metric alone (see [5]), in addition to the maximum LSP bandwidth and unreserved bandwidth, does not provide sufficient information to compute the best path between edges of the same region. This suggests that the inheritance mechanism of the TE metric for the FA link as defined in [4] has to be refined. The best example is a packet PSC-1 LSP nested into a PSC-2 LSP that lies over an LSC region. The TE metric of the FA link must not only take into account the packet boundary interface properties and TE attributes such as delay or bit-rate, but also, for instance, the distance over the region that the LSP will have to travel. For instance, the TE metric for the LSC LSP may be defined as a combination of the bit-rate and the distance, classically the bit-rate times the distance with some weighting factor(s). From this perspective, the main issue is that the joined path TE metric would not simultaneously tackle both packet and optical specifics.

This suggests the introduction of a more flexible TE metric definition, allowing a TE metric per ISCD (i.e. when multiple ISCDs attributes are associated to the corresponding TE link). This is because adjusting the TE metric of the FA link to (TE metric of the FA-LSP path - 1) is a valid approach between LSPs over the same region class (PSC-1, PSC-2, ... , PSC-N, for instance) but not necessarily between the PSC and the LSC region. The other TE attributes that need specific processing during inheritance are the Shared Risk Link Groups (SRLG), Resource Class, and Protection.

Recovery and FA abstraction

In MRN environments, the inheritance of protection and restoration TE link attributes must also be considered. For instance, consider a 1:1 end-to-end LSP recovery scheme; two FA-LSPs, a primary and a secondary, may be set up to form a single FA and to enhance its availability. The primary LSP is used to carry the normal traffic. Once a failure occurs affecting the primary LSP, the normal traffic is carried over the secondary LSP. From a routing perspective, there is no topological change during the recovery switching operation. Therefore, the two LSPs should be part of a single TE link advertisement, with the link protection type set to 1:1, and thus be processed by an upper layer as a span protected link.

Therefore, abstraction and summarization must be performed when advertising FA-LSPs as TE links (to an upper layer) but using the Link Protection Type flags and applying simple attribute inheritance might not be sufficient to distinguish different recovery schemes.

Protocol extensions

This section describes the extensions required to institute GMPLS routing and signaling to control MRN environments (see [8] for more details).

Interface Adaptation Capability Descriptor (routing)

In an MRN, some nodes, under the control of a single GMPLS instance, may combine multiple ISCs such as PSC and TDM or PSC and LSC or even LSC and WaveBand Switching Capability (WBSC). These nodes, hosting multiple switching capabilities, are required to hold and advertise resource information on link states and topology. They also may have to consider certain portions of their internal resources allowing termination of hierarchical LSPs, since circuit switch capable units such as TDMs and LSCs require rigid resources.

For example, an L2SC+TDM switching capable node can deliver connectivity for TDM LSPs but can never terminate the TDM LSP if there is no unused adaptation capability left between the L2SC and the TDM layers. Therefore, the advertisement of the so-called adaptation capability to terminate LSPs provides the critical information to be taken into account when performing multi-region path computation. This concept enables a local node to discriminate from remote nodes (and thus allows their selection during path computation) with respect to their adaptation capability e.g., to terminate TDM LSP at the L2SC level. This introduces the idea of discriminating the (internal) adaptation capability from the (interface) switching capability by considering an Interface Adaptation Capability Descriptor (IACD). This descriptor represents the link adaptation capability information for a given link described by its ISCDs.

By introducing such an additional descriptor, by analogy to the existing ISCD sub-TLV, the local GMPLS control plane can swiftly search which nodes can terminate a certain encoding type of LSP and successfully establish an LSP tunnel between two PSCs. Finally, integrated devices will not duplicate switching capacity at each SC and will not provide full capacity for interworking between the SC. The IACD is thus an enabler for GMPLS applications in those integrated situations.

Deterministic multi-layer signaling

Considering that path computation can take into account the richness of TE information regarding the SC available on node interfaces, it provides a means to deterministically indicate, during LSP signaling, the SC to be used on each link and where it is anticipated. Limiting extensions to the existing GMPLS RSVP signaling protocol, the MRN solution is expected to provide this strict indication of an SC. This functionality can be obtained by defining a new sub-object of the existing eXclude Route Object (XRO). This extension solves the ambiguous choice of SCs that are potentially used along a given path and makes it possible to optimize resource usage on a multi-layer basis.

Conclusion

Based on GMPLS and its intrinsic unified properties, the MRN approach contributes to satisfying major requirements for network evolution. By capitalizing on existing GMPLS protocol technology and fitting very common carrier network architectures, it takes only limited operational effort to build an integrated solution. MRN allows simplification, optimization of network usage, and ultimately reduced OPEX. Limited extensions of GMPLS protocols are sufficient to provide vertical control plane integration. The proposed MRN extensions do not introduce any interoperability issues, since they are backward compatible with the existing GMPLS installed base. These extensions can even be completely transparent, as the MRN architecture is complementary to the horizontal integration required for network partitioning purposes. Moreover, hitless upgrades are possible to gradually increase the number of switching capabilities that a single control plane instance can process. Finally, as carriers seek convergence, they should naturally go for integrated solutions based on standards, and in so doing solve interworking issues within their multi-technology networks.

Therefore, the converging evolution towards a unified control plane approach will deliver a ubiquitous set of tools for the control (including the engineering and the automated provisioning) of IP/MPLS packet, Ethernet, and circuit service. In turn, this will not only save costs related to daily and fastidious manual operation, planning, and configuration, but also deliver a long-term investment in inter-operable and upgradable control systems that carriers need to incorporate new switching layers as part of their network infrastructure.

Acknowledgment

The authors would like to thank authors and contributors of the Internet Draft “Generalized MPLS Architecture for Multi-Region Networks” as well as Bart Rousseau for the thorough review of this paper.

Glossary

AS	Autonomous Systems
DCN	Data Communication Network
FA	Forwarding Adjacency
GMPLS	Generalized Multi-Protocol Label Switching
IACD	Interface Adaptation Capability Descriptor
IETF	Internet Engineering Task Force
ITU-T	International Telecommunications Union - Telecommunication Standardization Sector
L2SC	Layer 2 Switching Capabilities
LSA	Link State Advertisement
LSP	Label Switched Path
LSR	Label Switching Routers
MPLS	Multi-Protocol Label Switching
MRN	Multi Region Networks
OPEX	Operating Expenses
OSPF	Open Shortest Path First
PSC	Packet Switching Capabilities
PSN	Packet Switched Networks
RSVP	Resource reSerVation Protocol
SC	Switching Capabilities
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
TDM-SC	Time Division Multiplexing Switching Capabilities
TE	Traffic Engineering
UNI	User Network Interface
VPNs	Virtual Private Network
XRO	Exclude Route Object

References

[1] "Generalized MPLS Architecture for Multi-Region Networks," K.Shiomoto, D.Papadimitriou, et al., Internet Draft (work in progress), draft-shiomoto-ccamp-gmpls-mrn-reqs-00.txt, October 2004.

[2] "Framework for Inter-domain MPLS Traffic Engineering," A.Farrel et al., Internet draft (work in progress), draft-ietf-ccamp-inter-domain-framework-00.txt, August 2004.

[3] "Demystifying GMPLS - A Technical Perspective," D.Papadimitriou and B.Rousseau, Alcatel Technical White Paper, No 02932, Ref. 3GQ 10001 0009 TQZZA Ed 01, December 2003.

[4] "LSP Hierarchy with Generalized MPLS TE," K.Kompella et al., Internet Draft (work in progress), draft-ietf-mpls-lsp-hierarchy-08.txt, September 2002.

[5] "Traffic Engineering (TE) Extensions to OSPF Version 2," D. Katz, et al., IETFRFC 3630, September 2003.

[6] "Link Bundling in MPLS Traffic Engineering," K.Kompella et al., Internet Draft (work in progress), draft-ietf-mpls-bundle-05.txt, October 2004.

[7] ''Generalized MPLS (GMPLS) RSVP-TE Signaling in support of Layer-2 Label Switched Paths (L2 LSPs)," D.Papadimitriou et al., Internet Draft (Work in progress), draft-papadimitriou-ccamlp-gmpls-l2sc-lsp-03.txt, October 2004.

[8] "Generalized Multi-Protocol Label Switching (GMPLS) Protocol Extensions for Multi-Region Networks (MRN)," D.Papadimitriou, K.Shiomoto, et al., Internet Draft (work in progress), draft-papadimitriou-ccamp-gmpls-mrn-extensions-00.txt, October 2004.

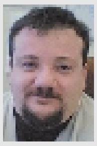

Emmanuel Dotaro (Emmanuel.Dotaro@alcatel.fr) (Emmanuel.Dotaro@alcatel.fr) received an M.S. degree in Computer Science from the University of Versailles, France in 1996. He spent three years spent at Institut National des Télécommunications Performance Evaluation lab while holding a teaching position at the University of Versailles. In 1999, he joined the Alcatel Research and Innovation lab at Marcoussis, France. He is now deputy project leader of the Packet Transport Infrastructure project in charge of the control and management plane activities His current research interests are network architecture and design, performance and dimensioning with a specific focus on GMPLS protocol suite, including standard bodies activities.

Martin Vigoureux (Martin.Vigoureux@alcatel.fr) received an M Sc. degree in high frequency electronics from the University Pierre et Marie Curie, France, in 2000. He then joined the Transport Network Systems and Architectures unit of Alcatel Research and Innovation where he is now study leader for software development of a flow based GMPLS network simulator. His current research interests are focused on GMPLS applicability to integrated systems architectures and advanced network environments.

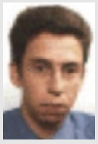

Dimitri Papadimitriou (dimitri.papadimitriou@alcatel.be) received an M.S. degree in physics from the Université Libre de Bruxelles and an M.S. degree in Computer Science from the Université de Liège. After one year of collaboration on Network Protocols R&D at the Research Center of Namur, he worked for 3 years in the design and evaluation of packet (IP/MPLS), Ethernet, and circuit (SDH/optical) transport networks as well as their compliance to IETF/ITU standards. In 2000, he joined Alcatel as a Network Analyst in the Network Technology and Analysis team of the Network Strategy Group. In 2004, he joined the Packet-Transport Interworking (PTI) team of Research and Innovation (R&I). His current areas of interest are focused on wavelength switched optical networks and their distributed control planes. He is also involved in standardization activities at IETF (in the Sub-IP Area) and OIF where he edited and co-authored numerous contributions. He has authored numerous standardization papers on Generalized MPLS (GMPLS)-based control planes and various technical papers on the aspects of outing and recovery techniques in wavelength switched optical networks.