A Service Plane over the PCE Architecture for Automatic Multidomain Connection-Oriented Services

Richard Douville, Alcatel-Lucent Bell Labs
Jean-Louis Le Roux, Orange Labs France Telecom
Jean-Louis Rougier, TELECOM ParisTech (ENST)
Stefano Secci, TELECOM ParisTech (ENST), Politecnico di Milano

ABSTRACT

In this article we concentrate on the automated provisioning of inter-AS services based on GMPLS-TE technology. We consider a provider alliance where TE connections are established between the members of the alliance. We propose an efficient and economically feasible architecture for the automatic provisioning of inter-AS GMPLS-TE, based on the introduction of a multidomain service plane coupled with the PCE-based architecture.

INTRODUCTION

Carriers usually deploy traffic engineering (TE) technologies such as generalized multi-protocol label switching (GMPLS)-TE to offer value-added services. Nevertheless, these technologies are restricted to intra-domain networks, narrowing the scope of potential service offers. Currently, there is a clear demand to extend guaranteed service offers beyond domain boundaries.

To support inter-domain services, one must rely on inter-provider quality of service (QoS) mechanisms, usually referred to as inter-autonomous system (inter-AS) QoS mechanisms. The Internet Engineering Task Force (IETF) has defined an extension to the GMPLS-TE technology, called inter-AS GMPLS-TE, which enables the establishment of inter-provider, explicitly routed, connections with stringent QoS and availability constraints.

In this article, we propose to enrich the current inter-AS GMPLS-TE technology to enable automatic provisioning of inter-domain TE services. We define the basis of a global architecture that enables an automatic provisioning of multidomain TE network services, defined within the French "Agence Nationale pour la Recherche — Approche Combinée de Technologies Réseaux Inter-domaine sous Contraintes Economiques" (ANR ACTRICE) project.

In the following section, we describe the inter-AS GMPLS-TE building blocks involved in our solution, and we indicate their current limitations in achieving our goal. We then propose to introduce a service plane coupled with the path computation element (PCE)-based architecture, and we show how this service plane allows for negotiations of provider service elements through successive phases of selection, instantiation, and activation. Finally, we define the required extensions to existing protocols and management elements, and we illustrate how they can be applied to offer automatic provisioning of inter-AS TE services.

INTER-AS GMPLS-TE TECHNOLOGY

Within provider boundaries, the GMPLS technology allows establishing connections, called label-switched paths (LSPs), based on any transport network solution (circuit- or packet-based).

The GMPLS-TE extensions add the possibility of routing an LSP explicitly, taking TE constraints into account; for instance, verifying resource availability, switching capability, and end-to-end or subpath protection possibility. As already mentioned, the IETF has extended the GMPLS-TE technology to support the configuration of inter-AS LSPs, meeting the requirements in [1].

INTER-AS PATH SIGNALING

A signaling protocol called Resource Reservation Protocol with Traffic Engineering Extensions (RSVP-TE) is used to establish GMPLS-TE LSPs. As explained in [2], an inter-AS LSP can be signaled in three ways:

LSP Nesting: In this mode a local high-level

0163-6804/08/$25.00 © 2008 IEEE

intra-domain LSP is used between domain border routers to transport many inter-domain LSPs sharing a common intra-domain subpath. For purely MPLS backbones, this corresponds to encapsulating an inter-domain tunnel into an intra-domain tunnel. For optical networks, this corresponds to grooming incoming inter-domain MPLS/GMPLS LSPs into lower-level intra-domain LSPs with coarser optical switching capabilities (fiber, waveband, or wavelength).

Contiguous LSP: In this mode, a single end-to-end LSP is signaled across the domains. The head-end router is connected to the tail-end router via a single signaling session. This means that single session and LSP identifiers are used across the inter-AS path. Hence, the reconfiguration of such an LSP is controlled by the head-end router, and intermediate domains should not modify their local subpath.

LSP Stitching: In this mode, intra-domain LSPs are signaled and then stitched at the boundaries to form a single inter-domain LSP perceived in the data-plane. From the control-plane standpoint, the local intra-domain LSPs are signaled separately, and every LSP has different source and destination (ingress and egress domain border routers). This signaling method would be applied particularly to the case in which some switching capabilities are not compatible with the nesting method. For instance, a lambda-LSP cannot be nested in another lambda-LSP.

Inter-AS Path Computation

An LSP is to be signaled over a pre-computed path. A head-end router has full topology visibility within its domain and hence, can compute alone an end-to-end intra-domain path but cannot compute alone an end-to-end inter-domain path. Two methods can be adopted for the inter-AS path computation:

- A per-domain path computation method, in which the source or ingress router determines the next domain and the ingress router in this next domain, and computes the corresponding subpath. Then, the path computation is moved to the ingress router of the next domain and so on, up to the tail-end router. This simple method does not allow computing a shortest inter-domain path and can lead to several crankbacks that might affect the stability of the control plane.
- An inter-domain path computation method that takes as input the AS chain (i.e., the succession of ASs to be used) and relies on computation servers present in each AS, called PCEs, to collaboratively compute an inter-AS shortest path along the given AS chain.

PCE-Based Architecture

The PCE architecture [3] consists in delocalizing the path computation from the head-end router to a PCE that computes paths on behalf of the head-end router. PCEs can collaborate together to compute constrained inter-AS paths, without being required to share any TE information with each other; thus, solving the topology visibility issue. PCEs are particularly useful when end-to-end constraints for protection or path diversity must be taken into account.

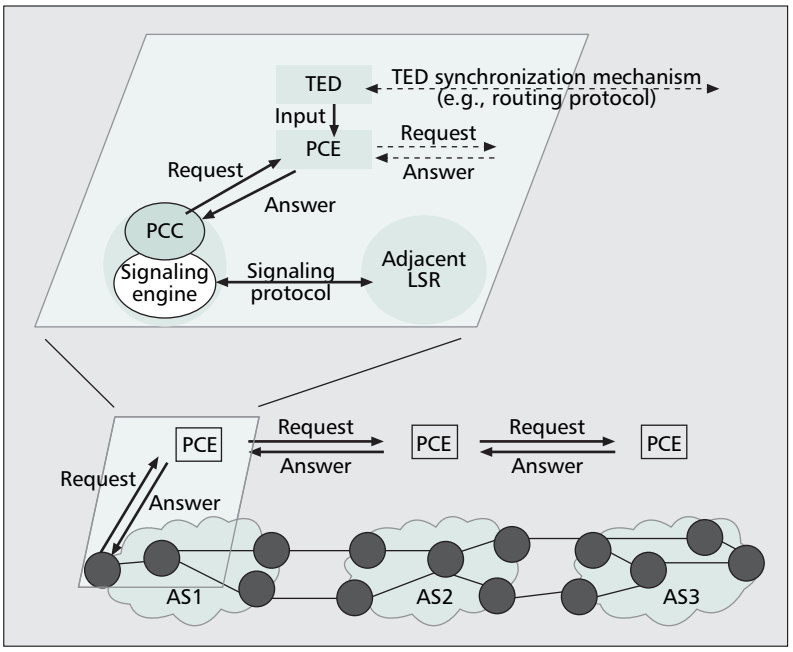

Figure 1. *Communications relating to PCE.*

As depicted in Fig. 1, at least one PCE is required per domain. A PCE serves requests sent by path computation clients (PCCs), for example, routers or switches, using information in a local TE database (TED). A PCE can query the PCEs of other domains to perform this computation, acting in turn as a PCC. A PCE Communication Protocol (PCEP) was defined to relay these requests and answer messages [4]. PCCs can dynamically discover external PCEs through extensions of existing routing protocols, meeting the requirements in [5].

As already mentioned, the PCE-based inter-AS path computation can be performed after the AS chain for the destination is known. Efficient distributed algorithms are required for the end-to-end inter-AS path computation, taking into account a pre-computed AS chain. A procedure called backward recursive path computation (BRPC) is the one that seems best, meeting the requirements of both operators and suppliers in terms of complexity and network information hiding [6]. It consists of computing recursively, at each PCE of the AS chain, an inverse tree of constrained shortest paths, with one branch for each ingress border router (and toward the destination), starting from the destination AS. Each path might be a loose path containing only the tail router, the border routers, and the cost of the corresponding shortest path. The tree is sent back to the previous AS, which does the same, and so forth back to the source AS.

Contributions

As we have mentioned, the IETF developed solutions for inter-AS LSP set up. However, we can outline the following open issues:

- For the PCE-based architecture, the standardization does not indicate how the input AS chain is calculated.

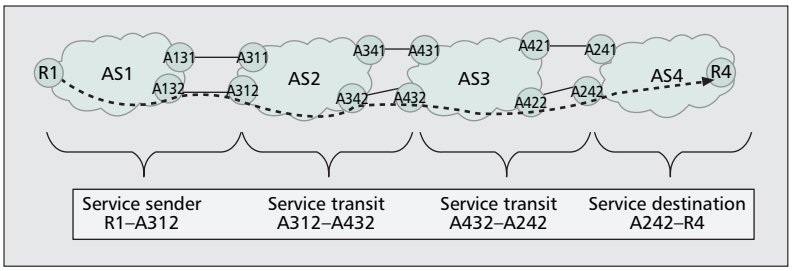

Figure 2. *Service elements.*

- The set up of an inter-AS tunnel is subject to strong business, security, and confidentiality aspects. Hence, the set up of an inter-AS must be performed only between trusted entities and requires a preliminary service instantiation and activation, to ensure billing and to manage routing and signaling requests at domain boundaries. Such procedures are beyond the scope of the IETF and are still not defined.

It is worth noting briefly how the AS chain selection is currently performed for connectionless services, via the inter-domain Border Gateway protocol (BGP). A cascade of criteria is employed to compare alternative paths. The first criterion is the "local preference," through which local policies, guided mainly by economic issues, can be applied: for example, a peering link (i.e., free transit) is preferred to a transit link (transit fees). The subsequent criteria incorporate purely operational network issues: smaller AS hop count, closer egress point, and so on. On the basis of all the criteria, the best path is selected, and it is the single advertised one. Hence, gathering one path per destination AS from BGP is not advisable because paths are chosen regardless of QoS, availability, and reliability constraints, and because the BGP decision process is too simple to model the complex AS relationships that exist today or that will emerge with the generalization of extended peerings (including valued-added services). However, the economical distinction of paths granted by the local preference criterion should not be lost, but on the contrary, enhanced.

Within the ACTRICE project, we have worked on solutions to overcome these issues. In particular, we have studied how the deployment of a multidomain service plane can support an automatic provisioning of multidomain LSPs and define a favorable business model. This service plane is to be adopted by an alliance of providers willing to collaborate for the delivery of multidomain TE services and willing to decrease the overwhelming operational efforts currently related to such a service (a chain of bilateral ad hoc agreements). Within the provider alliance, imprecise TE information is to be transmitted by means of service elements through which each provider advertises its transit capabilities and policies.

In this article, we focus on the service establishment within a provider alliance. It consists of a first phase of selection and validation of service elements to compose the AS chain and of a second phase of service configuration (computation and signaling). In the following, we characterize the required inter-AS service plane, and we detail inter-AS LSP provisioning steps.

The Notion of Inter-AS Network Service

As explained previously, we can benefit from the mechanisms defined in the IETF for the configuration of inter-AS GMPLS-TE LSPs. However, this should rely on a preliminary agreement between providers on a common service plane and should require the application of important TE and security policies at the provider boundaries. We tackle these fundamental service and policy aspects that are beyond scope of the IETF.

An inter-AS TE service is composed of one or more inter-AS LSPs between a head router in the source provider and a tail router in the destination provider, crossing a chain of transit providers. An inter-AS LSP can be unidirectional for content distribution or bidirectional for interactive services.

We characterize an inter-AS TE service by the following parameters:

- Address of the head and tail routers
- Source and destination AS numbers
- AS chain
- Direction: unidirectional or bidirectional
- Bandwidth
- Service level specifications (SLSs), containing performance parameters and cost
- Protection level: unprotected, global protection, local protection
- Reoptimization: enabled or disabled

An inter-AS TE service is the result of a composition of service elements offered by each operator. We introduce three service element categories:

- The *Edge Sender*, which assures the routing from the head router of the sender AS toward an ingress router of a neighboring AS
- The *Edge Destination*, which assures the routing from an ingress router of the destination AS toward the tail router
- The *Transit*, which assures the routing from an ingress router of the AS toward an ingress router of the next AS

Figure 2 illustrates an example of inter-AS TE service, unidirectional and unprotected, between the R1 router of the AS1 and the R2 router of the AS4 across AS2 and AS3. It is the composition of four service elements, two transit, one sender, and one destination. Each element indicates explicit edges as incoming and outgoing border routers.

Service Elements

To enable composition of service elements, every operator advertises to the other members of the alliance its service elements. It is worth mentioning that the IP Sphere Forum is currently specifying a framework that, among other features, can enable reliable multidomain advertisement of service data via Simple Object Access Protocol/ Extensible Mark-up Language (SOAP/XML)-based Web services.

We characterize a service element by the following parameters:

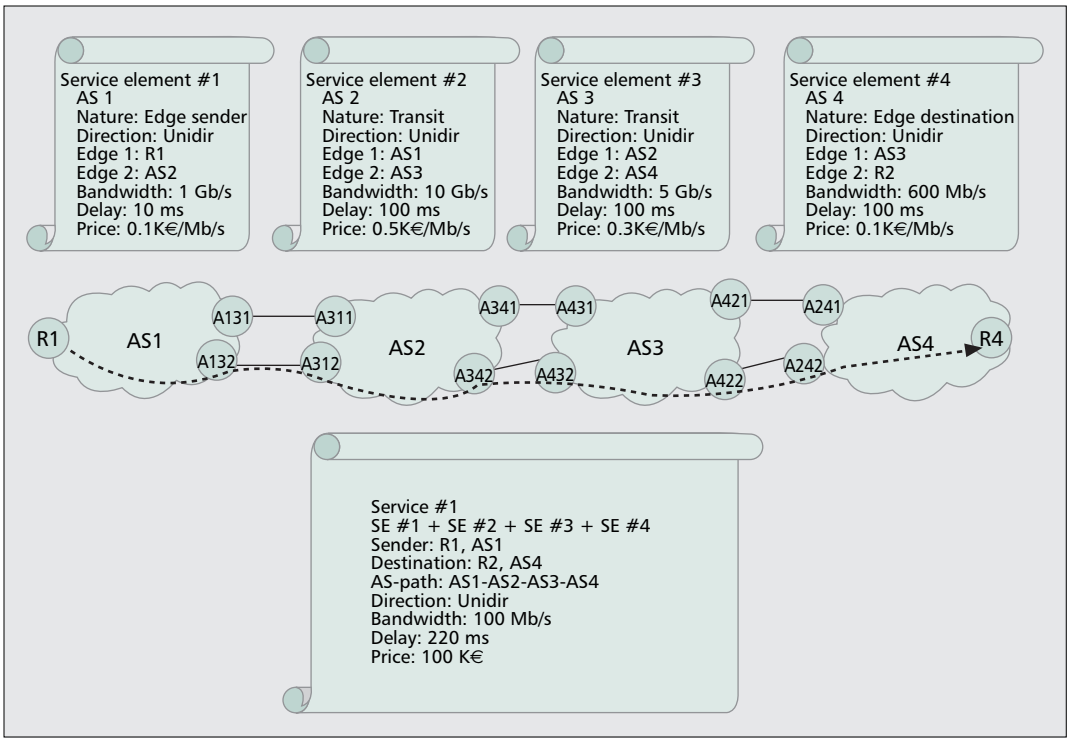

Figure 3. *Service elements composition.*

It is worth mentioning that the IP Sphere Forum is currently specifying a framework that, among other features, can enable reliable multidomain advertisement of service data via SOAP/XML-based Web services.

- Local AS number
- Nature of the service: Sender, Destination, or Transit
- Direction: unidirectional or bidirectional
- Ingress and egress edges
- Upper bounds of performance parameters
- Maximum bandwidth that can be reserved for a given session
- Protection level (unprotected, global, local)
- Transit cost, per bandwidth unit (the Mb/s) and/or per duration unit (the month)

An edge can be identified explicitly by the address of a router or a group of routers or implicitly by the AS number of the neighbor. Obviously, in the case of Edge Sender and Edge Destination service elements, one of the two edges would represent the head and the tail router (or group of routers), respectively.

In Fig. 3, a set of service elements that contributes to the composition of an end-to-end network service is displayed. In this example, each element indicates implicit edges as incoming and outgoing ASs. The choice of these four service elements is guided by the compliance with the request parameters and the minimization of the service cost. Coherently, the AS chain is built as a list of the corresponding ASs.

Within the provider alliance, business relationships are defined by the policy of the ASs in advertising service elements and in admitting requests. Existing transit or peering agreements for best effort IP routing may intervene in the alliance formation, in the service element definition, or they may be disregarded for connection-oriented services.

It is worth noting that the described framework is not restricted to pure MPLS networks, but it encompasses GMPLS-based optical network domains. As previously mentioned, LSP nesting and stitching methods can rely locally on GMPLS LSPs. The service plane signaling is agnostic on the used switching granularities, and the service parameters of the service element are not restricted for MPLS tunnels.

Requirements

The set up of an inter-AS LSP requires the following subsequent steps:

1 The discovery of the service elements offered by each provider.

2 The composition of the service elements to form an end-to-end inter-AS network service, namely, the computation of the cheapest constrained AS chain. This computation must handle:
- –The fact that transit service elements contain directional policies, from an ingress edge toward an egress edge. Thus, the AS graph is weighted by directional metrics and not simple metrics.
- –The presence of computation servers that would support pre-computation, which could allow decreasing the post-request complexity (i.e., the complexity of the sub-algorithm to execute after the request arrives).

3 The instantiation of the composed service. This should include a:
- –Connection admission control at the service plane to verify the availability of the service elements.
- –Confirmation of the SLS.

4 The activation of the service. This should include, in the following order:
- –Configuration of filters on the policy managers of each domain to validate inter-AS PCEP and RSVP-TE messages in function of the instantiated service.

The network layer encompasses the ASs with their core and border routers, and their PCEs. The network layer is guided by the management layer, where a network management server and a policy manager are required.

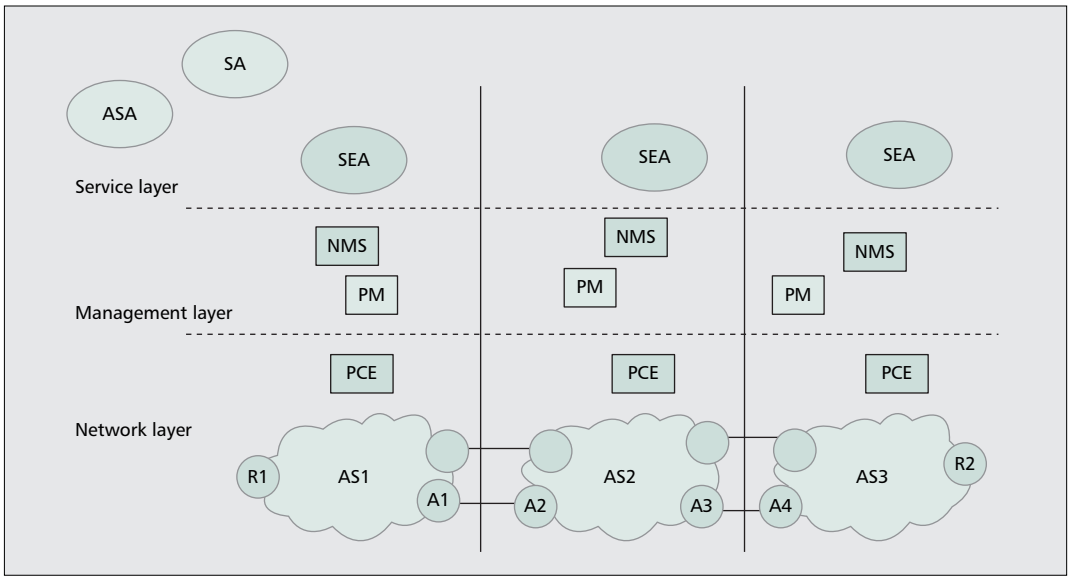

Figure 4. *Inter-AS multilayer service architecture.*

–Configuration of the LSP on the head router.

5 The inter-AS path computation along the composed AS chain, via the PCE-based architecture. Inter-AS PCEP messages must be filtered with:
–Translation of TE parameters (priority, class of service).
–Filtering of topological information to assure confidentiality.
–Rejection, if not compliant with the instantiated service (SLS, etc.).

6 The LSP signaling via RSVP-TE. Inter-AS RSVP-TE messages must be filtered similarly to PCEP messages as indicated previously.

7 The service maintenance. This should include:
–Accounting and measurement of the end-to-end and local performances.
–Fault detection, localization, and reporting.

Architecture for Automatic Provisioning of Inter-AS GMPLS Services

To reduce the operational costs and minimize service set up and restoration times, all the steps enumerated above must be automated. The set of required mechanisms forms the architecture for automatic provisioning of inter-AS GMPLS services.

The IETF defined mechanisms for steps 5 and 6. Nonetheless, extensions are required to support the transport of a service identifier in RSVP-TE and PCEP messages so as to apply the filters.

Steps 1 through 4 and step 7 require the introduction of a new plane called a service plane, which enables the exchange of service information among providers. This plane may rely on an adaptation of the IP Sphere Forum service plane, implemented via a SOAP/XML-based Web service platform.

Functional Elements

Figure 4 illustrates the elements required to automate the set up of an inter-AS GMPLS service. We can distinguish three layers.

The network layer encompasses the ASs with their core and border routers (ASBR), and their PCEs. The network layer is guided by the management layer, where a network management server (NMS) and a policy manager (PM) are required as explained hereafter.

The management layer is in turn guided by the novel layer that we introduce, the inter-AS service layer. This layer is composed of per-domain agents, called service element agents (SEAs); of end-to-end agents, called service agents (SAs); and of agents responsible for the service element composition, called AS selection agents (ASAs).

At the network layer, the ASBRs are linked by an IP/GMPLS link and interact via the RSVP-TE protocol. A router communicates with its local PCE via the PCEP protocol. The PCEs also communicate via the PCEP protocol for the end-to-end path computation.

Between the network and the management layers, the ASBR communicates with its NMS via a network management protocol (e.g., Simple Network Management Protocol [SNMP], XMLConf, Telnet command line interface [CLI]) to set up LSPs and to raise information about LSP status. An ASBR similarly communicates with its PM to perform the admission control of inter-AS RSVP-TE and PCEP messages (this may rely e.g., on a policy protocol such as common open policy service [COPS], Diameter, or a yet to be defined SOAP/XML solution).

Between the service and the management layers, the NMS receives inter-AS LSP set up, change, or deletion commands from the SEA and raises LSP status information to the SEA. The PM also communicates with its SEA to acquire filtering policies corresponding to the services activated at the service layer. The policies are to be indexed by a service ID. PM/SEA and NMS/SEA communications may be based on a SOAP/XML protocol.

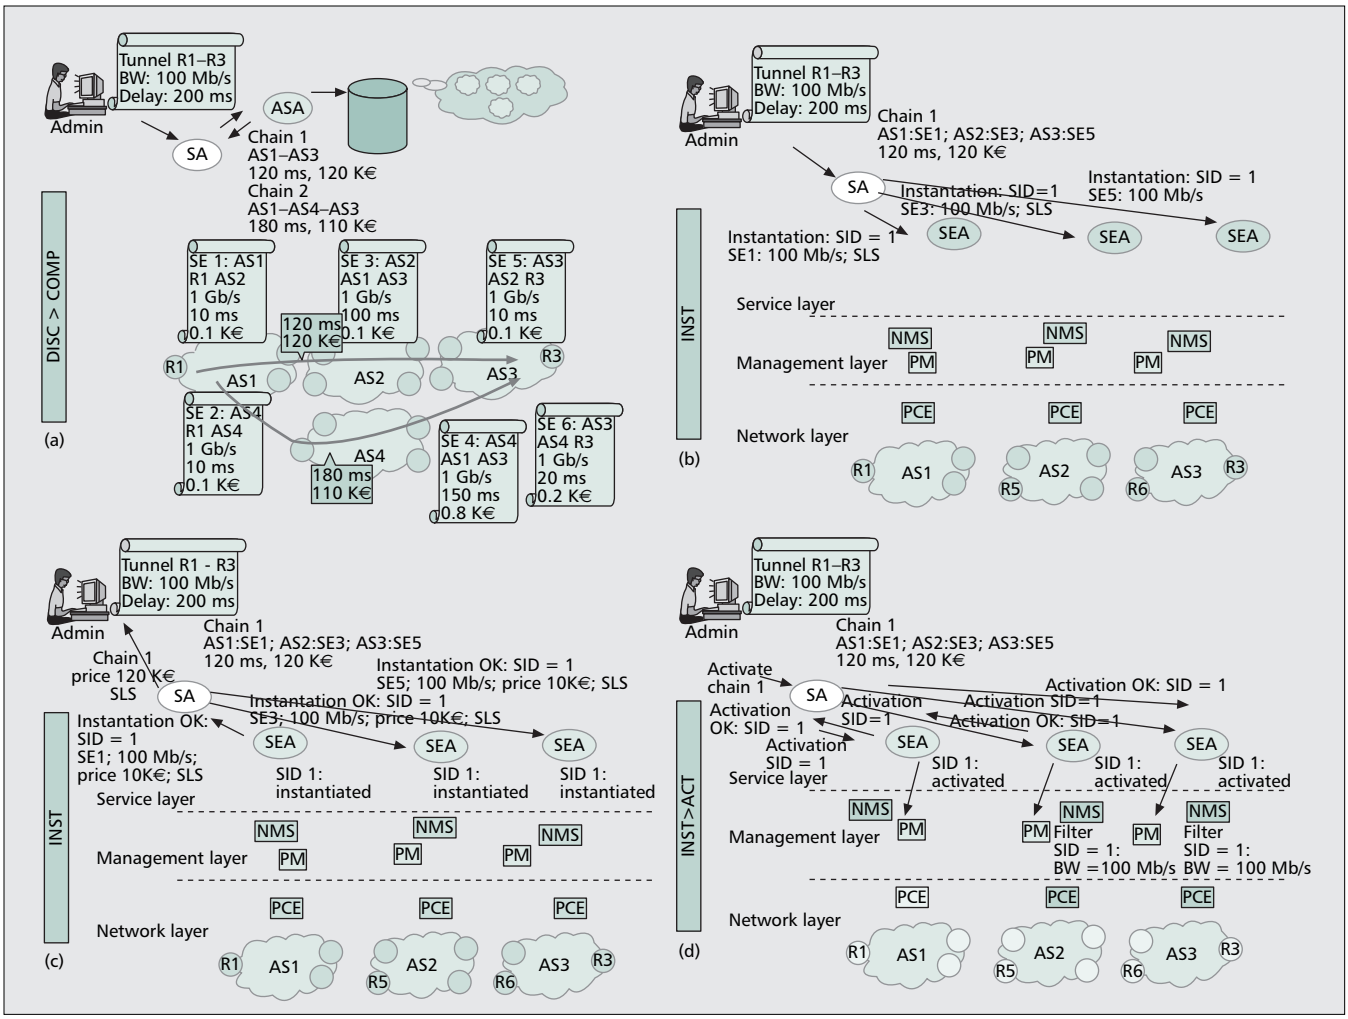

■ **Figure 5.** *Discovery, instantiation, and activation at the service layer.*

It is worth noting that there is no inter-AS communication at the management layer.

At the service layer, the SA is responsible for the construction of the end-to-end service. It receives inter-AS LSP set up, change, and deletion commands from an administrator. It queries the ASA to calculate an AS chain. Then, it communicates with the SEAs along the AS chain to instantiate and activate the service elements. SA/SEA and SA/ASA communications also can be based on SOAP/XML.

Functional Steps

We thus distinguish seven functional steps corresponding to the different phases of the lifecycle of an inter-AS GMPLS service at the service plane (Fig. 5) and at the management and network planes (Fig. 6).

Service Discovery (DISC) — This step consists of acquiring the inventory of all the service elements offered by the providers of the alliance.

Composition of Service Elements (COMP) — This step consists of determining the AS chain of the LSP (Fig. 5a). The administrator triggers the composition via the SA at the ASA, where constrained shortest path algorithms should be implemented. The ASA answers with one or more AS paths. Many diverse AS paths may be selected to increase the success of having at least one accepted. Indeed, a selected AS path can fail during the instantiation, the activation, the path computation, or the signaling phases.

The multi-constrained shortest path problem is known to be NP hard, and directional metrics and diversity constraints expand the complexity. By studying the current AS graph, the authors in [7] report that in a backbone AS graph (composed of ASs potentially interested in an alliance for automatic inter-AS service provisioning), well-known constrained shortest path algorithms, having at least $O(n^3)$ time complexity, would assume a complexity of at least $O(n^4)$ with directional metrics because these roughly multiply by a factor n the number of nodes in the graph to be employed. They conclude that breadth-first search algorithms with limited depth are advisable for this problem because they scale with directional metrics. Hence, they propose an ad hoc algorithm for a proactive quasi-optimal selection of diverse paths: proactive because it can improve the service acceptance in the following steps, without reacting to rejection (i.e., stateless computation is to be preferred to stateful computation); quasi-optimal because it often offers a 5 percent optimality gap, which is not

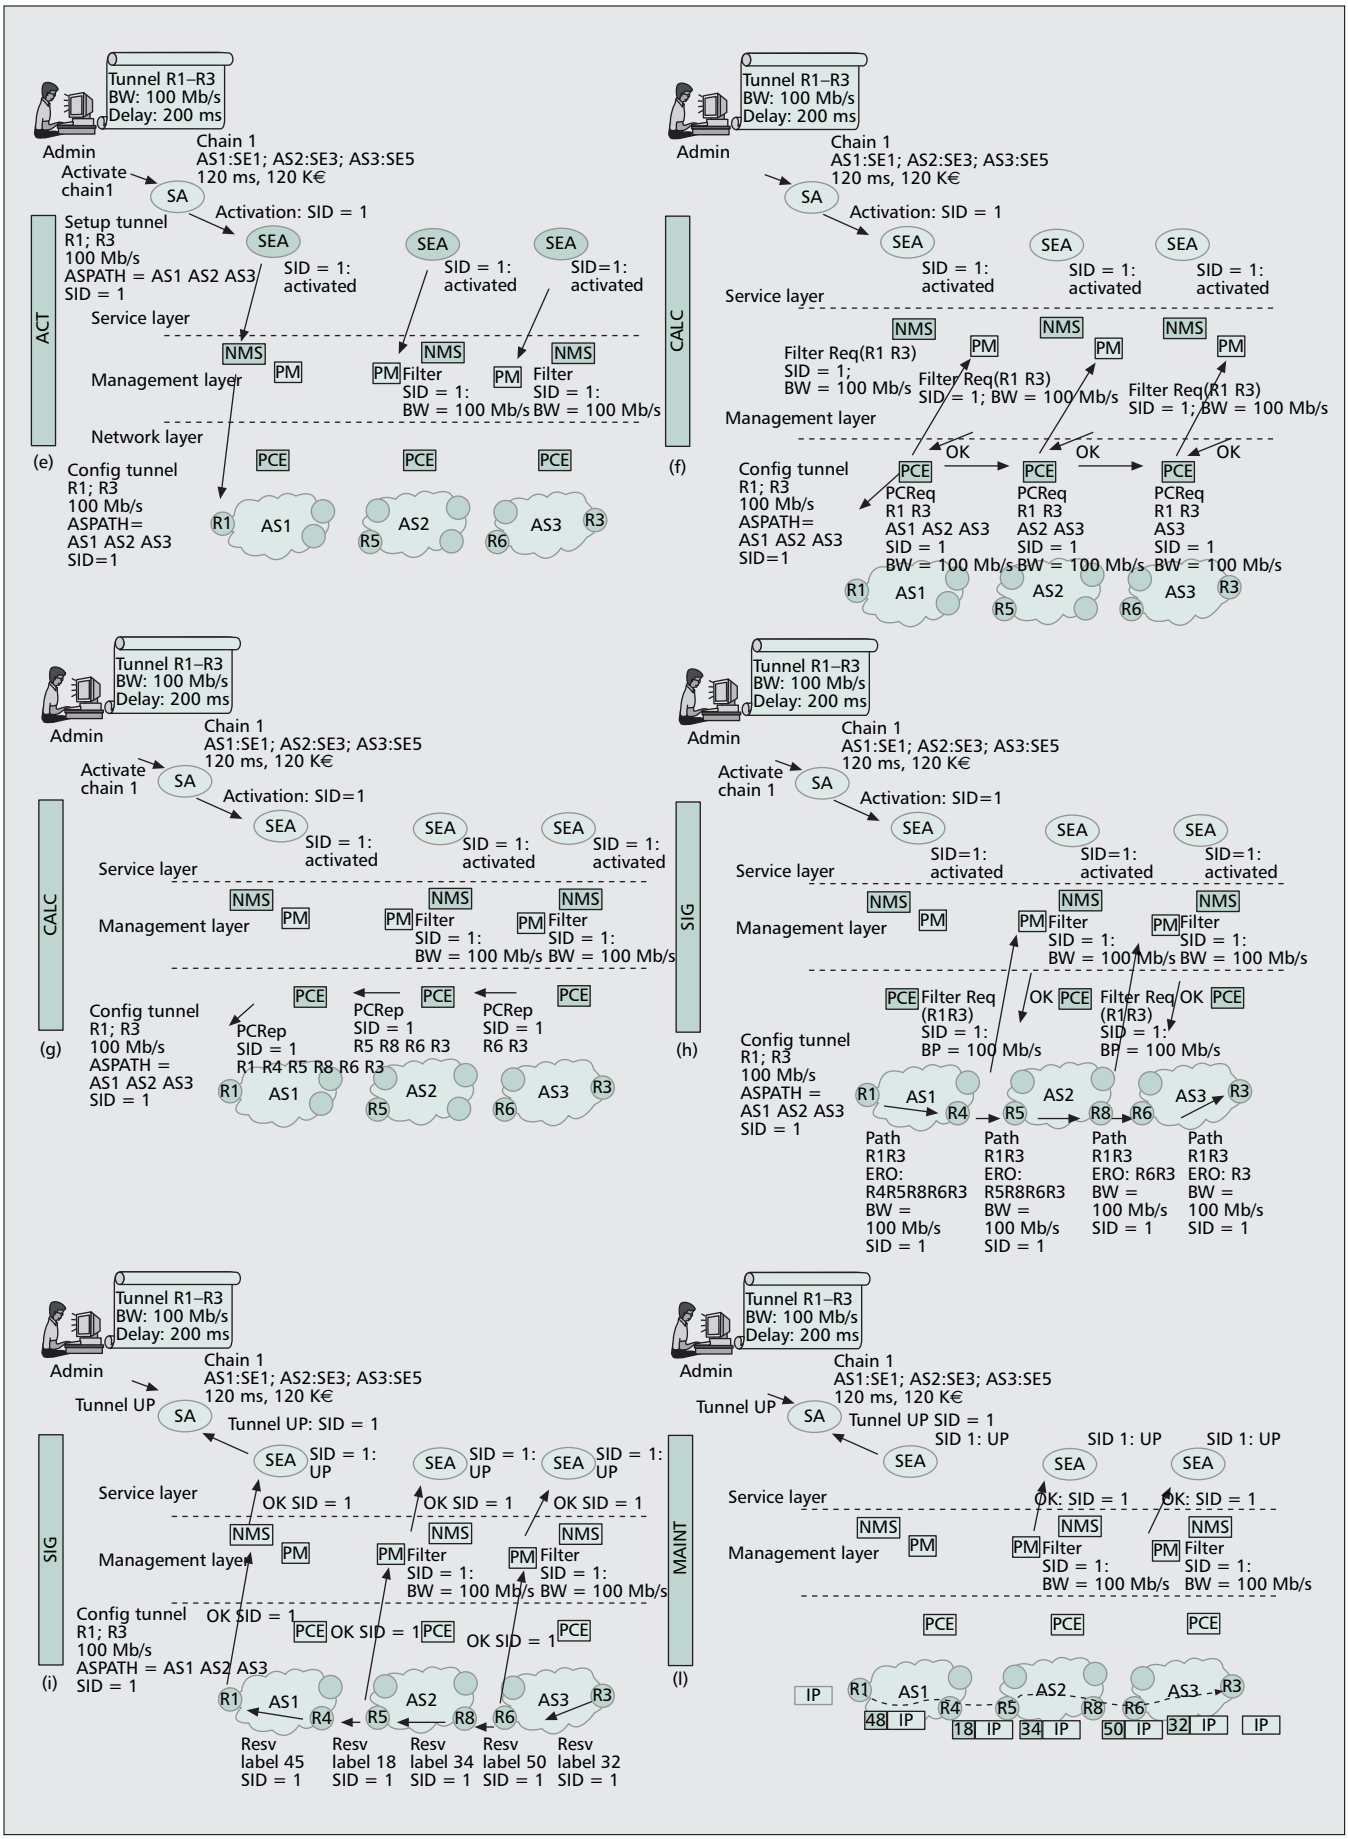

Figure 6. *Inter-AS computation and signaling at the management and network layers.*

poor because the price of the service element can be renegotiated during the instantiation. Two paths are considered diverse if at each AS-node they follow different directions; that is, if they rely on different service elements. Due to the possibility of pre-computing a part of the job independently of the request details, the algorithm proposed in [7] and its extension for the multipoint case [8] offer a $O(n^3)$ worst-case time complexity.

Service Instantiation (INST) — This action aims at verifying the availability of the service elements composing the AS chain and at agreeing upon the final SLS and cost (Fig. 5b–c). An `Instantiation` message is generated at the source SA and is sent to the SEAs of the involved ASs. The request contains a service Id (SID) that is produced to identify the service during all of its lifecycle. Then, in the case of availability, the SEAs send back an `Instantiation OK` message with the current SLSs (potentially changed), and current price; otherwise an instantiation not OK (NOK) message is sent. If an element is not available or if the SLS is not acceptable, the SA can test another AS chain.

Service Activation (ACT) — This step consists of triggering the service establishment (Fig. 5d, Fig. 6e). As a first action, the SA sends to all the SEAs an activation message with the SID. These SEAs send to the PMs a filtering policy associated to the SID, useful to filter inter-AS PCEP and RSVP-TE messages. If no error occurs, each SEA sends back an `Activation OK` message; otherwise an `Activation NOK` message. If all the responses are positive, the SA sends an `Activation` message to the source SEA, which then commands the LSP configuration to the local NMS with all the request details. If no error occurs, this SEA sends an `Activation OK` message to the SA. Then, the NMS configures the inter-AS LSP on the head router, passing the SID and the AS chain in addition to base TE parameters.

Path Calculation (CALC) — This step consists of computing the inter-AS path via the PCE-based architecture (Fig. 6f–g). Acting as PCC, the head router sends a path computation request (PCReq) message to its local PCE. This message is propagated along the PCEs of the AS chain up to the destination PCE, where the BRPC procedure can start. During the computation, a PCRep message is propagated backward toward the source PCE. A confidentiality key should be associated to this local information so as to retrieve the full intra-domain path during the signaling phase [9].

We introduce the following novelty in filtering the PCEP messages: by default an inter-AS PCEP message is to be rejected by a PCE for obvious reasons of security and confidentiality; it can be accepted only if it transports a SID corresponding to a service that has passed a preliminary activation at the service layer.

In [10], a PCEP extension was proposed to include a SID object. When a PCE receives a PCEP message, it transmits the request to its PM using a Filter Req message. The PM performs the following operations:

- It extracts the SID and the parameters of the request.
- It looks for a filter indexed by the SID; if there is not, a PCErr message is sent back to the source PCC.
- If a filter is found, its application entails the deletion of certain objects, the modification of others (e.g., the priority or the DiffServ class), or the rejection of the request if some parameters do not comply with the activated service.

Service Signaling (SIG) — This step consists of the final signaling of the inter-AS LSP (Fig. 6h-i). When the source PCE computes the final path to be employed, it sends a PCRep message to the source PCC containing an end-to-end path toward the destination router, as depicted in Fig. 6. This should be a loose path containing, for example, only the border routers to cross. Then, the signaling can proceed as explained previously, using this loose inter-AS path as an explicit route object (ERO), resolved locally via the confidentiality key.

The novelty we introduce is in filtering the inter-AS RSVP-TE messages. The RSVP-TE Path and Resv messages should be extended to transport the SID [2]. Employing the SID, the ingress router of each domain queries the local PM to perform the required filtering operations according to the instantiated service.

When an ASBR receives a Resv message, it sends an OK message to the PM with the SID. The PM then decrements the bandwidth allocated to the service (in Fig 6i, the LSP uses all the negotiated bandwidth, so there is no remaining bandwidth for the service).

Service Maintenance (MAINT) — After the inter-AS LSP is established (Fig. 6l), the events that could occur are the failure or the closing of the LSP. In case of failure, re-routing or re-provisioning operations should be executed. If a specific protection strategy was chosen at the corresponding service element, it should be implemented. Whether the failure happens on intra-domain links or routers, the recovery should not involve the service plane. If a failure on intra-domain equipment cannot be recovered or if a failure that occurs on inter-domain links cannot be recovered by rerouting the LSP on an alternative path between the two involved ASs, a status NOK message is sent to the service plane, and then the source SA is notified and should proceed with a new service request.

The novelty we introduce is in filtering the inter-AS RSVP-TE messages. The RSVP-TE Path and Resv messages should be extended to transport the SID. Employing the SID, the ingress router of each domain queries the local PM to perform the required filtering operations according to the instantiated service.

Dealing with Collateral Behaviors

It is evident that if a provider cannot guarantee the SLS, it should be penalized. Nonetheless, if a provider perturbs with unidentified inter-AS PCEP and RSVP-TE messages or advertises service elements that are revealed to be unavailable most of the time, the provider should be penalized.

Furthermore, it is worth noting that collateral business behaviors can appear within the alliance that can still be correct from an alliance standpoint. First, it is still possible to hide private bilateral agreements. Second, it is possible to prune or weight competitors' service elements.

The IETF has worked on extending existing protocols and architectures required to set up inter-domain connections. These extensions are referred to as inter-AS GMPLS-TE technology. However, some missing blocks are required to automate inter-AS services.

Third, an AS can avoid instantiating elements by blocking requests involving a specific AS. These behaviors can be detected locally, but this may not be sufficient. To automatically isolate them, some transaction statistics should be shared at the service plane. This data should contain the level at which the behavior is detected: at the service layer during the composition or instantiation, at the management layer during the message filtering, or at the network layer. By sharing this data, an SA can instruct its ASA with information on how to prune and weight some of the service elements. These transaction statistics can be collected by a shared service broker whose data is populated by the local SAs where they must report a proved incorrect behavior.

CONCLUSIONS

Currently, the GMPLS-TE technology is deployed mainly within provider boundaries to support real-time and interactive services. The extension of these services in the multidomain scope requires supporting inter-AS QoS guarantees between providers. The IETF has worked on extending existing protocols and architectures required to set up inter-domain connections. These extensions are referred to as inter-AS GMPLS-TE technology. However, some missing blocks are required to automate inter-AS services.

In this article, we first outlined these missing blocks and in particular, the importance of a service plane. In the context of a provider alliance, where TE connections are established between the members of the alliance, we defined the notion of inter-AS GMPLS-TE service as a composition of service elements. We then proposed a comprehensive architecture based on three planes: service, management, and network planes. We outlined the roles of each plane and showed how they can interact. We concluded by showing how inter-AS provisioning can be automated within this framework with an emphasis on service composition and activation.

REFERENCES

[1] R. Zhang and J.-P. Vasseur, "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements," RFC 4216, Nov. 2005.

[2] A. Farrel, A. Ayyangar, and J.-P. Vasseur, "Inter Domain MPLS and GMPLS Traffic Engineering — RSVP-TE Extensions," RFC 5151, Feb. 2008.

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

[4] J.-P. Vasseur and J.-L. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)," draft-ietf-pce-pcep-12, Mar. 2008.

[5] J.-L. Le Roux, "Requirements for Path Computation Element (PCE) Discovery," RFC 4674, Oct. 2006.

[6] J.-P. Vasseur *et al.*, "A Backward Recursive PCE-Based Computation (BRPC) Procedure to Compute Optimal Inter-Domain Traffic Engineering Label Switched Paths," draft-ietf-pce-bBpc-09, Apr. 2008.

[7] S. Secci, J.-L. Rougier, and A. Pattavina, "On the Selection of Optimal Diverse AS-Paths for Inter-Domain IP/GMPLS Tunnel Provisioning," *Proc. 4th IEEE IT-NEWS on QoS in Multiservice IP Networks*, Venice, Italy, Feb. 13–15, 2008.

[8] S. Secci, J.-L. Rougier, and A. Pattavina, "AS Tree Selection for Inter-Domain Multipoint MPLS Tunnels," *Proc. IEEE ICC '08*, Beijing, China, May 19–23, 2008.

[9] R. Bradford, J.-P. Vasseur, and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Key-Based Mechanism," draft-ietf-pce-path-key-02.txt, Feb. 2008.

[10] J. L. Le Roux, R. Jacob, and R. Douville, "Carrying a Contract Identifier in the PCE Communication Protocol (PCEP)," draft-leroux-pce-contract-id-01.txt, Mar. 2007.

BIOGRAPHIES

RICHARD DOUVILLE (richard.douville@alcatel-lucent.fr) received his M.Sc. degree in electrical engineering from the University of Versailles, France, in 2000. In 2000 he joined Alcatel Research Labs, where he has been involved in IP over optical internetworking issues, including network and system architectures, traffic engineering, resource dimensioning, and performance evaluation. His current research interests include control (GMPLS) and management architectures and solutions enabling the automation of end-to-end interdomain/layer service deployments.

JEAN-LOUIS LE ROUX (jeanlouis.leroux@orange-ftgroup.com) holds an engineering degree from TELECOM Bretagne. He is a senior MPLS architect at Orange Labs. He is working on short-term MPLS engineering activities and longer-term R&D projects. He contributes actively to the IETF, where he has edited and co-authored several RFCs and drafts in the area of MPLS/GMPLS. His current interests are fast rerouting, interdomain traffic engineering with PCEs, multicast MPLS, and GMPLS multiregion networks.

JEAN-LOUIS ROUGIER (rougier@telecom-paristech.fr) received his engineering diploma in 1996 and his Ph.D. in 1999 from TELECOM ParisTech (formerly called École Nationale Supérieure des Télécommunications). He joined the computer science and networks department of TELECOM ParisTech in 2000 as an associate professor. His research interests are performance evaluation and algorithms, traffic engineering, and routing in networks.

STEFANO SECCI (secci@telecom-paristech.fr) received an M.Sc. degree in telecommunications engineering in 2005 from Politecnico di Milano. Since June 2006 he has been working on a double Ph.D. degree at TELECOM ParisTech and Politecnico di Milano. His Ph.D. thesis is about multidomain service and transport architectures. He has worked as an assistant researcher at Politecnico di Milano and as a service infrastructure functional analyst at Fastweb S.p.A. His research interests are IP and optical network design, traffic engineering, and game theory.