Preserving Confidentiality in PCE-based Multi-domain Networks

Francesco Paolucci, Molka Gharbaoui, Alessio Giorgetti, Filippo Cugini, Barbara Martini, Luca Valcarenghi, and Piero Castoldi

***Abstract*—The path computation element (PCE) architecture has been proposed to effectively enable multi-domain traffic engineering (TE) in generalized multiprotocol label switching (GMPLS) networks while providing an adequate level of confidentiality among domains. However, a malicious utilization of the procedures defined within the PCE architecture might affect the confidentiality of network domain information in a multi-domain multi-carrier network scenario. This paper discusses the critical issues of the PCE architecture in terms of confidentiality. A two-step authorization scheme, named the behavior-based PCE authorization policy (BPAP), is proposed. The BPAP includes a novel add-on PCE component and a central authorization policy server to protect against confidentiality breaking. The scheme is based on the PCE protocol (PCEP) client behavior analysis and includes attack pattern detection procedures and possible partial information filtering of the reply message. The applicability of the BPAP scheme is validated in wavelength switched optical networks (WSONs) through simulations focusing on the exchange of a restricted set of available resources. Finally, a BPAP implementation is experimentally evaluated, showing the efficiency of the two-step scheme in terms of scalability, capability to limit the discovery of critical information, and reactivity to confidential attacks.**

***Index Terms*—Authorization policy; Confidentiality; Generalized multiprotocol label switching; Multi-domain; Path computation element; PCE protocol; Security.**

I. Introduction

The provisioning of quality-of-service- (QoS-) guaranteed applications has driven the introduction of traffic engineering (TE) solutions in multiprotocol label switching (MPLS) and generalized MPLS (GMPLS) networks. TE relies on constraint-based path computation, which essentially consists of finding a shortest path between a source and a destination node, subject to constraints such as reservable bandwidth, diversity, and resource class affinity.

Within a single administrative domain, TE techniques are nowadays successfully adopted [1]. However, when multiple domains and carriers are involved in a path computation, a number of significant issues arises. Above all, the need to preserve information confidentiality across domains controlled by different carriers prevents the open advertisement of detailed intra-domain network resources (e.g., topology, node capabilities, and bandwidth availability) [2]. This considerably complicates the constraint-based path computation and affects the inter-domain TE performance in terms of overall network resource utilization. As a matter of fact, network operators do not currently implement inter-domain TE techniques and the provisioning of QoS-guaranteed applications across multiple domains is performed manually, often requiring several weeks, and typically relies on sub-optimal solutions.

To overcome this issue, the Internet Engineering Task Force (IETF) has proposed a set of inter-domain TE techniques within the path computation element (PCE) architecture [3]. Such techniques, rather than exploiting advertised information, rely on a distributed path computation, performed in a cooperative way by entities (i.e., the PCEs) belonging to different domains. PCEs are enabled to communicate with each other by means of the PCE protocol (PCEP). Each PCE has the responsibility for the path computation only in its own domain. PCEs cooperate by sharing just the result of each (intra-domain) path computation expressed as, for example, border node(s) to traverse, encrypted intra-domain routes, metric values. The combination of these results provides the entire source–destination path, and no additional information is exchanged among different domains.

Concerning GMPLS networks, in the context of multi-area wavelength switched optical networks (WSONs), PCEP is extended with label set information, thus enabling effective end-to-end path computation with guaranteed wavelength continuity constraint [4]. However, confidentiality issues currently prevent the use of such an extension in the context of multi-domain multi-carrier WSONs, where poor TE performance is still experienced [2].

In MPLS networks, the PCE architecture potentially provides effective TE [5]. However, this potential might be jeopardized by the possibility for a PCE belonging to a different domain to maliciously perform bogus or false computation requests, aiming at discovering important confidential information inside other domains. Despite authentication and encryption on path segments (i.e., path key [6]), several parameters and patterns could be used to discover important confidential information inside other domains. Typical confidential information includes details on intra-domain network resources (e.g., available bandwidth), congested portions of the network, node architectural limitations and constraints, recovery schemes, the ability/inability to support advanced

Manuscript received December 22, 2010; revised March 16, 2011; accepted April 8, 2011; published April 29, 2011 (Doc. ID 140132).

Francesco Paolucci (e-mail: fr.paolucci@sssup.it), Molka Gharbaoui, Alessio Giorgetti, Luca Valcarenghi, and Piero Castoldi are with Scuola Superiore Sant'Anna, Via Moruzzi 1, 56124 Pisa, Italy.

Filippo Cugini and Barbara Martini are with CNIT, Via Moruzzi 1, 56124 Pisa, Italy.

Digital Object Identifier 10.1364/JOCN.3.000465

1943-0620/11/050465-10/$15.00

network services and QoS-guaranteed applications. In [3], the need for implementing effective access policies to avoid malicious utilizations of the PCEP procedures is indicated. However, no risk analysis, implementation details or solutions are provided. Indeed, confidentiality aspects of multi-domain multi-carrier networks are still undiscussed in the context of the PCE architecture.

In this paper, we propose a novel authorization scheme, named the behavior-based PCE authorization policy (BPAP), to be enforced in the context of multi-domain multi-carrier (hereafter referred to as multi-domain) networks. BPAP analyzes the sequence of requests coming from a PCEP peer and is able to either limit the exchange of information or block requests following pre-determined attack patterns over a given intra-domain resource. The BPAP scheme is applied on a two-step extended PCE architecture recently proposed in [7]. Such a solution offers a reasonable trade-off between two opposite requirements: the need to preserve strict intra-domain information and the need to effectively utilize network resources. Moreover, BPAP features dynamic procedures providing the reply filtering option, in order to preserve confidential information.

First, we show through simulations that BPAP can be efficiently applied in PCE-based multi-domain WSONs without excessively impacting network utilization by exchanging a restricted set of available resource information. In addition, we experimentally evaluate the proposed solution through a BPAP implementation and validate its scalability performance in a real GMPLS network testbed, showing in addition the BPAP capability to rapidly detect anomalous and malicious behavior of PCEP peers.

II. PCE Architecture

The PCE architecture relies on two functional components: the PCE and the path computation client (PCC). The PCE, possibly implemented on a dedicated server, is responsible for performing constraint-based path computation requested by PCCs, which are typically implemented on a network management system (NMS) or a network node. In the inter-domain scenario, a PCE may also behave as a PCC, requesting path computations from a PCE belonging to a different domain. Communication between the PCC and the PCE is guaranteed by the recently standardized PCE communication protocol (PCEP) [8]. To perform path computations, the PCC and PCE first open a PCEP session within a TCP session. A path computation request is then included within a PCReq message, specifying all the requested parameters and constraints. A reply (i.e., PCRep message) is provided by the PCE, specifying either the positive result (i.e., explicit path route) or negative result (i.e., no path found). Additional messages are also defined to close the PCEP session and to handle specific events and communication errors (e.g., error (PCErr) and notification (PCNtf) messages).

To perform inter-domain path computations, the PCE architecture defines two main procedures: the PCE-based per-domain (PPD) [9] and the backward recursive PCE-based computation (BRPC) [10]. They both exploit a backward recursive technique. The path computation request is first forwarded between PCEs, domain-by-domain, until the PCE responsible for the domain containing the destination node is reached. The PCE in the destination domain then computes either a single sub-path (as in PPD) or a tree of virtual sub-paths (as in BRPC) from one or more border nodes attached to the upstream PCE domain to the destination. The result is passed back to the previous PCE, which in turn expands the sub-path(s) and passes the result back until the source domain PCE completes the entire path computation. In the case of BRPC, the source PCE also selects the shortest path among those included in the final tree. The main path computation parameters defined in [8] and [10] consider end points (source and destination), connection bi-directionality, and requested bandwidth. Other important parameters include diversity (link, node, and/or shared risk link group (SRLG) disjointness), the need for local protection (i.e., fast reroute), and the application of the BRPC procedure. In addition, PCEP specifications allow the provision of information about failure in the path computation (i.e., NO-PATH information), to specify strict/loose sequences of hops to traverse or avoid, computed metric values, priority in the path computation, and information to perform re-optimization.

III. Confidentiality in Inter-PCE Communication

Inter-domain PCEP-based computations are performed upon general agreements between adjacent domains, which include technical (e.g., physical connectivity, interface switching capabilities) and economical specifications. The general agreement also encompasses confidentiality aspects, with the formalization of a set of rules and permissions aiming at defining the basic limitations in requests and replies due to confidentiality reasons. Above all, PCEs and PCCs do not exchange a strict explicit list of traversed intra-domain hops, and paths are expressed in the form of an encrypted key [6,11]. However, this basic level of trust agreement is not sufficient to fully guarantee the required level of confidentiality. In fact, the additional information exchanged to enable distributed path computation may disclose, explicitly or implicitly, intra-domain information that a network operator wants to keep private [2]. In [10], an overview of security considerations is provided. Requirements and possible solutions are indicated to address vulnerability aspects, including spoofing (PCC or PCE impersonation), snooping (message interception), falsification and denial of service. With reference to confidentiality aspects, [10] identifies the need to additionally define network policies aiming at preserving network information from bogus computation requests. Indeed, differently from connection requests triggered during signaling [12], PCEP-based computations do not imply the subsequent setup of the required connection, thus potentially enabling a malicious utilization of the PCE architecture.

A. PCEP Parameters

In this subsection the main PCEP parameters are discussed, highlighting their potential risk for a malicious utilization to break confidentiality.

1) Bandwidth (MPLS): Path computations requiring small values of bandwidth do not usually induce confidentiality issues. On the other hand, a request for a significant amount of bandwidth should require some careful treatment, or even an immediate rejection, since it might allow the discovery of bottlenecks, e.g., in the case of negative reply.

2) Label Set (WSON): Path computation in the context of WSON may require the exchange of available wavelength channels along the path in order to satisfy the wavelength continuity constraint. The detailed knowledge of the total amount of currently available wavelengths is considered strictly confidential [2]. Therefore specific policies (e.g., label set filtering) may be required in order to limit or partially hide the amount of replied information.

3) Diversity and Bi-directionality: Path computation diversity, performed through the synchronized vector (SVEC) object, local protection and bi-directionality imply the need to identify, within the requested domain, available resources along multiple disjoint routes or directions. This might increase the risk of discovering bottlenecks, topological limitations, or node architectural constraints, in particular when associated with requests with relatively high bandwidth values.

4) Metrics: Metric values returned to a PCC might be used to infer intra-domain topological information. For example, subsequent identical requests for which the returned metric value changes might indicate a variation in the intra-domain resource availability.

5) Backward Procedure: The backward nature of the PCEP procedures allows the requesting domain to retrieve information without providing any information about its own domain. This is particularly critical in the case of BRPC, where a tree between border nodes and the destination is returned together with the computed metric values.

B. Implication on Confidentiality: Correlation and Patterns

Correlations among different path computation requests including the aforementioned parameters might introduce additional risks of breaking confidentiality. For example, multiple apparently independent path computation requests, targeting destinations located in the same geographical area, might hide a confidentiality attack. In particular, positive replies under certain constraints and negative replies under different constraints (e.g., link disjointness and SRLG disjointness) or in different time periods could practically reveal lack/availability of intra-domain resources or intra-domain network performance (metrics).

Path computations performed by a transit domain (i.e., not including source and destination nodes) should not be typically considered critical for confidentiality. The information included in the reply is not related to a specific domain, but refers to the entire computed sub-path until the destination node. On the other hand, path computations returned from a destination domain should be taken into account for confidentiality attacks.

In the PCE architecture, path computation requests for which a PCE provided a positive reply might not be followed by the related setup procedure (i.e., signaling messages). On the one hand, this could refer to a truthful need to identify the optimal path along alternative routes controlled by different operators. In this case, just one route will be eventually set up, while the others will end up in the discard of the computed path upon the expiration of a pre-defined timeout (e.g., ten minutes [6]). On the other hand, expired path computations might be considered as an attempt to discover confidential information. Also the time period between a positive reply and the related connection setup or timeout should be carefully treated, since a burst of requests could take place without being eventually concluded by any setup.

Among the aforementioned events, the detection of *patterns* [13], i.e., sequences of requests with parameters matching specific criteria or with periodical behavior, may clearly reveal the possibility of being under attack. Pattern analysis is frequently utilized in the field of network security. For example, a sequence of requests targeting the same destination node and presenting values of bandwidth following periodical incremental step behavior may represent a possible attack candidate sequence to be classified under the suspected patterns. In this case it is likely that the client is periodically monitoring the resource availability toward a network node.

C. Previous Work on Authorization Policies

Confidentiality issues are typically addressed through the implementation of either static access lists or dynamic policies [14]. The former consist of simply associating a permission to a user to access a particular resource (by denying or granting the access). Although significantly simpler and faster, they are not able to effectively identify confidentiality attacks due to correlation among different events. Additional and more sophisticated dynamic schemes are then required. In [12] a policy-based authorization management is proposed in the context of GMPLS networks. Such a methodology offers a much finer granular access control during resource reservation procedures based on the resource reservation protocol (RSVP), which can enforce specific actions along with the permit or deny permission. Recent proposals in the context of network resource provisioning rely on extensible access control markup language (XACML), and policies are defined in order to authorize access to the resources and to enforce QoS parameters [15]. More specific applications (e.g., stringent QoS-based services) require policies extending the XACML standard to include the access control information, as in the case of either the RSVP policy control [12] or the next-step-signaling-protocol- (NSIS-) based authorization, authentication and service level agreement (SLA) enforcement procedures [16]. The only relevant proposal involving the PCE architecture is described in [17] and accounts for authorization, authentication, and request resource acknowledgment. However, [17] does not address the problem of the malicious utilization of PCEP performed by authenticated peers aiming at breaking confidentiality.

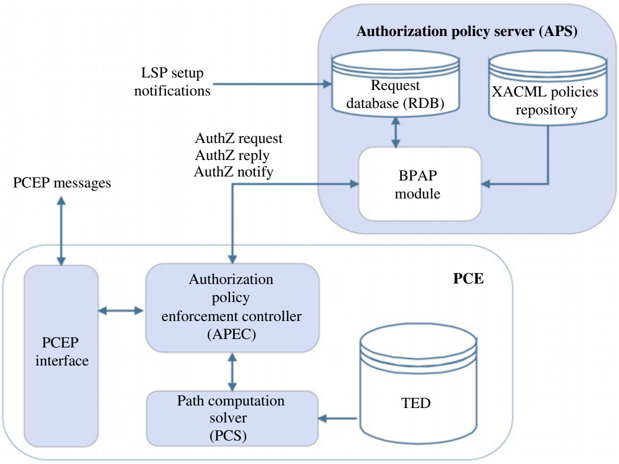
Authorization policy server (APS)
LSP setup notifications
Request database (RDB)
XACML policies repository
AuthZ request
AuthZ reply
AuthZ notify
BPAP module
PCEP messages
PCE
PCEP interface
Authorization policy enforcement controller (APEC)
Path computation solver (PCS)
TED

Fig. 1. (Color online) Extended PCE architecture.

IV. The Proposed Behavior-Based PCE Authorization Policy

A. Architecture

The extended architecture enabling BPAP in the context of the PCE architecture is shown in Fig. 1. It refers to a single domain which cooperates with adjacent domains to perform inter-domain path computations. As in [3], one or more PCCs per adjacent domain as well as one or more PCEs within the reference domain are considered. A PCE is equipped with a PCEP interface to handle PCEP communication and with a path computation solver (PCS) to perform path computations. The extended architecture follows the approach based on the policy decision point (PDP) and policy enforcement point (PEP) [15], encompassing two new additional elements: a new PCE add-on component, named the authorization policy enforcement controller (APEC), and a centralized authorization policy server (APS). The PCE acts as the PEP, while the APS acts as the PDP.

The APEC component is introduced to filter the inter-domain path computation requests and replies. Concerning input requests, the APEC performs basic request evaluations through simple permit/deny conditions specified in the form of access lists. Path computation results are also parsed by the APEC for policy enforcement operation, and specific output information filtering is performed, based on local confidentiality rules (e.g., optional reply objects, such as metrics and label set, may be totally or partially dropped).

The APS is introduced to run, when needed, more sophisticated authorization policies based on a detailed analysis of the behavior of the PCC with regard to past path computation requests and the risk to confidentiality related to incoming and previous requests. To accomplish the latter task, the APS utilizes a set of policy rules (e.g., XACML policies) stored in a local repository, a BPAP module devoted to the PCEP peer behavior analysis, and a request database (RDB) per domain. The RDB stores all the details (i.e., PCEP parameters, timestamp) of the completed path computations handled by the APS for each requesting domain, limited within a reasonable period of time (e.g., six months). RDB entries are also tagged with a status, based on the PCRep outcome and the possible related subsequent setup event: 1) *failure*: the requested path computation failed and the No-path object was included within PCRep; 2) *setup*: successful path computation with ERO included within PCRep, followed by the related LSP setup procedure (i.e., signaling); 3) *expired*: ERO included within PCRep, not followed by the related signaling, with setup timeout expired (typical value 10 min); 4) *pending*: ERO included within PCRep, not followed by the related signaling, with setup timeout not expired yet. The request status is dynamically updated based on the path computation outcome and the events occurring after path computation. In particular the pending state eventually changes into either expired or setup.

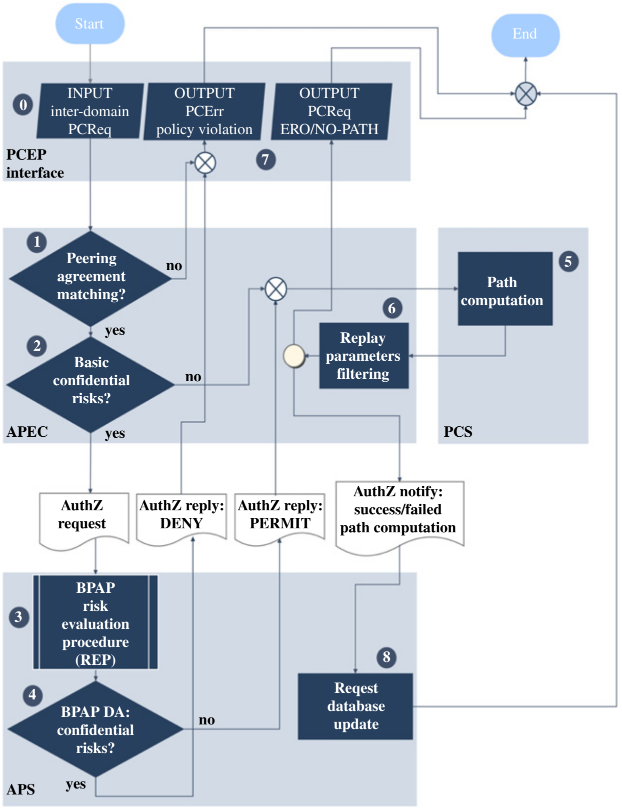
Start
End
0
INPUT inter-domain PCReq
OUTPUT PCErr policy violation
OUTPUT PCReq ERO/NO-PATH
7
PCEP interface
1
Peering agreement matching?
no
yes
2
Basic confidential risks?
no
yes
APEC
5
Path computation
6
Replay parameters filtering
PCS
AuthZ request
AuthZ reply: DENY
AuthZ reply: PERMIT
AuthZ notify: success/failed path computation
3
BPAP risk evaluation procedure (REP)
4
BPAP DA: confidential risks?
no
yes
APS
8
Reqest database update

Fig. 2. (Color online) PCE–APS authorization procedure workflow.

Communication between the APEC and the APS is achieved through the exchange of specific authorization messages. To maintain databases, the APS is notified with information on the final status (setup or expired) of the computed path, e.g., through simple network management protocol (SNMP) notifications from the network management system (NMS). The decoupling of the authorization evaluation performed by the APEC (involved in all inter-domain path computations) and possibly by the APS (only when complex evaluation is required) is introduced to better address the scalability requirements of the overall authorization scheme.

B. Authorization Procedure

Figure 2 shows the entire procedure performed upon the arrival of an inter-domain path computation request (step 0), forwarded from the PCEP interface to the local APEC.

Step 1: The APEC first evaluates the incoming request on the basis of the general trustworthiness agreement defined with the adjacent domain. The APEC decides whether to immediately reject the authorization request, tagged as *unacceptable* (e.g., excessive request burst, parameters not allowed, such as excessive bandwidth) or proceed with the next step. In the former case, the procedure goes to step 7.

Step 2: The APEC evaluates the request on the basis of a simple authorization policy. Requests not determining risks to confidentiality (e.g., negligible bandwidth requirement) are tagged as *risk-free* and then directly forwarded to the PCS (step 5), while the remaining are tagged as *critical* and passed to the APS for more careful authorization evaluation.

Step 3: The BPAP module evaluates the risk in terms of confidentiality of the incoming request, taking into account its constraints, its possible correlation with previous requests, and pattern detection. Such operation is referred to as the *risk evaluation procedure* (REP).

Step 4: The BPAP module decides to authorize or deny the request based on parameters computed at the previous step, through a *decision algorithm* (DA). The APS replies to the APEC with the following mutually exclusive options: a) Authorize the incoming request to be forwarded to the PCS for the subsequent path computation. The procedure continues at step 5. b) Deny the incoming request because of an excessive risk to confidentiality. The procedure continues at step 7.

Step 5: The PCS performs the required path computation and the result is passed to the APEC.

Step 6: The APEC performs path computation reply filtering (e.g., label set restriction), based on the current confidentiality rules. Specific policies are applied to the requesting PCEP peer to prevent information inferring. Results are returned to the PCEP interface. The procedure continues at steps 7 and 8, performed in parallel.

Step 7: The PCEP interface returns to the PCC the result of the path computation request in the form of: a) a PCRep message with path computation failure (i.e., No-path) or with the computed path; b) PCErr due to authorization failure. The procedure ends.

Step 8: The RDB is updated through a notification message from the APEC to the APS: the pending path computation request is now completed, with PCRep information provided to the client, and is inserted in the RDB, classified as either a *failure* state in the case of No-path or a *pending* state in the case of path computation success. The procedure ends.

If the requesting PCC receives a significant amount of authorization denials, the PCEP session between the domains may be eventually closed in order to verify and renegotiate the peering agreement. Warning messages could be defined to allow the requesting domain to become aware of its risky position.

C. Confidentiality Risk Evaluation and Decision Algorithm

The risk evaluation procedure performed at the APS (step 3) is the core of the proposed BPAP and is based on a detailed study of the client past behavior. The procedure accounts for all the received critical requests and related replies for each adjacent domain stored in the RDB. The evaluation is based on two main concepts: the identification of the attacked resource and the attack pattern detection. The evaluation considers, for each different PCEP client belonging to a set of authorized clients, a possible suspicious behavior toward a given resource, identified through a specific set of requests having a defined pattern trend. Upon a new critical PCReq being received, to identify possible confidential attacks, RDB entries having the same resource target (e.g., destination node, destination area, edge-to-edge transit segment) are selected, taking into account the status of each entry. Then, the procedure evaluates whether the sequence of the request subset not triggering LSP setup (e.g, failure, pending and expired requests) correlates with some standard or previously acquired confidentiality attack patterns. Attack patterns are defined by the network operator based on historical information about recent attack attempts or based on standard suspicious behavior. Attack patterns present specific characteristics typical of automatic processes, such as limited periodicity. As an example, a sequence of PCReq toward a given node with values of requested bandwidth following a periodical staircase function may be classified as a suspected pattern aimed at discovering intra-domain bandwidth bottlenecks. The RDB entry status is utilized to complete the peer behavior analysis (e.g., many expired entries might reveal confidential attacks). The evaluation results are then subject to the decision algorithm (step 4). For each attack class a, a vulnerability parameter ρ_a ($0 \le \rho_a \le 1$) is introduced to estimate the probability of being under attack. The parameter is computed by taking into account the number, the order, the status of the entries, and the possible suspicious pattern detection. A threshold T_a fixed by the network operator defines the decision between authorization and denial. If $\rho_a < T_a$, the request is authorized and is passed to the path computation procedure, otherwise it is refused. Note that, if parallel instances of the scheme run analyzing different attack classes, an updated attack vulnerability parameter vector is generated. In this case, the maximum vector value is utilized for the threshold-based authorization decision.

V. BPAP Applicability in WSON

The described BPAP procedures include, beyond the PCE behavior analysis, the policy enforcement stage (e.g., filtering), aimed at preserving confidential information within the PCRep messages. Since information restriction may potentially impact the overall network utilization, it needs to be carefully treated, and its applicability needs to be evaluated in detail.

A. PCE-Based Label Set Exchange Schemes

In PCE-based multi-domain WSONs the following four schemes are considered, among which the last two apply BPAP.

1) No Label Set (NoLS): NoLS is the currently available scheme. It does not apply advanced authorization policies, and the path computation is performed by assuming that it does not violate confidentiality. PCE-based path computation accounts for the inter-domain routing. Wavelength continuity is verified during path computation inside each domain, but, since no details are available from other domains, the risk of incurring blocking is high.

2) Full Label Set (FLS): Although not applicable in multi-carrier WSONs for confidentiality reasons, FLS can be considered as the reference bound. In FLS, by exploiting the PCEP label set (LS) extension proposed in [4], wavelength availability is exchanged between domains; therefore end-to-end wavelength continuity is verified during path computation.

3) BPAP-Based Dedicated Label Set (B-DLS): In B-DLS different domains agree to dedicate a fixed pool of P wavelengths to intra-domain requests, and the remaining $W-P$ wavelengths to inter-domain requests. During inter-domain path computation, PCEP LS is used only on the $W-P$ wavelengths. In this way, the intra-domain resources on the P wavelengths are completely hidden to other domains.

4) BPAP-Based Restricted Label Set (B-RLS): B-RLS adopts PCEP LS but arbitrarily manipulates the information included in the LS to partially hide the domain wavelength availability. In particular, the set of available wavelengths is restricted end-to-end by a specific percentage γ. Therefore, during inter-domain path computation, each PCE removes some available wavelength from the LS. In particular, if W_{ls}^i wavelengths are contained in the LS received by the PCE i, it removes the first

$$W_r^i = \left\lfloor \frac{\gamma \cdot W_{ls}^i}{h-1} \right\rfloor \qquad (1)$$

wavelengths from the LS, where h is the number of domains traversed by the LSP. The filtering operation is performed in step 6 at the APEC, as described in Subsection IV.B. In this way, similarly to B-DLS, only a subset of available wavelengths is made visible to other domains.

In both B-DLS and B-RLS, BPAP is utilized (i) to verify the incoming LS, (ii) to validate/restrict the outgoing LS, and (iii) to evaluate correlations among different requests and replies. Indeed, particularly in the case of B-RLS, correlation among replies might be exploited to break confidentiality. For example, if multiple requests targeting the same end point obtain different restricted subset information, correlation might be used to discover the end-to-end set W_{ls}.

The parameter $\mathscr{C}$ is used to assess the confidentiality degree provided by each scheme. $\mathscr{C}$ is defined as the ratio between the number W_r of hidden end-to-end wavelengths and the number W_{ls} of available end-to-end wavelengths:

$$\mathscr{C} = \frac{W_r}{W_{ls}}, \quad 0 \le \mathscr{C} \le 1. \qquad (2)$$

Considering the definition in Eq. (2), the NoLS scheme provides $\mathscr{C} = 1$, while the FLS provides $\mathscr{C} = 0$.

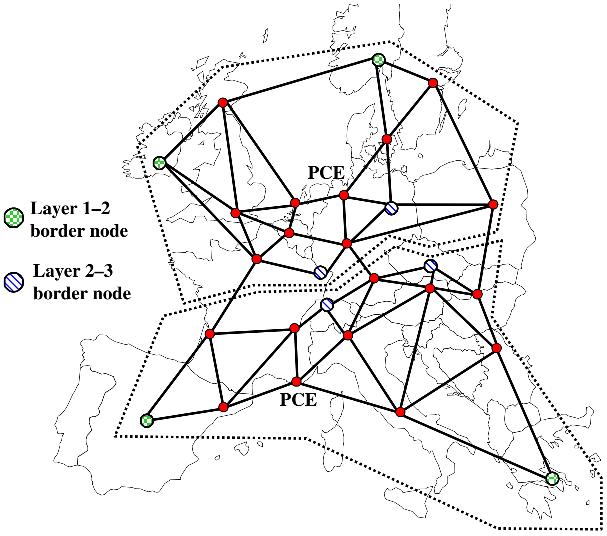

Fig. 3. (Color online) Pan-European triple layered 6-domain topology.

B. Simulation Results

The performance of the aforementioned schemes has been evaluated through simulations on a multi-domain WSON composed of six domains. The considered multi-domain topology is obtained by replicating on three layers the two-domain pan-European topology depicted in Fig. 3. Adjacent layers are connected through four inter-domain links between the border nodes highlighted in Fig. 3. The global topology comprises $N = 84$ nodes, $L = 156$ bi-directional links, each supporting $W = 32$ wavelengths. Lightpath requests, uniformly distributed between node pairs, are generated following a Poisson process with a fixed interarrival time of 100 s. An OSPF-TE instance is running in each domain, advertising detailed wavelength availability information (i.e., the status of each wavelength along every link). The PCE of each domain applies least fill routing among the set of shortest paths in terms of number of traversed links [18]. Wavelength assignment is first fit.

Two versions of each BPAP-based scheme have been evaluated. Since the considered traffic matrix is uniform, most of the lightpath requests are inter-domain. Thus B-DLS-10 and B-DLS-5 (i.e., with $P = 5$ and $P = 10$) respectively reserve 10 and 5 wavelengths for intra-domain traffic. B-RLS-50% and B-RLS-25% respectively use $\gamma = 50\%$ and $\gamma = 25\%$.

Figure 4 shows the overall blocking probability of the six considered schemes as a function of the offered network load. The plotted overall blocking is influenced by both blocking of intra-domain and blocking of inter-domain traffic. However, in the considered simulation scenario, the blocking of intra-domain traffic has resulted to be negligible. As expected, the NoLS scheme provides extremely poor performance with respect to the FLS, which represents the lower bound [4]. Conversely, BPAP-based schemes provide significant improvements with respect to NoLS. Among the BPAP-based schemes, B-RLS significantly outperforms B-DLS. Indeed, wavelength continuity on a single flexible set of resources provides higher network utilization with respect to two dedicated pools of resources, at any load. Moreover, the B-RLS block (especially

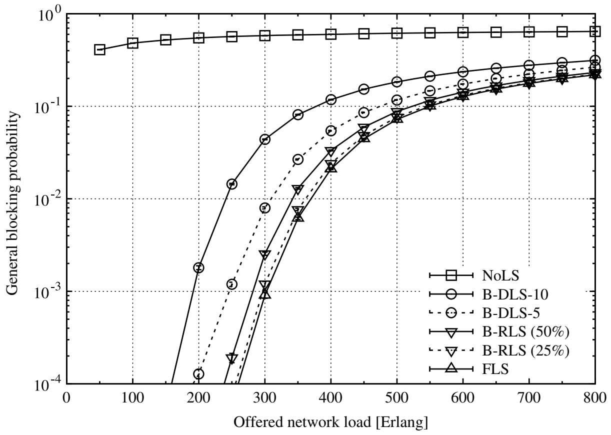

Fig. 4. BPAP simulation results: blocking probability.

TABLE I
BPAP SCHEMES' CONFIDENTIALITY DEGREE

Scheme	Load 250 [Er]	Load 350 [Er]	Load 450 [Er]	Load 550 [Er]
NoLS	1	1	1	1
B-DLS-10	0.42	0.5	0.53	0.55
B-DLS-5	0.19	0.24	0.28	0.31
B-RLS-50%	0.32	0.31	0.31	0.31
B-RLS-25%	0.15	0.13	0.12	0.12
FLS	0	0	0	0

using $\gamma = 25\%$) is very close to the FLS bound, confirming that filtering one fourth of the available LS induces a limited impact on the network resource utilization.

Table I shows the degree of confidentiality $\mathscr{C}$ of the considered schemes at different loads. As expected, the NoLS and the FLS schemes provide $\mathscr{C} = 1$ and $\mathscr{C} = 0$ for all loads, respectively. B-DLS-10 provides higher $\mathscr{C}$ with respect to B-DLS-5, resulting in a higher blocking (see Fig. 4). For both B-DLS schemes, $\mathscr{C}$ significantly increases for increasing loads because the number of hidden wavelengths (i.e., $W - P$) is fixed and the number of available wavelength decreases for increasing loads (see Eq. (2)). Conversely, B-RLS schemes provide $\mathscr{C}$ almost independent of the network load. Indeed, B-RLS hides a number of wavelengths which is proportional to the number of available wavelengths. In addition, while B-RLS-25% provides lower $\mathscr{C}$ with respect to B-DLS, B-RLS-50% provides higher $\mathscr{C}$ with respect to B-DLS-5 and, at the same time, it provides a lower blocking (see Fig. 4). This happens because, even if B-RLS-50%, on average, hides more wavelengths than B-DLS-5, it never removes all the available wavelengths from the LS. Conversely, with the B-DLS scheme, inter-domain lightpath requests are blocked if the only available wavelengths are in the intra-domain pool.

Finally, the obtained simulation results prove that the B-RLS provides a degree of confidentiality, almost independent on the network load (see Table I). This result is obtained without impacting the resource network utilization, thus reaching an overall blocking similar to the FLS lower bound (see Fig. 4).

VI. BPAP Experimental Implementation

In order to assess the BPAP capability to intercept anomalous/malicious behavior of PCEP peers, an experimental implementation has been evaluated in a real GMPLS domain testbed equipped with 9 commercial routers running OSPF-TE and RSVP-TE, forming a 3 × 3 grid topology, a C++ based PCE, and an external PCC performing PCEP requests. The BPAP module, including the risk evaluation procedure and decision, has been written in JAVA within a central APS in communication with the PCE through a dedicated XML-based socket, following the architecture in Subsection IV.A.

A. Considered Attack Classes and Patterns

The BPAP implementation considers a list of attack classes. Four attack classes are hereafter defined: the *bandwidth* monitor (*Bm*), the *wavelength* monitor (*Wm*), the *diversity* monitor (*Dm*), and the *topology* monitor (*Tm*).

Bm is the attack class attempting to collect updated information about the MPLS available bandwidth values toward or through a given resource, typically a node or a path. A typical example of Bm refers to a PCEP client sending PCReq messages periodically toward an egress PCE, requesting an LSP with the same destination end point and with increasing values of required bandwidth. Bm could also be used against a transit PCE by specifying the same transit border nodes to be traversed (e.g., by using exclude routing object (XRO) or include routing object (IRO) extensions).

Wm attempts to monitor the set or the amount of available wavelength channels toward a destination node and may be classified as the resource monitor in the context of WSONs.

Dm attempts to monitor the availability of reservable fully/locally protected paths toward a destination, together with information about diverse path node capabilities (e.g., node architectural constraints, nodal degree).

Tm attempts to monitor a specific portion of intra-domain network topology, aiming at inferring the network graph through cross-analysis of the intra-domain link metrics.

The pattern analysis of each monitor attack considers specific PCReq object values, as specified in Table II. RDB requests entries matching the target, containing the PCEP objects tagged as *mandatory* (M), and at least one of the objects indicated in Table II are selected for pattern analysis. The patterns considered for Wm apply only to input LS and are the full LS occurrence (i.e., the whole wavelength set is available), the XOR-based correlation (LS sequence having a high amount of transitions), and the LS contiguous subset (LS sequence having large subsets of available wavelengths and following fixed/mobile window behavior). The considered Bm patterns are the constant bandwidth function (e.g., PCC aims at monitoring the temporal availability of a certain bandwidth value), the incremental and decremental staircase function, the sawtooth function, and the generic monotone function. The patterns considered for Dm and Tm are the generic sequence presenting parameters with constant values and the generic sequence presenting parameters with periodical variation occurrences.

TABLE II
BPAP ATTACK CLASSES DETECTION FEATURES

Attack class	Target	Pattern parameters	Patterns
Bm (MPLS)	Destination end point	Bandwidth value (M)	Constant, inc/dec staircase, sawtooth, monotonic
Wm (WSON)	Destination end point	Label set (M)	Full LS, XOR correlation, LS contiguous subset
Dm	Destination end point	SVEC, bi-directional flag, LSPA, BRPC flag	Alternate parameters
Tm	Destination area	SVEC, BRPC flag, Metric (M)	Alternate parameters

B. Implemented Risk Evaluation Procedure

The decision algorithm applies on the set $\mathscr{L}$ of the RDB entries having the same target, provided by the risk evaluation procedure. For each attack class, the vulnerability parameter ρ is defined as $\rho = \alpha\rho_p + (1-\alpha)\rho_s$, where ρ_p accounts for the detection of one or more patterns, ρ_s accounts for the request entries (each tagged with status s_l) collected from the RDB and α is a (0, 1) tunable weight that enhances or reduces the impact of the pattern discovery on the authorization decision. The sub-parameter ρ_s is defined as

$$\rho_s = \frac{N_{\mathscr{L}} - N_{\mathscr{L}}^{setup}}{N_{\mathscr{L}}^2} \sum_{\mathscr{L}} w_l, \qquad (3)$$

where $N_{\mathscr{L}}$ is the number of entries of the set $\mathscr{L}$ selected for analysis, $N_{\mathscr{L}}^{setup}$ is the number of setup tagged entries, and w_l is the weight of the lth entry, dependent on the status s_l:

$$w_l = \begin{cases} 0 & s_l = setup \\ 0.5 & s_l = failure \\ 1 & s_l = expired \\ [0.5, 1) & s_l = pending. \end{cases} \qquad (4)$$

The w_l weights are introduced in order to tune the risk contribution given by each entry status. In particular, expired request entries are considered probable attack candidates, failure requests are considered intermediate potential attacks, pending requests worsen as the setup time increases, while setup requests are risk-free. ρ_s equals 0 if all the requests are setup (minimum alert state) and equals 1 if all requests are expired (maximum alert state). The parameter is designed also to smooth the alert state in case the PCC performs requests followed by setup.

The sub-parameter ρ_p is defined as a discrete case function depending on a set of flag parameters:

$$\rho_p = \begin{cases} 0 & f_D = 0, f_F = 0 \\ 0.25 & f_D = 0, f_F = 1 \\ 0.5 & f_D = 1, f_F = 0, f_M = 0 \\ 0.75 & f_D = 1, XOR(f_F, f_M) = 1 \\ 1 & f_D = 1, f_F = 1, f_M = 1. \end{cases} \qquad (5)$$

Each of these flag parameters specifies an additional risk level related to the pattern detection. In particular, $f_D \in \{0,1\}$ accounts for the *detection* of at least one pattern (0.5 weighted), $f_F \in \{0,1\}$ accounts for the detection of periodical (suspected) request value occurrences, which reveals a malicious (*frequency*-active) process, $f_M \in \{0,1\}$ specifies whether the incoming request values *match* at least one of the detected patterns. Multi-level ρ_p values are chosen to describe possible intermediate scenarios and reduce false positive events.

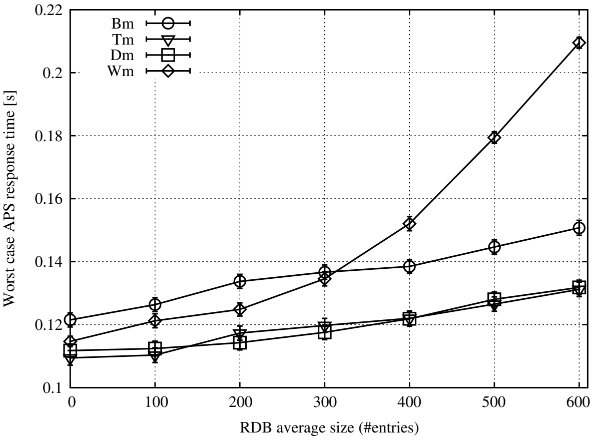

Fig. 5. Experimental results: APS response time.

C. Experimental Results

In Fig. 5, the time required by the APS to authorize/deny a request is plotted with the confidence interval at 90% of confidence level, as a function of the RDB size, assuming the worst case, i.e., all the entries are selected for pattern analysis. Times range from 100 to 210 ms considering up to 600 entries, showing good scalability performance. Dm and Tm curves present a similar trend below 140 ms. Bm requires additional time due to specific analysis of the bandwidth value sequence. Wm, based on the analysis of the label set, requires a greater additional amount of time with respect to the other monitors, due to its computationally intensive pattern identification procedures based on vectors of size $W = 40$.

For each attack class, a pre-defined attack pattern of 30 PCReq messages has been submitted to the BPAP to test reactivity to incoming attacks. In Fig. 6 the Wm attack is tested and the number of PCReq messages required to trigger the first PCErr is reported as a function of the number of initial consecutive PCE failure replies (i.e., No-path PCRep). The same benchmark attack (constant label set) has been run with different PCReq interarrival times: the *slow* test triggers requests with interarrival times greater than the path-key validity timeout (10 min) [6], while the *fast* test generates requests with interarrival times below the timeout (in the test, 1 min). Results show that using low α values the reactivity is slower and significantly dependent on the amount of failures, while with high α values the pattern

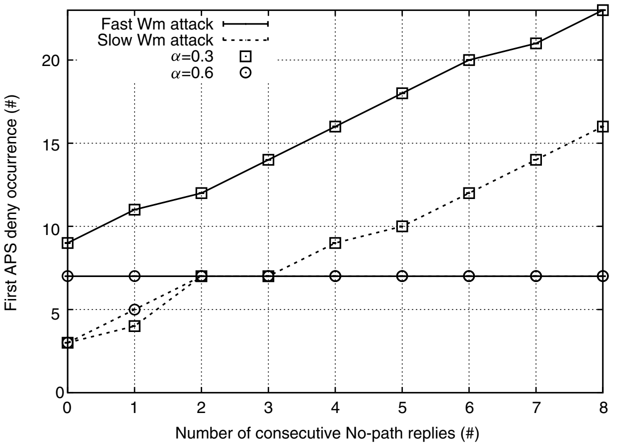

Fig. 6. Experimental results: reactivity during Wm attack ($T = 0.8$).

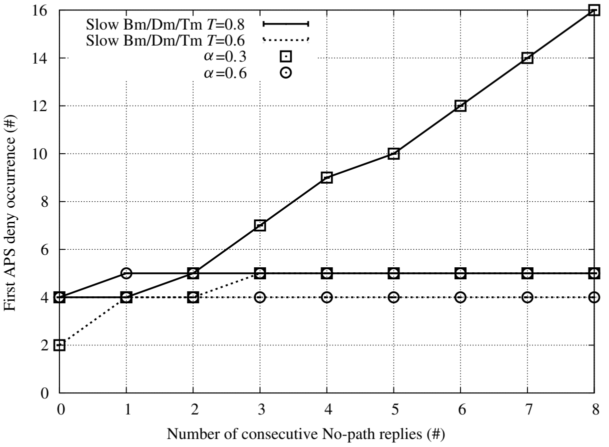

Fig. 7. Experimental results: reactivity during Bm/Dm/Tm slow attacks.

identification is given priority and PCErr is triggered long in advance, as soon as the pattern is discovered. The implemented pattern detection reveals the attack process after a sequence of 4 PCReq messages. Furthermore, the fast tests introduce a reactivity delay with respect to the slow tests. This is due to the dynamic update of the weights w_l in the pending state, assuming an intermediate value between failure and expired status values, as defined in Eq. (4). Indeed, pending requests are considered to be more suspected as time increases and expiration due to timeout is more likely to occur. However, while the failure state is considered uncertain and the expired state is considered risky, an excessive setup delay may reveal suspicious activity, but it should be considered licit as well. For these reasons reactivity to fast pending requests is slower.

In Fig. 7 Bm, Dm, and Tm are tested only with the slow benchmark attack sequences using two different values of the decision threshold T. The results are qualitatively similar to the Wm test, showing that the reactivity level is the same for all the considered attacks. The plot also shows the impact of the threshold T. As T decreases, the reactivity increases and the BPAP behavior flattens, regardless of the value of α.

Reactivity tests show that the BPAP may be flexibly tuned (e.g., by setting T and α) in order to provide the required level of security and to determine the level of tolerated suspicious PCEP behavior. It is worthwhile to note that a tradeoff may be required to both preserve confidentiality and address high network utilization. Reactivity tests also suggest that either giving excessive emphasis to pattern identification (i.e., high α values) or setting a low threshold may lead to frequent false positive detection.

VII. Conclusion

This paper discussed the confidentiality issues arising when the PCE architecture is utilized for enabling TE in multi-domain multi-carrier GMPLS networks. The study showed that the PCEP might be potentially used by authenticated peers to maliciously monitor the intra-domain network, thus breaking confidentiality.

The proposed BPAP scheme, enforced in an extended PCE architecture, enables the decoupling of basic access-list-based authorization evaluation involved in all inter-domain path computations from the complex-pattern-detection-based evaluation. The two-step structure assures the scalability requirements of the overall authorization scheme.

Simulation results showed that BPAP applicability in the WSON scenario is feasible using dynamic availability information restriction, inducing a limited impact on the overall network resource utilization.

The experimental evaluation carried out on a real GMPLS testbed showed that the BPAP implementation successfully prevents four different attack classes, showing good scalability performance in terms of response time and tunable flexibility concerning detection reactivity toward PCEP attacks.

Acknowledgment

This work is supported by the STRONGEST project, funded by the European Community. Part of this work has been presented in [7].

References

[1] A. Farrel and I. Bryskin, *GMPLS: Architecture and Applications*. Morgan Kaufmann Publishers Inc., 2005.

[2] Y. Lee, G. Bernstein, J. Martensson, T. Takeda, and T. Tsuritani, “PCEP Requirements for WSON Routing and Wavelength Assignment,” *Draft-ietf-pce-wson-routing-wavelength-03*, IETF, Nov. 2010.

[3] A. Farrel, A. J. Bruce, and J. P. Vasseur, “A path computation element (PCE)-based architecture,” *IETF RFC 4655*, Aug. 2006.

[4] R. Casellas, R. Martinez, R. Munoz, and S. Gunreben, “Enhanced backwards recursive path computation for multi-area wavelength switched optical networks under wavelength continuity constraint,” *J. Opt. Commun. Netw.*, vol. 1, no. 2, pp. A180–A193, July 2009.

[5] S. Dasgupta, J. de Oliveira, and J.-P. Vasseur, “Path-computation-element-based architecture for interdomain

MPLS/GMPLS traffic engineering: overview and performance," *IEEE Network*, vol. 21, no. 4, pp. 38–45, 2007.

[6] R. Bradford, J.-P. Vasseur, and A. Farrel, "Preserving topology confidentiality in inter-domain path computation using a key-based mechanism," *IETF RFC 5520*, Apr. 2009.

[7] F. Paolucci, F. Cugini, B. Martini, M. Gharbaoui, L. Valcarenghi, and P. Castoldi, "Preserving confidentiality in PCEP-based inter-domain path computation," in *Proc. ECOC'10*, Sept. 2010.

[8] J. Vasseur and J. Le Roux, "Path computation element (PCE) communication protocol (PCEP)," *IETF RFC 5440*, Mar. 2009.

[9] N. Bithar, R. Zhang, and K. Kumaki, "Inter-AS requirements for the path computation element communication protocol (PCECP)," *IETF RFC 5376*, Nov. 2008.

[10] J. Vasseur, R. Zhang, N. Bitar, and J. Le Roux, "A backward-recursive PCE-based computation (BRPC) procedure to compute shortest constrained inter-domain traffic engineering label switched paths," *IETF RFC 5441*, Apr. 2009.

[11] R. Casellas, R. Martinez, R. Munoz, T. Tsuritani, L. Liu, and M. Tsurusawa, "Lab-trial of multi-domain lightpath provisioning with PCE path computation combining BRPC and path-key topology confidentiality in GMPLS translucent WSON networks," in *Proc. ECOC'10*, Sept. 2010.

[12] E. Toktar, E. Jamhour, and E. Maziero, "RSVP policy control using XACML," in *Proc. POLICY'04*, June 2004.

[13] L. Fang, C. Wang, and G. Ma, "A framework for network security situation awareness based on knowledge discovery," in *Proc. ICCET'10*, Apr. 2010.

[14] N. Nordbotten, "XML and web services security standards," *IEEE Commun. Surv. Tutorials*, vol. 11, no. 3, pp. 4–21, Mar. 2009.

[15] Y. Demchenko, M. Cristea, and C. de Laat, "XACML policy profile for multidomain network resource provisioning and supporting authorisation infrastructure," in *Proc. POLICY'09*, July 2009.

[16] S. Polito, D. Gebbers, M. Chamania, and A. Jukan, "A new NSIS application for LSP setup with security features," in *Proc. ICC'10*, May 2010.

[17] S. Polito, M. Chamania, and A. Jukan, "Extending the inter-domain PCE framework for authentication and authorization in GMPLS networks," in *Proc. ICC'09*, June 2009.

[18] A. Giorgetti, N. Sambo, I. Cerutti, and P. Castoldi, "Impact of link-state advertisement in GMPLS-based wavelength-routed networks," in *Optical Fiber Communication Conf. and Expo. and the Nat. Fiber Optic Engineers Conf.*, 2008, JWA98.