Behavior-based Authorization Policy for Multi-domain PCE-based MPLS and WSON

M. Gharbaoui[1], F. Paolucci[1], A. Giorgetti[1], F. Cugini[2], B. Martini[2], P. Castoldi[1]
1: Scuola Superiore Sant'Anna, Pisa, Italy, e-mail: m.gharbaoui@sssup.it
2: CNIT, Pisa, Italy

Abstract: A novel path computation authorization policy based on PCEP peer behavior analysis and attack pattern detection is presented. Applicability is demonstrated in multi-domain PCE-based WSON. Experimental validation in MPLS network testbed is provided.

OCIS codes: (060.0060) Fiber optics and optical communications; (060.4253) Networks, circuit-switched.

1. Introduction

In multi-domain multi-carrier networks, confidentiality issues inhibit the advertisement of accurate network information across domains. This impacts the Traffic Engineering (TE) performance, i.e. the overall network resource utilization. To address this issue, the Path Computation Element (PCE) architecture specifies the PCE Protocol (PCEP [1]) and path computation procedures performing inter-domain path computation through the cooperation between multiple PCEs, each having full visibility of the resources inside its controlled domain. However, this is currently not sufficient to always provide effective TE and guarantee an adequate confidentiality.

In [2], multi-*area* Wavelength Switched Optical Networks (WSONs) are considered: PCEP is extended with Label Set information thus enabling effective end-to-end path computation with wavelength continuity constraint guaranteed. However, confidentiality issues currently prevent the use of such extension in the context of multi-domain multi-carrier WSONs and poor TE performance are still experienced in this context.

In MPLS networks, the PCE architecture potentially provides effective TE. However, as indicated in [1], such potential might be jeopardized by the possibility for a PCE belonging to a different domain to maliciously perform bogus or false computation requests. Despite authentication [1] and encryption on path segments (i.e., Path-key), several parameters and patterns could be used to discover important confidential information inside other domains. For example, inter-domain PCEP requests for a significant amount of bandwidth might allow the discovery of bottlenecks towards that end-point in case of negative reply. Same for constraints on diversity (e.g., SRLG disjointness), local protection and bi-directionality which imply the need to identify, within the requested domain, available resources along multiple disjoint routes or directions. The backward nature of the PCEP procedures allows the requesting domain to retrieve information without providing any information about its own domain. In addition, differently from connection requests triggered during signaling, PCEP-based computations do not imply the subsequent setup of the required connection. This facilitates correlations among different malicious path computation replies, significantly increasing the risks to break confidentiality.

This paper proposes an authorization policy, named *Behavior-based PCE Authorization Policy* (BPAP), that 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. BPAP applies on a two-step architecture and offers a reasonable trade-off to two opposite requirements: the need of preserving strict intra-domain information and the need of effectively utilizing network resources. We show through simulations that BPA can be applied in PCE-based inter-domain WSONs. In addition, we experimentally validate the BPAP performance in the context of inter-domain MPLS networks.

2. Proposed Behavior-based PCE Authorization Policy (BPAP)

To provide authorization functionalities and guarantee confidentiality, the proposed BPAP exploits a two-step procedure [3]. At the first step, authentication [1] and basic authorization evaluations are performed. Through simple permit/deny conditions specified in the form of access lists, the PCEP Request (PCReq) message is tagged either: *unacceptable*, *risk-free* or *critical*. If unacceptable (e.g., upon authentication failure), the request is denied and a Error (PCErr) message is returned. If risk-free (e.g., authentication is successfully performed and requested parameters fall within acceptable ranges), the request is accepted for path computation and a Reply (PcRep) message is provided. If critical, (e.g., when bandwidth exceed a predefined significant threshold), a second step with more sophisticated authorization policies is performed. To this extent, BPAP accounts for all the received critical requests and related replies for each adjacent domain (stored in a Request Database (RDB)). Critical requests are tagged with a status, based on the PCRep outcome and the possible related subsequent setup event: 1) *failure*: the requested path computation failed and *NO_PATH* 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 minutes); 4) *pending*: ERO included within PCRep, not followed by the related signaling, with setup timeout not expired yet.
The requests 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*.
Upon a new critical PCReq is received and evaluated, to identify possible confidential attacks, BPAP first selects the RDB entries having the same resource target (e.g., destination node, destination area, edge-to-edge transit segment). Then, BPAP evaluates whether the sequence of the selected requests correlates some standard or previously acquired confidentiality attack patterns. For each attack type *a,* a correlation parameter ρ_a ($0 \leq \rho_a \leq 1$) is introduced to estimate the probability to be under attack. The parameter is computed by taking into account the number, the order, the status of the entries and the possible pattern likelihood detection. A threshold T_a defines the decision between authorization and deny. If $\rho_a < T_a$ the request is authorized and is passed to path computation procedure, otherwise it is refused and a PCErr containing a proper Error object (Policy Violation: Confidentiality) is sent back to the client. Note that, if parallel instances of the scheme run analyzing different attack types, an updated attack correlation parameter vector is generated. In this case, the maximum vector value is utilized for the threshold-based authorization decision.

3. BPAP application in multi-domain WSON

In PCE-based multi-domain WSONs the following four schemes are considered. (1) *No Label Set (NoLS)*. NoLS is the currently available scheme: advanced authorization policies are not applied and path computation is performed not violating confidentiality. PCE-based path computation accounts for the multi-domain routing. Wavelength continuity is verified just at the destination domain up to the border node. Since no details are available within other domains, the risk to incur into blocking is significantly high. (2) *Full Label Set (FLS)*. Although not applicable in multi-carrier WSON for confidentiality reasons, FLS can be considered as the reference bound. In FLS, by exploiting the PCEP LS extension proposed in [2], wavelength availability is exchanged between domains, as confidentiality would not represent an issue. In FLS, end-to-end wavelength continuity is verified during path computation. Two alternative schemes are then proposed based on BPAP.
(3) *BPAP-based Dedicated Label Set (B-DLS)*. In B-DLS different domains agree to dedicate to multi-domain requests a fixed pool *P* of wavelengths and to intra-domain requests the remaining *W-P* wavelengths. During multi-domain path computations, PCEP LS is used just on the considered *P* resources. In this way, the intra-domain resources on the *W-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 LS to partially hide the full set of available information. In particular, the set *W'* of available wavelengths is restricted by removing the first *P'* wavelengths which satisfy the wavelength continuity within the domain. In this way, similarly to B-DLS, only a subset of available intra-domain resources 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, (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 apparently independent requests targeting the same end-point obtain different *(W'-P')* information (extracted from the same *W'* available wavelengths), correlation might be used to discover the whole set *W'*.
The performance of the aforementioned schemes have been evaluated through simulations on a two-domain WSON with, overall, N=28 nodes, L=55 bidirectional links each supporting W=40 wavelengths. In particular N_1=N_2=14 nodes and L_1=L_2=25 links per domain, with the two domains connected by 5 inter-domain links. Lightpath requests are generated following a Poisson process and are uniformly distributed between node pairs and, as a consequence, between intra- and inter-domain requests. Least fill routing is applied among the set of shortest paths in terms of number of traversed hops. Wavelength assignment is first fit. Fig. 1 shows the overall blocking probability of the four considered schemes as a function of the offered network load. As expected [2], NoLS provides extremely poor performance, in particular with reference to FLS which represents the (inapplicable) lower bound. BPA-based schemes provide significant improvements with respect to NoLS. Among the BPA-based schemes, B-RLS (using P'=10) outperforms B-DLS (P=20). Indeed, wavelength continuity on a single flexible set of resources provides higher network utilization with respect to two dedicated pools of resources. Fig. 2 shows the blocking probability contributions due to either inter- and intra-domain requests. B-DLS provides higher blocking for both contributions. Conversely, in B-RLS, performance are affected just in terms of inter-domain blocking, while intra-domain requests are able to achieve a comparable blocking with respect to the bound FLS. Indeed, B-RLS efficiently exploits for intra-domain requests the resources not used by inter-domain requests.

4. Experimental BPAP implementation and results

To assess the BPAP effectiveness not only in the context of multi-domain WSON, an experimental implementation has been evaluated in a real MPLS testbed equipped with commercial routers, a C++ based PCE and an external

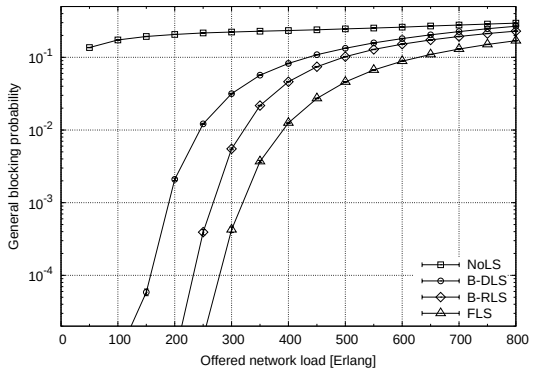

Fig. 1. Blocking probability vs. network load

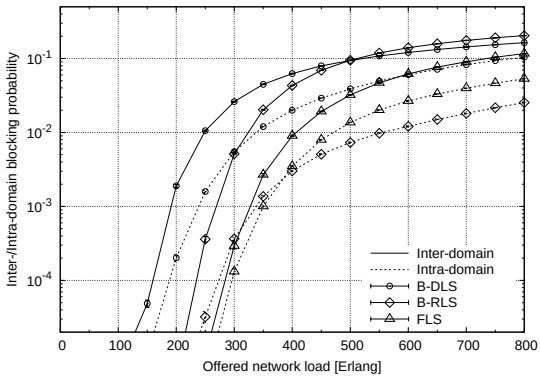

Fig. 2. Intra- and Inter-domain blocking vs. network load

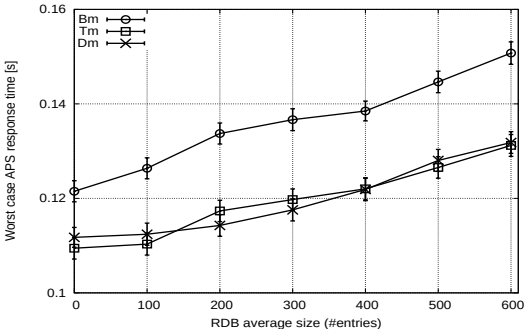

Fig. 3. Experimental testbed: BPAP response time

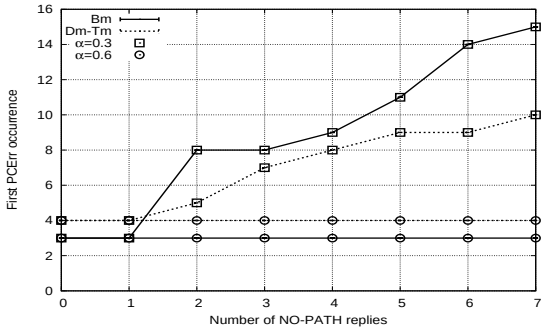

Fig. 4. Experimental testbed: BPAP reactivity (T_a=0.8)

PCC performing PCEP requests. BPAP has been written in JAVA within a central Authorization Policy Server (APS) in communication with the PCE through a dedicated XML-based socket, following the architecture in [3].

The considered attack types are: *Bandwidth monitor* (Bm), *Diversity monitor* (Dm) and *Topology monitor* (Tm). Bm and Dm attacks target a destination node in case of egress target PCE or a remote border node in case of transit target PCE. Bm and Dm attacks aim at discovering network information in terms of total available bandwidth and path diversity capabilities, respectively. Tm attacks target a specific network area and attempt to detect portions of an intra-domain topology and its variations. The considered domain is egress-type, thus the target is the destination node/area. The periodical trend, the time sequence of optional parameters occurrences and of LSP bandwidth values are the possible patterns. The 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 requests samples *status* 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. In Fig. 3, the time required by the APS to authorize/deny a request is plotted with a confidence interval at 90% 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 150 ms considering up to 600 entries, showing very good scalability performances. Bm requires additional time due to specific analysis of the bandwidth values sequence, while Dm and Tm curves present a similar trend.

For each attack type, a benchmark attack pattern of 20 PCReqs has been submitted to BPAP to test reactivity to incoming attacks. In Fig. 4 the number of PCReqs to trigger the first PCErr is reported as a function of the number of initial PCE failure replies (i.e., NO-PATH PcRep). Using low α values the detection reactivity is slower and significantly dependent on the amount of failures, while with high α values the pattern identification is given priority and PCErr is triggered in large advance, as soon as the pattern is discovered. However, giving excessive emphasis to pattern identification may lead to frequent false positive detection.

5. Conclusion

In this study, the Behavior-based PCE Authorization Policy (BPAP) is proposed to address PCEP confidentiality in multi-domain multi-carrier networks. Path computation schemes exploiting BPAP are proposed and evaluated through simulation in WSONs showing good blocking probability performance and the capability to preserve confidential a subset of the exchanged Label Set. BPAP experimental validation in a MPLS network successfully prevents three different PCEP attack types, showing good scalability performance in terms of response time.

Acknowledgements: This work was partially supported by the STRONGEST project.

References

[1] J-P. Vasseur and J. LeRoux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, Mar 2009.
[2] R Casellas, R Martínez, R Muñoz, S Gunreben, J. Optical Communication Networking (JOCN), Jul 2009
[3] F. Paolucci *et al.*, "Preserving confidentiality in PCEP-based inter-domain path computation", ECOC 2010.