Inter-domain Path Provisioning with Security Features: Architecture and Signaling Performance

S. Greco Polito[1], S. Zaghloul[2], M. Chamania[3], and A. Jukan[3]
[1]Università degli Studi di Enna "Kore", [2]Red Knee, Inc, [3]Technische Universität Carolo-Wilhelmina zu Braunschweig

***Abstract*—Significant research and standardization efforts are underway to enable an automated computation and setup of connection-oriented paths across multiple domains. Despite the technological advances in transport network technologies, such as Ethernet PBB-TE, MPLS-TP, WDM, security features related to connection origination, monitoring and accounting remain an open challenge. In absence of secure authentication and authorization, carriers will continue to provision connections manually, which may lead configuration errors and large setup delays. Carriers also lack mechanisms to meter connection quality during the service lifetime and typically do not exchange accounting information for established connections for auditing and billing purposes. In this paper, we address the challenge for automatic multi-domain path provisioning with authentication, authorization and accounting system in carrier-grade transport networks. The designed solution secures computation and reservation for path provisioning and allows providers to choose different AAA features on the basis of domain-internal management. Our framework also incorporates a standard accounting model, for which we provide an analytical performance model to assess its signaling performance in the proposed system. We verify the analysis by simulations and quantify the feasibility of our model in terms of signaling load and delay in a wide range of scenarios.**

***Index Terms*—AAA, connection-oriented networks, PCE, RSVP, Diameter, peering agreements, inter-domain routing**

I. Introduction

The connection-oriented networking across multiple domains is gaining momentum due to the advances in circuit-switching technologies, such as carrier-grade Ethernet, MPLS-TP, and optical wavelength division multiplexing (WDM). While the research community has widely embraced the concept of automated multi-domain path computation and reservation in carrier-grade transport networks, the path computation and setup in commercial networks remains manual. This is primarily due to the static nature of service level agreements (SLA) which are negotiated off-line, but also to the lack of automatic authentication and authorization features in the current provisioning frameworks. Carriers also lack mechanisms to meter connection quality during the service lifetime and typically do not exchange accounting information for established connections for auditing and billing purposes.

Adopting the authentication and authorization mechanisms matured in packet-switched IP networks [1] [2] carries however challenges for connection-oriented networks since the current IP networks support primarily the so-called *cascaded peering model*. Whereas in the cascaded peering model the service provisioning is driven by peering agreements between neighboring domains, the connection-oriented networks are expected to use multiple peering models, i.e., cascaded and alliance peering models, and their combinations [3]. In the alliance peering model, for instance, the SLAs exist between the requester of circuit and all the domains along the inter-domain path. In addition, circuit provisioning encompasses discrete path computation and reservation procedures [4], which makes it essential that the management framework also protects the system from configuration failures and even malicious attacks derived from these procedures. To illustrate this, let us consider the example in Fig. 1. The path computation system of domain A asks domains B and C for computation of local paths to a destination in domain C and a reservation request is then sent from the source in A to build a circuit to a destination in C. A possible misconfiguration can come from the "middle" domain B. First, domain B can erroneously initiate reservation procedures from its own node along this computed path, either due to attack while pretending to be A, or by a configuration error. In addition, domain B can route the incoming computation request toward domains untrusted by the source, while trusted by itself, and intentionally or unintentionally affect the path integrity.

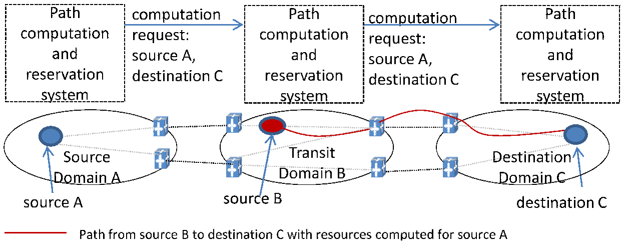
Path computation and reservation system
computation request: source A, destination C
Path computation and reservation system
computation request: source A, destination C
Path computation and reservation system
Source Domain A
Transit Domain B
Destination Domain C
source A
source B
destination C
Path from source B to destination C with resources computed for source A

Fig. 1. Current path computation and setup and security challenges

In this paper, we address the issue of inter-domain authentication, authorization and accounting (AAA) for carrier grade networks and propose an architectural and signaling solution that is flexible so as to adapt to different peering models. Moreover, we propose a security solution via coupling of computation and reservation procedures, which is novel. In our architecture, an AAA server interacts with the management system during path computation to enable secure path provisioning, and collects accounting information for accounting, monitoring or metering purposes during the connection lifetime. As such, our proposal is fully compatible with the IETF proposals for multi-domain computation and reservation, such as the emerging third-party management sub-system PCE (Path Computation Element) [4], and RSVP (Resource Reservation Protocol) [5]. We design four-way authentication mechanisms integrated with the computation signaling to secure path computation and use authentication tokens to ensure that path reservation signaling traverses the securely computed paths for authenticated and authorized domains. Authentication for path computation is performed using static symmetric/asymmetric keys agreed upon a-priori, while dynamic session keys are derived to secure the path reservation signaling. For the proposed architecture, we provide a new analytical model to characterize its signaling load and delay, and verify the results by simulations in a range of realistic test scenarios, including the mentioned peering models. We also show the stability and robustness of our signaling mechanisms in response to variations in connection duration statistics.

The rest of this paper is organized as follows: Section II presents the related work. Section III describes the provisioning scenario under study, while Section IV presents the proposed architecture including the path computation, reservation and accounting, respectively. Section V presents the analytical model for the proposed architecture and signaling. The results are presented in Section VI. Section VII concludes the paper.

II. Related Work and our Contribution

The IETF framework [1] presents a general purpose AAA architecture with inter-working AAA servers in multiple domains. Here, verification of identity and authorization rights of incoming service requests can be performed according to three main authorization models, i.e., agent, pull and push [6]. In this paper, authorization of path computation requests is performed using the pull model, during which tokens are computed and later used for authorization of reservation requests.

Unlike in [1] where the agent model is used to transfer service requests between domains, in our model, the inter-domain computation and reservation messages remain under the control of PCE and RSVP servers, according to the current IETF standards. We define extensions to the PCE protocol to include authentication and authorization, and we use of RSVP policy object [1] to carry the authentication and authorization token [7]. We also include mechanisms for the domains to establish shared dynamic session keys similarly to what done within the IKE (Internet Key Exchange) protocol [10].

Note that in a cascaded peering model, and given that the PCE protocol (PCEP) is implemented over TCP, an authentication scheme based on TLS may be a valid alternative. However, the same is not applicable when using alliance peering which requires authentication and authorization between non-neighboring domains. On the other hand, transport and network layer encryption mechanisms discussed within [11] can be used in our architecture to secure communication between neighboring domains for any kind of peering. Our solution is compliant with the security specifications of the IP/MPLS forum in [12] stating that border routes of Autonomous Systems must support MD5 authentication for all protocols using TCP and for the RSVP Integrity object [9], and exchange of signaling over IPsec tunnels. Implementation in border routers of mechanisms to filter and rate limit signaling are suggested in [12] against denial of service (DoS) attacks. These mechanisms implemented in PCE nodes can be used to block DoS attacks activated via PCE signaling.

Only a few past research papers addressed the issue of AAA integration in multi-domain connection-oriented service provisioning. In [13] a pull-like model is proposed with inter-domain path setup requests controlled by AAA agents deployed in each domain. These AAA servers authenticate and authorize incoming path requests, ask local servers for computation of intra-domain paths, and forward the requests to the AAA servers of the next-hop domains. This approach, like ours, promotes strong integration of authentication and authorization signaling with path setup signaling. However, to facilitate ubiquitous adoption, our goal is to secure provisioning over existing and standardized service provisioning protocols, namely PCE and RSVP, which also allows for easy integration with other legacy systems [14] and eliminates the need for new inter-domain interfaces. From the later perspective, our work carries significant novelty.

Papers [15] [16] propose the use of an inter-domain service plane. The approach associates a unique identifier to each negotiated service which is included in both the path computation and reservation requests for service authorization, effectively coupling the computation and reservation procedures. In our approach, we endorse the concept of coupling computation and reservation functions. However, we do not use a service identifier negotiated on the service plane, but the PCE path-key element representing paths computed according to the authorization profile of the requester. Enforcing utilization of path-keys for setup requests guarantees provisioning of resources authorized during the path computation. We use a secure token to transfer the path-key element within setup signaling with mechanisms to verify its origin and secure its transfer between domains. Utilization of tokens for authorization in multi-domain contexts is described in [17] in which the push model from [6] is extended to provide the local node with the authorization context associated with the token. Finally, paper [18] describes an authorization system for combined network and Grid resources that makes use of tokens. In our model, we use an approach similar to [17] for the specific PCE/RSVP architecture but we also take advantage of the existing interface between PCE and RSVP node to transfer the authorization context to the serving RSVP node.

In our preliminary work [19], we proposed security features for mutual authentication between the source and the other domains along an inter-domain path. In this paper, we extend the authentication mechanisms with those that can guarantee higher robustness to security attacks. This is achieved via a combination of 4-way mechanisms for mutual authentication and utilization of dynamic keys to secure the reservation process. In this paper, we also offer higher flexibility to providers as they can select among multiple authentication mechanisms based on domain-internal management and security requirements. Finally, this paper also proposes that the standard Diameter protocol for accounting is used in connection-oriented networks, for which we also develop an analytical model for the AAA signaling rate. This is based on our past research in [20], where we presented an analytical model to evaluate the AAA signaling rate for various components of the AAA

[1] The RSVP policy object is part of a set of extensions to the RSVP protocol [8] [7] [9] defined to facilitate authentication and authorization and integrity check of messages.

infrastructure in platforms for IP cellular networks. From an accounting perspective, the works in [21] and [20] have investigated performance issues due to the frequent exchange of accounting records between AAA servers and metering nodes that can be located in different domains. In [21], the authors observe how the IETF Diameter protocol [22], can result in an heavy signaling overhead because of the short length of accounting records (1-4 bytes for IP flows). They propose extensions for the accounting Diameter application to reduce the overhead for transfer of accounting records.

III. Provisioning Scenario and our Objectives

In our provisioning scenario, we use the PCE/RSVP system, a de-facto standard within IETF. In this system, path computation is performed by a stand alone PCE server [4]; this server computes domain-internal paths satisfying the user's QoS constraints and can implement policy mechanisms [23] to make computation dependent on the authorization profile of requesters. If the requested path spans multiple domains, see example in Fig. 2, the PCE of the source domain issues a computation request (PCEreq) which is forwarded downstream by each PCE along the path till the destination domain [24]. Here, the PCE issues a path computation response (PCEresp) carrying the available paths from the destination to its edge nodes toward the upstream domain. PCEs of transit domains compute local paths to extend the paths carried in the PCE response from the edge nodes corresponding to the downstream domain to their edge nodes corresponding to the upstream domain. In this fashion, a PCE computes the so-called Virtual Shortest Path Tree (VSPT) [25] describing the available optimal paths from source to destination.

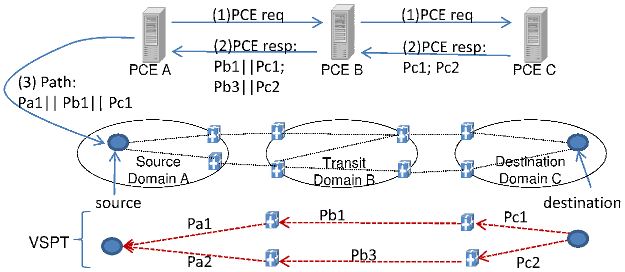

Fig. 2. Current architecture with PCE-based inter-domain computation

Upon completion of signaling, the PCE of the source domain computes the optimal segment to finish the VSPT and provides the requesting node with the computed path. The requesting node then asks for a label switched path (LSP) via the *RSVP Path* message [5] carrying a label request which is forwarded downstream till the destination. Labels assigned to the path are communicated by each node to their upstream neighbor via the *RSVP Resv* message which is issued by the destination and sent upstream along the path described by the *RSVP Path* message. As RSVP is a soft-state protocol, connection state is refreshed using periodic *RSVP Path* and *RSVP Resv* messages having the same format as the initial signaling which act as keep-alive messages along the connection path. We use the path-key mechanism as introduced in RFC 5520 [26] in order to prevent disclosure of domain-internal topology information to other domains. A path key is a token associated with the intra-domain path computed inside a domain and is used to obtain the path description during the reservation phase. Path keys are embedded inside the RSVP requests, and upon arrival at the domain ingress, the corresponding path-key is resolved by the PCE.

In carrier-grade networks, path computation and reservation signaling, along with routing and any other inter-domain signaling, can be manipulated to pose security threats for the provisioning system [11]. Such manipulations can lead to different kinds of attacks such as resource stealing, malicious hijack of inter-domain routes and cross connects, disclosure of private management information as well as Denial of Services attacks. In this paper, we focus on the security threats posed via manipulation of inter-domain computation and reservation signaling for unauthorized access to resources, resource stealing or path mis-configurations. Our approach to address these threats is integrating AAA features in the path computation/reservation system covering requirements of both cascaded and alliance peering models. We leave issues about intra-domain mechanisms for SLA verification [27] as well as SLA negotiation and key management for future work.

A. Objectives

The design of the proposed AAA architectural and signaling solution for multi-domain circuit provisioning is driven by the following objectives:

1) Secure access to multi-domain provision of network resources in multi-domain network scenarios, i.e., to authenticated and authorized entities only.
2) Secure provisioning from inter-domain path changes or resource stealing made by malicious (unauthenticated) domains along the inter-domain path.
3) Flexible AAA model to adapt to different peering models (such as cascaded and alliance peering models) as well as different cryptographic mechanisms (symmetric key and asymmetric key, characterized by different computational costs and key management requirements).
4) Backward compatibility by deploying existing standards for inter-domain signaling and management.

To achieve the objectives introduced above, we propose a PCE/RSVP management eco-system as illustrated in Fig. 3. From an architectural perspective, a "standard-based" AAA server, such as a DIAMETER or RADIUS server, is added to the PCE/RSVP components. These components, namely the PCE server and the RSVP border switching node (BSW), communicate with the AAA server via intra-domain interfaces (interfaces (a) and (b) in Fig. 4) to trigger authentication, authorization and accounting procedures. The AAA server refers to SLA repositories with information about QoS and authentication keys to perform its functions.

Keys negotiated in SLAs are called –SLA keys– and can be symmetric or asymmetric, i.e., shared keys and public certificates. From an inter-domain signaling perspective, as shown in Fig. 4, both PCE and RSVP signaling are enhanced with content for inter-domain authentication and authorization (AA), and the AAA server exchanges accounting information across domains facilitated by standard AAA protocols [22].

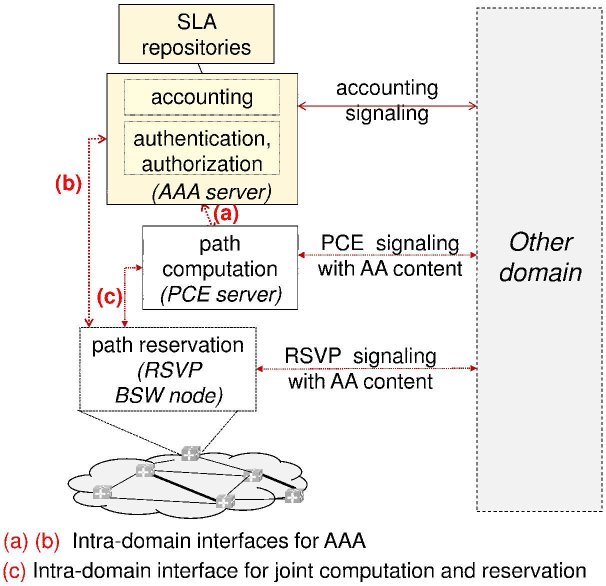
SLA repositories
accounting
authentication, authorization
(AAA server)
accounting signaling
(b)
(a)
path computation (PCE server)
PCE signaling with AA content
Other domain
(c)
path reservation (RSVP BSW node)
RSVP signaling with AA content
(a) (b) Intra-domain interfaces for AAA
(c) Intra-domain interface for joint computation and reservation

Fig. 3. The proposed PCE-based management eco-system

The AA content for path computation signaling includes authentication requests/responses and also contributes for generation of dynamic session keys. These keys are derived from Diffie-Hellman (DH) negotiation [28] and referred to as DH keys in the rest of the paper; these keys are used to secure authentication tokens for the RSVP setup signaling. Therefore, per session DH keys are used for authentication of reservation requests, while the SLA keys are used for authentication of computation requests.

IV. The Proposed Architecture

In this section, we describe implementation of the security features for path computation and reservation, and we also propose methods to include standardized mechanisms for accounting. In our descriptions, we will assume a cascaded peering model between domains for easier understanding and later on present the alliance model.

A. Path computation with security features

Our approach migrates from a single request/response signaling to a four-way signaling mechanism for path computation. As shown in the example in Fig. 4 (see path computation phase), the four-way signaling is composed of the *PCEreq(1)*, *PCEauth-resp(2)*, *PCEreq(3)* and the *PCEresp(4)* messages.

PCEreq(1): This message starts the path computation signaling and carries an authentication request from the source to the transit domain, and from the transit to the destination domain. The authentication request is encoded in the *AuthReq object*.

PCEauth-resp(2): This message is issued by the destination domain in response to the PCEreq(1). Destination and transit domains include the authentication responses for their upstream domains in this message along with requests for them to authenticate. Authentication responses and requests are encoded in the *AuthResp/Req object*.

PCEreq(3): This message carries the authentication responses as well as DH contributes from the upstream domains to the downstream domains. This message, upon successful authentication triggers the computation of the VSPT at the destination domain.

PCEresp(4): This message carries the computed paths in the VSPT along with the DH contribute issued by the downstream domains to their neighboring upstream domains.

Therefore each domain challenges its neighbors for authentication and provides authentication responses via the four-way PCE signaling. In each domain the PCE server interacts with the AAA server to facilitate AA using the AA content embedded inside PCE messages. This is described in Fig. 4 (Steps (S:a)-(S:c)). We now present in detail the PCEP AA information in the following subsection.

1) PCE authentication content: In general, authentication requests have the content shown in (1). This includes a domain ID for identification with the other party, a session ID which is an identifier for the session and an authentication random nonce to challenge the other party. The session ID is defined by the AAA server of the requesting domain and it is included in all subsequent authentication messages.

$$\text{Authentication Request= [DomainID, sessionID, local-nonce]} \quad (1)$$

Authentication responses typically depend on the kind of the key negotiated in SLAs: Keyed hashes are used for authentication if the parties share symmetric keys, while digital signatures are used if the parties share asymmetric keys. The content of the authentication responses for the two types of keys is shown in (2)-(3), respectively, with the sign "||" meaning concatenation.

$$\text{Authentication Response}_{Sym-key}\text{= [DomainID, sessionID,}$$
$$\text{Hash(otherDomain-nonce||key||sessionID||DomainID]} \quad (2)$$

$$\text{Authentication Response}_{Asym-key}\text{= [DomainID, sessionID,}$$
$$\text{Sig(otherDomain-nonce||sessionID||DomainID)]} \quad (3)$$

Domain ID and session ID are included in the Authentication Response and in the content for keyed-hash/signature computation for secure identification of issuer and authentication session, respectively. The nonce received from the other domain is included to guarantee protection against replay attacks. Authentication responses to upstream domains also include the DH contributes which, according to the specifications in [28], are computed as g^{secret} mod p with the prime number p and g a primitive root of p. Both p and g are public values, while $secret$ is a secret value issued by each party for its DH contribute. p and g can be agreed upon by neighboring domains or can be provided to downstream neighbors within authentication requests. Protection against man-in-the-middle attacks to DH keys is inherited by the hash or signature of the authentication response. The DH contributes for the downstream domains are also secured with keyed hashes or signatures.

2) Dealing with unsuccessful authentication: If a transit domain fails to authenticate its downstream neighbor, it continues the computation but includes a failure notification field in the *PCEauth-resp(2)*. This field is for the source domain which is the entity interested in knowing the security robustness of the inter-domain chain. Upon receiving the *PCEauth-resp(2)* message, the source can evaluate the robustness of

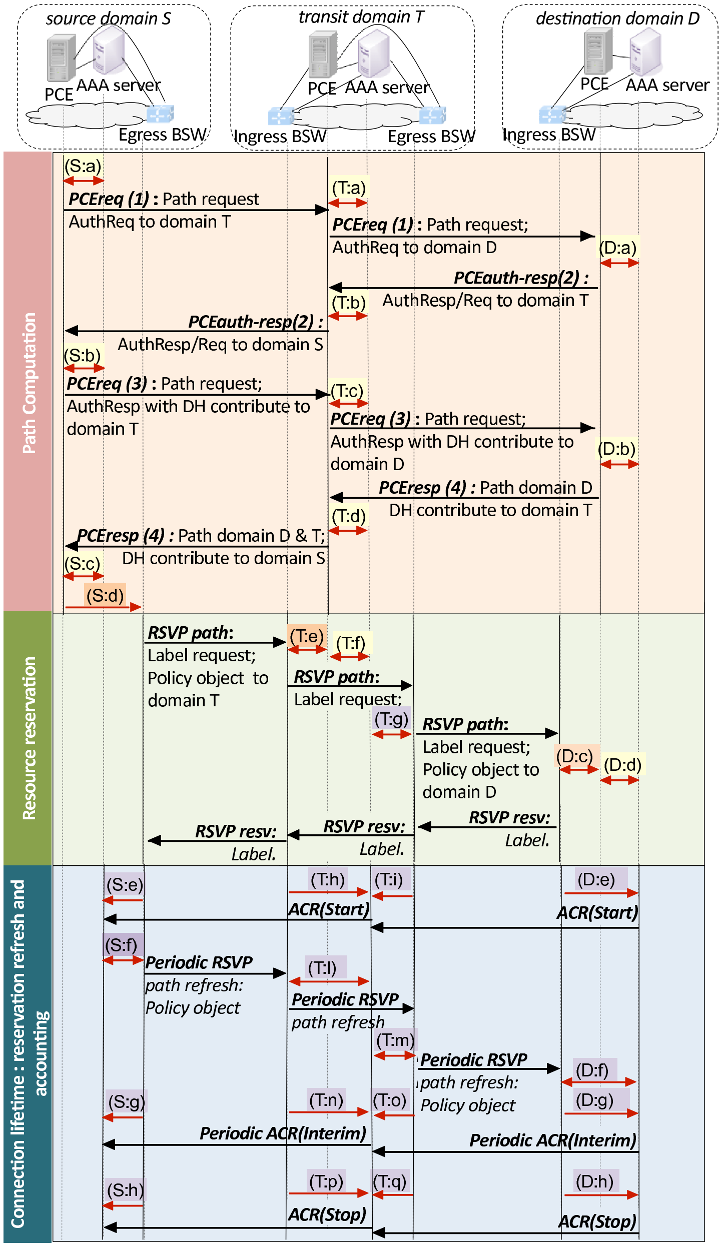

(a)

Intra-domain signaling for path computation with security features	
(S:a)	-Provisioning of AuthReq for domain T
(T:a)	-Computation of AuthResp/req for domain S -Provisioning of AuthReq for domain D
(D:a)	-Provisioning of AuthResp/req for domain T
(T:b)	-Verification of AuthResp from domain D and computation of DH contribute for domain D -Provisioning of AuthResp/req for domain S
(S:b)	-Verification of AuthResp/req of domain T and provisioning AuthResp with DH contribute for domain T
(T:c)	-Verification of AuthResp from domain S, computation of DH contribute and session key derivation for domain S -Provisioning of domain S authorization profile -Provisioning of AuthResp with DH contribute for domain D
(D:b)	-Verification of AuthResp from domain T, provisioning of DH contribute and computaion of DH key with domain T -Provisioning of domain T authorization profile
(T:d)	-Computation of DH key with domain D -Provisioning of DH contribute for domain S
(S:c)	-Computation of DH key with domain T -Provisioning of policy object for domain T
(S:d)	-Provisioning of computed path and policy object for domain T

(b)

Intra-domain signaling for resource reservation with security features	
(T:e)	-Provisioning of path-key & policy object
(T:f)	-Verification of policy object
(T:g)	-Provisioning of policy object for domain D
(D:c)	As T:e
(D:d)	-As T:f

(c)

Intra-domain signaling for connection lifetime functions with security features	
(S:e), (T:h), (T:i), (D:e)	-Notification of starting of metering
(S:f), (T:m)	-Provisioning of policy objects for RSVP refresh
(T:l), (D:f)	-Verification of policy object
(S:g), (T:n), (T:0), (D:g)	-Provisioning of accounting interim records
(S:h), (T:p), (T:q), (D:h)	-Notification of ending of metering

(d)

🟨 PCE server -AAA server interface

🟧 PCE-BSW interface

🟪 BSW-AAA server interface

Fig. 4. (a) Inter-domain signaling for path computation, reservation and connection lifetime procedures; (b) Intra-domain signaling for computation; (c) intra-domain signaling for reservation; (d) intra-domain signaling for connection lifetime procedures

the domain chain and can decide to terminate/continue path computation procedures based on its security requirements. We assume here that authenticated domains can be trusted and therefore will not activate path integrity attacks. Note however that while authentication of downstream domains is introduced for checking the integrity of the domain chain, authentication of upstream domains is essential for billing purposes. This ensures provisioning of resources to identified and authorized requesters. Therefore, if authentication of the upstream domain with its downstream neighbor fails on receiving the *PCEreq(3)*, the computation request is rejected with a *PCEError message*. If authentication of upstream domains is successful, the AAA servers provide the local PCE with information about the authorization profile of the authenticated upstream domain described in its SLA. Authorization policies are thus enforced during path computation.

3) Four-way signaling in the alliance model: Here, authentication request/responses and DH contributes are exchanged between the source and all domains along the path. To facilitate that exchange, we take advantage of the cascaded nature of PCE signaling: The source issues a PCEreq(1) message which contains an authentication request addressed to all downstream domains along the inter-domain path. Each domain issues authentication requests and responses that reach the source using the PCEauth-resp(2) message. Similarly, the PCEreq(3) and PCEresp(4) are used to carry authentication responses and DH contributes between source and all other domains. In this scenario, although the authentication content exchanged between source and a generic domain can be read by domains in the middle, they cannot perform impersonation, man-in-the-middle or replay attacks. This is due to the robustness of the implemented authentication mechanism.

B. *Resource reservation with security features*

As shown in Fig. 4 (see resource reservation phase), authentication and authorization of reservation request is performed via a RSVP policy object issued by each domain. In this object, each domain describes authentication and authorization assertions that include the path-key received by the downstream domain during path computation. The path-key acts as correlation token for authorization between computation and reservation procedures. This is because the path-key is resolved in the computed path by each domain [26] and therefore setup is enforced on resources computed according to the requester authorization profile. For authentication purposes, each domain pair can decide to use either SLA keys or DH keys. DH keys, if used provide the advantage of lower computation cost of the cryptographic algorithms compared with the cost for encryption with asymmetric SLA keys, and enforce protection against encryption attacks on symmetric SLA keys during the reservation process. Utilization of DH keys also ensures that the entity asking for reservation is the same that was authenticated during computation as authentication is made with keys derived during the computation phase. In our model, not using DH keys means *trusting* each authenticated domain asking for reservation, since there is no mechanism to verify if computation was made for the same. We will later provide results to quantify the higher overhead of the asymmetric-based authentication mechanism in terms of computation cost and its effect on signaling delay. This to support the decision of domains with regard to key usage.

1) Policy object computation and handling: The format of the policy object is presented in (4) and (5). Keyed hashes are used to compute the policy object with symmetric SLA keys or DH keys as in (4). Digital signatures are in stead used for authentication with asymmetric SLAs keys as in (5).

$$\text{Policy object}_{Sym/DH-key}\text{= [DomainID, sessionID, Hash(path-key}||\text{session-key/SLAkey}||\text{sessionID}||\text{DomainID)]} \quad (4)$$

$$\text{Policy object}_{Asym-key}\text{= [DomainID, sessionID, Sig(path-key}||\text{sessionID}||\text{DomainID)]} \quad (5)$$

Regardless of the format, the policy object (see Fig. 4, (S:c)-(D:d) steps), in the source domain, is provided to the egress BSW node along with the computed path. In a generic transit domain, the ingress BSW node provides the PCE with the policy object while asking for resolution of the path key (Step (T:e)). The PCE asks the AAA server to verify the policy object (Step (T:e)) and if verification is successful, provides the ingress BSW node with the extended intra-domain path. Once intra-domain reservation is performed, the egress BSW node asks the local AAA server for a policy object for authentication with the destination domain (Step (T:g)).

2) RSVP AA in alliance model: In the alliance peering model, the source domain includes policy objects in the RSVP reservation request towards all the domains along the path. If DH or symmetric SLA keys are used for authentication, the source will issue a policy object computed as in (4) for each domain along the path. If SLA asymmetric keys are used, as digital signatures can be verified by any domain along the path, the source issues a single policy object to authenticate with all the domains. Such object, which is called cumulative policy object, includes the path keys provided by each domain and is digitally signed. The decision to use a single cumulative policy object instead of a list of policy objects per each domain is to reduce the computational cost.

C. *Connection lifetime: reservation refresh and accounting*

We now investigate the mechanisms activated during the connection lifetime to secure signaling used to refresh reservation states and to provide accounting records related to the multi-domain services (see Fig. 4: connection lifetime phase).

1) Reservation refresh: The procedures with policy objects for authentication during resource reservation are also applied to *refresh* reservations with policy objects added to the periodic RSVP path refresh messages; this is in both cases of cascaded and alliance models. Minor modifications, however, are required in the format of the policy object in which the path-key element (see (4) and (5)) is substituted by the $path - key + n$ element. n is an integer value that represents the sequence number of the RSVP path reservation request and is used to verify the freshness of the policy object; it is included in the RSVP-TE message itself. The edge node keeps memory of the sequence number of any received message and only accepts reservation requests with a sequence number higher than that in the last processed message. Other changes are related to the intra-domain procedures for provisioning of the policy token: egress and ingress BSW nodes directly ask the AAA server for the policy object (Steps (S:e), (T:m)) or for verification of a received policy object (Steps (T:l), (D:f)). In the initial reservation this was made via the PCE that had to be contacted by the border switching nodes for path-key resolution. We recall that once the path is reserved, path key resolution from PCEs is not required for subsequent refresh messages.

2) Accounting: Accounting signaling for inter-domain paths is used for: 1) Start and Stop notifications for accounting and 2) facilitating intermediate accounting record generation. For a cascaded model, destination and transit domains provide accounting information for billing or QoS monitoring to their upstream customer domains. To this end, we adopt the standard accounting Diameter application [22], with incorporates message formats similar to its predecessor RADIUS. As shown in Fig. 4 (connection lifetime phase), the BSW nodes of transit and destination domains - which implement metering functions - inform their AAA server about the start of the accounting procedures (Steps (T:n),(T:o),(D:g)). These servers provide such information to the AAA servers of their customer domains using Diameter *ACR (Start)* messages. Similarly, the BSW nodes send periodic accounting information related to active connections (Steps (T:p),(T:q),(D:h)) to their AAA servers which are pushed to the customers AAA using Diameter *Interim* messages. *ACR(Stop)* messages are used to indicate the termination of the accounting procedures and are pushed in a similar fashion.

In the alliance model, provisioning of accounting records is the same as in the cascaded model with the difference that such records are provided to the source domain. In this case, the source has a customer-provider relationship with each domain along the path. As a consequence transit and destination domains issue, but do not receive accounting records.

V. Analysis

In this section, we develop analytical models for the mean signaling load and mean delay in the proposed architecture and compare these metrics with the standard mechanism. We consider generic multi-domain path provisioning scenarios with variable number of transit domains. We denote the sets of domains that originate and terminate inter-domain connections as $\mathbb{O}$ and $\mathbb{D}$ respectively. We also denote the set of services (e.g., video, VPNs) supported by an arbitrary domain k as $\mathbb{S}_k$. In the following analysis we focus on the first order statistics of the signaling load and assume reasonable signaling intervals of the order of few seconds during which all connection related signaling procedures will mostly likely finish.

A. Signaling load analysis

In our model, domains can act as source, transit and destination domains for different connections. We categorize connections that result in inter-domain signaling into two types, namely: (1) *source connections* which represent inter-domain connections generated from within the domain, (2) *transit connections* which represent inter-domain connections generated in other domains and transit or terminate in the domain under consideration. At an arbitrary domain k, the signaling due to *source connections* is primarily a function of the connection arrival rate λ from the supported services $\mathbb{S}_k$ (i.e, $\sum_{i\in\mathbb{S}_k}\lambda_{k,i}$). On the other hand, the signaling rate due to *transit connections* at domain k is equal to the sum of the product of the arrival rates from all other domains (i.e., $j \in \mathbb{O}$ where $j \neq k$), and the likelihood of passing through or terminating at domain k as $\sum_{\forall j\neq k}\sum_{i\in\mathbb{S}_k}\lambda_{j,i}P_T(k,j)$. The term $P_T(k,j) = Pr\{k \textit{ is transit or destination } \mid j \textit{ is source}\}$ can be calculated as the ratio of the number of all possible connections originating in domain j and passing through domain k (a.k.a., $N^*_{k,j}$) and the total number of possible connections from domain j to all domains (a.k.a., T_j) as $\frac{N^*_{k,j}}{T_j}$.

In the following, we study the signaling load at each network component for the standard and the proposed architecture in the three distinctive phases of path computation, resource reservation, and connection lifetime.

1) Path computation: In both the standard and the proposed frameworks, the signaling load arriving from interior switches at the PCE is equal to the source connections rate of $\sum_{i\in\mathbb{S}_k}\lambda_{k,i}$ from all services. The difference is in the inter-domain signaling between PCEs. In the standard framework, the signaling load on the interface between PCEs is proportional to the sum of the arrival rates from source and transit connections as,

$$\zeta^{\text{PCE-PCE}}_{\text{std}}(k) = \sum_{i\in\mathbb{S}_k}\lambda_{k,i} + \sum_{\forall j\in\mathbb{O},j\neq k}\sum_{i\in\mathbb{S}_j}\lambda_{j,i}P_T(k,j) \tag{6}$$

In the proposed method, the corresponding signaling load is double that of the standard model since it includes two rounds (i.e., 4-way handshake) as,

$$\zeta^{\text{PCE-PCE}}_{\text{prop}}(k) = 2\zeta^{\text{PCE-PCE}}_{\text{std}}(k) \tag{7}$$

In addition, on the interface between the PCE and the AAA server in the proposed model, the signaling load corresponds to three interactions with the AAA server for source connections (Steps (S:a),(S:c),(S:b) in Fig. 4). For transit connections that do not terminate in the domain, the AAA server is contacted for each received message in both upstream and downstream directions to establish security association with both the neighboring domains. Therefore four interactions are incurred in this case (Steps (T:a)-(T:d) in Fig. 4). For connections that terminate in the domain under consideration, two interactions with the AAA server are required (Steps (D:a),(D:b), in Fig. 4). Thus, the signaling load on the PCE-AAA interface is,

$$\begin{aligned}\zeta^{\text{PCE-AAA}}_{\text{prop}}(k) = 3\sum_{i\in\mathbb{S}_k}\lambda_{k,i} + \\ 4\sum_{\forall j\in\mathbb{O},j\neq k}\sum_{i\in\mathbb{S}_j}\lambda_{j,i}(P_T(k,j) - P_{dst}(k,j)) \\ + 2\sum_{\forall j\in\mathbb{O},j\neq k}\sum_{i\in\mathbb{S}_j}\lambda_{j,i}P_{dst}(k,j)\end{aligned} \tag{8}$$

Where $P_{dst}(k,j)$ represents the probability of domain k being the destination domain.

2) Resource reservation: Path reservation involves the use of RSVP signaling to reserve a computed path. Our framework almost keeps the signaling load of the RSVP signaling between switches identical to that of the standard framework. In our analysis, we focus on the RSVP signaling at border switches, ρ, as we are mostly interested in inter-domain scenarios. Similar to (6), ρ is given as the sum of the proportions of connections due to source connections leaving from the border switch and those entering and leaving from the border switch for transit connections, as,

$$\begin{aligned}\rho^{\text{BSW}}_{\text{std}}(\nu,k) = \rho^{\text{BSW}}_{\text{prop}}(\nu,k) = \sum_{i\in\mathbb{S}_k}\lambda_{k,i}P_{eg}(\nu,k) \\ + \sum_{\forall j\in\mathbb{O},j\neq k}\sum_{i\in\mathbb{S}_j}\lambda_{j,i}P_T(k,j)\left(P_{in}(\nu,k,j) + P_{eg}(\nu,k,j)\right)\end{aligned} \tag{9}$$

The term $P_{eg}(\nu,k)$ denotes the likelihood of leaving from border switch ν (i.e., the switch under consideration). The terms $P_{in}(\nu,k,j)$ and $P_{eg}(\nu,k,j)$ denote entering and leaving from border switch ν in domain k given that the source domain is j. These terms are estimated similar to $P_T(k,j)$ but for each BSW ν in transit domain k. Note however that $P_{eg}(\nu,k,j)$ excludes connections that terminate in domain k. The signaling load between the border switch and the PCE is the same for the standard and the proposed frameworks ($\rho^{\text{BSW-PCE}}_{\text{std}}(\nu,k) = \rho^{\text{BSW-PCE}}_{\text{prop}}(\nu,k)$) and is given as,

$$\rho^{\text{BSW-PCE}}_{\text{std}}(\nu,k) = \sum_{\forall j\in\mathbb{O},j\neq k}\sum_{i\in\mathbb{S}_j}\lambda_{j,i}P_T(k,j)P_{in}(\nu,k,j) \tag{10}$$

Due to the added security in the proposed method, when the ingress border switch contacts the PCE for path key resolution (Steps (T:e),(D:c) in Fig. 4), the PCE contacts the AAA server to verify the policy object (Steps (T:f),(D:d) in Fig. 4) in the incoming RSVP request.

$$\rho^{\text{PCE-AAA}}_{\text{prop}}(k) = \sum_{\forall j \in \mathbb{O}, j \neq k} \sum_{i \in \mathbb{S}_j} \lambda_{j,i} P_T(k,j) \tag{11}$$

Note that the PCE is not contacted in the source domain as the token is already created at the end of path computation.

3) During connection lifetime: During the connection lifetime, RSVP refresh messages are sent periodically every Δ_R time units to ensure the liveliness of the connection. These messages cause border switches to contact the AAA server to include a fresh policy object (Steps (S:f),(T:e) in Fig. 4). Since the RSVP signaling load depends on the mean number of RSVP refreshes during a service connection and since the service connection may have any arbitrary distribution (i.e., short or long tailed), we need to investigate their number in general. To this end, we define the function $\psi(E_{S_i}, \Delta) = E\left[\lfloor \frac{S_i}{\Delta} \rfloor\right]$ where E_{S_i} represents the mean duration of a connection and Δ is an arbitrary constant. Thus, the mean load of RSVP refresh messages at a given border switch is given by the product of their numbers from source and transit connections as,

$$\begin{aligned}\gamma^{\text{BSW}}_{\text{std}}(\nu,k) = \gamma^{\text{BSW}}_{\text{prop}}(\nu,k) &= \sum_{i \in \mathbb{S}_k} \lambda_{k,i} \psi(E_{S_i}, \Delta_{R_i}) P_{eg}(\nu,k) \\ &+ \sum_{\forall j \in \mathbb{O}, j \neq k} \sum_{i \in \mathbb{S}_j} \lambda_{j,i} P_T(k,j) \psi(E_{S_i}, \Delta_{R_i}) \\ &\times (P_{in}(\nu,k,j) + P_{eg}(\nu,k,j)) \end{aligned} \tag{12}$$

While the standard mechanism only entails RSVP signaling, the proposed framework requires that border switches contact the AAA server to secure RSVP refresh messages by including policy objects and to report accounting interim records. The signaling load from border switch ν that pertains to securing RSVP messages is denoted as $\gamma^{\text{BSW-AAA}}_{\text{prop, Auth}}(\nu,k)$ and includes three components: (i) one from source connections as the egress border switch contacts the AAA to obtain the policy object for outgoing RSVP refresh messages (Step (S:f) in Fig. 4), (ii) one from transit/terminating connections as the ingress border switch contacts the AAA for verification of the incoming policy object (Steps (T:i) (D:f) in Fig. 4), (iii) and the last from transit connections as the egress border switch contacts the AAA to get the policy object for the next domain hop (Step (T:g) in Fig. 4). Thus, the authentication signaling load from border switch ν to the AAA server is given as,

$$\begin{aligned}\gamma^{\text{BSW-AAA}}_{\text{prop, Auth}}(\nu,k) &= \sum_{i \in \mathbb{S}_k} \lambda_{k,i} \psi(E_{S_i}, \Delta_{R_i}) P_{eg}(\nu,k) \qquad (13) \\ &+ \sum_{\forall j \in \mathbb{O}, j \neq k} \sum_{i \in \mathbb{S}_j} \Bigg[\lambda_{j,i} \psi(E_{S_i}, \Delta_{R_i}) P_T(k,j) \times \\ &\Big(P_{in}(\nu,k,j) + P_{eg}(\nu,k,j)\Big) \Bigg] \end{aligned}$$

For accounting signaling, we consider the load on the interface between border switches and the domains' AAA server and the inter-domain interface between AAA servers of different domains to report metering information to downstream domains. We assume that accounting records are only reported by egress border switches in the source domain, by ingress border switches in the destination domain, and by both ingress and egress border switches in transit domains. Since accounting messages consist of start, interim, and stop records, their number during a service connection ($i \in \mathbb{S}$) is given as,

$$\Omega(E_{S_i}, \Delta_{T_i}) = 2 + \psi(E_{S_i}, \Delta_{T_i}) \tag{14}$$

where the '2' indicates the count of accounting start and stop messages, and the function $\psi(E_{S_i}, \Delta_{T_i})$ gives the number of accounting interim reports during the connection every accounting interim interval (Δ_{T_i}). Hence, signaling rate from a border switch ν to the AAA server, denoted as $\gamma^{\text{BSW-AAA}}_{\text{prop, Acct}}(\nu)$, is given as,

$$\begin{aligned}\gamma^{\text{BSW-AAA}}_{\text{prop, Acct}}(\nu,k) &= \sum_{i \in \mathbb{S}_k} \lambda_{k,i} \Omega(E_{S_i}, \Delta_{T_i}) P_{eg}(\nu,k) + \\ &\sum_{\forall j \in \mathbb{O}, j \neq k} \sum_{i \in \mathbb{S}_j} \Bigg[\lambda_{j,i} \Omega(E_{S_i}, \Delta_{T_i}) P_T(k,j) \times \\ &\Big(P_{in}(\nu,k,j) + P_{eg}(\nu,k,j)\Big) \Bigg] \end{aligned} \tag{15}$$

In addition, the AAA server receives and sends accounting records to the AAA servers in its downstream and upstream neighbor domains respectively. Hence, the signaling load for receiving accounting records from downstream domains is,

$$\begin{aligned}\gamma^{\text{AAA-AAA}}_{\text{prop, Recv Acct}}(k) &= \sum_{i \in \mathbb{S}_k} \lambda_{k,i} \Omega(E_{S_i}, \Delta_{T_i}) \qquad (16) \\ &+ \sum_{\forall j \in \mathbb{O}, j \neq k} \sum_{i \in \mathbb{S}_j} \lambda_{j,i} (P_T(k,j) - P_{dst}(k,j)) \Omega(E_{S_i}, \Delta_{T_i}) \end{aligned}$$

Since AAA servers only send/proxy accounting records to upstream domains, only transit and terminating connections are considered. Hence, the signaling rate is given as,

$$\gamma^{\text{AAA-AAA}}_{\text{prop, send Acct}}(k) = \sum_{\forall j \in \mathbb{O}, j \neq k} \sum_{i \in \mathbb{S}_j} \lambda_{j,i} P_T(k,j) \Omega(E_{S_i}, \Delta_{T_i}) \tag{17}$$

4) The Derivation of $\psi(E_{S_i}, \Delta)$: The mean number of RSVP refresh messages and accounting interims is defined by the mean of the floor of the ratio connection duration S and the interval Δ which represents either the refresh time or the accounting interim interval. Let J be a random variable defined as $J = \lfloor \frac{S}{\Delta} \rfloor$. Then the probability density function of J is given as [20],

$$\begin{aligned} f_J(j) &= \int_{j\Delta}^{(j+1)\Delta} f_S(x)\, dx = F_S((j+1)\Delta) - F_S(j\Delta) \\ &= \bar{F}_S(j\Delta) - \bar{F}_S((j+1)\Delta) \end{aligned} \tag{18}$$

Using (18), the mean value of J (i.e., $E[\lfloor \frac{S_i}{\Delta_{T_i}} \rfloor]$) is given as,

$$\begin{aligned} E[\lfloor \frac{S}{\Delta} \rfloor] &= \sum_{j=0}^{\infty} j f_J(j) = \sum_{j=0}^{\infty} j \left[\bar{F}_S(j\Delta) - \bar{F}_S((j+1)\Delta) \right] \\ &= \sum_{j=1}^{\infty} j \bar{F}_S(j\Delta) - \sum_{j=1}^{\infty} (j-1) \bar{F}_S(j\Delta) = \sum_{j=1}^{\infty} \bar{F}_S(j\Delta) \end{aligned} \tag{19}$$

For exponentially distributed session durations, the infinite sum in (19) can be written in closed form as $\left(e^{\frac{\Delta}{E[S]}} - 1\right)^{-1}$.

B. Delay Analysis

In this section, we derive the signaling delay for time sensitive procedures - path computation and path reservation - for the proposed method and compare its performance with the standard procedures. Since transmission delays are insignificant for the message sizes and interfaces under consideration, we only consider processing and propagation delays.

1) Path Computation: For the standard PCE mechanism, the processing delay of PCE requests is attributed to their end-to-end propagation delay (D_{prop}), packet processing delays within the servers including decapsulation, validation and forwarding (D_{proc}), and path computation delay at the PCE (D_{cmp}). If we denote the number of hops on a path between a source and destination pair i, j as $h_{(i,j)}$ ($h_{(i,j)}$ = *number of domains in the path from* i *to* j -1), then the total PCE signaling delay is given as,

$$D^{\text{PCE,std}}_{(i,j)} = (h_{(i,j)} + 1)(2D_{\text{proc}} + D_{cmp}) + 2D_{prop(i,j)} \quad (20)$$

Notice that the mean delay $E[D^{\text{PCE}}_{(i,j)}]$ is given by averaging over all possible sources and destinations.

For the proposed signaling model, the total delay is,

$$\begin{aligned} D^{\text{PCE,prop}}_{(i,j)} &= (h_{(i,j)} + 1)(4D_{\text{proc}} + D_{cmp}) + 4D_{prop(i,j)} \\ &+ 2(2h_{(i,j)} + 3)D_{\text{PCE-AAA}} + D_{auth} + 4D_{\text{exp}} \quad (21) \end{aligned}$$

where the processing and propagation delays are double of those in the standard solution due to the extra round in the 4-way handshake. The delay also includes the propagation delay between PCE and AAA server ($D_{\text{PCE-AAA}}$), the delay for generating DH contributes due to exponentiation (D_{exp}), and the delay for processing of authentication contents (D_{auth}). The latter delay (D_{auth}) depends on whether symmetric or asymmetric keys are used for authentication. If symmetric key authentication is used, $D_{auth} = 6h_{(i,j)}D_{\text{hash}}$ as hashes are generated both for inclusion in authentication responses and for verification of authentication responses and DH objects from downstream domains. If asymmetric authentication is used, $D_{auth} = 3h_{(i,j)}\,(D_{\text{sig}} + D_{\text{ver}})$ as signatures are generated and verified for authentication responses and DH objects.

2) Resource reservation: To evaluate the RSVP signaling delay in *the standard mechanism*, we sum the intra-domain and inter-domain reservation delays. For the intra-domain delay component, we assume that for a given domain k, the average *intra-domain* router hops is η_k and that the reservation delay per router is given by D_{Res}. For the inter-domain component, we consider latencies pertaining to the propagation delay between ingress border switches and the PCE in transit domain ($D^{\text{BSW-PCE}}$), PCE request processing in the PCE ($D^{\text{PCE}}_{\text{proc}}$), and path key resolution by the PCE (D^{PCE}_{Key}). Let us denote the propagation delay between interior routers in domain k as $D^{\text{Intra}}_{\text{prop}}(k)$, and the inter-domain path between domain i to domain j by the set $\mathbb{P}_{i,j} = \{i, ..., j\}$. Then, given a hop count of ($h_{i,j}$) between domains i and j, the RSVP signaling delay in the standard PCE framework is obtained as,

$$\begin{aligned} D^{\text{RSVP,std}}_{(i,j)} &= h_{(i,j)}(2D^{\text{BSW-PCE}} + D^{\text{PCE}}_{Key} + D^{\text{PCE}}_{proc}) \quad (22) \\ &+ \sum_{k \in \mathbb{P}_{i,j}} 2D^{\text{Intra}}_{\text{prop}}(k) + (\eta_k + 1)\,(D_{\text{Res}} + 2D_{\text{Proc}}) \end{aligned}$$

In the *proposed method*, the AAA signaling delay should be added to the delay in (22) and proportionally to the number of hops in the path between domains i, j (i.e, $h_{(i,j)}$) as,

$$\begin{aligned} D^{\text{RSVP, prop}}_{(i,j)} &= D^{\text{RSVP,std}}_{(i,j)} \\ &+ h_{(i,j)}(2D_{\text{PCE-AAA}} + D_{\text{proc}}) + D_x \quad (23) \end{aligned}$$

The term D_x depends on the type of the used authentication key. If symmetric or DH keys are used, then D_x corresponds to the delay incurred by comparing the hash from the downstream domain and computing hash for the upstream domain except at the destination, i.e., $D_x = 2(h_{(i,j)} - 1)D_{\text{hash}} + D_{\text{hash}}$. If asymmetric keys are used instead, the term D_x corresponds to the delay incurred for verifying the signature of the downstream domain hop and for signing the reservation request for the upstream domain hop except at the destination, i.e., $D_x = (h_{(i,j)} - 1)(D_{\text{ver}} + D_{\text{sig}}) + D_{\text{ver}}$.

VI. Numerical Results

In this section, we study the performance of the proposed signaling framework with respect to the standard PCE mechanism. In our evaluation, we use the BRITE topology generation libraries [29] to generate random domain topologies of sizes ranging from 120 to 200 domains. We use the Barabasi Albert (BA) algorithm to generate topologies that exhibit a hierarchical (power law) structure in terms of the number of peerings per domain (i.e., nodal degree). The BA algorithm balances structure and randomness, and generates topologies that match the observed domain-level structure of the internet [30]. For each generated topology, we categorize domains in the network into three types based on their nodal degrees: (i) *core*: degree $\geq$ 16; (ii) *intermediary*: 16 $>$ degree $>$ 8; (iii) *stub*: degree $\leq$ 8. This is similar to existing transport networks in which stub, intermediary and core domains represent metro ring networks, national backbone networks, and continental and transcontinental backbone networks respectively.

TABLE I
Topology, service, and protocol parameters

Network Parameters	
Domain Types	Core domains with degree $\geq$ 16 , intermediary domains with 8 $<$ degree $<$ 16, and stub domains with degree $\leq$ 8
Intra domain Parameters	Stub: typical metro rings with 10 switches and diameter of 40 kms. Intermediary: Germany17 [32] topology. Core: NSFNet backbone.
Services and Protocol Parameters	
Types	Broadcast Video (BV), Content Distribution (CD), Leased Line Service (LS), Whole Sale Backhaul (WB)
Arrival Rates	BV (2.9/hr), CD (14.5/hr), LS(0.2/hr), WB(0.1/hr)
Mean Duration	BV (1 hr), CD (10 min), LS(15 days), WB(1 month)
Protocol Parameters	Δ_R = 10 min (all services), Δ_T = {BV (30 min), CD (5 min), LS(12 hr), WB(12 hr)}
Computational and Roundtrip Delays	
Cryptographic Delays	D_{ver} = 369μs, D_{sig} = 5.31 ms, D_{Hash} = 121μs, D_{exp} = 962μs
Processing Delays	D_{cmp} = 0.5ms, D_{proc} = 1 ms, D_{Resv} = 1 ms
Propagation Delays	$D^{\text{PCE-AAA}}$ = $D^{\text{PCE-BSW}}$ = $D^{\text{BSW-AAA}}$ = 5 ms

We use the shortest path algorithm to compute the inter-domain paths using equal weights for the inter-domain links.

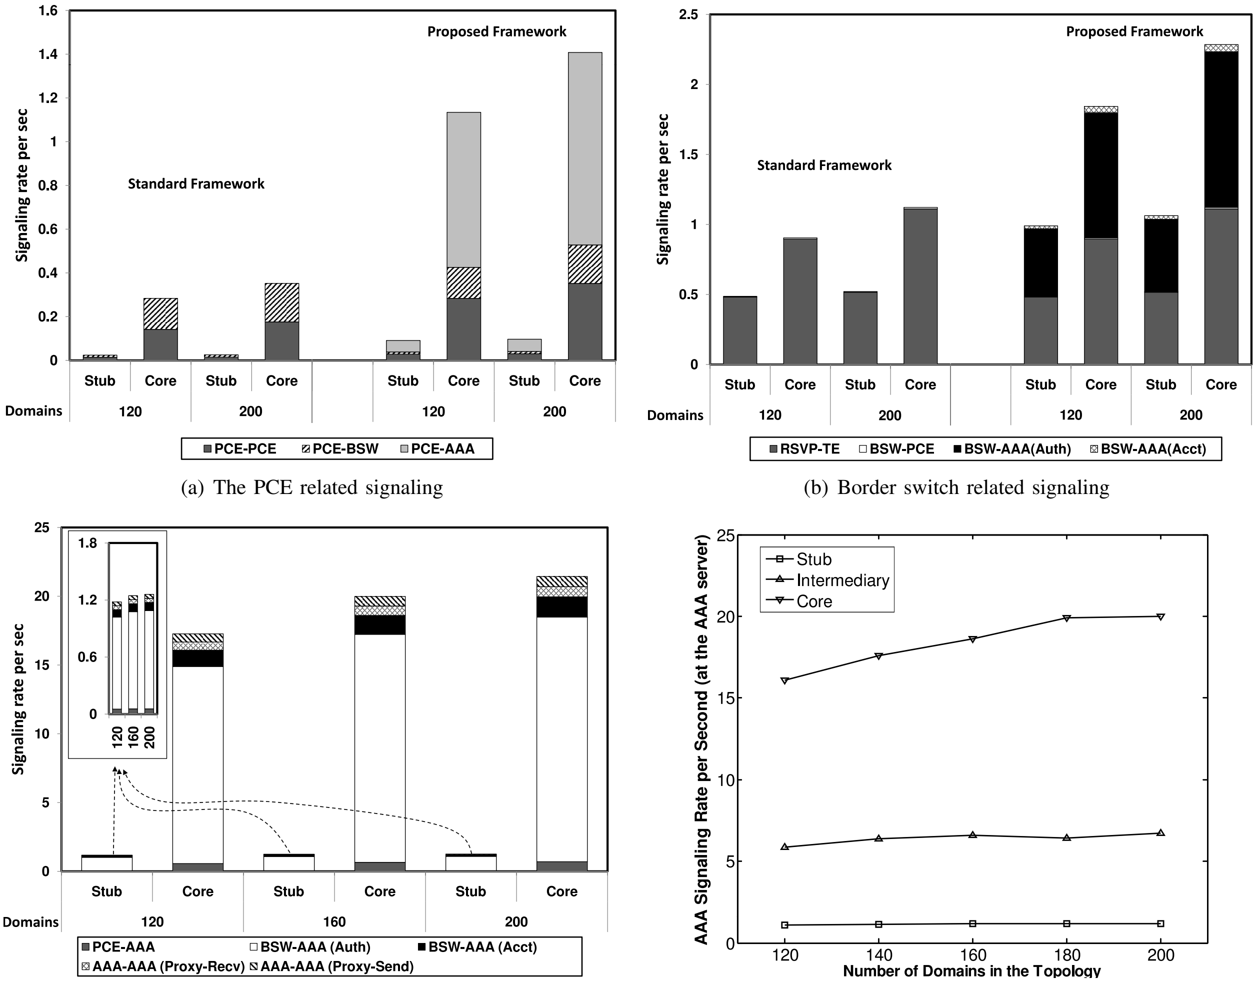

(a) The PCE related signaling

(b) Border switch related signaling

(c) AAA related signaling

(d) AAA signaling vs topology size

Fig. 5. Scalability and signaling load on interfaces [200 random network topologies, 90% confidence levels]. In Fig. 5(b), the BSW-PCE signaling barely shows between the RSVP-TE and the BSW-AAA (Auth) bars.

We use this data to estimate parameters of the average hop count, propagation delay for each topology, and the probabilities of transiting domains (P_T (see Section V-A)). For the topologies investigated, we have verified that the average domain hop count is around 3 and ranges from 2 to 5. Furthermore, the average hop counts are 3.3, 2.8, and 2.4 for connections initiated from stub, intermediary, and backbone domains respectively. For the domain-internal parameters (i.e., propagation delay and mean hop count), we assume typical topologies for each of the three domain types, summarized in Table I. In our analysis, we adopt four service types from the ETNA project [3] including broadcast video, content distribution, leased line service, and whole sale backhaul (Table I). We assume that core domains do not generate traffic towards other domains and only offer transit services.

To obtain the cryptographic processing delays, which are shown in Table I, we used the OPENSSL library (version 0.9.8) and C compiler (GCC 4.3.2) under Linux Fedora Core 10 [31], and averaged delays over 1000 runs. The measurements were carried out on a HP ProLiant DL160 Generation 5 Rack server (Intel Xenon Quad Core 5400 series 3.00 GHz processor and 16 GB of RAM and 12MB L2 Cache).

In the following subsections, we first show basic results on signaling load and delay for deployments implementing cascaded peerings; Afterwards, we study the impact of protocol and service parameters on system performance, validate the analytical framework by event-driven simulations, and conclude the section by showing results relevant to peering models.

A. Signaling Load

In this section, we consider the signaling load for connection provisioning at the protocol interfaces of the network components in the system. We start with the load on the PCE and then discuss similar results for the BSW and the AAA server.

1) PCE load: In Fig. 5(a), we show the signaling load on the different interfaces of the PCE server for the standard PCE mechanism and the proposed framework. From Fig. 5(a), we see that the *total* signaling load on all interfaces pertaining to our method is about 3-4 times that of the standard mechanism. Specifically, the load on the PCE-PCE interface is twice that of the standard mechanism due to the additional round trip signaling of the 4-way handshake and is the same on the PCE-BSW interface as the proposed and standard mechanisms follow a similar signaling flow. The PCE-AAA

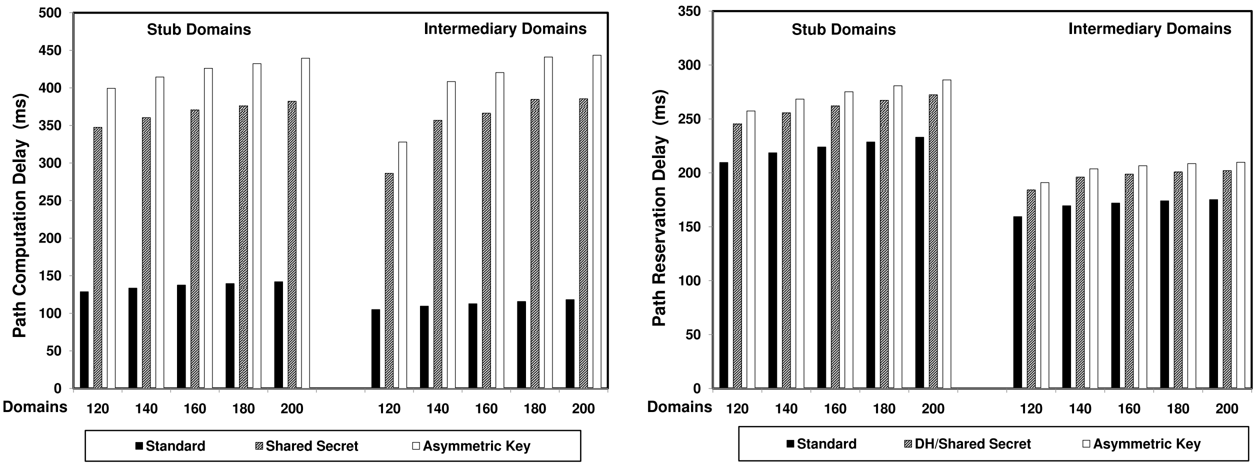

(a) Mean signaling delay during path computation

(b) Mean signaling delay during path reservation

Fig. 6. Signaling delay comparison [200 random network topologies, 90% confidence levels]

interface carries the largest signaling traffic as the PCE server triggers the AAA server for almost each PCE message during path computation and for each request to resolve path-keys during path reservation. We also observe that the signaling load on core domains is significantly larger than that at stub domains (about 10 to 15 times) which is a consequence of having hierarchical domain topology structures (as in practice). Interestingly, increasing the topology size from 120 to 200 does not impact the signaling load significantly due to the hierarchical structure of the network.

2) RSVP border switching node load: In Fig. 5(b), we analyze the signaling load on the interfaces of the border switch (BSW) with other border switches (for RSVP signaling), the PCE (for path-key resolution signaling), and the AAA server (to provision and verify policy objects for RSVP refresh messages and to report accounting records). First, we see that RSVP signaling is almost equal in the standard and the proposed mechanisms for both stub and core domains. The signaling between the BSW and the PCE triggered during connection setup is negligible compared with the others as it only happens once during the lifetime of the connection (barely shows). The BSW-AAA signaling for authentication (BSW-AAA (auth)) is less than or equal to that of the RSVP-TE signaling because border switches trigger the AAA server for initial and refresh RSVP messages (compare black to gray bars). The accounting signaling load is lower than of authentication because the accounting interim intervals, which determine the accounting rate, are much larger than RSVP refresh intervals, which determines the authentication rate.

3) AAA server load: In Fig. 5(c), we show the AAA signaling observed at each interface of the AAA server. We observe that the authentication signaling on the interface with the BSW constitutes the majority of the load on the AAA server. This is because the frequent RSVP refresh messages are secured by the AAA server. The signaling on the PCE-AAA interface is low as it is only triggered during connection setup - to secure path computation and path reservation signaling. The BSW-AAA interface receives signaling from all border switches in the domain and carries less load than that pertaining to authentication traffic due to the lower frequency of accounting records. The inter-domain AAA-AAA interface allows the domain to receive accounting records from downstream neighbors (a.k.a., Proxy-Recv) and send accounting reports to its upstream neighbors (a.k.a., Proxy-Send). Due to the cascaded peering in our case study, both streams are almost equal. We also see that these trends apply for a range of topology sizes.

In Fig. 5(d), we investigate the impact of the number of domains per topology on the AAA signaling load in all domain types (stub, intermediary, and core). We see that the topology size barely impacts the signaling load on stub domains as they are unlikely to transit traffic. On the other hand, intermediary and core domains exhibit a slight linear growth of the signaling load as function of the number of domains.

The results discussed above, show that the scalability of the proposed scheme is similar to the standard method for path computation and reservation. Even though the security offered by the proposed framework comes on the cost of higher signaling load, the signaling rate with the AAA server is still quite modest (below 25 exchanges/sec). This is significantly low load for currently available commercial AAA servers which are able to handle around one thousand exchanges/sec. Hence, we conclude that introducing AAA mechanisms to the standard PCE framework is unlikely to pose scalability issues.

4) Validation Via Simulation: Table II shows the verification of the analytical model by simulations. To this end, we use a custom event-driven simulator in JAVA, which emulates the process signaling as described in the paper. Note that the diversity in the arrival rates for different services as shown in Table I implies that simulating a sufficient number of connections for all services in an event-driven simulation would require an extremely long time. We therefore compare analytical and simulation results for a randomly generated 120 domain topology with only the Content Distribution (CD) service in operation in the network. We also do not simulate the actual intra-domain signaling to reduce the number of events and make the simulation manageable, and instead use the same average values for hop-count and signaling delays inside the domain as those used for the analytical model.

The event-driven simulation was allowed to run over a time unit of one year, and measurements for average signaling rates were taken over a duration of 1 day. We saw a very small variation ($< 2\%$) in the average signaling rates, and it was seen that the signaling rates measured via simulation were also very close to the signaling rates as obtained by analysis. Table II presents a comparison of the analytical and the simulation signaling load as the BSW-AAA interface when varying the accounting interim interval. As can be seen, the variation in the values is very small, indicating that the analytical model can provide us with a good estimate for the signaling load. Note however that we did not take into consideration blocking of connections in the simulation model, which can affect the routing of connections and in turn affect the probabilities used by the analytical model.

TABLE II
SIGNALING LOAD AT THE BSW-AAA INTERFACE AS A FUNCTION OF ACCOUNTING INTERIM INTERVAL (Δ_T). [AN: ANALYTICAL, SIM: SIMULATION, NETWORK SIZE = 120 DOMAINS]

Msg/sec	**0.1**Δ_T	**0.3**Δ_T	**0.5**Δ_T	**0.7**Δ_T	**0.9**Δ_T	**1.1**Δ_T
AN	13.01	4.89	3.27	2.82	2.57	2.36
SIM	12.73	5.23	3.47	3.13	2.49	1.92

B. Signaling Delay

In this section, we investigate the signaling delay for path computation and reservation which impacts the connection setup time. We consider the different authentication models based on symmetric and asymmetric keys. Fig. 6(a), shows the end-to-end delay for connections originated in stub and intermediary domains. We do not consider core domains as we assumed that core domains do not originate connections but only transit traffic (e.g., Tier-1). If core domains were considered, they would exhibit shortest delays due to their high degree of connectivity leading to low propagation delays for their originated traffic. In the same vein, traffic generated in the intermediary domains generally incur lower signaling delay than stub domains do as intermediary domains have higher nodal degrees. In addition, the delay in the proposed method is about 3 to 4 times that of the standard scheme. This is primarily due to the extra rounds with the AAA server rather than the cryptographic computations. The variation of the signaling delay depends on the implemented cryptographic scheme for securing the signaling messages (i.e., symmetric or asymmetric). For path reservation (see Fig. 6(b)), we also observe that intermediary domains incur lower path reservation delays than stub domains. Interestingly, we see that when implementing the proposed scheme, the delay increases approximately by 20-35% depending on whether symmetric or asymmetric key schemes are used. Clearly, the path computation phase incurs a higher delay penalty than path reservation - due to the higher number of interactions with the AAA server.

Since the round trip delay between the PCE and the AAA server can play a significant role in the delay performance of connection setup, the results suggest that implementations may consider collocating the AAA functions within the PCE platform to reduce the overall connection setup latency. The results in Table III show that co-location reduces the *path computation delay* to around 2-2.5 times that of the standard mechanism. It also makes the delay for the procedures to resolve the path-key and verify the policy object during resource reservation almost identical to that of the standard method. It should be noted that co-locating AAA functions within the PCE system does not lead to unacceptable loads at the hosting platform. In fact, as seen from Figs.5(a),5(d), the signaling load at the PCE is in the vicinity of 1 exchange/sec and that at the AAA is below 25 exchange/sec.

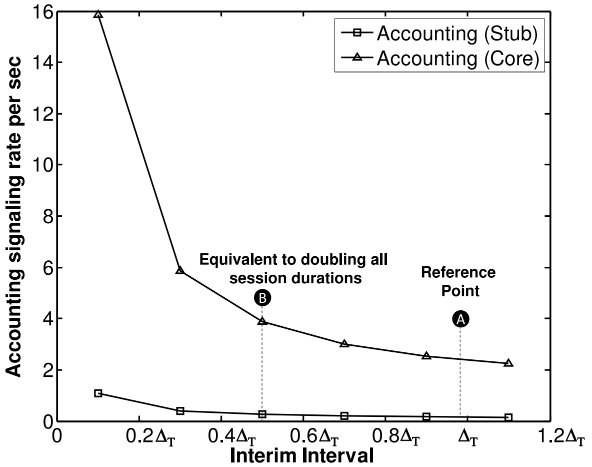

Fig. 7. Signaling load at the BSW-AAA interface as function of the accounting interim interval [120 domains, 200 random topologies, 90% confidence levels]

TABLE III
THE DELAY WHEN THE PCE AND THE AAA SERVER ARE COLLOCATED IN THE SAME PLATFORM [90% CONFIDENCE, 200 RANDOM TOPOLOGIES]

	Stub			**Intermediary**		
Number of Domains	120	160	200	120	160	200
Path Computation Phase						
Standard	129	137	142	105	113	118
Shared Secret	254	272	281	204	220	231
Asymmetric key	306	327	339	245	264	278
Path Reservation Phase						
Standard	210	224	233	159	172	175
Shared Secret	213	228	237	162	175	178
Asymmetric key	226	241	251	169	182	186

C. Impact of protocol parameter settings and service statistics

In this section, we first evaluate how protocol parameters including RSVP refresh and accounting interim intervals affect the signaling load on the protocol interfaces in the system. Then, we study an exemplary mix of two connections and investigate how fluctuations of arrival rate and connection duration of one service affects the load on the system. Finally, since the mean signaling load depends on the *distribution* of the connection duration rather than its average value (see (19)), we see whether knowing the distribution is necessary for the load analysis for shortly and heavily tailed connection durations.

1) Protocol parameters: In our study, we vary the interim intervals (Δ_T) for all services - which are defined in Table I - by scaling them all by a linear factor. As shown in Fig. 7, we see that increasing the accounting interim interval quickly reduces the accounting signaling load at the AAA server due to the reduced frequency of accounting messages from border switches. We see that when the interim interval is largely increased, the load does not change significantly since no

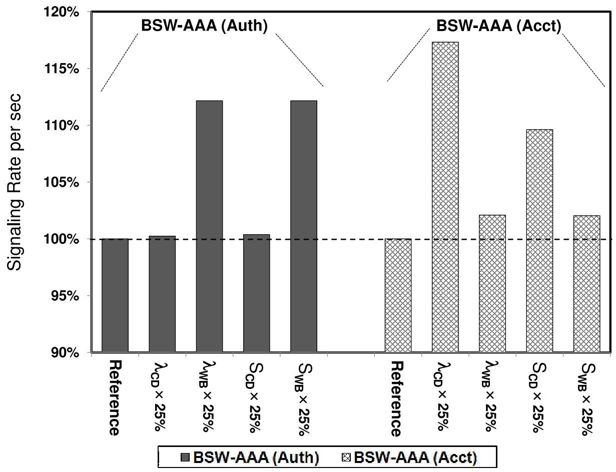

Fig. 8. Impact of arrival rate and session duration on AAA signaling load [120 domains, 200 random network topologies, 90% confidence levels]

interim messages are sent during the connection lifetime. The opposite is true when the interim interval is reduced. From (19), we know that it is the ratio of the connection duration and the session duration that determines the signaling rate rather than their individual values - approximately doubling the mean session duration leads to the same signaling rate of when the interim interval is halved. This implies that one should choose interim settings for operation in the "linear" portion of the curve (e.g., say between $0.5\Delta_T$-$1.0\Delta_T$) to maintain stable accounting signaling rate as the connection duration statistics fluctuate over time. For instance, if the mean connection duration doubles for all services over time, the accounting signaling load changes by less than 50% (see points A and B). The same reasoning applies for choosing the RSVP refresh intervals (not shown here for bevity).

2) Service parameters: Let us now get into more details on the impact of the perturbation of the arrival rates and session durations on the signaling load on the BSW-AAA interface for authentication and accounting traffic during the connection lifetime. In the analysis, we consider two services characterized by long and short durations, i.e., Content Distribution (CD) and the Wholesale Backhaul (WB). Arrival rates and durations are described in Table I. To see the impact of the arrival rate perturbations, we increase the arrival rate of each service by 25%. From Fig. 8, we observe that this increase barely impacts the total authentication signaling rate for CD services, but results in an increase in the accounting signaling rate. For the WB service, the opposite behavior is observed. The reason for such effect is that most of the authentication signaling traffic comes from long duration services due to the large number of RSVP-TE refreshes. On the other hand, for accounting signaling, short duration services produce considerable amount of accounting signaling compared to long duration services. This is because of the relatively high arrival rate of short session services and the relatively low frequency of interims during long duration connections in practice.

It is noteworthy to state that the behavior in Fig. 8 depends on the behavior of the $\psi(E_{S_i}, \Delta_{R_i})$ and $\psi(E_{S_i}, \Delta_{T_i})$ functions (see Eqs. (13), (15)). When the connection duration is very high compared to the interval under consideration, the ψ function becomes linear of the mean connection duration. This explains why increasing the arrival rate - a linear operation - leads to almost the same rate when increasing the connection duration for WB services. It also explains why the same increase is not observed for accounting signaling in CD services as the ratio of the $E_{S_{CD}}$ to $\Delta_{T_{CD}}$ is not high enough and hence behaves non-linearly.

3) Session statistics: We now evaluate the dependence of the signaling rate on the connection duration statistics. To this end, we estimate the signaling rate by assuming exponentially distributed connection durations and then compare it with estimates for connection durations with Erlang and lognormal probability distributions. We choose Erlang distributions as they are shortly tailed while lognormal distributions are heavily tailed distributions. In Table IV, we show the signaling load at the AAA interfaces with the PCE, BSWs, and other AAA servers. We see that for both Erlang and lognormal distributions, the estimation error (ϵ) of the signaling load with respect to the exponential estimate is below 2% for all domain types. This result is interesting since the exponential estimate only requires the knowledge of the mean connection duration to obtain the mean signaling rate (see below (19)), and is quite accurate regardless of the actual underlying distribution.

TABLE IV
THE EFFECT OF THE SESSION DISTRIBUTION [EX: EXPONENTIAL DIST., ER: ERLANG DIST. (COEFF. VARIATION = 0.5), AND LN: LOGNORMAL DIST. (COEFF. VARIATION = 3), NETWORK SIZE = 120 DOMAINS]

	Stub			**Intermediary**			**Core**		
	EX	$\epsilon\%$ **ER**	$\epsilon\%$ **LN**	**EX**	$\epsilon\%$ **ER**	$\epsilon\%$ **LN**	**EX**	$\epsilon\%$ **ER**	$\epsilon\%$ **LN**
PCE-AAA	**0.05**	0.00	0.00	**0.22**	0.00	0.00	**0.58**	0.00	0.00
BSW-AAA (Auth)	**0.97**	0.15	0.18	**5.12**	0.16	0.18	**14.70**	0.16	0.18
BSW-AAA (Acct)	**0.08**	1.12	1.35	**0.43**	1.05	1.46	**1.22**	1.04	1.46
AAA-AAA	**0.08**	1.12	1.47	**0.43**	1.05	1.46	**1.22**	1.04	1.47

D. The impact of the peering model

Table V shows the signaling load (message/sec) in the alliance model with respect to the cascaded model at the core domains. First, we see that deploying the alliance model does not alter the load on the PCE interfaces in general. This is because neither the number of messages nor the amount of cryptographic operations change during the path computation phase. Second, the alliance peering results in lower signaling load on the BSW-AAA interface as only ingress RSVP signaling is authenticated in transit domains due to their single peering with the source. The rest of the BSW interfaces incur similar signaling as the BSW-AAA interface (see Table V). Third, we observe that the AAA interfaces with other systems incur less load in the alliance model than the cascaded model. This is because only ingress RSVP signaling is secured by the AAA server. Also, AAA servers only need to send accounting messages to the source domain and never receive accounting messages for transit connections from its neighboring domains as in the cascaded case. Finally, as shown in Table VI, the delay performance (in milliseconds) of path computation and reservation phases is almost identical to that observed in the cascaded deployments and is only (slightly) smaller when asymmetric keys are used. This is because in alliance models the source domain issues a single authentication response to all the other domains using one signature, while in cascaded models each domain signs and verifies signatures for its upstream and downstream neighbors.

TABLE V
SIGNALING LOAD IN ALLIANCE (A) AND CASCADED DEPLOYMENTS (C) [200 DOMAINS, 200 RANDOM TOPOLOGIES, 90% CONFIDENCE LEVELS]

Peering Model	**Protocol Interfaces**							
	PCE-PCE	**PCE-AAA**	**PCE-BSW**	**BSW-BSW (RSVP)**	**BSW-AAA (Auth)**	**BSW-AAA (Acct)**	**AAA-AAA (Proxy-Recv)**	**AAA-AAA (Proxy-Send)**
Alliance	0.4	0.5	0.2	1.1	8.9	1.5	0.0	0.7
Cascaded	0.4	0.9	0.2	1.1	17.8	1.5	0.7	0.7

TABLE VI
SIGNALING DELAY IN ALLIANCE (A) AND CASCADED DEPLOYMENTS (C) [200 DOMAINS, 200 RANDOM TOPOLOGIES, 90% CONFIDENCE LEVELS]

Peering Model	**Computation**		**Reseravation**	
	Shared secret	**Asymmetic key**	**Shared secret**	**Asymmetric key**
Alliance	382.2	426.2	272.1	273.0
Cascaded	382.2	439.5	272.4	286.2

VII. CONCLUSIONS

In this paper, we proposed an architectural and signaling framework that facilitates authentication, authorization, and accounting (AAA) in connection-oriented multi-domain networks with security features. The solution builds on the recent IETF PCE/RSVP standards and introduces security mechanisms which enforce the verification of identity and authorization rights of incoming computation and reservation requests in transit and destination domains. The proposed framework can guarantee path integrity to the source domains and also allows accounting information to flow between domains for auditing and billing purposes. As performance is one of the primary concerns when extending protocols, we proposed analytical models which estimate the signaling load and delay metrics, and model verification by simulations.

The results indicate that the additional signaling load in the proposed framework scales almost linearly with the number of domains. Although the connection setup delay is naturally higher due to the introduced multi-round security mechanisms, it is still acceptable and is largely a function of the domain-internal propagation delay between the PCE system and the AAA server. This indicates that network providers can select authentication mechanisms based on their management and security requirements and that systems can have a high measurable benefit from the security features in the PCE servers.

The results on signaling rate showed that a proper choice of the RSVP refresh and accounting interim intervals can largely minimize the load perturbations as the average connection duration fluctuates around the planned values. Interestingly, the results also demonstrated that although the developed model of the mean signaling load requires knowledge of the connection duration distribution, only average values are sufficient to get an accurate estimate for the signaling load. Finally, we also showed that the proposed framework applies to cascaded and alliance peering models with comparable performance.

REFERENCES

[1] C. de Laat et al., *Generic AAA Architecture*, RFC 2903, Aug 2000.

[2] P. Morand et al, *Management of End-to-end Quality of Service Across the Internet*, MESCAL report, http://www.mescal.org/deliverables/MESCAL-D14-final-v2.pdf, Jan 2004.

[3] ETNA, *Ethernet Transport Networks, Architectures of Networking (ETNA), Network Architecture WP*, D2.1 Issue 2.1, Dec 2008.

[4] A. Farrel, et al, *A PCE-Based Architecture*, RFC 4655, Aug 2006

[5] D. Awduche et al., *RSVP-TE*,RFC 3209, Dec 2001.

[6] J. Vollbrecht et al., *AAA Authorization Framework*, RFC 2904, Aug 2000.

[7] S. Herzog, *RSVP Extensions for Policy Control*, RFC 2750, Jan. 2000.

[8] H. Tschofenig et al., *RSVP Security Properties*, RFC 4230, Dec 2005.

[9] F. Baker et al., *RSVP Cryptographic Authentication*, RFC 2747, Jan 2000.

[10] C. Kaufman, *Internet Key Exchange Protocol*, RFC 4306, Dec 2005.

[11] L. Fang, *Security Framework for MPLS/GMPLS*,RFC 5920, Jul 2010.

[12] N. Bitar, *MPLS InterCarrier Interconnect Technical Specification*, IP/MPLS Forum 19.0.0, Apr 2008.

[13] L. Gommans et al., *Applications drive secure lightpath creation across heterogeneous domains*, IEEE Communications Magazine, Vol. 44, No. 3, pp. 100-106, Mar 2006.

[14] OIF, *E-NNI Signaling Specification*, IA OIF-E-NNI-Sig-02.0

[15] R. Douville et al., *A Service Plane over the PCE Architecture for Automatic Multidomain Connection-Oriented Services*, IEEE Communications Magazine , Vol. 46, No. 6, pp. 94-102, Jun 2008.

[16] A. P. Bianzino et al., *Testbed Implementation of Control Plane Extensions for Inter'Carrier GMPLS LSP Provisioning*, International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Washington (USA), Apr 2009.

[17] L. Gommans et all, *Multi-domain Lightpath Authorization using Tokens*, Future Generation Computer Systems Journal, Vol. 25, Issue 2, Feb 2009.

[18] Y. Demchenko et al., *Authorization Infrastructure for On-Demand Network Resource Provisioning*, IEEE/ACM International Conference on Grid Computing, Tsukuba, Sept 2008.

[19] S. Greco Polito, M. Chamania and A. Jukan, *Extending the Inter-domain PCE Framework for Authentication and Authorization in GMPLS Networks*, IEEE ICC, Dresden, Germany, Jun 2009.

[20] S. Zaghloul, A. Jukan, *Signaling Rate and Performance for Authentication, Authorization, and Accounting (AAA) Systems in all-IP Cellular Networks*, IEEE Transactions on Wireless Communications.

[21] P. Racz, B. Stiller , *IP Flow Accounting Application for Diameter*, IEEE Transactions on Network and Service Management, Vol.5, n. 4, Dec 2008.

[22] P. Calhoun et al., *Diameter Base Protocol*, RFC 3588, September 2003.

[23] I. Bryskin, et al, *Policy-Enabled PCE Framework*, RFC 5394, Dec 2008.

[24] J.P. Vasseur et al, *PCE Comm. Protocol (PCEP)*, RFC 5440, Mar 2009.

[25] J.P. Vasseur et al., *A Backward Recursive PCE-based Computation (BRPC) Procedure To Compute Shortest Constrained Inter-domain Traffic Engineering Label Switched Paths*, RFC 5441, Apr 2009.

[26] R. Bradford, JP. Vasseur and A. Farrel, *Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism*, RFC 5520, Apr 2009.

[27] I. Aib and R. Boutaba, *On Leveraging Policy-Based Management for Maximizing Business Profit*, IEEE Transactions on Network and Service Management, Vol.4, n. 3, Dec 2007.

[28] W. Diffie and M.E. Hellman, *New direction in cryptography*, IEEE Trans. Inform Theory, IT-22, 6, 1976.

[29] BRITE: Boston University Representative Internet Topology gEnerator, http://www.cs.bu.edu/brite/

[30] A.L. Barabasi and R. Albert, *Emergence of Scaling in Random Networks*,Science, pp 509-512, Oct 1999.

[31] OPENSSL, http://www.openssl.org/

[32] Germany 17 Node Reference Network, http://www.bmbf.de/de/6103.php

Silvana Greco Polito is Assistant Professor at the University of Enna "Kore". She received MSc and PhD degrees from the University of Palermo (Italy) and was a visiting PhD student at the Internet RealTime Lab of the Columbia University (USA). She has worked as researcher at the University of Palermo and the Technical University Carolo-Wilhelmina of Brunswick (Germany). Her research interests include security models and protocols, carrier-grade control-plane, p2p technologies, cloud computing.

Said Zaghloul is currently with Redknee Corporation in Canada where he leads the design of data policy and real-time rating and charging solutions, and represents Redknee in 3GPP. He received his PhD and MSc degrees from the Technical-University of Braunschweig and the University of Kansas respectively. Prior to joining Redknee, he was with Sprint-Nextel, USA and Siemens Germany. Said is a Fulbright Alumnus and author of several refereed journal and conference articles, and industry patents.

Mohit Chamania is currently pursuing his PhD at the Technical University Carolo-Wilhelmina of Braunschweig, Germany. Prior to his PhD studies, he pursued his B.Tech and M.Tech in Mechanical Engineering at the Indian Institute of Technology, Bombay. His research interests include hybrid IP-over-optical networks, inter-domain routing, Ethernet and Optical technologies.

Admela Jukan Admela Jukan (SM) is Chair Professor of Communication Networks in Electrical and Computer Engineering Department at the Technische Universität Carolo-Wilhelmina zu Braunschweig (Brunswick) in Germany. Prior to coming to Brunswick, she was research faculty at the Institut National de la Recherche Scientifique (INRS), University of Illinois at Urbana Champaign (UIUC) and Georgia Tech (GaTech). From 2002-2004, she served as Program Director in Computer and Networks System Research at the National Science Foundation (NSF) in Arlington, VA. She received her M.Sc. degree in Information Technologies from the Politecnico di Milano, Italy, and the Ph.D. degree (cum laude) in Electrical and Computer Engineering from the Technische Universität Wien, in Vienna, Austria. She received her B.Sc. (Dipl. -Ing.) degree from the Faculty of Electrical Engineering and Computing (FER), in Zagreb, Croatia. Dr. Jukan has chaired and co-chaired several international conferences, including IFIP ONDM, IEEE ANTS, IEEE ICC and IEEE GLOBECOM. She serves as Associate Technical Editor for IEEE Communications Surveys, IEEE Communications Magazine and IEEE Network. She is elected Secretary of the IEEE Optical Network Technical Committee (Vice Chair in 2012, Chair in 2014). She currently coordinates the EU-Project ONE, with focus on next generation network management systems in optical networks.

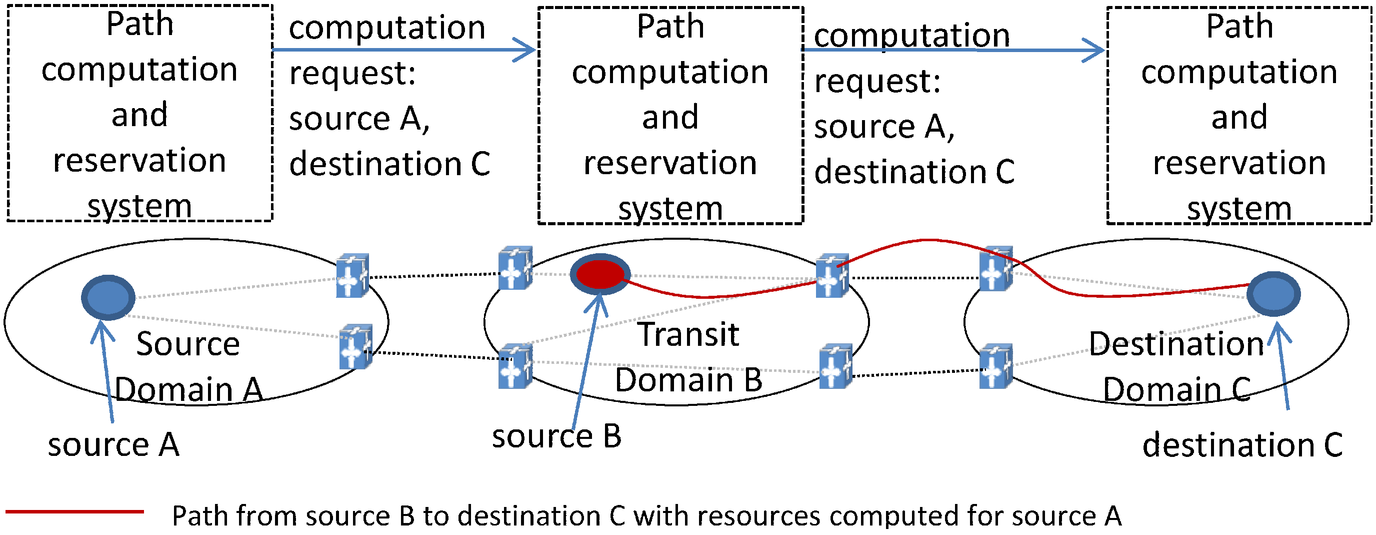
Path computation and reservation system
computation request: source A, destination C
Path computation and reservation system
computation request: source A, destination C
Path computation and reservation system
Source Domain A
Transit Domain B
Destination Domain C
source A
source B
destination C
Path from source B to destination C with resources computed for source A

Fig. 1. Current path computation and setup and security challenges

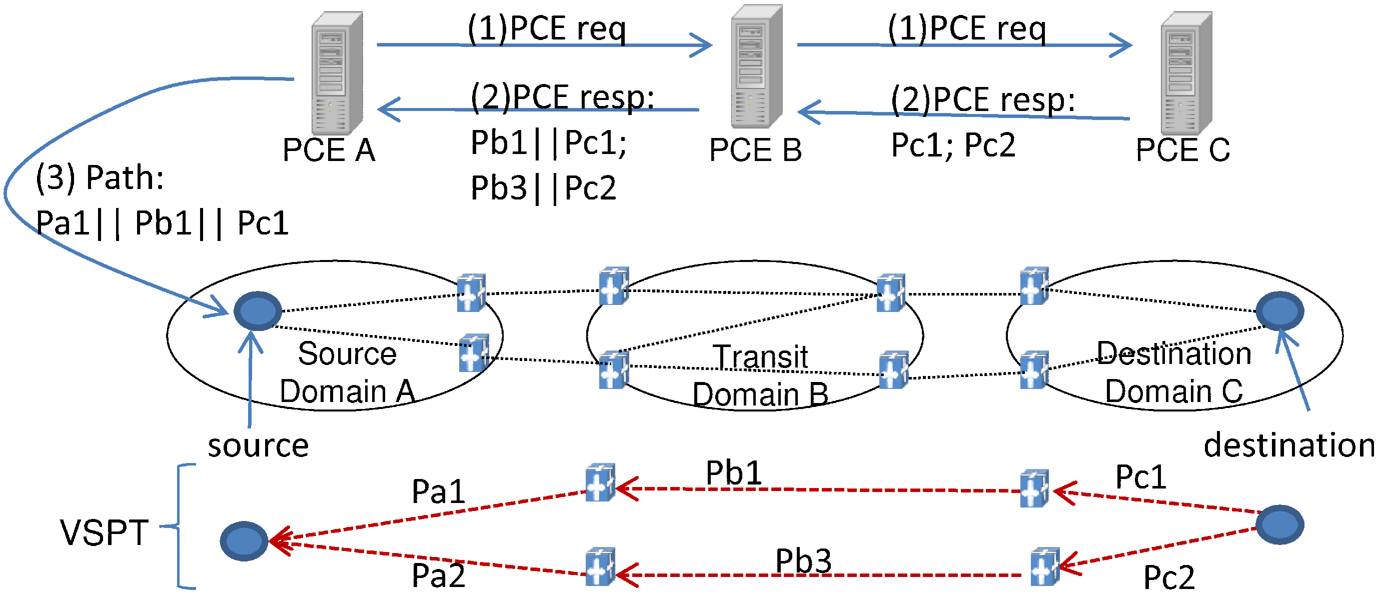

Fig. 2. Current architecture with PCE-based inter-domain computation

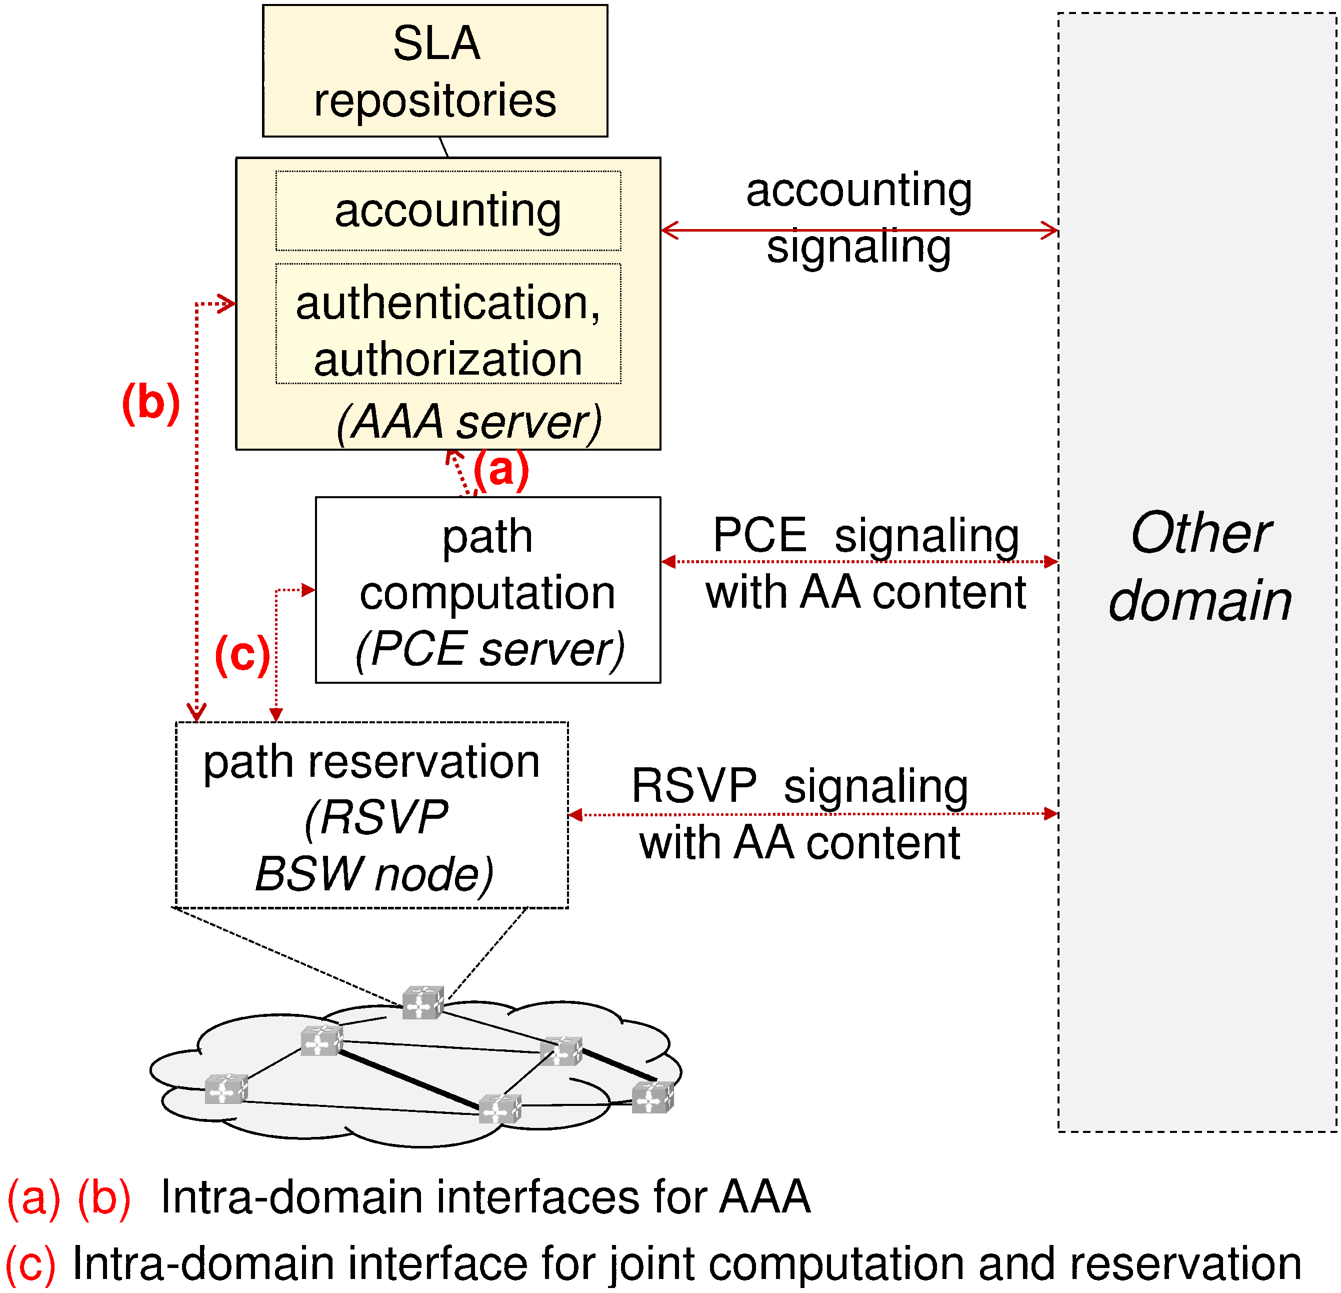
SLA repositories
accounting
authentication, authorization
(AAA server)
accounting signaling
(b)
(a)
path computation (PCE server)
PCE signaling with AA content
(c)
path reservation (RSVP BSW node)
RSVP signaling with AA content
Other domain
(a) (b) Intra-domain interfaces for AAA
(c) Intra-domain interface for joint computation and reservation

Fig. 3. The proposed PCE-based management eco-system

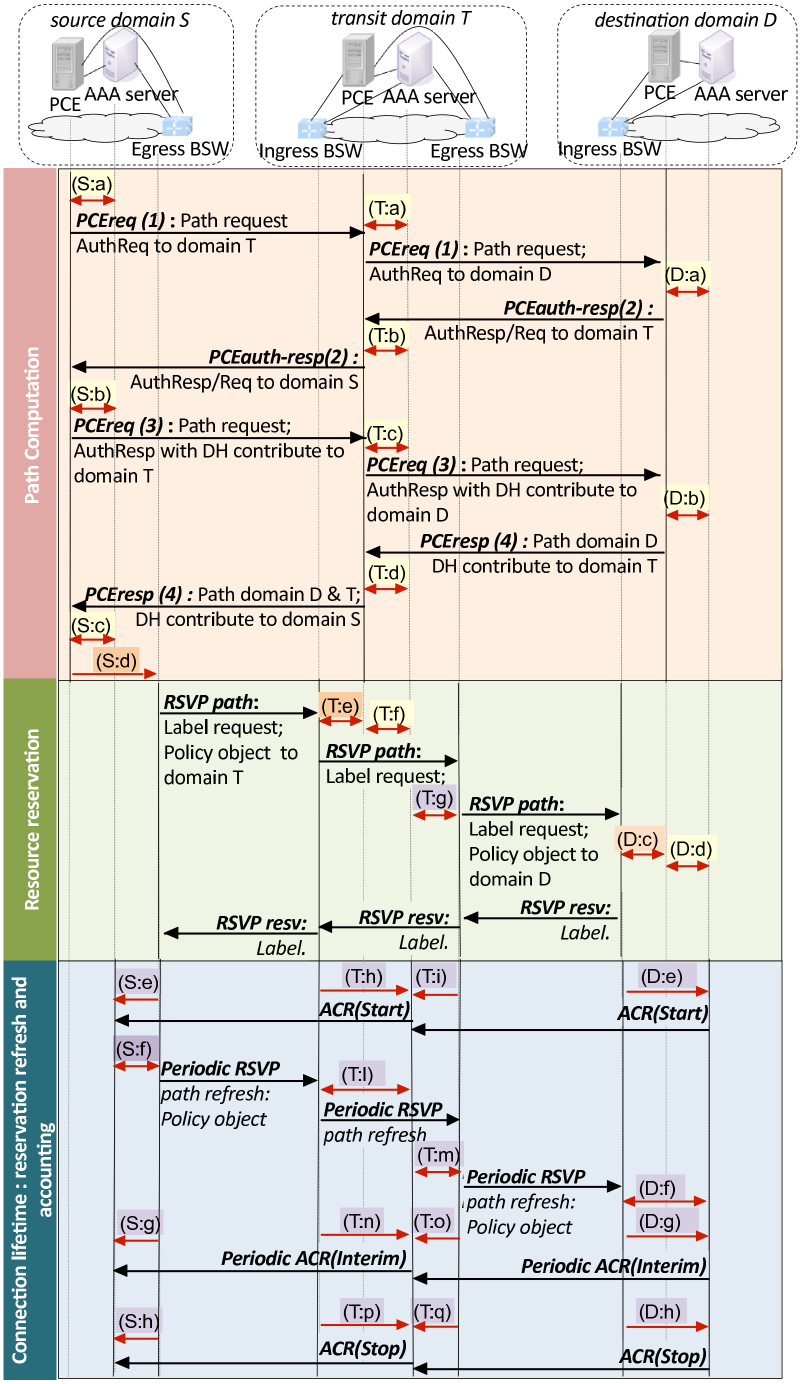

(a)

Intra-domain signaling for path computation with security features

(S:a)	-Provisioning of AuthReq for domain T
(T:a)	-Computation of AuthResp/req for domain S -Provisioning of AuthReq for domain D
(D:a)	-Provisioning of AuthResp/req for domain T
(T:b)	-Verification of AuthResp from domain D and computation of DH contribute for domain D -Provisioning of AuthResp/req for domain S
(S:b)	-Verification of AuthResp/req of domain T and provisioning AuthResp with DH contribute for domain T
(T:c)	-Verification of AuthResp from domain S, computation of DH contribute and session key derivation for domain S -Provisioning of domain S authorization profile -Provisioning of AuthResp with DH contribute for domain D
(D:b)	-Verification of AuthResp from domain T, provisioning of DH contribute and computaion of DH key with domain T -Provisioning of domain T authorization profile
(T:d)	-Computation of DH key with domain D -Provisioning of DH contribute for domain S
(S:c)	-Computation of DH key with domain T -Provisioning of policy object for domain T
(S:d)	-Provisioning of computed path and policy object for domain T

(b)

Intra-domain signaling for resource reservation with security features

(T:e)	-Provisioning of path-key & policy object
(T:f)	-Verification of policy object
(T:g)	-Provisioning of policy object for domain D
(D:c)	As T:e
(D:d)	-As T:f

(c)

Intra-domain signaling for connection lifetime functions with security features

(S:e), (T:h), (T:i), (D:e)	-Notification of starting of metering
(S:f), (T:m)	-Provisioning of policy objects for RSVP refresh
(T:l), (D:f)	-Verification of policy object
(S:g), (T:n), (T:0), (D:g)	-Provisioning of accounting interim records
(S:h), (T:p), (T:q), (D:h)	-Notification of ending of metering

(d)

□ PCE server -AAA server interface

□ PCE-BSW interface

□ BSW-AAA server interface

Fig. 4. (a) Inter-domain signaling for path computation, reservation and connection lifetime procedures; (b) Intra-domain signaling for computation; (c) intra-domain signaling for reservation; (d) intra-domain signaling for connection lifetime procedures

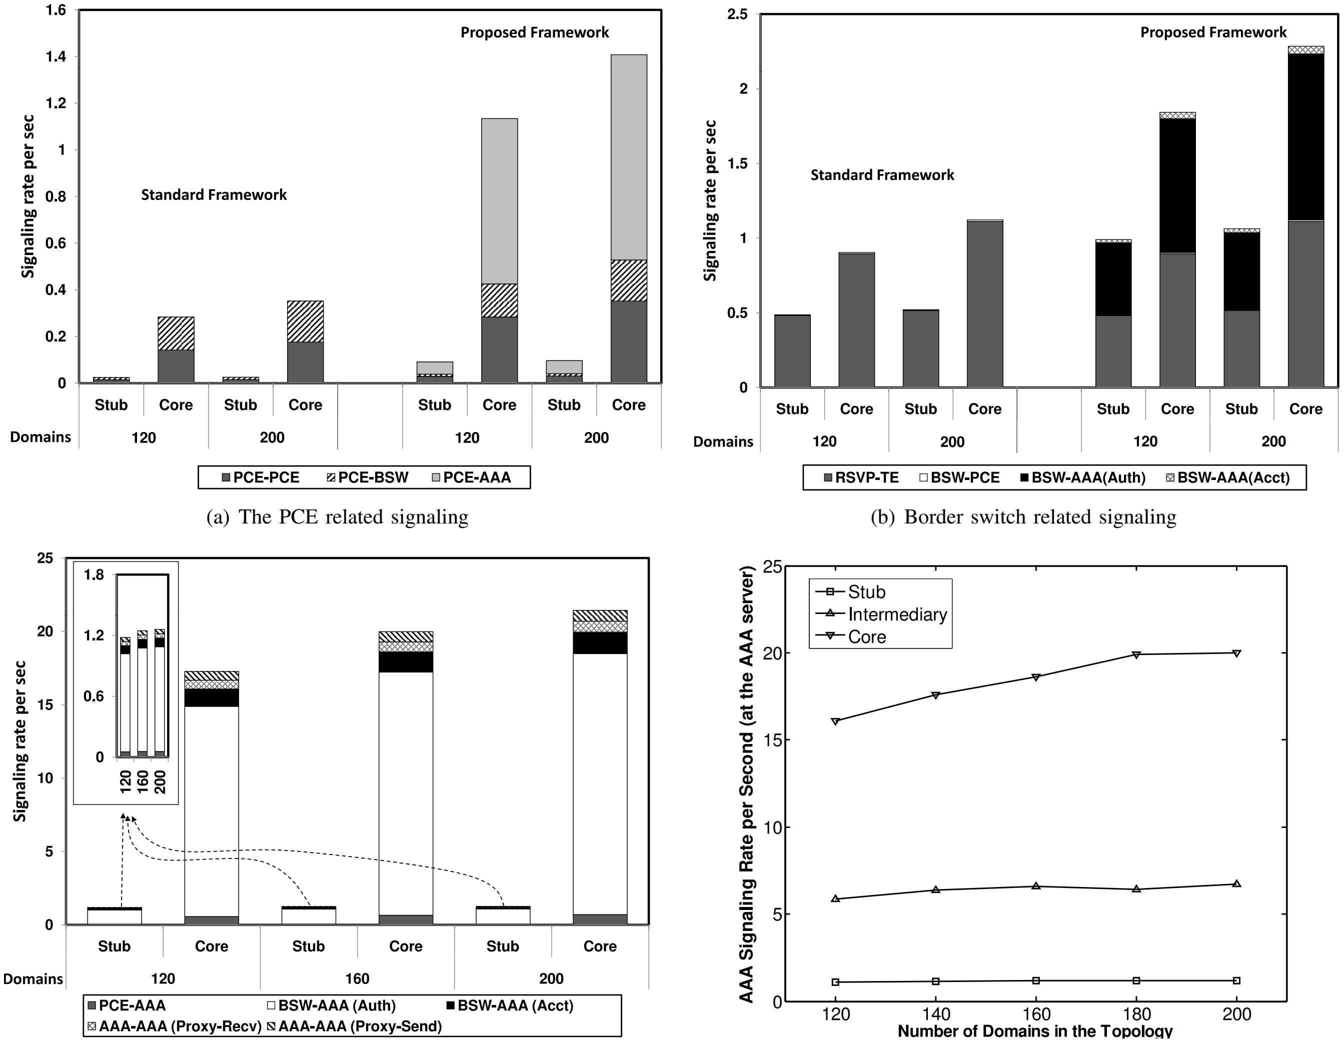

(a) The PCE related signaling

(b) Border switch related signaling

(c) AAA related signaling

(d) AAA signaling vs topology size

Fig. 5. Scalability and signaling load on interfaces [200 random network topologies, 90% confidence levels]. In Fig. 5(b), the BSW-PCE signaling barely shows between the RSVP-TE and the BSW-AAA (Auth) bars.

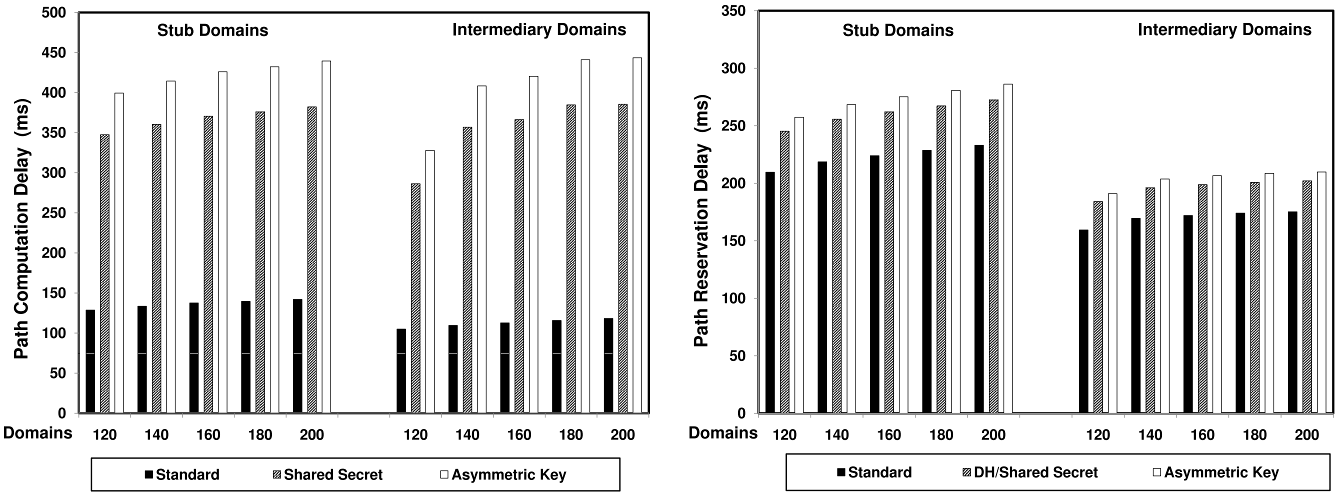

(a) Mean signaling delay during path computation

(b) Mean signaling delay during path reservation

Fig. 6. Signaling delay comparison [200 random network topologies, 90% confidence levels]

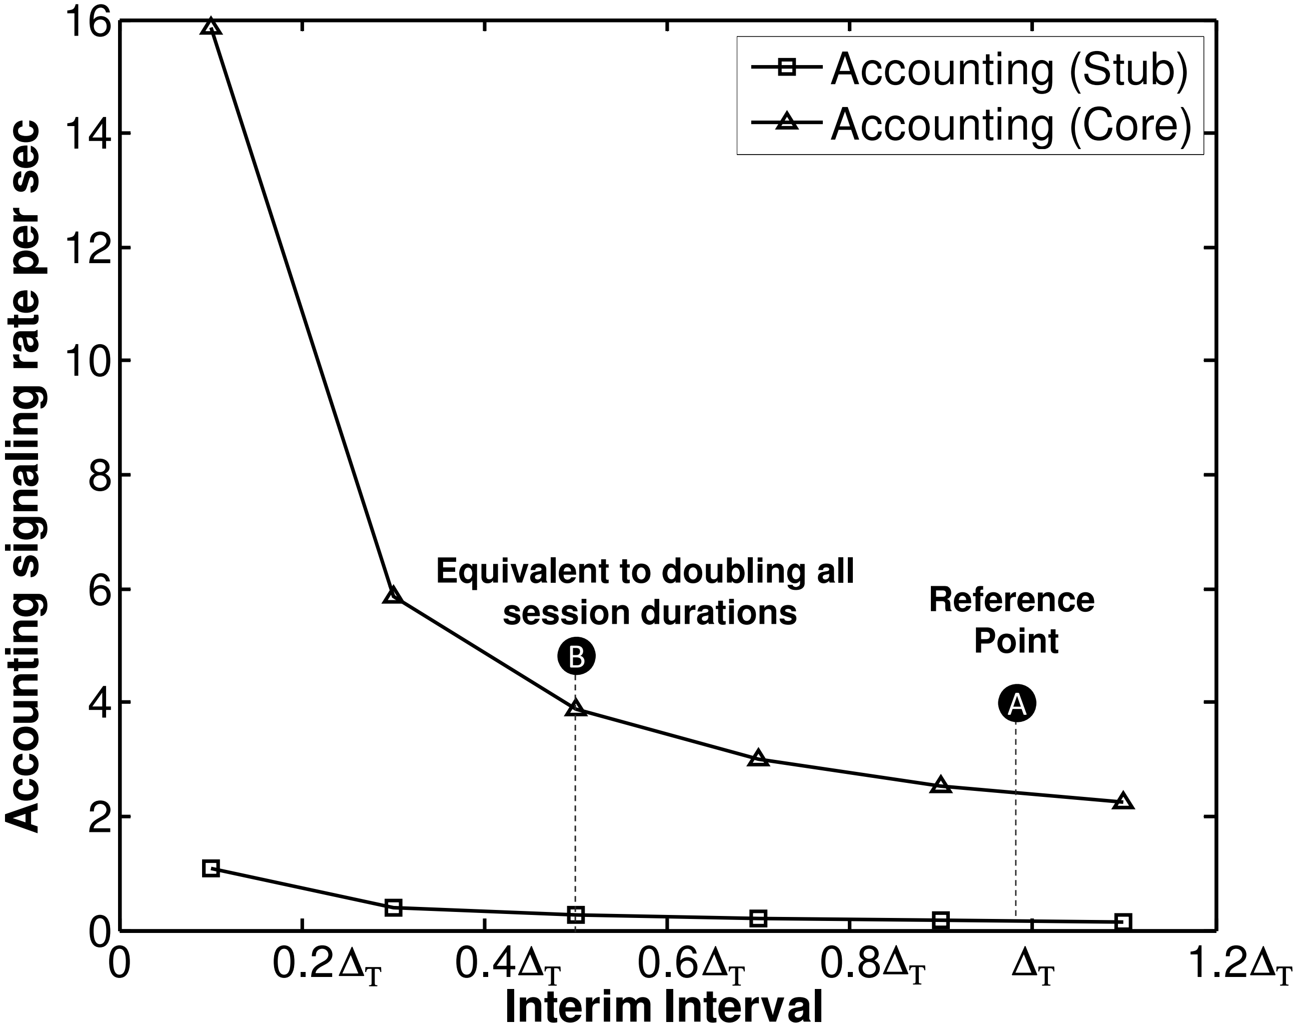

Fig. 7. Signaling load at the BSW-AAA interface as function of the accounting interim interval [120 domains, 200 random topologies, 90% confidence levels]

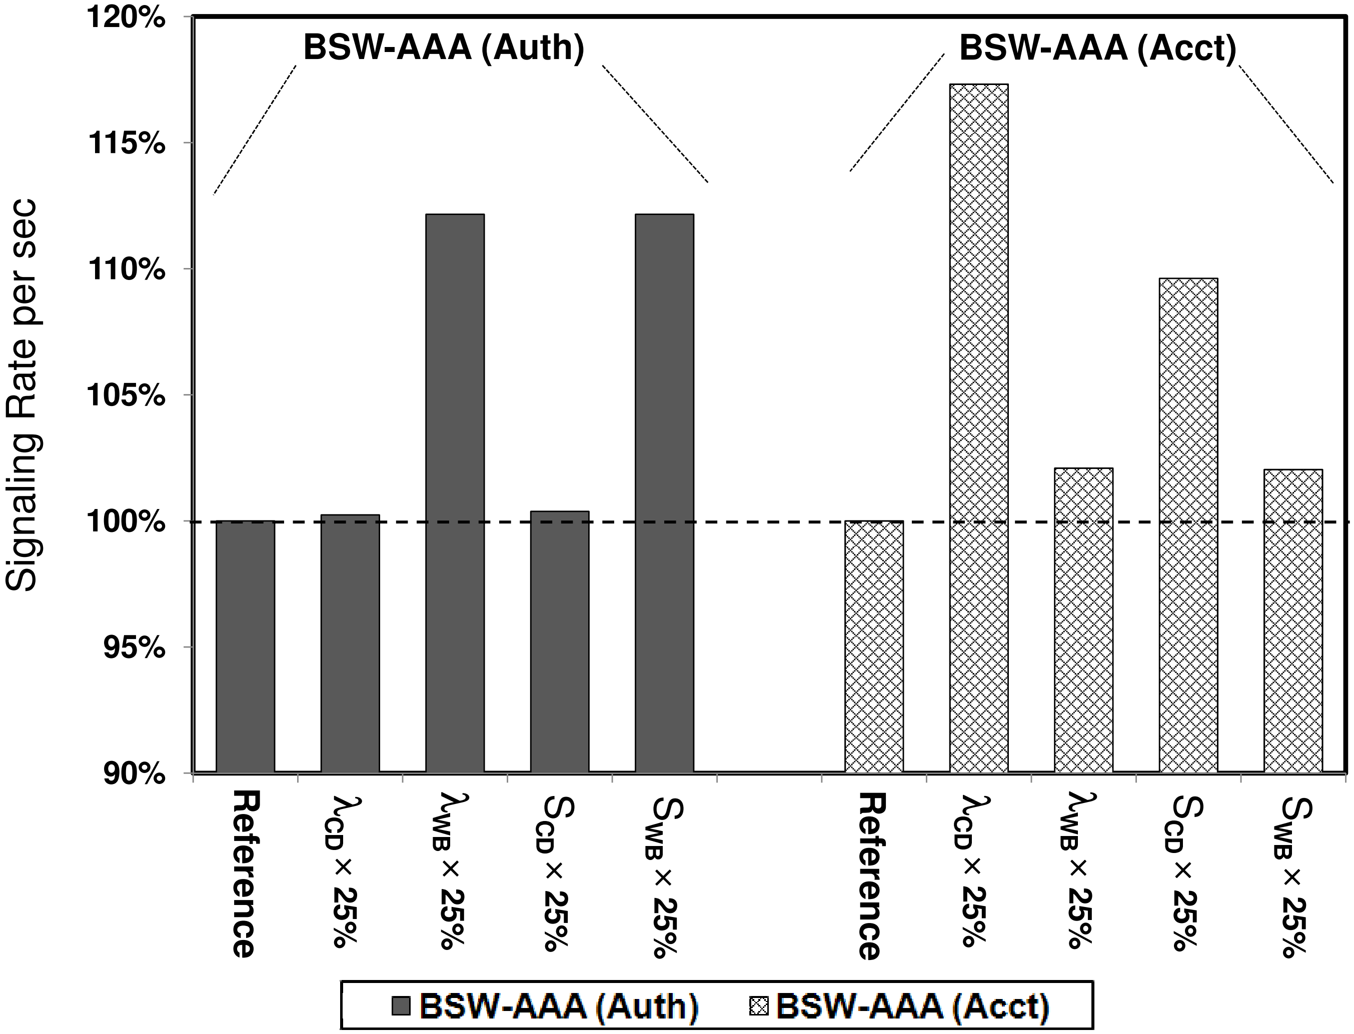

Fig. 8. Impact of arrival rate and session duration on AAA signaling load [120 domains, 200 random network topologies, 90% confidence levels]