

Path Computation Element
Internet-Draft
Intended status: Experimental
Expires: August 18, 2014

D. Lopez
O. Gonzalez de Dios
Telefonica I+D
Q. Wu
D. Dhody
Huawei
February 14, 2014

Secure Transport for PCEP
draft-lopez-pce-pceps-02

Abstract

The Path Computation Element Communication Protocol (PCEP) defines the mechanisms for the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE), or among PCEs. This document describe the usage of Transport Layer Security (TLS) to enhance PCEP security, hence the PCEPS acronym proposed for it. The additional security mechanisms are provided by the transport protocol supporting PCEP, and therefore they do not affect its flexibility and extensibility.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 18, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Requirements Language	3
3. Applying PCEPS	4
3.1. TCP ports	4
3.2. TLS Connection Establishment	4
3.3. Peer Identity	6
3.4. Connection Establishment Failure	7
4. Discovery Mechanisms	7
4.1. DANE Applicability	8
5. Backward Compatibility	8
6. IANA Considerations	8
7. Security Considerations	8
8. Acknowledgements	9
9. References	9
9.1. Normative References	9
9.2. Informative References	10
Authors' Addresses	11

1. Introduction

PCEP [RFC5440] defines the mechanisms for the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between two PCEs. These interactions include requests and replies that can be critical for a sustainable network operation and adequate resource allocation, and therefore appropriate security becomes a key element in the PCE infrastructure. As the applications of the PCE framework evolves, and more complex service patterns emerge, the definition of a secure mode of operation becomes more relevant.

[RFC5440] analyzes in its section on security considerations the potential threats to PCEP and their consequences, and discusses several mechanisms for protecting PCEP against security attacks, without making a specific recommendation on a particular one or defining their application in depth. Moreover, [RFC6952] remarks the importance of ensuring PCEP communication privacy, especially when PCEP communication endpoints do not reside in the same AS, as the interception of PCEP messages could leak sensitive information related to computed paths and resources.

Among the possible solutions mentioned in these documents, Transport Layer Security (TLS) [RFC5246] provides support for peer authentication, and message encryption and integrity. TLS supports the usage of well-know mechanisms to support key configuration and exchange, and means to perform security checks on the results of PCE discovery procedures via IGP ([RFC5088] and [RFC5089]).

This document describes a security container for the transport of PCEP requests and replies, and therefore it will not interfere with the protocol flexibility and extensibility.

This document describes how to apply TLS in securing PCE interactions, including the TLS handshake mechanisms, the TLS methods for peer authentication, the applicable TLS ciphersuites for data exchange, and the handling of errors in the security checks. In the rest of the document we will refer to this usage of TLS to provide a secure transport for PCEP as "PCEPS".

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Applying PCEPS

3.1. TCP ports

Since PCEP can operate either with or without TLS, it is necessary for the PCEP speaker to indicate whether it wants to set up a TLS connection or not. There are two main ways of achieving this:

- o One option is to use a different port number for TLS connections (for example, the port 443 used for HTTPS)
- o The other is to use the regular port number and have the PCEP speaker request that the PCE switch the connection to TLS using a protocol-specific mechanism (for example, the STARTTLS for mail and news protocols)

To avoid requiring a specific PCEP extension to request TLS, this document proposes the usage of the former solution to implement PCEPS.

The default destination port number for PCEPS is TCP/XXXX.

NOTE: This port has to be agreed and registered as PCEPS with IANA.

3.2. TLS Connection Establishment

PCEPS has no notion of negotiating TLS in an established connection. PCEP peers MAY either discover that the other PCEP endpoint supports PCEPS or can be preconfigured to use PCEPS for a given peer (see section Section 4 for more details). The connection establishment SHALL follow the following steps:

1. After completing the TCP handshake, immediately negotiate TLS sessions according to [RFC5246]. The following restrictions apply:
 - * Support for TLS v1.2 [RFC5246] or later is REQUIRED.
 - * Support for certificate-based mutual authentication is REQUIRED.
 - * Negotiation of mutual authentication is REQUIRED.
 - * Negotiation of a ciphersuite providing for integrity protection is REQUIRED.
 - * Negotiation of a ciphersuite providing for confidentiality is RECOMMENDED.

- * Support for and negotiation of compression is OPTIONAL.
 - * PCEPS implementations MUST, at a minimum, support negotiation of the TLS_RSA_WITH_3DES_EDE_CBC_SHA, and SHOULD support TLS_RSA_WITH_RC4_128_SHA and TLS_RSA_WITH_AES_128_CBC_SHA as well. In addition, PCEPS implementations MUST support negotiation of the mandatory-to-implement ciphersuites required by the versions of TLS that they support.
2. Peer authentication can be performed in any of the following two REQUIRED operation models:
- * TLS with X.509 certificates using PKIX trust models:
 - + Implementations MUST allow the configuration of a list of trusted Certification Authorities (CAs) for incoming connections.
 - + Certificate validation MUST include the verification rules as per [RFC5280].
 - + Implementations SHOULD indicate their trusted CAs. For TLS 1.2, this is done using [RFC5246], Section 7.4.4, "certificate_authorities" (server side) and [RFC6066], Section 6 "Trusted CA Indication" (client side).
 - + Peer validation always SHOULD include a check on whether the locally configured expected DNS name or IP address of the peer that is contacted matches its presented certificate. DNS names and IP addresses can be contained in the Common Name (CN) or subjectAltName entries. For verification, only one of these entries is to be considered. The following precedence applies: for DNS name validation, subjectAltName:DNS has precedence over CN; for IP address validation, subjectAltName:ipAddr has precedence over CN.
 - + NOTE: Consider here whether peer validation MAY be extended by means of the DANE procedures, including its specs as informative references.
 - + Implementations MAY allow the configuration of a set of additional properties of the certificate to check for a peer's authorization to communicate (e.g., a set of allowed values in subjectAltName:URI or a set of allowed X509v3 Certificate Policies)

- * TLS with X.509 certificates using certificate fingerprints:
Implementations MUST allow the configuration of a list of trusted certificates, identified via fingerprint of the Distinguished Encoding Rules (DER) encoded certificate octets. Implementations MUST support SHA-256 as the hash algorithm for the fingerprint.

3. Start exchanging PCEP messages.

To support TLS re-negotiation both peers MUST support the mechanism described in [RFC5746]. Any attempt of initiate a TLS handshake to establish new cryptographic parameters not aligned with [RFC5746] SHALL be considered a TLS negotiation failure.

3.3. Peer Identity

Depending on the peer authentication method in use, PCEPS supports different operation modes to establish peer's identity and whether it is entitled to perform requests or can be considered authoritative in its replies. PCEPS implementations SHOULD provide mechanisms for associating peer identities with different levels of access and/or authoritativeness, and they MUST provide a mechanism for establish a default level for properly identified peers. Any connection established with a peer that cannot be properly identified SHALL be terminated before any PCEP exchange takes place.

In TLS-X.509 mode using fingerprints, a peer is uniquely identified by the fingerprint of the presented client certificate.

There are numerous trust models in Public-Key Infrastructure (PKI) environments, and it is beyond the scope of this document to define how a particular deployment determines whether a client is trustworthy. Implementations that want to support a wide variety of trust models should expose as many details of the presented certificate to the administrator as possible so that the trust model can be implemented by the administrator. As a suggestion, at least the following parameters of the X.509 client certificate should be exposed:

- o Peer's IP address
- o Peer's fully qualified domain name (FQDN)
- o Certificate Fingerprint
- o Issuer

- o Subject
- o All X509v3 Extended Key Usage
- o All X509v3 Subject Alternative Name
- o All X509v3 Certificate Policies

In addition, a PCC MAY apply the procedures described in [RFC6698] (DANE) to verify its peer identity when using DNS discovery. See section Section 4.1 for further details.

3.4. Connection Establishment Failure

In case the initial TLS negotiation or the peer identity check fail according to the procedures listed in this document, the peer MUST immediately terminate the session. It SHOULD follow the procedure listed in [RFC5440] to retry session setup along with an exponential back-off session establishment retry procedure.

4. Discovery Mechanisms

A PCE can advertise its capability to support PCEPS using the IGP advertisement and discovery mechanism. The PCE-CAP-FLAGS sub-TLV is an optional sub-TLV used to advertise PCE capabilities. It MAY be present within the PCED sub-TLV carried by OSPF or IS-IS. [RFC5088] and [RFC5089] provide the description and processing rules for this sub-TLV when carried within OSPF and IS-IS, respectively. PCE capability bits are defined in [RFC5088]. A new capability flag bit for the PCE-CAP-FLAGS sub-TLV that can be announced as attribute to distribute PCEP security support information is proposed in [I-D.wu-pce-discovery-pceps-support]

NOTE: A new bit must be added here to advertise the PCEPS capability.

When DNS is used by a PCC (or a PCE acting as a client, for the rest of the section, PCC refers to both) willing to use PCEPS to locate an appropriate PCE [I-D.wu-pce-dns-pce-discovery], the PCC as initiating entity chooses at least one of the returned FQDNs to resolve, which it does by performing DNS "A" or "AAAA" lookups on the FQDN. This will eventually result in an IPv4 or IPv6 address. The PCC SHALL use the IP address(es) from the successfully resolved FQDN (with the corresponding port number returned by the DNS SRV lookup) as the connection address(es) for the receiving entity.

If the PCC fails to connect using an IP address but the "A" or "AAAA" lookups returned more than one IP address, then the PCC SHOULD use

the next resolved IP address for that FDQN as the connection address. If the PCC fails to connect using all resolved IP addresses for a given FDQN, then it SHOULD repeat the process of resolution and connection for the next FQDN returned by the SRV lookup based on the priority and weight.

If the PCC receives a response to its SRV query but it is not able to establish a PCEPS connection using the data received in the response, as initiating entity it MAY fall back to lookup a PCE that uses TCP as transport.

4.1. DANE Applicability

DANE [RFC6698] defines a secure method to associate the certificate that is obtained from a TLS server with a domain name using DNS, i.e., using the TLSA DNS resource record (RR) to associate a TLS server certificate or public key with the domain name where the record is found, thus forming a "TLSA certificate association". The DNS information needs to be protected by DNSSEC. A PCC willing to apply DANE to verify server identity MUST conform to the rules defined in section 4 of [RFC6698].

5. Backward Compatibility

Since the procedure described in this document describes a security container for the transport of PCEP requests and replies carried on a newly allocated TCP port there will be no impact on the base PCEP and/or any further extensions.

6. IANA Considerations

NOTE: PCEPS has to be registered as TCP port XXXX.

No new PCEP messages or other objects are defined.

7. Security Considerations

While the application of TLS satisfies the requirement on privacy as well as fine-grained, policy-based peer authentication, there are security threats that it cannot address. It is advisable to apply additional protection measures, in particular in what relates to attacks specifically addressed to forging the TCP connection underpinning TLS. TCP-AO (TCP Authentication Option [RFC5925]) is fully compatible with and deemed as complementary to TLS, so its usage is to be considered as a security enhancement whenever any of

the PCEPS peers require it, especially in the case of long-lived connections. The mechanisms to configure the requirements to use TCP-AO and other lower-layer protection measures, as well as the association of the required crypto material (MKT in the case of TCP-AO) with a particular peer are outside the scope of this document. [I-D.chunduri-karp-using-ikev2-with-tcp-ao] defines a method to perform such association.

Since computational resources required by TLS handshake and ciphersuite are higher than unencrypted TCP, clients connecting to a PCEPS server can more easily create high load conditions and a malicious client might create a Denial-of-Service attack more easily.

Some TLS ciphersuites only provide integrity validation of their payload, and provide no encryption. This specification does not forbid the use of such ciphersuites, but administrators must weight carefully the risk of relevant internal data leakage that can occur in such a case, as explicitly stated by [RFC6952].

When using certificate fingerprints to identify PCEPS peers, any two certificates that produce the same hash value will be considered the same peer. Therefore, it is important to make sure that the hash function used is cryptographically uncompromised so that attackers are very unlikely to be able to produce a hash collision with a certificate of their choice. This document mandates support for SHA-256, but a later revision may demand support for stronger functions if suitable attacks on it are known.

8. Acknowledgements

This specification relies on the analysis and profiling of TLS included in [RFC6614].

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.
- [RFC5089] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element

(PCE) Discovery", RFC 5089, January 2008.

- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.
- [RFC5746] Rescorla, E., Ray, M., Dispensa, S., and N. Oskov, "Transport Layer Security (TLS) Renegotiation Indication Extension", RFC 5746, February 2010.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.
- [RFC6066] Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, January 2011.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, August 2012.

9.2. Informative References

- [I-D.chunduri-karp-using-ikev2-with-tcp-ao]
Chunduri, U., Tian, A., and J. Touch, "A framework for RPs to use IKEv2 KMP",
draft-chunduri-karp-using-ikev2-with-tcp-ao-06 (work in progress), February 2014.
- [I-D.wu-pce-discovery-pceps-support]
Lopez, D., Wu, Q., Dhody, D., and D. King, "IGP extension for PCEP security capability support in the PCE discovery", draft-wu-pce-discovery-pceps-support-00 (work in progress), February 2014.
- [I-D.wu-pce-dns-pce-discovery]
Wu, W., Dhody, D., King, D., and D. Lopez, "Path Computation Element (PCE) Discovery using Domain Name System(DNS)", draft-wu-pce-dns-pce-discovery-04 (work in progress), January 2014.

- [RFC6614] Winter, S., McCauley, M., Venaas, S., and K. Wierenga, "Transport Layer Security (TLS) Encryption for RADIUS", RFC 6614, May 2012.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, May 2013.

Authors' Addresses

Diego R. Lopez
Telefonica I+D
Don Ramon de la Cruz, 82
Madrid, 28006
Spain

Phone: +34 913 129 041
Email: diego@tid.es

Oscar Gonzalez de Dios
Telefonica I+D
Don Ramon de la Cruz, 82
Madrid, 28006
Spain

Phone: +34 913 129 041
Email: ogondio@tid.es

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: sunseawq@huawei.com

Dhruv Dhody
Huawei
Leela Palace
Bangalore, KA 560008
India

Email: dhruv.ietf@gmail.com

