

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2012

O. Gonzalez de Dios, Ed.
Telefonica I+D
R. Casellas
CTTC
C. Margaria
Nokia Siemens Networks
Y. Lee
F. Zhang
Huawei
March 9, 2012

PCEP Extensions for Temporary Reservation of Computed Path Resources and
Support for Limited Context State in PCE
[draft-gonzalezdedios-pce-reservation-state-01](#)

Abstract

The Path Computation Element (PCE) provides path computation functions in support of traffic engineering in Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

A limited form of statefulness is useful to improve PCE functionality in situations in which the local TED might not be up to date, or in the case of concurrent requests where most of the LSPs are computed before the end of the set-up of the LSPs, when the TED is updated. The PCC is responsible to setup or not the TE-Path computed by the PCE. By providing an indication that it intends to use the resources on the TE-Path a PCC can help the PCE to get a more accurate dynamic TED view and thus the PCE can avoid suggesting the use of the same resources for subsequent TE LSPs.

This document proposes an extension to the PCEP protocol to allow the PCC to indicate to the PCE to block or reserve the resources computed in a path request of a TE LSP for subsequent requests for a certain time and can help to reduce the number of crankbacks.

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

Internet-Draft PCEP Ext for Reserv of Resources in PCE

March 2012

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 September 10, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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	4
2.	PCEP Requirements	6
3.	PCEP Extensions (Encoding)	8
3.1.	Requesting a Reservation of Resources	8
3.2.	Replying a reservation status	10
3.3.	Cancelling a Reservation	10
3.4.	RESERVATION object format	11
3.5.	RESERVATION_CONF object format	12
3.6.	RESERVATION_ID TLV	13
4.	Procedures	13
5.	Use cases	14
5.1.	Multiple LSP restoration in a WSON network	14
5.2.	Domain path selection	15
5.3.	Multidomain path computation	16
6.	Manageability Considerations	16
6.1.	Control of Function and Policy	16
6.2.	Information and Data Models	16
6.3.	Liveness Detection and Monitoring	17
6.4.	Verifying Correct Operation	17
6.5.	Requirements for Other Protocols and Functional Components	17
6.6.	Impact on Network Operation	17
7.	Security Considerations	17
8.	IANA Considerations	18
8.1.	RESERVATION object	18
8.2.	RESERVATION_CONF object	18
8.3.	RESERVATION_ID TLV	18
8.4.	PCEP Errors	18
9.	Contributing Authors	18
10.	Acknowledgements	18
11.	Normative References	19
	Authors' Addresses	19

1. Introduction

According to [[RFC4655](#)], a PCE can be either stateful or stateless. In the former case, there is a strict synchronization between the PCE, the network state (in terms of topology and per link aggregated resource information such as unreserved bandwidth), and also the set of computed paths, active Label Switched Paths (LSPs) and resources in use in the network.

In other words, a stateful PCE utilizes information from the TED as well as information about existing paths (for example, TE LSPs) in the network when processing new requests. However, the maintenance and synchronization of a stateful LSP database (LSP-DB) can be non-trivial, not only because it should verify the actual establishment of computed paths, and because it might not be the unique element to compute paths.

In addition, it can be argued that maintaining such a stateful database is not a function of the PCE, but rather of a Network Management System (NMS).

On the other hand, a stateless PCE does not typically keep track of computed paths, and each set of request(s) is processed independently of each other, typically using a local copy of the TED. Since a stateless PCE typically operates on a graph with computation constraints and without tracking the current state of paths, independent requests will be processed on the same TED graph, until the graph is updated.

With a stateless PCE, there is a 'potential window of TED inaccuracy', where a stateless PCE may compute paths based on current TED information, which could be out-of-sync with the actual or potential network state changes given other recent PCE-computed paths.

For example, some sources for this potential TED inaccuracy are:

- o Control Plane link latencies may increase due to several factors such as: a) the time required for a PCC to obtain the paths after a successful computation, requiring several Round-Trip-Times (RTT) as per TCP; b) the setup delay and c) the time it takes for the PCE to update the local TED given IGP update times.
- o The routing and topology dissemination protocol (i.e. OSPF-TE), which may operate with timers for LSA updates, to avoid excessive control plane overhead.

- o Concurrent requests that arrive during the time window, between a response is sent and the LSP is setup and the topology changes flooded. Even for very fast networks with low latency, there may be 'batched' requests: several path computation requests within a PCReq message or, in dynamic restoration without pre-planning, several LSPs that need to be rerouted avoiding a failed link.
- o Local PCE contention, where the PCE needs to concurrently serve path computation requests and update the LSA (e.g. parsing OSPF-TE LSA updates). A PCE implementation may need to find a trade-off, when synchronizing access to the local TED: favor OSPF-TE parsing which means that some path computations are slightly delayed to allow an 'update' to be processed, or give strict priority to computation requests.

In consequence, a stateless PCE may assign the same (or a subset of the same) resources to several requests, which may result in contention and degraded network performance. The effects are detected late, typically during path signaling, causing path blocking and excessive crank-backs and retries.

Note that, as per [RFC 5440](#) [[RFC5440](#)], a PCC may include a set of

previously computed paths in A given request, in order to take them into account, for instance, to avoid double bandwidth accounting or to try to minimize changes (minimum perturbation problem).

[Section 6.8 of RFC 4655](#) [RFC4655] suggests that a limited form of statefulness might be applied within an otherwise stateless PCE. The PCE may retain some context from paths it has recently computed so that it avoids suggesting the use of the same resources for other TE LSPs, using heuristics / statistic or forecasting for improved resource (i.e. wavelength) allocation. In other words, a given PCE implementation may decide to perform additional book-keeping and management of resources, deploying policies that prevent sub-optimal allocations. For instance a PCE may compute the mean time used to update the TED based on the previous calculated TE-LSPs and TED updates. Those kinds of mechanisms may reduce the TED inaccuracy but in all cases they cannot infer the PCC use of the TE-path.

This document proposes a set of extensions to the PCEP protocol to allow a PCC to request a PCE to block or reserve the resources associated with a path computation for a given path request. By reservation, it is implied that a set of resources which have been associated to such computation are excluded for subsequent path computations for a given time period. This time-based temporary reservation PCE system would be a compromise between a full-blown stateful PCE and a stateless PCE to achieve efficiency without costing and excessive resource commitment associated with the

maintenance and synchronization of a stateful PCE system. Due to the fact that the PCC is explicitly indicating this reservation, the PCE can get a more accurate view of the dynamic of the TED and thus increase the accuracy of its routing. In addition having an explicit input may simplify how a PCE take into account the fact that the TED may be outdated.

In the cases where the resource is uniquely identified in the topology update (such as receiving an OSPF-TE TE LSA with a bitmap encoded wavelength availability reflecting a change in the link status)), the reservation can still hold after a topology update, as there is a correspondence between the resource in both reservation and traffic engineering update, and the PCE can infer whether a given reserved resource has actually been committed. Otherwise, when the traffic engineering update reaches the PCE, there is no way to

distinguish the resource in the reservation among the resources shown in the TE update. Thus, to assure a coherent behavior, the general rule is that as soon as the PCE gets updated traffic engineering information, all the reservations are deleted, save the the cases where the resource is uniquely identified and the PCE can infer whether a given reserved resource has actually been committed.

Examples of resources potentially subject to reservations are: the bandwidth computed for the path in PSC or L2SC layers, a specific time slot (SDH) or tributary slot (OTN ODU-k) in TDM networks or a given wavelength or regenerator (WSON or OTN OCh).

This document also presents some illustrative use cases where the PCC would want the PCE to retain some context or state, like multiple LSP restoration, and counterexamples where the PCC does not have the intention to immediately set up the LSP, i.e., multidomain cases where the PCE is probing different paths to decide the sequence of domains.

[2.](#) PCEP Requirements

This section provides a set of requirements, both for PCCs and PCEs, to support context awareness.

When requesting a path computation (PCReq) to a PCE, a PCC should be able to indicate:

- o Whether the resources computed in the request should be made unavailable for further requests.
- o The amount of time the resources should be committed/reserved for the current computation request so that keeps subsequent requests

from taking.

- o The type and granularity of the resources to be blocked in the request. The type refers to the actual resources blocked such as path bandwidth or timeslot, wavelength, fiber... The granularity refers to the possibility of not only reserving the resource computed for the path but whether the associated links/nodes/SRLGs may need to be reserved too.

A PCE should be able to:

- o Apply policies whether a reservation request can be applied or not.
- o Compute one or more paths according to the request parameters and, based on the PCC indications, prevent (part of) the resources committed in the computed route from being used for subsequent computation requests for a given period.
- o If the request is allowed, the given reservation period SHOULD be no less than the requested period by the PCC (e.g. for the cases where the PCE is only able to reserve for multiples of a given value). This does not preclude the fact that, if configured by policy, a PCE MAY limit the period to a lower period. Alternatively, a PCE MAY be configured to reply with a PCEP_ERROR stating the cause of the failed computation/reservation.
- o The PCE MAY decide to apply a different granularity for the reservation request (e.g. block a given Time Slot or wavelength but not the TE links). In this case, the PCE MUST reply with the actual reservation.

Note that, the means by which a PCE can perform the reservation/commitment of the resources are out of the scope of the present document but could include, for example, marking the resources as 'reserved', applying internal exclude objects etc.

A PCE should be able to respond (PCRep) to the PCC the following:

- o If the resources have been effectively locked, and the effective allocated reservation period (if different from the requested one).
- o The granularity of the reservation, which may be different from the requested one.
- o Provide a means to allow a PCC to request the cancellation of an active reservation (for example an identification of the

reservation to allow its cancellation).

The PCC should be able to request the cancellation of an active resource reservation.

[3.](#) PCEP Extensions (Encoding)

[3.1.](#) Requesting a Reservation of Resources

EDITORS NOTE: OPTION WITH PCC-ID-REQ

A PCC that wants to indicate a PCE to temporarily reserve or block resources does so by including a RESERVATION object along with a client PCC_ID_REQ in a request within a PCReq message.

Analogously to [[RFC5886](#)] the PCC-ID-REQ object is used to specify the IP address of the requesting PCC. The PCC-ID-REQ MUST be inserted within a PCReq message to specify the IP address of the requesting PCC. In [[RFC5886](#)] two PCC-ID-REQ objects (for IPv4 and IPv6) are defined.

EDITORS NOTE: OPTION WITHOUT PCC-ID-REQ

A PCC that wants to indicate a PCE to temporarily reserve or block resources does so by including a RESERVATION object in a request within a PCReq message.

A PCE that processes a PCEP request with a RESERVATION object MUST act according to the P-bit in the object header: if the P-bit is set, the object MUST be treated as mandatory and the request must either be processed using the contents of the object or rejected as defined in [[RFC5440](#)]. If the P-bit is clear, the object MAY be used by the PCE or MAY be ignored.

The RESERVATION object is optional in a PCEP request. Multiple instances of the object MUST NOT be used on a single PCEP request and if a PCE finds multiple instances of the object it MUST use the first one and discard the rest (Editors note: alternatively, it could send a PCErr, OR it could allow several RESERVATION objects, and let the PCE choose which one will be used). The RESERVATION object may appear either at an individual request level or within a SVEC. The latter means that the RESERVATION object applies to all requests involved in the SVEC object.

The PCReq ([[RFC5440](#)][[RFC5541](#)][[RFC5557](#)]) message is

<PCReq Message> ::= <Common Header>

[svec_list]

<request-list>

where

<svec-list> ::= <SVEC>

[<OF>]

[<GC>]

[<XRO>]

[<metric-list>]

[<PCC-REQ-ID> <RESERVATION>]

[<svec-list>]

<metric-list> ::= <METRIC>

[<metric-list>]

<request-list> ::= <request>

[<request-list>]

<request> ::= <RP>

<END-POINTS>

[<LSPA>]

[<BANDWIDTH>]

[<metric-list>]

[<OF>]

[<PCC_REQ_ID> <RESERVATION>]

[<RRO> [<BANDWIDTH>]]

[<LOAD-BALANCING>]

3.2. Replying a reservation status

If the PCE that receives the request applies the reservation, it indicates so using a RESERVATION_CONF object in the PCRep message.

EDITOR'S NOTE: Alternatively a RESERVATION object can be used in the PCReq message

The PCRep message is extended with regard to the one defined in [\[RFC5440\]](#) as follows:

```
<attribute-list>::=[<LSPA>
```

```
    [<BANDWIDTH>
```

```
    [<metric-list>
```

```
    [<IRO>
```

```
    [<RESERVATION_CONF>]
```

Note that the reservation applies at PATH level, and a RESERVATION_CONF object is included within every path in a given PCEP response. This means distinct reservations for each path, which can be cancelled independently (Editor's Note: TDB, the PCC could indicate whether to have a single reservation or multiple reservation).

It is RECOMMENDED that the RESERVATION_CONF object appears the last attribute for a Path (or as an optional object in the attribute-list associated to a NO_PATH object.

3.3. Cancelling a Reservation

A PCC that wishes to cancel a reservation may send an unsolicited notification to the PCE, including the identifier of the reservation.

The PCNtf message used for one or more cancellations has no RP object. As with [[RFC5440](#)], the PCNtf message MUST carry at least one NOTIFICATION object and MAY contain several NOTIFICATION objects should the PCE or the PCC intend to notify of multiple events:

<PCNtf Message>::=<Common Header>

<notify-list>

<notify-list>::=<notify> [<notify-list>]

<notify>::= <notification-list>

<notification-list>::=<NOTIFICATION>[<notification-list>]

NOTIFICATION objects used for the purposes of cancelling an active reservation MUST include the RESERVATION_ID TLV. It is RECOMMENDED to use dedicated PCNtf messages for the purposes of cancelling reservations.

Both the Notification-type and Notification-value are TBD by IANA

The following Notification-type and Notification-value values are currently defined:

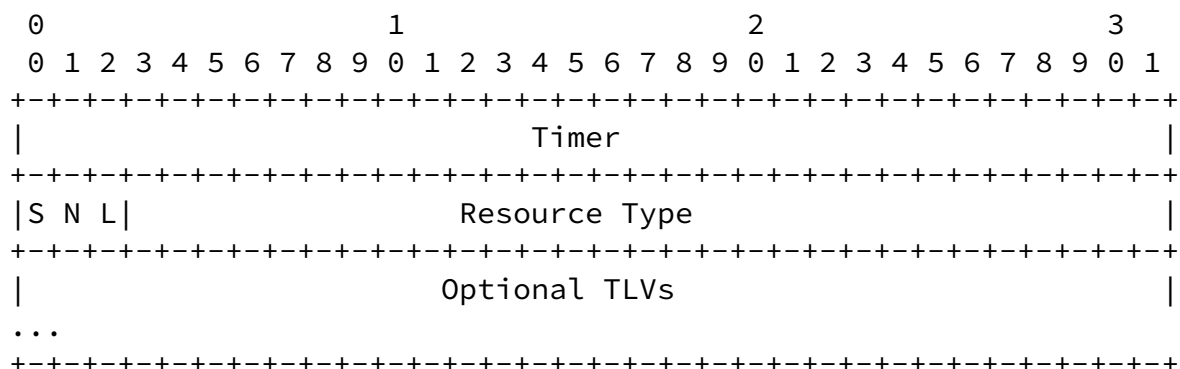
- o Notification-type=TBD: Pending Reservation cancelled
- o Notification-value=TBD (sug 1): PCC cancels a set of reservation requests.

[3.4.](#) RESERVATION object format

RESERVATION Object-Class is to be assigned by IANA.

RESERVATION Object-Type is to be assigned by IANA (recommended value=1)

The RESERVATION object indicates the intention of the PCC to set up the requested path and request the PCE to reserve the resources of the computed path to avoid being used by other requests.



- o Timer is the value in ms of the time that the resources should be blocked, encoded as a 32 bit unsigned integer.
- o Resource Type indicates the type of resource to be reserved. A value of 0 means the default resource type:
 - * Bandwidth (PSC, L2SC, ...)
 - * Time Slot (Sonet/SDH TDM)
 - * Tributary Slot (G709 OTN ODU-k TDM)
 - * Wavelength (G709 OTN OCh or WSON LSC)
- o Bit L: if set, TE Links should be part of the reservation, and excluded from subsequent request.
- o Bit N: if set, Nodes should be part of the reservation.
- o Bit S: if set, the set of SRLG (Shared Risk Link Group) deduced from the associated resources (i.e., the union of SRLGs of the links) should be part of the reservation.

Currently no TLVs are defined.

3.5. RESERVATION_CONF object format

The RESERVATION_CONF object is optional. The RESERVATION_CONF object indicates that the PCE has reserved the resources of computed path to avoid being used by other requests. The RESERVATION_CONF object is sent in the PCRep.

The RESERVATION_CONF Object-Class is to be assigned by IANA.

The RESERVATION_CONF Object-Type is to be assigned by IANA
(recommended value=1)

The format of the RESERVATION_CONF object body is:

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reservation ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reservation timer                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S N L|                               Reservation Type                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Timer is the value in ms of the time that the resources are blocked. The PCE May decide to apply a different value that the one requested by the PCC.

A PCC MUST NOT send a RESERVE_RESPONSE object if the client has not requested a RESERVATION in the PCReq message. A PCE MAY apply reservations as a means of internal policy and/or operation.

3.6. RESERVATION_ID TLV

The TLV indicates the reservation ID (Type TBA by IANA).

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reservation ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

4. Procedures

A client that wishes to request a path computation with reservation shall:

- o Include a PCC_REQ_ID and RESERVATION objects in the involved Request within the PCReq message.
- o Specify what level of reservation to apply after the request.

Upon receiving a PCReq with a Resource Reservation object, the PCE may:

- o Perform the Path Computation using the local Traffic Engineering Database which has been extended to account for resources that have been marked reserved or blocked and which SHOULD not be used while blocked. This includes both synchronized / dependent path computations via SVEC or individual Path Computations requested in the request_list.
- o For the successful path computations, and for all paths corresponding to a given Request, determine the type of resources to be blocked (marked as reserved) with the granularity requested by the client once mapped to PCE policies.
- o It will start a local timer associated with this blocking action.

- o Send the Responses (successful or not) using PCRep message(s) and, where appropriate, indicate the level of reservation and associated period.
- o For subsequent requests, perform path computation as detailed above, updating the local TED with potential new reservations.

Whenever a timer expires, the PCE will:

- o Remove the reservation status / blocking that affected the reservation (e.g. add the previously subtracted unreserved

bandwidth, mark the label, wavelength or time slot as available, etc).

- o Delete any data related with this blocking action.

Whenever a traffic engineering update reaches the PCE, the PCE will:

- o If the reserved resource can be uniquely identified in the traffic engineering update, keep the reservation
- o If the reserved resource cannot be uniquely identified in the traffic engineering update, delete the reservation

[5.](#) Use cases

This section aims to show the use cases of the proposed possibility to activate the limited context awareness.

[5.1.](#) Multiple LSP restoration in a WSON network

One of the most challenging scenarios for a PCE-based architecture is the one of PCE-based dynamic multiple LSP restoration in a WSON network without pre-planning. In the event of a network failure affecting a high number of LSPs (e.g. a fiber cut), a PCE could potentially receive a significant amount of restoration requests in a short period of time from different PCCs.

One of the various challenges in this scenario is the fact that the PCE needs to sequentially perform multiple independent path computations including routing and wavelength assignment. In this scenario, a stateless PCE would rely on TED information, which could potentially be up-to-date before the first incoming request (e.g. in case the routing algorithm has disseminated the failure event), but will definitely be outdated for subsequent requests.

It might be expected that the paths calculated for different

connections would rely on the same nodes, TE links or even labels (lambdas). It might occur at the signaling phase that multiple connection requests are contending for the same resources. After the eventual failure in the establishment of some of the connections,

subsequent requests to the PCE would be triggered. After a number of loops, the PCE-based restoration would be eventually solved, but the potential number of retries could be significantly high.

The main issue is that the stateless PCE relied on an outdated TED to perform path computation. As the subsequent connection request is expected to be computed immediately, there is either no time for the routing algorithm to update the TED after a successful signaling or for the signaling process to successfully finish.

In this context, the availability of a limited context aware PCE could potentially solve the issue in a graceful fashion. Each of the restoration path requests will have an associated Resource Reservation object, which will state the kind of resources and the amount of time they should be blocked.

The PCE will then temporarily 'mark' the resources as blocked, so as not to consider them in subsequent connection requests, and thus avoiding the contention at the signaling phase. The timer should be in line with the LSP set up time and TED time to update.

This use case might be solved in the PCE by having a policy to implicitly pre-reserve the resources for a given time, which can be based on the mean time between a PCRep and a TED update indicating that the labels are not available. The drawback of this implicit reservation is that path establishment time may depend on a variety of factors that may be strongly dependent on the chosen path and technology used (e.g. power equalization algorithms). In this case the PCC has a better view on those aspects and can provide more accurate view on when the TED will be updated.

[5.2.](#) Domain path selection

When selecting the set of domains of a multidomain path, a PCE may request paths to several PCEs of different domains. Thus, the intention of the request is not to establish a LSP, but to obtain a hint on the domain path. Thus, in this case, no Reservation Object would be sent.

Here implicit policies in PCE will be inaccurate as they cannot determine if the PCC will setup the path or not.

[5.3.](#) Multidomain path computation

Once the domain path is known, when computing the actual path, the reservation object can be used. Note that multidomain paths may take a long time to be established, as it involves several AS or domains with different behavior and policies. Thus, it is a way to guarantee the availability of resources.

[6.](#) Manageability Considerations

Standard PCEP [[RFC5440](#)] describes various manageability considerations in PCEP, and most of the manageability requirements are already covered there. Specific aspects are detailed in this section.

[6.1.](#) Control of Function and Policy

In addition to PCE configuration parameters listed in [[RFC5440](#)], the following additional parameters might be required:

- o The ability to enable or disable reservations on the PCE.
- o The ability to retrieve a list of reservations currently active in the PCE.
- o The ability to configure which PCCs are allowed to perform reservations and the ability to configure limits on the timer periods requested. This includes, for example, the configuration of IP based access lists for PCCs.
- o The ability to configure which PCCs are allowed to perform reservations for single-domain and multi-domain scenarios, typically according to pre-defined agreements.
- o The ability to configure which reservation granularity a given PCC group is able to request, and the associated action (error or downgrade).
- o TDB: Advertisements of capabilities via IGP and configurability

[6.2.](#) Information and Data Models

A number of MIB objects have been defined for general PCEP control and monitoring of P2P computations in [PCEP-MIB]. For the time being, no extra models are considered although it could be possible that current means to retrieve information from the PCE be extended

to include eventual resource reservations.

[6.3.](#) Liveness Detection and Monitoring

Other than the considerations expressed in [[RFC5440](#)], a PCE could provide extensions to [MONITORING] to verify reservation status, and to obtain statistics on the system.

[6.4.](#) Verifying Correct Operation

There are no additional requirements for verifying the correct operation of the PCEP sessions. Future MIB objects could facilitate verification of correct operation and reporting of reservations and errors.

[6.5.](#) Requirements for Other Protocols and Functional Components

The method for the PCC to obtain information about a PCE capable of reservation may include extensions to IGP protocols.

[6.6.](#) Impact on Network Operation

It is expected that the use of PCEP extensions specified in this document will not significantly increase the level of operational traffic. However, mis-configured, excessive reservation requests, excessive reservation periods, or excessive granularity may increase the number of failed requests or cause the PCE to provide sub-optimal routes due to existing reservations. Coarse reservations may also limit the resources that are available for a PCE in order to serve requests.

An excessive number of reservation requests and reservation cancellations may degrade server performance. A PCE SHOULD provide a means to control the rate of messages with reservation, extending the proposed mechanism of [[RFC5440](#)].

[7.](#) Security Considerations

In the event of an unauthorized path computation request with mandatory resource reservation, or in case of a (distributed) denial of service attack, the subsequent state/context managed within the

PCE may be disruptive to the network, resulting in performance degradation or sub-optimal computed routes. Implementations should conform to the relevant security requirements of [[RFC5440](#)] that specifically help to control unauthorized requests. These mechanisms include securing the PCEP session requests and responses using TCP security techniques, authenticating the PCEP requests and responses to ensure the message is intact and sent from an authorized node, providing policy control by explicitly defining which PCCs are

allowed to perform resource reservations to the PCE and disallowing reservation requests that may block an excessive amount of resources.

[8.](#) IANA Considerations

IANA maintains a registry of PCEP parameters. A number of IANA considerations have been highlighted in previous sections of this document.

[8.1.](#) RESERVATION object

[8.2.](#) RESERVATION_CONF object

[8.3.](#) RESERVATION_ID TLV

[8.4.](#) PCEP Errors

For the RESERVATION object, the default error procedures regarding supported unknown objects defined in [[RFC5440](#)] apply

- o Unsupported Reservation Option
- o Reservation Forbidden by Policy
- o Unknown Reservation Request

[9.](#) Contributing Authors

Telefonica I+D

Victor Lopez

Don Ramon de la Cruz

email: vlopez@tid.es

Francisco Javier Jimenez Chico

10. Acknowledgements

The authors thank Meral Sherazipur for the discussions and suggestions in the topic.

Gonzalez de Dios, et al. Expires September 10, 2012

[Page 18]

Internet-Draft PCEP Ext for Reserv of Resources in PCE

March 2012

11. Normative References

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", [RFC 5541](#), June 2009.
- [RFC5557] Lee, Y., Le Roux, JL., King, D., and E. Oki, "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", [RFC 5557](#), July 2009.
- [RFC5886] Vasseur, JP., Le Roux, JL., and Y. Ikejiri, "A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture", [RFC 5886](#), June 2010.

Authors' Addresses

Oscar Gonzalez de Dios (editor)

Telefonica I+D
Don Ramon de la Cruz
Madrid, 28006
Spain

Phone: +34 913328832
Email: ogondio@tid.es

Ramon Casellas
CTTC
Av. Carl Friedrich Gauss n7
Castelldefels, Barcelona 08860
Spain

Phone:
Email: ramon.casellas@cttc.es

Gonzalez de Dios, et al. Expires September 10, 2012

[Page 19]

Internet-Draft PCEP Ext for Reserv of Resources in PCE

March 2012

Cyril Margaria
Nokia Siemens Networks
St Martin Strasse 76
Munich, 81541
Germany

Phone: +49 89 5159 16934
Email: cyril.margaria@nsn.com

Young Lee
Huawei
1700 Alma Drive, Suite 100
Plano, TX 75075
U.S.

Phone: (972) 509-5599
Email: leeyoung@huawei.com

Fatai Zhang
Huawei
F3-5-B RD Center
Bantian, Longgang District, Shenzhen 518129
P.R.China

Phone:
Email: zhangfatai@huawei.com