

COPS usage for Path Computation Servers (COPS-PCS)
<[draft-brunner-mpls-cops-pcs-00.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

This memo proposes to use COPS for the communication between a Label Switching Router (LSR) and a Path Computation Server (PCS). Path computation is in much regard a complex function and might be out-sourced. For this reason a protocol between an LSR and a Path Computation Server is needed. This memo proposes to use COPS as a base protocol for that task.

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 [RFC 2119](#).

1 Introduction

Path computation in MPLS and GMPLS might be a computationally complex function especially if several constraints need to be taken into account. Therefore a protocol between an LSR and a path

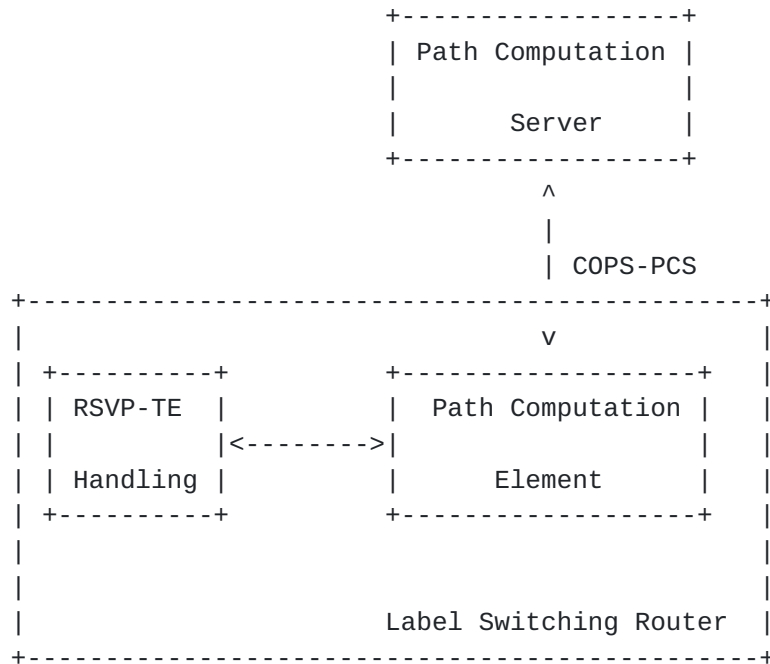
computation server is needed.

M. Brunner

[Page 1]

1.1 Framework

The following figure shows the basic framework, where COPS-PCS is applied. The path computation server can be located at a central location or on LSR itself. The interaction between the local path computation element and the RSVP-TE protocol handling is out of scope of this document.



1.2 Motivation

[3] proposes to use RSVP extensions [1][2]. The advantage of using RSVP lies in the easiness of reusing the objects from RSVP and RSVP-TE for that particular communication. On the other hand, it is not natural to use RSVP for client server communication. Additionally, new RSVP messages need to be defined.

Therefore this memo proposes to use COPS for the LSR to PCS communication. COPS is a protocol which might already be used for admission control for RSVP and therefore also for RSVP-TE. For inter-domain use of RSVP-TE implementing authentication and authorization, COPS or similar mechanisms must be used anyway. Additionally, COPS as a protocol already has the notion of running clients on routers and a server somewhere on the network. Additionally, it has been design to be used together with RSVP, which makes it easy to extend it for RSVP-TE. Whether the server runs on an LSR itself or on separate entity does not matter. Definitely the path computation server needs topology information in order to perform its task. But how to get that information is out of scope of this document.

The basic operation of COPS nicely covers the message types used. Basically, the COPS request message is used to request path computation and a COPS decision message replies the computed paths.

To incorporate RSVP objects into COPS requests and decisions has already been foreseen.

Note that this memo does not define any policy-based admission control. Nor does it define an RSVP-TE extension to the Usage of COPS for RSVP [4]. However, such an RSVP-TE extension might include the semantic of this memo.

Actually, [RFC2749](#) COPS usage for RSVP might be used directly for path computation, because it specifies that all RSVP object in an RSVP message are sent to Policy Decision Point (PDP), in our case RSVP-TE messages sent to the path computation server.

However, since policy decisions for admission control and path computation are inherently different tasks, we propose to add a new COPS client type with restricted functionality not including policy decisions. But the proposal takes advantage of the COPS features for synchronizing states in case the PCS is a statefull implementation.

[2](#) RSVP-PCS values for COPS objects

[2.1](#) Client Type

RSVP-PCS is client-type [0x03, IANA]

[2.2](#) Context Object

In COPS-PCS, only R-Type 0x01 = Incoming-Message request is used. R-Type 0x01 MUST be implemented; all other R-Types MAY be implemented.

The semantics of the context object is as follows:

Incoming-Message request

This context is used when a PEP gets a RSVP-TE Path message in order to get the path computed.

[2.3](#) Client-Specific Information

The client specific information contains all the required information about path computation request and decisions. Since [4] already defines that all RSVP object are sent from the PEP to the PDP (in our case called Path Computation Server), also the base specification of COPS usage for RSVP would work. However, see Section below on the RSVP objects, which MUST be included and supported.

[2.4](#) Decision Object

For COPS-PCS only two commands apply.

Install: the decision contains a positive answer, meaning the path has been computed.

M. Brunner

[Page 3]

Remove: the negative answer; the PCS was not able to compute the path with the given constraints.

If the Trigger Error flag is set, RSVP-TE SHOULD schedule a PathErr in response to a path message.

3 Client-specific Information objects

In order to simplify Path Computation Server (PDP) implementations, we list the RSVP-TE object, which MUST be send to a PDP with client-type COPS-PCS. Every other RSVP object encapsulated within a COPS request is skipped and not evaluated in any regard. If the listed objects are not contained in the request message the path computation server MUST return an <Error> in the decision message indicating, "Mandatory client-specific info missing".

3.1 The RSVP-TE objects in a request message

The request MUST contain the Session object, when C-Type == LSP_TUNNEL_IPv4 or LSP_TUNNEL_IPv6.

It MUST contain the PCS object as defined below.

It might contain the Explicit Route object if included in RSVP-TE.

It might also contain the session attribute object. SESSION_ATTRIBUTE object (class=207, C-type=1) allows carrying setup and holding priorities, resource affinities, etc.

It might contain the sender TSPEC if bandwidth is constraint.

3.2 The RSVP-TE objects in a decision message

Explicit Route object as computed by the Path Computation Server. If no PCS object is contained, the Explicit Route object is copied to the RSVP-TE message and the message is sent towards the next hop in the ERO object.

Additionally, the PCS object might be contained if special handling was requested.

3.3 New COPS object for COPS-PCS

The PCS object is a new object encapsulated in a client specific information object (clientSI) (C-Num =9, C-Type = 2 (named)).

We currently only define one object encapsulated in the named client-specific information. Therefore, no TLV type of object structure is defined.

0	1	2	3	
+	+	+	+	+
ETC	T-Type	Prot-Elem	Flags	

+-----+-----+-----+-----+

Element-to-compute (ETC) : The type of path to be computed.

0x00 default, computes one primary path to the destination

0x01 p+b, computes the primary and the backup (type of backup depends on the protection element type see below)

0x02 backup, computes the backup for a given primary. If this is set, the primary path needs to be in the request as ERO.

Topology-Type (T-Type): Since especially for GMPLS several topologies are possible, this identifies the topology the PCS should calculate the path.

Protection-Element (Prot-Elem): The element, which needs to be protected in case a backup path needs to be computed (Element-to-compute set to 1 or 2.

0x00 default, no backup to be computed

0x01 link, protect against the next link failure

0x02 node, protect against next node failure

0x03 path, compute backup path up to the destination

Flags: A set of bits controlling the path computation.

0x1: Re-optimization: the field defines that the request as well as the decision is a re-optimization. The re-optimization could be triggered by the PCS or the LSR.

Other objects of parameters in the COPS-PCS object are for further study.

4 Statefull versus Stateless PCS

A PCS can be implemented statefull or stateless, which means the PCS can store all the paths (primary and backup) it has computed, and take them into account for future path computation. This means the state between PCS and the LSR needs to be synchronized upon state change.

Statefull PCS implies that if the LSR receives a RSVP PathTear or ResvTear message, it needs to communicate this fact to the PCS. According to [RFC 2749](#) [4], PathTear and ResvTear are not valid message types in the M-Type of the Context Object. Similarly, PathErr or ResvErr must be reported.

Therefore, the LSR MUST send a Delete Request State (DRQ) message to the PCS on receipt of PathTear or ResvTear. The DRQ contains a reason object as defined in [RFC 2748](#) [5]. No client specific sub-code is defined. For RSVP tear down messages the reason is Tear (4). If the LSP with that particular route is not refreshed, reason

Timeout (5) is used.

Statefull PCS MUST be notified about the failure or success of setting up the LSP tunnel with the computed ERO. Upon successful

M. Brunner

[Page 5]

receipt of the RSVP Resv message, the LSR MUST send a Report State message to the PCS. The report type is set to success. The message SHOULD contain the record route object of the RSVP message (RRO), if available. The RRO is used by the PCS in case the path computes was a loose one, then it must update the state for future computation.

Even so COPS is well supporting statefull PCS, the whole implementation gets much easier with stateless PCS. However, stateless PCS must get information about the allocation of resource by other means, when bandwidth constraints are taken into account.

5 Security Considerations

The security considerations have been handled in the Security Considerations section of [5]. The same considerations apply here.

6 Reference

- [1] Braden, R., Zhang, L., Berson, S., Herzog, A., Jamin, S., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [2] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [3] J.P. Vasseur et al., "RSVP Path computation request and reply messages", [draft-vasseur-mpls-computation-rsvp-03.txt](#), June 2002.
- [4] S. Herzog et al., "COPS usage for RSVP", [RFC 2749](#), January 2000.
- [5] D. Durham et al., "The COPS (Common Open Policy Service) Protocol", [RFC 2748](#), January 2000.

7 Author's Addresses

Marcus Brunner
NEC Europe Ltd.
Network Laboratories
Adenauerplatz 6
D-69115 Heidelberg
Germany
E-Mail: brunner@ccrle.nec.de

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it

or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this

document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

