

Internet Draft

June 2002

Network Working Group

Internet Draft

Document: [draft-vasseur-mpls-computation-rsvp-03](#)

JP Vasseur (editor)

Carol Iturralde

Cisco Systems, Inc

IETF Internet Draft

Expires: December, 2002

Raymond Zhang

Infonet Services

Corporation

Xavier Vinet

Equant

Satoru Matsushima

Japan Telecom

Alia Atlas

Avici System

June 2002

RSVP Path computation request and reply messages

[draft-vasseur-mpls-computation-rsvp-03](#)

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/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

Abstract

This document describes extensions to RSVP-TE to support a new message type called a "Path computation" message. This message is to

Vasseur et al.

1

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

be used between an LSR and a Path Computation Server, which may be an LSR or a centralized path computation tool. An RSVP Path Computation Request message is used by the head-end LSR to send its request to the Path Computation Server. The Path Computation Server in turn sends an RSVP Path Computation Reply message containing either:

- a positive reply, containing one or more paths, if the request can be satisfied.
- a negative reply if no path obeying the requested constraints can be found. The Path Computation Server may also optionally suggest new constraint values for which one or several paths could be found.

There are many situations where a Path Computation Server may be used. A typical example is in the context of Inter-area MPLS TE. A head-end LSR could request that a Path Computation Server compute one or more paths obeying a specified set of constraints for a TE LSP spanning multiple areas. The path computation server could be a centralized path computation server or an ABR. Another example is the use of a Path Computation Server to compute diversely routed paths between two end points. This may be useful in the context of MPLS TE LSP Path protection or GMPLS LSP Path protection. The computation of Multi-constraints paths requires intensive CPU resources, and may be yet another usage example. Lastly, those protocol extensions could be used as a "UNI" like protocol between a CE (Customer Edge equipment) and a PE (Provider Edge equipment) where the CE is not part of the PE's IGP domain.

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

Contents

1. Introduction -----	4
2. Terminology -----	4
3. RSVP Path computation Request/Reply messages -----	5
3.1 RSVP Path Computation Request message format -----	6
3.2 RSVP Path Computation Reply message format -----	7
3.3 New RSVP objects used in Path computation messages -----	8
3.3.1 REQUEST-ID Object -----	8
3.3.2 METRIC-TYPE -----	10
3.3.3 PATH_COST -----	12
3.3.4 NO-PATH-AVAILABLE object -----	14
3.3.5 NB-PATH object -----	15
3.3.6 PATH-CORRELATED Object -----	22
3.3.7 MIN-BW object -----	22
3.3.8 LSP-BANDWIDTH object -----	23
3.3.9 EXCLUDE-ELEMENT object -----	25
3.3.10 OPAQUE object -----	26

4. Definition of the "closest possible solution"-----	27
5. Path Computation Server discovery -----	27
6. Acknowledgment -----	27
7. Security Considerations -----	28
8. References -----	28
9. Authors' Addresses -----	29

1. Introduction

As mentioned in the abstract, there are various situations where a Path Computation Client may need to compute one or more paths obeying

a specified set of constraints, and may ask a Path Computation Server to perform this task. This exchange does not allocate any resources, it is simply a mechanism by which a client may send a path computation request to a server and get in return a reply (positive or negative). Note also this is not related to policy.

Let's briefly describe a typical sequence of events:

1) the Path Computation Client (an LSR) sends a request to the Path Computation Server (LSR, centralized path computation tool,...). A Path Computation Request message will be sent containing:

- a) already specified objects defined in [9] characterizing the request, and
- b) new objects defined in the present draft related to the request.

2) the Path Computation Request message is sent to the Path Computation Server

3) the Path Computation Server processes the request and sends either:

- a) a positive reply to the client containing one or more computed paths that obey the requested constraints,
- b) a negative reply to the client, which, if requested by the PCC, can optionally contain alternate constraints values for which the request would have been positive and the computed paths which meet the alternate constraints values.

4) the Path Computation Client can in turn:

- a) If the reply is positive
 - i) If the client has sent the same request to several servers in parallel
 - 1. Compare the reply with other replies it received from other servers.
 - 2. Select the preferred path.
 - Otherwise, select the returned path.
 - ii) Establish the LSP using RSVP with extensions as defined in [9].
- b) If the reply is negative
 - i) If the alternative constraints values given by the PCS are acceptable, consider the computed path sent with the replies from other servers to select the best path.
 - ii) Otherwise, send another request to the Path Computation Server with the new constraints (potentially taking into account the returned suggested constraints values by the server, if any).
 - iii) Wait for an answer from other servers, if any.
 - iv) Go to 4).

[2. Terminology](#)

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

Terminology used in this draft

PCS: Path Computation Server (may be any kind of LSR (ABR, ASBR, ...) or a centralized path computation server

PCC: Path Computation Client (any LSR) requesting a path computation of the Path Computation Server.

[3.](#) RSVP Path computation Request/Reply messages

As defined in [rfc2205](#), an RSVP message consists of a common header followed by a body consisting of a variable number of variable-length, typed "objects". As a reminder, the common header format is:

0	1	2	3
+-----+	+-----+	+-----+	+-----+
Vers	Flags	Msg Type	RSVP Checksum
+-----+	+-----+	+-----+	+-----+
Send_TTL	(Reserved)	RSVP Length	
+-----+	+-----+	+-----+	+-----+

See [RFC2205](#) for details.

One new RSVP message type is defined in this draft: a Path Computation Message. The message type is [TBD by IANA].

The Flags field is used to identify whether the message is a path computation request or a reply. Each has different contents, defined below. Flags 0x01-0x8 are reserved. A request has its flag set to [TBD by IANA] and a reply has its flag set to [TBD by IANA].

An RSVP Path Computation Request message is sent by an LSR to request one or more path computations of a Path Computation Server obeying a set of specified constraints. The objects carried in this message may include those defined in [\[8\]](#) and [\[9\]](#), new objects defined in this draft, as well as other objects that may be defined in the future characterizing, for instance, the constraints of the LSP for which one or several paths should be computed. The IP source address is the

IP address of the requesting LSR and the IP destination address is the IP address of the Path Computation Server.

An RSVP Path Computation Reply message is sent by the PCS to the requesting LSR (PCC) to return one or more computed Paths, if any. The object(s) carried will include one or more EROs (Explicit Route Objects), as defined in [9], plus additional objects defined in this draft. If no path can be found, the PCS reply will be negative. An ERROR-SPEC object will be carried in the reply, and optionally additional information as defined in [section 3.3.4](#). The source IP address will be the IP address of the PCS and the destination IP address will be the IP address of the PCC.

RSVP Path computation messages are sent without the router alert option. Path computation messages should be sent in reliable mode as
Vasseur, et al 5

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

defined in [RFC2961](#). This allows an acknowledgment message to be used to acknowledge the receipt of a Path computation message (request and reply). In case of message loss, the message will be fast retransmitted as defined in [RFC2961](#). Note that the DSCP field of the IP packet carrying an RSVP Path Computation message may be set appropriately to provide the appropriate quality of service delivery to the packet.

The same Path Computation Request may be sent to several PCSs. In this case, the decision process used by the PCC to select from among (possibly) multiple replies is a local decision and is beyond the scope of this document.

[3.1](#) RSVP Path Computation Request Message Format

The Path Computation Request message format is as follows:

The Flags field of the common header must be set to [TBD] by IANA to identify a request.

```
<Path computation request> ::= <Common Header> [ <INTEGRITY> ]  
                                [[<MESSAGE_ID_ACK> |  
                                <MESSAGE_ID_NACK>] ... ]  
                                [<MESSAGE_ID>]  
                                <SESSION>
```

```

<REQUEST-ID>
[<NB-PATH>]
[<EXPLICIT_ROUTE>]
[<METRIC-TYPE>]
[<MIN-BW>]
[<EXCLUDE-ELEMENT>]
[<SESSION_ATTRIBUTE>]
[<POLICY_DATA> ... ]
<sender descriptor>

```

```

<sender descriptor> ::= <SENDER_TEMPLATE> <SENDER_TSPEC>
                        [ <ADSPEC> ]
                        [ <RECORD_ROUTE> ]

```

One new (mandatory) object is defined: REQUEST-ID to identify the request. This object is also used to indicate the request's priority and LSP type. See 3.3.1 for details.

The SESSION_ATTRIBUTE object (class=207, C-type=1) allows carrying setup and holding priorities, resource affinities, etc. Other constraints may also be carried in the Path Computation message. Note that any other constraints that could be defined in the future can be expressed as new objects.

The ERO object may also be present in the RSVP Path Computation Request. The reason is that the head-end may want to specify some LSR(s) that the LSP must traverse. At least one sub-object in the ERO must have its L flag bit set to 1, referring to a loose hop. Note that if the path computation request is related to a reoptimization, a second ERO object specifying the existing TE LSP path will be

Vasseur, et al

6

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

inserted to avoid double bandwidth accounting. This ERO will be easily differentiated from the ERO related to the path constraint as it just contains strict hops subobjects. It should be placed after the path constraints ERO.

The optional METRIC-TYPE object allows the PCC to specify the metric the PCS must use in its CSPF to select the "best" path obeying the requested constraints. See the METRIC-TYPE object definition in 3.3.2 for more details.

The optional NB-PATH object (defined in 3.3.5) allows the PCC to specify the number of requested paths to the PCS for the specific set of constraints specified in the RSVP Path Computation Request message.

[3.2](#) RSVP Path Computation Reply message format

The Path Computation Reply message format is as follows:

The Flags field of the common header must be set to [TBD by IANA] to identify a Path Computation Reply.

```
<Path Computation Reply>::=<Common Header> [ <INTEGRITY> ]
                                     [ <MESSAGE_ID_ACK> | <MESSAGE_ID_NACK> ] ... ]
                                     [ <MESSAGE_ID> ]
                                     <REQUEST-ID>
                                     [ <NB-PATH> ]
                                     [ <EXPLICIT_ROUTE> [ <PATH-COST> ] [ <LSP-
                                     BANDWIDTH> ] ] ...
                                     [ <ERROR_SPEC> ]
                                     [ <NO-PATH-AVAILABLE> ]
                                     [ <POLICY_DATA> ... ]
```

The REQUEST-ID is the same REQUEST-ID (contains the same value) as the one contained in the Path Computation Request to which the reply corresponds.

In case of a negative reply (no path obeying the constraints can be found), the PCS must send a reply containing an ERROR_SPEC object with:

- ERROR_CODE [TBD by IANA] and ERROR_VALUE [TBD by IANA]
- "Ipv4 Error Node address" ("Ipv6 Error Node address") set to the Ipv4 (Ipv6) address of the PCS. This must be the same IP address as was used in the Path Computation Request message.

There are various reasons why the Path Computation Server may not be able to satisfy the request:

- the Path Computation Request message was not valid ("unknown object class (error_code=23, "unknown object C-type (error_code=14), "Routing problem" (error_code=24), ...), ... See [\[8\]](#) and [\[9\]](#) for details. In such a situation, the PCS must send the Path Computation Reply message without any ERO objects and without NO-PATH-AVAILABLE object.

- No path can be found obeying the set of requested constraints. If no path can be found by the PCS for the specified constraints, and only in this situation, a NO-PATH-AVAILABLE object may be inserted into the RSVP Path Computation Reply message sent by the PCS. This object (defined in [section 3.3.4](#)) is optional and may specify the constraint(s) that explain(s) why no path has been found. In addition, the PCS may use the NO-PATH-AVAILABLE object to suggest new constraint values for which a path can be found.

- a) If the PCC did not specify that a less-constrained path is of interest (L flag of the REQUEST-ID object set to 0), the Path Computation Reply message must not contain any ERO objects.
- b) If the PCC specified that a less-constrained path is of interest (L flag of the REQUEST-ID object set to 1), then the PCS can optionally use the NO-PATH-AVAILABLE object to indicate new constraint values for which the included path could be found. This is a circumstance where the Path Computation Reply message can contain one or more EROs.

Note that a Path Computation Reply message may contain several EROs if and only if several paths have been requested by the LSR in its Path Computation Request message using the NB-PATH object (see 3.3.5).

Moreover, multiple replies may be combined in the same Path Computation Reply message. This is done using a list of EROs, each following its corresponding REQUEST-ID as shown in the example below.

A reply should not be delayed in order to bundle several path computation results for requests whose priority REQUEST-ID field (see 3.3.1) has been set to "high".

As an example, if a PCC sends several requests:

- L low priority requests with REQUEST-ID = R(i) I=1 ... [L](#)
- P high priority requests with REQUEST-ID = R'(i) I=1 ... P

Then the PCS MUST reply to every high priority request as soon as the computation result is completed. On the other hand, the low priority request results could be bundled in the SAME Path Computation Reply message using the following format: <REQUEST-ID R(1)> <ERO(1)> <REQUEST-ID R(2)> <ERO(2)> <ERO (2)> ... <REQUEST-ID R(L)> <ERO(L)>.

If no path can be found for a specific request, an individual negative Path Computation Reply message must be sent for the corresponding request.

concerns a reoptimization (a new path computation for a TE LSP already in place). This requires for the PCC to provide the path of the TE LSP in place in the path computation request to avoid double bandwidth counting. The ERO must be formed of strict hops only. If the PCC has previously requested a Partial ERO (T bit cleared), a reoptimization can still be requested by the PCC but this implies for the PCS to be statefull (keep a trace of the previously computed path with the associated list of strict hops).

If a set of TE LSPs are in place and a reoptimization is triggered:

- If all the elements of the set share the same constraints, then the PCC should send a single path computation request message containing the list of the EROs for the existing TE LSPs to avoid double bandwidth counting.
- If the TE LSPs of the set do not share the same set of constraint, the PCC should send N path computation requests (where N is the number of TE LSPs) with the R bit of their REQUEST-ID object set and the corresponding ERO for the existing TE LSPs.

B (Bi-directional): when set, the PCC specifies the TE LSP is bi-directional. When cleared, the TE LSP is unidirectional.

L (less-constrained path): when set in the Path Computation Request message, the PCC indicates to the PCS that a less-constrained path is of interest in case of a negative reply (the request cannot be fully satisfied). If set in a Path Computation Reply message, this indicates both that the PCS is capable of computing alternate constraint values for which a path could be found and that the PCC requested information on such a less-constrained path. This flag should not be set in a Path Computation Reply message unless the corresponding Path Computation Request message had it set. In this case, the associated constraints must be provided in the path computation reply making use of the NO-PATH-AVAILABLE object. The PCS will in this case provide the closest possible solution (see [section 4](#)).

Pri (Priority): 2 bits

This field may be used for the requesting LSR to specify to the PCS the urgency of this request. The decision of which priority should be used for a specific request is a local matter and must be set to 0 when unused. A possible use of this field is when several computations may be requested, each having different timing requirements: typically a request for a reoptimization would have a lower priority than a re-routing request.

0x0: normal. No timing requirement specified.

0x3: high. Urgent request to be served as soon as possible.

Epoch: 24 bits

Epoch is as described in [RFC2961](#) and can be the same number.

Request-ID-number: 32 bits

This value (combined with the IP address of the PCC) uniquely identifies the Path Computation Request the message refers to and is incremented every time a new request is sent to the PCS. If no Path computation reply is received from the PCS, the request may be resent with the same Request-ID-number.

The same Request-ID-number may be sent to different PCS's. The Path Computation Reply will be identified by the IP source address of the sender.

The presence of the REQUEST-ID object is mandatory in every Path Computation Request and Reply message.


[3.3.2](#) METRIC-TYPE

The METRIC-TYPE object may be used in Path Computation Request message. This object is optional.

When computing the path(s) obeying a set of specified constraints, the PCS will run a CSPF and will select the "shortest" path from the subset of the topology which meets the constraints. The shortest Path is defined as the path having the lowest cost for a specific

Vasseur, et al

10



The PCC sends a request for the computation of one path obeying the set of specified constraints. The returned path must be the shortest path using the TE metric and must specify the associated cost.

Reply

REQUEST-ID=a

<ERO 1>

<PATH-COST>: metric-type="TE", metric-value=C1 (sum of the TE metric of the links along the path for ERO 1)

Note: if the "C" bit is cleared in the REQUEST-ID object of the path computation request, the PCC may (but is not required to) return the computed path(s) with PATH-COST objects.

Example 2

Request

<SESSION>

<REQUEST-ID>=a, "C" bit=0x1

<METRIC-TYPE>: IGP metric & TE metric

<sender descriptor>

The PCC sends a request for the computation of one path obeying the set of specified constraints. The returned path must be the shortest path using the TE metric.

Reply

REQUEST-ID=a

<ERO 1>

<PATH-COST>: metric-type="TE", metric-value=C1 (sum of the IGP metric of the links along the path for ERO 1)

<ERO 2>

Vasseur, et al

13

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

<PATH-COST>: metric-type="IGP", metric-value=C2 (sum of the TE metric of the links along the path for ERO 2)

[3.3.4](#) NO-PATH-AVAILABLE object

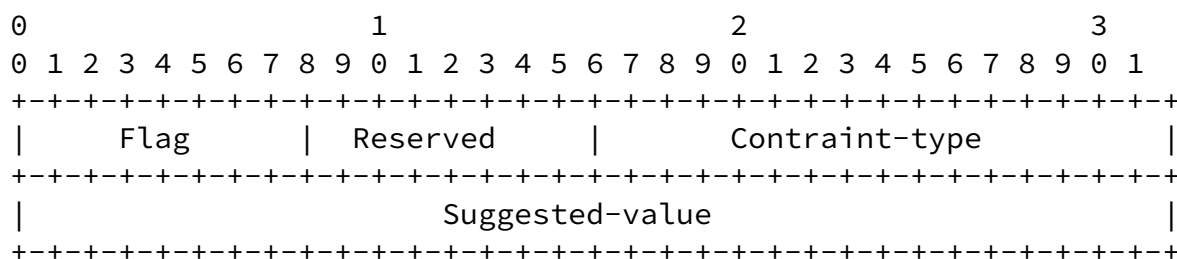
The NO-PATH-AVAILABLE object may be used in Path Computation Reply

message. This object is optional.

When present, it allows the PCS to indicate the unsatisfied constraint(s) (the reason why no path can be found).

- if the L bit in the REQUEST-ID of the path computation request has been set (less-constrained path is of interest) and the PCS is capable of suggesting new constraint values, then the NO-PATH-AVAILABLE object allows the PCS to specify the alternate constraint values for which a path could be found. These new constraint values were used to compute the ERO included in the case where the request failed and an ERROR-SPEC is included in the Path Computation Reply message.
- if the L bit in the REQUEST-ID of the path computation request has not been set and the request failed, a negative path computation reply is returned to the PCC with an ERROR-SPEC. No NO-PATH-AVAILABLE object is included in the reply.

NO-PATH-AVAILABLE Class-Num is [TBD by IANA] - C-Type is [TBD by IANA]



Flags: 8 bits

0x00: the PCS indicates the constraint for which no path can be found but does not suggest any other value for the constraint for which a path could be found.

0x01: the PCS indicates the constraint for which no path can be found and suggests another value for this constraint (as close as possible to the original requested constraint) for which a path could be found. This value is indicated in the suggested-value field.

0x02: the PCS indicates that no path can be found with the requested constraints but an unconstrained path could be found. In this case both the Constraint-type and Suggested-value fields must be set to 0.

Constraint-type: 16 bits

Defines the constraint for which no path has been found by the Path Computation Server.

0x0001 = no path can be found with the requested bandwidth

0x0002 = no path can be found with the requested protection
0x0003 = no path can be found with the requested class affinity
attribute
0x0004 = the path-correlation requested cannot be satisfied
0x0005 = no path can be found since the LSP was requested as bi-
directional

Suggested-value: 32 bits

The PCS may, for each constraint, suggest a value (potentially the closest to the requested constraint in the original Path computation request) for which a path could be found. In this case, the flag must be set to 0x01. For example, if a bandwidth of X is requested by the head-end LSR and a path may be found but with a bandwidth of Y (with $Y < X$), Y may be mentioned by the Path Computation Server in the Suggested-value field.

The suggested-value field must be set to 0x00000000 if the flag field is set to 0x00.

Note that several NO-PATH-AVAILABLE objects may be included in the Path Computation Reply message so that the PCS may indicate the list of constraints for which no path has been found. In the specific case where just an unconstrained path can be found, the Path Computation Server will return a single NO-AVAILABLE-PATH object with the Flags=0x02 (this simplifies the RSVP message in this particular case, instead of having to include one NO-AVAILABLE-PATH per constraint).

A PCS is not required to support the capability of identifying the constraint(s) that cannot be satisfied, and suggesting other values for the constraint(s) and will not do so unless the L bit was set in the request's REQUEST-ID.

[3.3.5](#) NB-PATH object

The NB-PATH object may be used in Path Computation Request message and Path Computation Reply messages. This object is optional in the Path Computation Request message. When present in the Path Computation Request message, it must be also present in the corresponding Path Computation Reply message.

There are many situations where a PCC may desire to get more than one

LSP (and so more than one path) between two points:

- to perform load balancing between two LSRS in order to reduce the traffic disruption impact in case of link failure,
- when a single TE LSP with the requested bandwidth cannot be found due to lack of bandwidth. In this case, the PCC could request a set of N TE LPS such that the sum of their bandwidth is equal to a known value X:
 - without specifying N,
 - without specifying N but a maximum for N,

Vasseur, et al

15

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

Also, in this case the PCC may want specify the minimum of bandwidth for each TE LSPs in the set (a set of 10000 TE LSPs of 10K bandwidth each is likely to be undesirable). This type of request is called a global bandwidth request.

- with path protection to compute both the primary and backup TE LSP (in this case they must be diversely routed),
- ...

Also:

- the PCS requests the computation of N (potentially correlated) paths sharing the same set of constraints. By correlated, we mean link/node/SRLG diverse.
- the PCS requests the computation of N correlated paths each path having a different set of constraints. Such a request is called a global request and is made of individual requests. For each individual request, a specific path computation request message is sent to the PCS. A global path computation request is said satisfied if each individual request can be satisfied with their corresponding correlation.

When the NB-PATH object is not present, exactly one path must be returned (provided one path satisfying the constraints exists).

NB-PATH Class-Num is [TBD by IANA] C-Type is [TBD by IANA]

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Res|N|S|           Path-correlation           |           number-path           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Flags: 4 bits

S bit:

- S=0: the PCS requests the computation of (potentially correlated) number-path paths that share the same set of constraints.
- S=1: the PCS requests the computation of number-path correlated paths having different set of constraints. The set of constraints of each path is defined in a different Path computation request message identified by the REQUEST-ID_number. The list of requests is specified in the PATH-CORRELATED object (defined in [section 3.3.6](#)). Note that in this case, the PCS starts the path computation after having received the number-path Path computation request.

N bit:

- N=0: the PCS requests the computation of exactly number-path paths
- N=1: the PCS requests the computation of a maximum of number-path paths.

Note

Vasseur, et al

16

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

In some situation, the PCC may want to request the computation of a set of n TE LSPs such that the sum of their bandwidth is equal to a specified value X. In this case, the sender-descriptor included in the path computation request must reflect this value X for the bandwidth. n may be considered as a strict value (N bit =0 in the NB-PATH object and number-path=n), a maximum (N bit =1 in the NB-PATH object and number-path=n) or may be unknown (number-path=0xFFFFF). Then in order to specify that each element of the set of TE LSPs must have a minimum bandwidth, the MIN-BW should be inserted (see 3.3.7).

Path-correlation

Defines the correlation between the various requested paths.

0x000: no path correlation

0x001: the various paths must be diversely routed (link diverse)
0x002: the various paths must be diversely routed (node diverse)
0x003: the various paths must be diversely routed (SRLG disjoint)
0x004: the various paths must be diversely routed (link and SRLG)
0x005: the various paths must be diversely routed (node and SRLG)

number-path

Defines the number of requested paths obeying the specified constraints to the Path Computation Server. The number-path must always be strictly greater than 1. If just one path is required no NB-PATH object must be inserted in the request.

When the value 0xFFFF is specified, the PCC requires the computation of a number of TE LSPs which is not known a priori: typically, when the PCC requires the computation of a set of TE LSPs such that the sum of their bandwidth is equal to X but does not know the number of TE LSPs that will be required and does not impose any bound. When number-path=0xFFFF, the N bit must be set to 1.

It may of course not be possible to compute number-path paths that obey the specified constraints; this will be more likely the case if diverse routing is requested.

A NB-PATH object must be used in every Path Computation Reply message if the corresponding Path Computation Request message also contained an NB-PATH object.

If the PCS requests the computation of (potentially correlated) number-path paths sharing the same set of constraints:

- The S bit flag of the NB-PATH object in the Path computation request must be set to 0,

If the reply is positive:

If exactly N EROs have been requested (number-path=N and N bit=0 in the NB-PATH object of the request), then the number-path field of NB-PATH in the reply is N. If up to N EROs have been requested (number-path=N and N bit=1 in the NB-PATH object of the request), then the number-path field of NB-PATH in the reply should be M

($M < N$), the number of paths which were found and necessary to meet the request.

If the reply is negative:

The number-path field of NB-PATH in the Path computation reply contains the number of paths that could be computed by the PCS. If N EROs have been requested (number-path=N and N bit=0 in the NB_PATH object of the request) but the PCS can only find M EROs obeying the constraints ($M < N$), the PCS will return a Path Computation Reply message with an NB-PATH object where number-path=M. If up to N EROs have been requested (number-path=N and N bit=1 in the NB_PATH object of the request) but the PCS cannot meet the constraints with the M paths that it can find ($M > N$), then the PCS will return Path Computation Reply message with an NB-PATH object where number-path=M.

If the request desired information on less-constrained paths in the event of a failure (L bit=1 in REQUEST-ID), then one or more EROs may be included along with one or more NO-PATH-AVAILABLE objects. The reply will also contain an ERROR_SPEC object.

If the request did not desire information on less-constrained paths in the event of a failure (L bit=0 in REQUEST-ID), no ERO object will be inserted in the reply.

The reply will also contain an ERROR_SPEC object.

If the PCS requests the computation of number-path correlated paths having different set of constraints, the NB-PATH object must just be present in the first of number-path requests.

- The S bit flag of the NB-PATH object in the Path computation request must be set to 1,
- In a request for number-path sharing different set of constraints, the N bit must always be set to 0,
- The path computation request must contain a PATH-CORRELATED object listing the number-path REQUEST-ID. The set of constraints of each path is defined in number-path path computation requests,

If the reply is positive:

If N correlated paths have been requested (number-path=N

in the NB-PATH object of the request), the number-path field of NB-PATH in the reply is N.

If the reply is negative:

Vasseur, et al

18

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

The number-path field of NB-PATH in the Path computation reply contains the number of correlated paths that could be computed by the PCS. If N path computation have been requested (number-path=N in the NB-PATH object of the request) but the PCS can only find M paths obeying the constraints ($M < N$), the PCS will return a Path Computation Reply message with an NB-PATH object where number-path=M. The PCS will also include a PATH-CORRELATED object listing the set of N-M requests that could not be satisfied (identified by their REQUEST-ID-number).

If the request desired information on less-constrained paths in the event of a failure (L bit=1 in REQUEST-ID) and the PCS is capable of computing alternate constraint values for which a path could be found, then:

- the EROs for the satisfied individual requests must be provided,
- one or more EROs may be included along with one or more NO-PATH-AVAILABLE objects for the individual requests that could not be satisfied as requested.

If the request did not desire information on less-constrained paths in the event of a failure (L bit=0 in REQUEST-ID), no ERO object will be inserted in the reply.

The reply will also contain an ERROR_SPEC object.

Example 1 (Path computation of 5 paths sharing the same set of constraints). Less constrained path is of interest.

The PCC sends a request to the PCS with the following constraints:

Request

<SESSION>

<REQUEST-ID>=a, L bit=1

<NB-PATH>: Flag: N=0, S=0, number-path=5, Path-correlation=0x02 (Node diversely routed paths)

sender descriptor: bw=x, ...

If just M=3 ($M < N=5$) paths can be found by the PCS with the requested constraints but 5 such paths can be found if the requested bandwidth is y ($y < x$), then the PCS may send the path computation reply of this form:

Reply

REQUEST-ID=a

<ERROR-SPEC> (specifying a negative reply)

<NB-PATH>: Flag: N=0, S=0 number-path=3, Path-correlation=0x02

<ER01>

<ER02>

Vasseur, et al

19

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

<ER03>

<NO-PATH-AVAILABLE>: Flags=0x01, Constraint-type=0x0001 (bandwidth), Suggested-value=y

Which means:

- the reply is negative (the request cannot be satisfied with the specified constraints)
- the unsatisfied constraint is "Bandwidth"
- Exactly M=3 EROs ($M < N$) could be found with the requested constraints. Their corresponding paths are ER01, ER02 and ER03.
- N=5 node diversely routed path could be found if the bandwidth is set to y.

Then the PCC could decide to resend a path computation request for N EROs but with the new suggested value for the bandwidth (y).

Example 2 (Path computation of N=4 paths sharing different set of constraints). Less constrained path is not of interest.

The PCC sends a request to the PCS with the following constraints:

Request 1

<SESSION>

<REQUEST-ID>=a, L bit=0

<NB-PATH>: Flag: N=0, S=1, number-path=4, Path-correlation=0x02 (Node diversely routed paths)

<PATH-CORRELATED>=a,b,c,d

sender descriptor: bw=x, ...

The PCS receives a Global path computation request for N=4 paths sharing different set of constraints. The PCS will start the computation after having received the N=4 path computation requests having the request-ID-number a, b, c and d.

Request 2

<SESSION>

<REQUEST-ID>=b

sender descriptor: bw=y, ...

Request 3

<SESSION>

<REQUEST-ID>=c

sender descriptor: bw=z, ...

Request 3

<SESSION>

<REQUEST-ID>=d

sender descriptor: bw=w, ...

Reply 1

REQUEST-ID=a

<ERROR-SPEC> (specifying a negative reply)

<NB-PATH>: Flag: N=0, S=1, number-path=2, Path-correlation=0x02

<PATH-CORRELATED>: b,c

Vasseur, et al

20

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

Which means:

- the reply is negative (the global request cannot be satisfied with the specified constraints)
- just 2 requests could be satisfied: a and d

- the requests whose request-ID-number are b and c could not be satisfied

Example 3 (Path computation of N=4 paths sharing different set of constraints). Less constrained path is of interest.

The PCC sends a request to the PCS with the following constraints:

Request 1

<SESSION>

<REQUEST-ID>=a, L bit=1

<NB-PATH>: Flag: N=0, S=1, number-path=4, Path-correlation=0x02 (Node diversely routed paths)

<PATH-CORRELATED>=a,b,c,d

sender descriptor: bw=x, ...

The PCS receives a Global path computation request for N=4 paths sharing different set of constraints. The PCS will start the computation after having received the N=4 path computation requests having the request-ID-number a, b, c and d.

Request 2

<SESSION>

<REQUEST-ID>=b

sender descriptor: bw=y, ...

Request 3

<SESSION>

<REQUEST-ID>=c

sender descriptor: bw=z, ...

Request 3

<SESSION>

<REQUEST-ID>=d

sender descriptor: bw=w, ...

Reply 1

REQUEST-ID=a, L bit=1

<NB-PATH>: Flag: N=0, S=1, number-path=2, Path-correlation=0x02

<PATH-CORRELATED>: b,c

<ERO>

Reply 2

REQUEST-ID=d, L bit=1

<ERO>

Reply 3

REQUEST-ID=b

<ERROR-SPEC> (specifying a negative reply)

```
<NO-PATH-AVAILABLE>: Flags=0x01, Constraint-type=0x0001 (bandwidth),
Suggested-value=y'
<ERO>
```

```
<NO-PATH-AVAILABLE>: Flags=0x01, Constraint-type=0x0001 (bandwidth),
Suggested-value=z'
<ERO>
```

- the reply is negative (the global request cannot be satisfied with the specified constraints)
- just 2 requests could be satisfied: a and d and their corresponding ERO are provided.
- the requests whose request-ID-number are b and c could not be satisfied as specified,
- the unsatisfied constraint for request b and c is "Bandwidth"
- the global request can be satisfied if the request bandwidth for request b is y' and the requested bandwidth for c is z' . The corresponding EROs are provided in the reply since the L bit in the REQUEST-ID object of the path computation request had been set.

The PATH-CORRELATED object may be used in Path Computation Request message. This object is optional.

It allows the PCC to request to the PCS the computation of N correlated paths having different set of constraints and contains the list of the REQUEST-ID of those paths.

PATH-CORRELATED C-Type is [TBD]

```
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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

```

|
//                               (Subobjects)                               //
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Subobjects

The contents of a PATH-CORRELATED object is a series REQUEST-ID objects.

3.3.7 MIN-BW object

The MIN-BW object is used in Path Computation Request messages. This object is optional and used by the PCC to request a set of TE LSPs

Vasseur, et al 22

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

such that the sum of their bandwidth is determined, the number of elements in the set may or not be specified, and the minimum bandwidth of each element in the set is specified thanks to the MIN-BW object.

Note that the TE LSPs may or not share the same constraints.

MIN-BW Class-Num is [TBD]
 MIN-BW C-Type is [TBD]

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               MIN-BW-LSP                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

MIN-BW-LSP: Minimum bandwidth of any element of the backup tunnel set.

3.3.8 LSP-BANDWIDTH object

The LSP-BANDWIDTH object is used in Path Computation reply messages and is included in any positive path computation reply to specify the bandwidth of a computed TE LSP path when the request was a global

bandwidth request (see the definition in [section 3.3.5](#)). It immediately follows the ER0.

LSP-BANDWIDTH Class-Num is [TBD]

LSP-BANDWIDTH C-Type is [TBD]

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               LSP-Bandwidth                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

LSP-Bandwidth: Actual bandwidth determined for the associated path

Example 1: Path computation of a set diversely routed paths sharing the same set of constraints such that:

- number of elements in the set is not specified and not bounded by the PCC,
- the sum of their bandwidths is X,
- each element of the set has a minimum bandwidth of m

Request

<SESSION>

<REQUEST_ID>=a

<NB_PATH>: Flag: N=1, S=0, number-path=0xFFFFF,

Path_correlation=0x02 (Node diversely routed paths)

sender descriptor: bw=X, ...

<MIN-BW>: MIN-BW-LSP=m

Vasseur, et al

23

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

Supposing the PCS can find such a solution with 4 node diverse TE LSPs: LSP1, LSP2, LSP3, LSP4 such that:

- TE LSP 1 : ER01, Bw=x1, x1>m
- TE LSP 2 : ER02, Bw=x2, x2>m
- TE LSP 3 : ER03, Bw=x3, x3>m,
- TE LSP 4 : ER04, Bw=x4, x4>m,
- x1+x2+x3+x4=X

Reply

REQUEST_ID=a

<NB_PATH>: Flag: N=1, S=0, number-path=4, Path_correlation=0x02

<ER01>, <LSP-BANDWIDTH> (Bw=x1)
<ER02>, <LSP-BANDWIDTH> (Bw=x2)
<ER03>, <LSP-BANDWIDTH> (Bw=x3)
<ER04>, <LSP-BANDWIDTH> (Bw=x4)

Example 2: Path computation of a maximum of n diversely routed paths sharing the same set of constraints such that:

- the sum of their bandwidths is X,
- each element of the set has a minimum bandwidth of m

Request

<SESSION>

<REQUEST_ID>=a

<NB_PATH>: Flag: N=1, S=0, number-path=n, Path_correlation=0x02 (Node diversely routed paths)

sender descriptor: bw=X, ...

<MIN-BW>: MIN-BW-LSP=m

Supposing the PCS can find such a solution with 3 node diverse TE LSPs: LSP1, LSP2, LSP3 such that

- TE LSP 1 : ER01, Bw=x1, $x1 > m$
- TE LSP 2 : ER02, Bw=x2, $x2 > m$
- TE LSP 3 : ER03, Bw=x3, $x3 > m$,
- $x1+x2+x3=X$

Reply

REQUEST_ID=a

<NB_PATH>: Flag: N=1, S=0, number-path=n, Path_correlation=0x02

<ER01>, <LSP-BANDWIDTH> (Bw=x1)

<ER02>, <LSP-BANDWIDTH> (Bw=x2)

<ER03>, <LSP-BANDWIDTH> (Bw=x3)

Example 3: Path computation of exactly 2 diversely routed paths not sharing the same set of constraints such that the sum of their bandwidths is X.

Request

<SESSION>

<REQUEST_ID>=a

<NB_PATH>: Flag: N=0, S=1, number-path=2, Path_correlation=0x02 (Node diversely routed paths)

sender descriptor: bw=X, ...

<PATH_CORRELATED>=a,b

Vasseur, et al

...

```
<REQUEST_ID>=b
<NB_PATH>: flag: N=0, S=1, number_PATH=2, Path_correlation=0x02 (Node
diversely routed paths)
sender descriptor: bw=X, ...
...
```

Supposing the PCS can find such a solution with 3 node diverse TE LSPs: LSP1, LSP2, LSP3 such that

- TE LSP 1 : ER01, Bw=x1, x1>m
- TE LSP 2 : ER02, Bw=x2, x2>m
- x1+x2=X

```
Reply
REQUEST_ID=a
<NB_PATH>: Flag: N=0, S=1, number-path=2, Path_correlation=0x02
<ER01>, <LSP-BANDWIDTH> (Bw=x1)
<ER02>, <LSP-BANDWIDTH> (Bw=x2)
```

[3.3.9](#) EXCLUDE-ELEMENT object

The EXCLUDE-ELEMENT object may be used in Path Computation Request message. This object is optional.

It allows the PCC to specify to the PCS another constraint related to the computed path: the exclusion of one or more network elements in the computed path. A network element may be a link, an entire node or even an Autonomous System. The EXCLUDED-ELEMENT object contains the list of network elements to exclude from the computed path. Each network element is represented as a subobject.

EXCLUDE-ELEMENT Class-Num is [TBD]
EXCLUDE-ELEMENT C-Type is [TBD]

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|//                               (Subobjects)                               //
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```


Subobjects

3.3.10 OPAQUE object

The OPAQUE object may be present in both Path Computation Request and Reply message types. This object is optional.

The OPAQUE object may be used by the PCC to transfer information to the PCS (in a Path Computation Request message) or by the PCS to transfer information to the PCC (in a Path Computation Reply message). Opaque object may be used for future extensions and the exact content of the OPAQUE object is beyond the scope of this draft.

OPAQUE Class-Num is [TBD by IANA] - C-Type is [TBD by IANA]



Where

Type: identifies the TLV type

Length: length of the value field in octets

The OPAQUE object may also contain sub-TLV to be defined in the future.

4. Definition of the "closest possible solution"

At several places in this draft the term "closest possible solution" is mentioned. This refers to the solution (the set of paths) the Path Computation Server could find if the request cannot be fully satisfied (negative reply). The definition of "closest possible solution" is determined by the PCS (local decision).

In a future version of this draft, this could be defined by the PCC

in its request (detailing the set of constraints that should be relaxed by order of priority): this is for further study.

Example

- Request for 3 TE LSPs sharing the same set of constraint for 10M of bandwidth each,
- The PCS can find three solutions:
 - Solution 1: LSP1=10, LSP2=10, LSP3=10, LSP4=5, LSP5=5
 - Solution 2: LSP1=8, LSP2=8, LSP3=8, LSP4=8, LSP5=8
 - Solution 3: LSP1=11, LSP2=11, LSP3=11, LSP4=11

Which solution is the closest of the initial request depends on the definition of "closest". It could be:

- Solution 1: 3 LSPs could be found with 10M
- Solution 2: exactly 5 LSPs with the same bandwidth could be found
- Solution 3: 4 LSPs could be found but the total of their bandwidth is 44M closest to the initial request for 50M.

5. Path Computation Server discovery

There are several possibilities for the PCC to learn the PCS(s) location (IP addresses) and capabilities:

- by static configuration: the list of PCS can be manually configured on each LSR, optionally:
 - o with an order of priority,
 - o their respective capabilities,
 - o ...
- using IGP extensions for an automatic PCS discovery (see [14] and [15]).

6. Acknowledgment

The authors would like to thank Bob Thomas, Francois Le Faucheur, Rob Goguen, Anna Charny and Ashok Naranayan for their very valuable comments.

Vasseur, et al

27

7. Security Considerations

No new security issues are raised in this document. See [8] for a general discussion on RSVP security issues.

8. References

- [1] Katz, D., Yeung, D., "Traffic Engineering Extensions to OSPF", [draft-katz-yeung-ospf-traffic-04.txt](#)
- [2] Smit, H. and T. Li, "ISIS Extensions for Traffic Engineering," [draft-ietf-isis-traffic-02.txt](#), work in progress.
- [3] Kompella, K., and Rekhter, Y., "LSP Hierarchy with MPLS TE", Internet Draft, [draft-ietf-mpls-lsp-hierarchy-02.txt](#), Feb., 2001.
- [4] Kompella, K., Rekhter, Y., "Signalling Unnumbered Links in RSVP-TE", Internet Draft, [draft-ietf-mpls-rsvp-unnum-01.txt](#), February 2001
- [5] Ashwood-Smith, P. et al, "Generalized MPLS Signaling - CR-LDP Extensions", Internet Draft, [draft-ietf-mpls-generalized-cr-ldp-01.txt](#), February 2001.
- [6] Ashwood-Smith, P. et al, "Generalized MPLS - Signaling Functional Description", Internet Draft, [draft-ietf-mpls-generalized-signaling-02.txt](#), February 2001.
- [7] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," [RFC 2119](#).
- [8] Braden, R. Ed. et al, "Resource ReserVation Protocol-- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [9] Awduche, D.O., Berger, L., Gan, D.-H., Li, T., Swallow, G., and Srinivasan, V., "RSVP-TE: Extensions to RSVP for LSP Tunnels," Internet Draft, [draft-ietf-mpls-rsvp-lsp-tunnel-09.txt](#), August 2001.
- [10] Berger L., Gan D., Swallow G., Pan P., Tommasi F., Molendini S., "RSVP Refresh Overhead Reduction Extensions", [RFC 2961](#), April 2001.
- [11] Le faucheur, "Use of IGP Metric as a second TE Metric", Internet draft, [draft-lefaucheur-te-metric-igp-00.txt](#).
- [12] Fedyk D., Ghanwani A., Ash J., Vedrenne A. "Multiple Metrics for Traffic Engineering with IS-IS and OSPF", Internet draft, [draft-fedyk-isis-ospf-te-metrics-01.txt](#)
- [13] Vasseur "IS-IS Path Computation Server discovery TLV", [draft-vasseur-mpls-isis-pcsd-discovery-00.txt](#), work in progress.

[14] Vasseur, Psenak, "OSPF Path Computation Server discovery",
Vasseur, et al

28

Internet [draft-vasseur-mpls-path-computation-rsvp-te-03.txt](#) June 2002

[draft-vasseur-mpls-ospf-pcsd-discovery-00.txt](#), work in progress.

9. Author's Addresses

JP Vasseur
CISCO Systems, Inc.
11, rue Camille Desmoulins
92782 Issy les Moulineaux Cedex 9
FRANCE
Email: jpv@cisco.com

Carol Iturralde
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
USA
Email: cei@cisco.com

Raymond Zhang
Infonet Services Corporation
2160 E. Grand Ave.
El Segundo, CA 90025
USA
Office: +310-335-1039
Email: raymond_zhang@infonet.com

Xavier VINET
EQUANT
9 rue du Chêne Germain - BP 80
35512 Cesson Sevigne cedex
FRANCE
Email : xavier.vinet@equant.com

Satoru Matsushima
Japan Telecom
4-7-1, Hatchobori, Chuo-ku
Tokyo, 104-8508 Japan

Phone: +81-3-5540-8214
Email: satoru@japan-telecom.co.jp

Alia Atlas
Avici Systems
101 Billerica Avenue
N. Billerica, MA 01862
email: aatlas@avici.com
phone: +1 978 964 2070