Network Working Group
Internet Draft
Category: Informational
Expires: December 2006

J.L. Le Roux (Editor)
France Telecom

June 2006

Requirements for Path Computation Element (PCE) Discovery

draft-ietf-pce-discovery-reqs-05.txt

Status of this Memo

Abstract

This document presents a set of requirements for a Path Computation Element (PCE) discovery mechanism that would allow a Path Computation Client (PCC) to discover dynamically and automatically a set of PCEs along with certain information relevant for PCE selection. It is intended that solutions that specify procedures and protocols or extensions to existing protocols for such PCE discovery satisfy these requirements.

Conventions used in this document

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.

Table of Contents

1. Contributors

The following are the authors that contributed to the present document:

Jean-Louis Le Roux (France Telecom)
Paul Mabey (Qwest Communications)
Eiji Oki (NTT)
Richard Rabbat (Fujitsu)
Ting Wo Chung (Bell Canada)
Raymond Zhang (BT Infonet)

2. Terminology

Terminology used in this document

LSR: Label Switch Router

TE-LSP: Traffic Engineered Label Switched Path

PCE: Path Computation Element: an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph, and applying computational constraints.

PCC: Path Computation Client: any client application requesting a path computation to be performed by a Path Computation Element.

IGP Area: OSPF Area or ISIS level/area

ABR: IGP Area Border Router (OSPF ABR or ISIS L1L2 router)

AS: Autonomous System

ASBR: AS Border Router

Intra-area TE LSP: A TE LSP whose path does not cross IGP area boundaries.

Inter-area TE LSP: A TE LSP whose path transits through two or more IGP areas.

Inter-AS MPLS TE LSP: A TE LSP whose path transits through two or more ASs or sub-ASs (BGP confederations).

Domain: any collection of network elements within a common sphere of address management or path computational responsibility. Examples of domains include IGP areas and Autonomous Systems.

3. Introduction

The PCE-based network Architecture [PCE-ARCH] defines a Path Computation Element (PCE) as an entity capable of computing TE-LSP paths based on a network graph, and applying computational constraints. A PCE serves path computation requests sent by Path Computation Clients (PCC).
A PCC is a client application requesting a path computation to be performed by a PCE. This can be, for instance, an LSR requesting a path for a TE-LSP for which it is the head-end, or a PCE requesting a path computation of another PCE (inter-PCE communication). The communication between a PCC and a PCE requires a client-server protocol whose generic requirements are listed in [PCE-COM-REQ].

The PCE based architecture requires, that a PCC be aware of the location of one or more PCEs in its domain, and also potentially of some PCEs in other domains, e.g. in case of inter-domain path computation.

In that context it would be highly desirable to define a mechanism for automatic and dynamic PCE discovery, which would allow PCCs to automatically discover a set of PCEs, to determine additional information required for PCE selection, and to dynamically detect new PCEs or any modification of the PCEs' information. This includes the discovery by a PCC of a set of one or more PCEs in its domain, and potentially in some other domains. The latter is a desirable function in the case of inter-domain path computation, for example.

This document lists a set of functional requirements for such an automatic and dynamic PCE discovery mechanism. Section 4 points out the problem statement. Section 5 illustrates an application scenario. Finally, section 6 addresses detailed requirements.

It is intended that solutions that specify procedures and protocols or protocol extensions for PCE discovery satisfy these requirements. There is no intent either to specify solution-specific requirements or to make any assumption on the protocols that could be used for the discovery.

Note that requirements listed in this document apply equally to PCEs that are capable of computing paths in MPLS-TE-enabled networks and PCEs that are capable of computing paths in GMPLS-enabled networks (and PCEs capable of both).

It is also important to note that the notion of a PCC encompasses a PCE acting as PCC when requesting a path computation of another PCE (inter-PCE communication). Hence, this document does not make the distinction between PCE discovery by PCCs and PCE discovery by PCEs.

4. Problem Statement and Requirements Overview

4.1. Problem Statement

A routing domain may, in practice, contain multiple PCEs:

- The path computation load may be balanced among a set of PCEs to improve scalability;
- For the purpose of redundancy, primary and backup PCEs may be used;
- PCEs may have distinct path computation capabilities (multi-constrained path computation, backup path computation, etc.);
- In an inter-domain context there can be several PCEs with distinct inter-domain functions (inter-area, inter-AS, inter-layer), each PCE being responsible for path computation in one or more domains.

In order to allow for effective PCE selection by PCCs, that is to select the appropriate PCE based on its capabilities and perform efficient load balancing of requests, a PCC needs to know the location of PCEs in its domain, along with some information relevant to PCE selection, and also potentially needs to know the location of some PCEs in other domains, for inter-domain path computation purpose.
Such PCE information could be learnt through manual configuration, on each PCC, of the set of PCEs along with their capabilities. Such a manual configuration approach may be sufficient, and even desired in some particular situations, (e.g. inter-AS PCE discovery, where manual configuration of neighbor PCEs may be preferred for security reasons), but it obviously faces several limitations:

- This may imply a substantial configuration overhead;
- This would not allow a PCC to dynamically detect that a new PCE is available, that an existing PCE is no longer available, or that there is a change in the PCE's information.

Furthermore, as with any manual configuration approach, there is a risk of configuration errors.

As an example, in a multi-area network made up of one backbone area and N peripheral areas, and where inter-area MPLS-TE path computation relies on multiple-PCE path computation with ABRs acting as PCEs, the backbone area would comprise at least N PCEs, and the configuration of PCC would be too cumbersome (e.g. in existing multi-area networks, N can be beyond fifty).

Hence, an automated PCE discovery mechanism allowing a PCC to dynamically discover a set of PCEs is highly desirable.

4.2. Requirements overview

A PCE discovery mechanism that satisfies the requirements set forth in this document MUST allow a PCC to automatically discover the location of one or more of the PCEs in its domain.
Where inter-domain path computation is required and policy permits, the PCE discovery method MUST allow a PCC to automatically discover the location of PCEs in other domains that can assist with inter-domain path computation.

A PCE discovery mechanism MUST allow a PCC to discover the set of one or more domains where a PCE has TE topology visibility and can compute paths. It MUST also allow the discovery of the potential inter-domain path computation functions of a PCE (inter-area, inter-AS, inter-layer, etc.).

A PCE discovery mechanism MUST allow the control of the discovery scope, that is the set of one or more domains (areas, ASs) where information related to a given PCE has to be disclosed.

A PCE discovery mechanism MUST allow PCCs in a given discovery scope to dynamically discover that a new PCE has appeared or that there is a change in PCE's information.

A PCE discovery mechanism MUST allow PCCs to dynamically discover that a PCE is no longer available.

A PCE discovery MUST support security procedures. In particular, key consideration MUST be given in terms of how to establish a trust model for PCE discovery.

OPTIONALLY a PCE discovery mechanism MAY be used so as to disclose a set of detailed PCE capabilities so that the PCC may make advanced and informed choices about which PCE to use.

5. Example of application scenario

```
   <----------------AS1-------------------->           <----AS2---
    Area 1           Area 0        Area 2
  R1---------R3-----R5-------R6-----------R9----------R11----R13
  |          |               |             |           |
  |          |               |             |           |
  R2---------R4-----R7-------R8-----------R10---------R12----R14
       |
       |
       --
      |S1|
       --
```

Figure 1

Figure 1 illustrates a multi-area/AS network with several PCEs:

- The ABR R3 is a PCE that can take part in inter area path computation. It can compute paths in area 1 and area 0;
- The ABR R6 is a PCE that can take part in inter-area path computation. It can compute paths in area 0 and area2;
- The ASBR R9 is a PCE that can take part in inter-AS path computation. It is responsible for path computation in AS1 towards AS2;
- The ASBR R12 is a PCE that can take part in inter-AS path computation. It is responsible for path computation in AS2 towards AS1;
- The server S1 is a PCE that can be used to compute diverse paths and backup paths in area 1.

By meeting the requirements set out in this document, the PCE discovery mechanism will allow:

- each PCC in areas 1 and 0 to dynamically discover R3, as a PCE for inter-area path computation, and that R3 can compute paths in area0 and area1;
- each PCC in areas 0 and 2 to dynamically discover R6, as a PCE for inter-area path computation, and that R6 can compute paths in area2 and area0;
- each PCC in AS1 and one or more PCCs in AS2 to dynamically discover R9 as a PCE for inter-AS path computation in AS1 towards AS2;
- each PCC in AS2 and one or more PCCs in AS1 to dynamically discover R12 as a PCE for inter-AS path computation in AS2 towards AS1;
- each PCC in area 1 to dynamically discover S1, as a PCE for intra-area path computation in area1, and optionally to discover its path computation capabilities (diverse path computation and backup path computation).

6. Detailed Requirements

6.1. PCE Information to be disclosed

We distinguish two levels of PCE information to be disclosed by a PCE discovery mechanism:

- General information. Disclosure MUST be supported by the PCE discovery mechanism.
- Detailed information. Disclosure MAY be supported by the PCE discovery mechanism.

The PCE discovery mechanism MUST allow disclosure of general PCE information that will allow PCCs to select appropriate PCEs. This comprises discovery of PCE location, PCE domains supported by the PCEs, and PCE inter-domain functions.

The PCE discovery mechanism MAY also allow disclosure of detailed PCE information. This comprises any or all information about PCE path computation capabilities and alternate PCEs. This information is not part of PCE discovery; this is additional information that can facilitate the selection of a PCE by a PCC. Support of the exchange of this information is optional in the context of the PCE discovery mechanism itself. This does not mean that the availability of this information is optional in the PCE-based architecture, but such information could also be obtained by other mechanisms, such as the PCC-PCE communication protocol.

6.1.1. General PCE Information (Mandatory support)

6.1.1.1. Discovery of PCE Location

The PCE discovery mechanism MUST allow the discovery, for a given PCE, of the IPv4 and/or IPv6 address to be used to reach the PCE. This address will typically be an address that is always reachable, if there is any connectivity to the PCE.

This address will be used by PCCs to communicate with a PCE, through a PCC-PCE communication protocol.

6.1.1.2. Discovery of PCE Domains and Inter-domain Functions

Inter-domain path computation is a key application of the PCE architecture. This can rely on a multiple-PCE path computation, where PCEs in each domain compute a part of the end-to-end path and collaborate with each other to find the end-to-end-path. Inter-domain path computation can also rely on a single-PCE path computation where a PCE has visibility inside multiple domains and can compute an entire end-to-end inter-domain path (that is a path from the inter-domain TE-LSP head-end to the inter-domain TE-LSP tail end).

Hence the PCE discovery mechanism MUST allow the discovery of the set of one or more domains where a PCE has visibility and can compute paths. These domains could be identified using a domain identifier: For instance, an IGP area can be identified by the Area ID (OSPF or ISIS), and an AS can be identified by the AS number.

Also the PCE discovery mechanism MUST allow discovery of the inter-domain functions of a PCE, i.e. whether a PCE can be used to compute or to take part in the computation of end-to-end paths across domain borders. The inter-domain functions include non exhaustively: inter-area, inter-AS and inter-layer path computation. Note that these functions are not mutually exclusive.

Note that the inter-domain functions are not necessarily inferred from the set of domains where a PCE has visibility. For instance a PCE may have visibility limited to a single domain, but may be able to take part into the computation of inter-domain paths, by collaborating with PCEs in other domains. Conversely, a PCE may have visibility in multiple domains but the operator may not want that the

PCE be used for inter-domain path computations.

The PCE discovery mechanisms MUST also allow discovery of the set of one or more domains toward which a PCE can compute paths. For instance in an inter-AS path computation context, there may be several PCEs in an AS, each one responsible for taking part in the computation of inter-AS paths toward a set of one or more destination ASs, and a PCC may have to discover the destination ASs each PCE is responsible for.

6.1.2. Detailed PCE Information (Optional support)

6.1.2.1. Discovery of PCE Capabilities

In the case where there are several PCEs with distinct capabilities available, a PCC has to select one or more appropriate PCEs.

For that purpose the PCE discovery mechanism MAY support the disclosure of some detailed PCE capabilities.

For the sake of illustration this could include the following path computation related PCE capabilities:

- The link constraints supported: e.g. bandwidth, affinities.
- The path constraints supported: maximum IGP/TE cost, maximum hop count;
- The objective functions supported: e.g. shortest path, widest path;
- The capability to compute multiple correlated paths: e.g. diverse paths, load balanced paths;
- The capability to compute bidirectional paths;
- The GMPLS technology specific constraints supported: e.g. the supported interface switching capabilities, encoding types.

And this could also include some specific PCE capabilities:

- The capability to handle request prioritization;
- The maximum size of a request message;
- The maximum number of path requests in a request message;
- The PCE computation power (static parameters to be used for weighted load balancing of requests).

Such information regarding PCE capabilities could then be used by a PCC to select an appropriate PCE from a list of candidate PCEs.

Note that the exact definition and description of PCE capabilities is out of the scope of this document. It is expected that this will be described in one or more separate documents which may be application specific.

6.1.2.2. Discovery of Alternate PCEs

In the case of a PCE failure, a PCC has to select another PCE, if one is available. It could be useful in various situations, for a PCE to indicate a set of one or more alternate PCEs that can be selected in case the given PCE fails.

Hence the PCE Discovery mechanism MAY allow the discovery, for a given PCE, of the location of one or more assigned alternate PCEs.

The PCE Discovery mechanism MAY also allow the discovery, for a given PCE, of the set of one or more PCEs for which it acts as alternate PCE.

6.2. Scope of PCE Discovery

The PCE Discovery mechanism MUST allow control of the scope of the PCE information disclosure on a per PCE basis. In other words it MUST allow control of to which PCC or group of PCCs the information related to a PCE may be disclosed.

The choice for the discovery scope of a given PCE MUST include at least the followings settings:

- All PCCs in a single IGP area

- All PCCs in a set of adjacent IGP areas

- All PCCs in a single AS

- All PCCs in a set of ASs

- A set of one or more PCCs in a set of one or more ASs

In particular, this also implies that the PCE Discovery mechanism MUST allow for the discovery of PCE information across IGP areas and across AS boundaries.

The discovery scope MUST be configurable on a per PCE basis.

It MUST be possible to deactivate PCE discovery on a per PCE basis.

6.2.1. Inter-AS specific requirements

When using a PCE-based approach for inter-AS path computation, a PCC in one AS may need to learn information related to inter-AS capable PCEs located in other ASs. For that purpose, and as pointed out in the previous section, the PCE discovery mechanism MUST allow disclosure of information related to inter-AS capable PCEs across AS boundaries.

Such inter-AS PCE discovery must be carefully controlled. For security and confidentiality reasons, particularly in an inter-provider context, the discovery mechanism MUST allow the discovery scope to be limited to a set of ASs and MUST also provide control of the PCE information to be disclosed across ASs. This is achieved by applying policies (See also section 6.4). This implies the capability to contain a PCE advertisement to a restricted set of one or more ASs, and to filter and translate any PCE parameter (PCE domains, PCE inter-domain functions, PCE capabilities, etc.) in disclosures that cross AS borders. For the sake of illustration, it may be useful to disclose detailed PCE information (such as detailed capabilities) locally in the PCE's AS but only general information (such as location and supported domains) in other ASs.

6.3. PCE Information Synchronization

The PCE discovery mechanism MUST allow a PCC to discover any change in the information related to a PCE that it has previously discovered. This includes changes to both general information (e.g. a change in the PCE domains supported), and detailed information if supported (e.g. a modification of the PCE's capabilities).

In addition, the PCE discovery mechanism MUST allow to dynamically discover new PCEs in a given discovery scope.

Note that there is no requirement for real-time detection of these changes, the PCE Discovery Mechanism SHOULD rather allow discovery of these changes in an order of magnitude of 60 seconds, and the operator should have the ability to configure the Discovery delay.

Note that PCE information is relatively static, and is expected to be fairly stable and to not change frequently.

6.4. Discovery of PCE deactivation

The PCE discovery mechanism MUST allow a PCC to discover when a PCE that it has previously discovered is no longer alive or is deactivated. This may help reducing or avoiding path computation service disruption.

Note that there is no requirement for real-time detection of PCE failure/deactivation, the PCE Discovery Mechanism SHOULD rather allow such discovery in an order of magnitude of 60 seconds, and the operator should have the ability to configure the Discovery delay.

6.5. Policy Support

The PCE Discovery mechanism MUST allow for policies to restrict the discovery scope to a set of authorized domains, to control and restrict the type and nature of the information to be disclosed, and also to filter and translate some information at domains borders. It MUST be possible to apply these policies on a per PCE basis.

Note that the Discovery mechanisms MUST allow disclosing policy information so as to control the disclosure policies at domain boundaries.

Also, it MUST be possible to apply different policies when disclosing PCE information to different domains.

6.6. Security Requirements

The five major threats related to PCE discovery mechanisms are:
- Impersonation of PCE;
- Interception of PCE discovery information (sniffing);
- Falsification of PCE discovery information;
- Information disclosure to non-authorized PCCs (PCC spoofing).
- DoS Attacks

Note that security of the PCE Discovery procedures is of particular importance in an inter-AS context, where PCE discovery may increase the vulnerability to attacks and the consequences of these attacks.

Hence mechanisms MUST be defined to ensure authenticity, integrity, confidentiality, and containment of PCE discovery information:
- There MUST be a mechanism to authenticate discovery information;
- There MUST be a mechanism to verify discovery information integrity;
- There MUST be a mechanism to encrypt discovery information;
- There MUST be a mechanism to restrict the scope of discovery to a set of authorized PCCs and to filter PCE information disclosed at domain boundaries (as per defined in 6.5).

A PCE and PCC MUST be identified by a globally unique ID, which may be for instance a combination of AS number an IP address.

Mechanisms MUST be defined in order to limit the impact of a DoS attack on the PCE discovery procedure (e.g. filter out excessive PCE information change and flapping PCEs). Note also that DOS attacks may be either accidental (caused by a mis-behaving PCE system) or intentional. As discussed in [PCE-COM-REQ] such mechanisms may include packet filtering, rate limiting, no promiscuous listening, and where applicable use of private addresses spaces.

Also, key consideration MUST be given in terms of how to establish a trust model for PCE discovery. The PCE discovery mechanism MUST explicitly support a specific set of one or more trust models.

6.7. Extensibility

The PCE discovery mechanism MUST be flexible and extensible so as to easily allow for the inclusion of additional PCE information that could be defined in the future.

6.8. Scalability

The PCE discovery mechanism MUST be designed to scale well with an increase of any of the following parameters:

- Number of PCCs discovering a given PCE;
- Number of PCEs to be discovered by a given PCC;
- Number of domains in the discovery scope.

The PCE discovery mechanism MUST NOT have an adverse effect in the performance of other protocols (especially routing and signaling) already operating in the network.

Note that there is no scalability requirement with regards to the amount of information to be exchanged.
Information disclosed in the PCE discovery mechanism is relatively static. Changes in PCE information may occur as result of PCE configuration updates, PCE deployment/activation or PCE deactivation/suppression, and should not occur as a result of the PCE activity itself. Hence, this information is quite stable and will not change frequently.

6.9. Operational orders of magnitudes

This section gives minimum order of magnitude estimates of what the PCE discovery mechanism should support.

- Number of PCCs discovering a given PCE: 1000
- Number of PCEs to be discovered by a given PCC: 100

6.10. Manageability considerations

Mechanisms are REQUIRED to manage PCE discovery operations. This includes the configuration of PCE Discovery functions and policies, as well as, the monitoring of the discovery protocol activity.

6.10.1. Configuration of PCE Discovery parameters

It MUST be possible to enable and disable the PCE discovery function at a PCC and at a PCE.

On the PCC it MUST be possible for an operator to activate/deactivate automatic PCE discovery. The activation of automatic discovery MUST not prevent static configuration of PCE information that may supplement discovered information.

On the PCE it MUST be possible for an operator to control the application of discovery policies by which the specific PCE is discovered. As described in Section 6.5, this control MUST include the ability to:
- restrict the discovery scope to a set of authorized domains;
- define the type and nature of the information disclosed;
- specify the filtering and translation to be applied to the PCE information disclosed at domain borders.

These configuration options MAY be supported through an implementation-specific local configuration interface, or MAY be supported via a standardised interface (such as a MIB module, as below).

6.10.2. PCE Discovery MIB modules

PCE Discovery MIB modules MUST be specified for the control of the function on PCCs and PCEs.

6.10.2.1. PCC MIB module

The MIB module that will run on PCCs MUST include at least:
- A control to disable automatic discovery by the PCC;
- The set of known PCEs;
- The number of known PCEs, and the number of discovered PCEs.

For each PCE reported in the MIB module, the following information MUST be available:
- Information advertised by the PCE (i.e., discovered information);
- Information locally configured about the PCE;
- The time since the PCE was discovered;
- The time since any change to the discovered information for the PCE;

Note that when a PCE is no longer alive (see section 6.4), it SHOULD no longer be reported in the PCC MIB module.

The MIB module SHOULD also provide the average and maximum rates of arrival, departure and modification of PCE discovery to enable effective analysis of the operation of the protocols. Further, the MIB module SHOULD report on the operation of the discovery protocol by counting the number of unacceptable and incomprehensible information exchanges.

The PCC MIB module SHOULD also be used to provide notifications when thresholds (e.g. on the maximum rate of change, on the number of unacceptable messages) are crossed, or when important events occur (e.g. the number of discovered PCEs decreases to zero).

6.10.2.2. PCE MIB module

The MIB module that will run on PCEs MUST include at least:
- A control to disable automatic discovery announcements by the PCE;
- Information to be advertised by the PCE, although this information MAY be present as read-only;
- The discovery policies active on the PCE, although this information MAY be present as read-only.

The MIB module SHOULD also include:
- The time since the last change to the advertised PCE information;
- The time since the last change to the advertisement policies;
- Control of on which interfaces the PCE issues advertisements where this is applicable to the protocol solution selected.

Note that a PCE MAY also be configured to discover other PCEs. In this case it SHOULD operate the MIB module described in section 6.10.2.1 as well as the module described here.

6.10.3. Monitoring Protocol Operations

It MUST be possible to monitor the operation of any PCE discovery protocol. Where an existing protocol is used to support the PCE discovery function, this monitoring SHOULD be achieved using the techniques already defined for that protocol, enhanced by the MIB modules described above. Where, those techniques are inadequate, new techniques MUST be developed.

Monitoring of the protocol operation demands support for at least the following functions:
- Correlation of information advertised against information received;
- Counts of dropped, corrupt, and rejected information elements;
- Detection of 'segmented' networks. That is, the ability to detect and diagnose the failure of a PCE advertisement to reach a PCC.

6.10.4. Impact on network operations

Frequent changes in PCE information may have a significant impact on PCCs that receive the advertisements, might destabilise the operation of the network by causing the PCCs to swap between PCEs, and might harm the network through excessive advertisement traffic. Hence it MUST be possible to apply at least the following controls:

- Configurable limit on the rate of announcement of changed parameters at a PCE;
- Control of the impact on PCCs such as through discovery messages rate-limiting;
- Configurable control of triggers that cause a PCC to swap to another PCE.

7. Security Considerations

This document is a requirement document and hence does not raise by itself any particular security issue.

A set of security requirements that MUST be addressed when considering the design and deployment of a PCE Discovery mechanism have been identified in section 6.6.

8. Acknowledgments

We would like to thank, in chronological order, Benoit Fondeviole, Thomas Morin, Emile Stephan, Jean-Philippe Vasseur, Dean Cheng, Adrian Farrel, Renhai Zhang, Mohamed Boucadair, Eric Gray, Igor Bryskin, Dimitri Papadimitriou, Arthi Ayyangar, Andrew Dolganow, Lou Berger, Nabil Bitar, and Kenji Kumaki.

Thanks also to Ross Callon, Ted Hardie, Dan Romascanu, Russ Housley and Sam Hartman for their review and constructive discussions during the final stages of publication.

9. References

9.1. Normative references

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[PCE-ARCH] Farrel, A., Vasseur, J.P., Ash, J., "Path Computation Element (PCE) Architecture", draft-ietf-pce-architecture, work in progress.

9.2. Informative references

[PCE-COM-REQ] Ash, J., Le Roux, J.L., "PCE Communication Protocol Generic Requirements", draft-ietf-pce-comm-protocol-gen-reqs, work in progress.

10. Editor Address

Jean-Louis Le Roux (Editor)
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: jeanlouis.leroux@francetelecom.com

11. Contributors' Addresses

Paul Mabey
Qwest Communications
950 17th Street,
Denver, CO 80202,
USA
Email: pmabey@qwest.com

Eiji Oki

NTT
Midori-cho 3-9-11
Musashino-shi, Tokyo 180-8585,

JAPAN
Email: oki.eiji@lab.ntt.co.jp

Richard Rabbat
Fujitsu Laboratories of America
1240 East Arques Ave, MS 345
Sunnyvale, CA 94085
USA
Email: richard@us.fujitsu.com

Ting Wo Chung
Bell Canada
181 Bay Street, Suite 350
Toronto, Ontario, M5J 2T3
CANADA,
Email: ting_wo.chung@bell.ca

Raymond Zhang
BT Infonet
2160 E. Grand Ave.
El Segundo, CA 90025
USA
Email: raymond_zhang@infonet.com

12. Intellectual Property Statement

Disclaimer of Validity

Copyright Statement