

Internet Draft
Category: Standards Track
Expiration Date: December 2006

Igor Bryskin (Movaz Networks)
Dimitri Papadimitriou (Alcatel)
Lou Berger (LabN Consulting, LLC)

June 2006

Policy-Enabled Path Computation Framework

[draft-bryskin-pce-policy-enabled-path-comp-02.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

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/1id-abstracts.html>

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

Abstract

The PCE architecture [[PCE-ARCH](#)] introduces the concept of policy in the context of path computation. This document provides additional details on policy within the PCE Architecture and also provides context for the support of PCE Policy. This document introduces the use of the Policy Core Information Model (PCIM) as a framework for supporting path computation policy. This document also provides representative scenarios for the support of PCE Policy.

Contents

1	Terminology	3
2	Introduction	3
3	Background	4
3.1	Motivations	4
3.2	Representative Policy Scenarios	6
3.2.1	Scenario: Policy Configured Paths	6
3.2.2	Scenario: Provider Selection Policy	7
3.2.3	Scenario: Policy Based Constraints	8
4	Requirements	10
5	Path Computation Policy Information Model (PCPIM)	11
6	Policy-Enabled Path Computation Framework Components	13
7	Policy Component Configurations	15
7.1	PCC-PCE Configurations	15
7.2	Policy Repositories	17
7.3	Cooperating PCE Configurations	18
7.4	Policy Configuration Management	20
8	Inter-Component Communication	20
8.1	Policy Communication	20
8.2	PCE Discovery Policy Considerations	22
9	Path Computation Sequence of Events	22
9.1	Policy-enabled PCC, Policy-enabled PCE	23
9.2	Policy-ignorant PCC, Policy-enabled PCE	24
10	Introduction of New Constraints	25
11	Security Considerations	26
12	Acknowledgements	27
13	IANA Considerations	27
14	References	27
14.1	Normative References	27
14.2	Informative References	28
15	Authors' Addresses	28
16	Full Copyright Statement	29
17	Intellectual Property	29

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 [\[RFC2119\]](#).

1. Terminology

The reader is assumed to be familiar with the following terms:

CSPF: Constraint-based Shortest Path First, see [\[RFC3630\]](#).
LSP: Label Switched Path, see [\[RFC3031\]](#).
LSR: Label Switching Router, see [\[RFC3031\]](#).
PCC: Path Computation Client, see [\[PCE-ARCH\]](#).
PCE: Path Computation Element, see [\[PCE-ARCH\]](#).
TE LSP: Traffic Engineering MPLS Label Switched Path, see [\[RFC3209\]](#) and [\[RFC3473\]](#).
CIM: Common Information Model, see [\[DMTF\]](#).
PCIM: Policy Core Information Model, see [\[RFC3060\]](#).
PC: Path Computation.
PCCIM: Path Computation Core Information Model.
QPIM: QoS Policy Information Model, see [\[RFC3644\]](#).
PBM: Policy-based Management, see [\[RFC3198\]](#).
PEP: Policy Enforcement Points, see [\[RFC2753\]](#).
PDP: Policy Decision Points, see [\[RFC2753\]](#).
COPS: Common Open Policy Service, see [\[RFC2748\]](#).
COPS-PR: COPS Usage for Policy Provisioning, see [\[RFC3084\]](#).

2. Introduction

The PCE architecture is introduced in [\[PCE-ARCH\]](#). This document describes the impact of policy on the PCE architecture and provides additional details on, and context for, policy within the PCE Architecture.

Policy-based Management (PBM) enables network administrators to operate in a high-level manner through rule-based policies; the latter are translated automatically into individual device configuration directives, aimed at controlling a network as a whole. Two IETF Working Groups have considered policy networking in the past: The Resource Allocation Protocol working group and the Policy Framework working group.

A framework for policy-based admission control [\[RFC2753\]](#) was defined and a protocol for use between Policy Enforcement Points (PEP) and Policy Decision Points (PDP) was specified: Common Open Policy Service (COPS) [\[RFC2748\]](#). This document uses the terms PEP and PDP to

refer to the functions defined in the COPS context. This document makes no assumptions nor requires that the actual COPS protocol be used.

The IETF has also produced a general framework for representing, managing, sharing, and reusing policies in a vendor independent, interoperable, and scalable manner. It has also defined an extensible information model for representing policies, called the Policy Core Information Model (PCIM) [RFC3060], and an extension to this model to address the need for QoS management, called the QoS Policy Information Model (QPIM) [RFC3644]. However, additional mechanisms are needed in order to specify policies related to the path computation logic as well as its control.

This document introduces PCIM as a core component in a framework for providing policy-enabled path computation. This document also covers the other aspects of the framework. It also discusses related background to provide a context for this work. Specific PCIM definition to support path computation will be discussed in a separate document.

3. Background

This section provides some general background on the use of policy within the PCE architecture. It presents some rationale behind the use of policy, as well as representative policy usage scenarios. This information is intended to provide context for the presented policy framework. This section does not attempt to present an exhaustive list of rationale or scenarios.

3.1. Motivations

The PCE architecture is introduced in [PCE-ARCH]. It includes policy as an integral part of the PCE architecture. This section presents some of the rationale for this inclusion.

Network operators require a certain level of flexibility to shape the TE path computation process, so that the process can be aligned with their business and operational needs. Many aspects of the path computation may be governed by policies. For example, a PCC may use policies configured by the operator to decide which optimizations, constraints, diversities and their relaxation strategies to request while computing path(s) for a particular service. Depending on SLAs, TE and cost/performance ratio goals, path computation requests may be built differently for different services. Service A, for instance, may require two SRLG-disjoint paths for building end-to-end recovery

scheme, while for service B link-disjoint paths may be sufficient. Service A may need paths with minimal end-to-end delay, while service B may be looking for shortest (minimal-cost) paths. Different constraint relaxation strategies may be applied while computing paths for service A and for service B, and so forth.

Likewise, a PCE may apply policies to decide which algorithm or algorithms to use while performing path computations requested from a particular PCC or for a particular domain, see [PCECP-IA-REQ]; whether to seek the cooperation of other PCEs to satisfy a particular request or to handle a request on its own (possibly responding with non explicit paths); or how the request should be modified before being sent to other member(s) of a group of cooperating PCEs, etc.

Additional motivations for supporting policy within the PCE architecture can be described as follows. Historically, a path computation entity was an intrinsic part of an LSR's control plane and always co-located with the LSR's signaling and routing subsystems. This approach allowed for unlimited flexibility in providing various path computation enhancements, such as: adding new types of constraints, diversities and their relaxation strategies, adopting new objective functions and optimization criteria, etc. All that had to be done to support an enhancement was to upgrade the control plane software of a particular LSR (and no other LSRs or any other network elements).

With the introduction of the PCE architecture, the introduction of new PCE capabilities becomes more complicated: it isn't enough for a PCE to upgrade its own software. In order to take advantage of a PCE's new capabilities, new advertising and signaling objects may need to be standardized, all PCCs may need to be upgraded with new software, and new interoperability problems may need to be resolved, etc.

Within the context of the PCE architecture, it is therefore highly desirable to find a way to introduce new path computation capabilities without requiring modifying either the discovery/communication protocols or PCC software. One way to achieve this objective is to consider path selection constraints, their relaxations and objective functions, as path computation request specific policies. Furthermore that such policies may, on one hand, be configured and managed by an operator as any other policies or, on the other hand, may be interpreted in real time by PCCs and PCEs.

There are a number of advantages and useful by-products of such an approach:

- New path computation capabilities may be introduced without

changing PCE-PCC communication and discovery protocols or PCC software. Only the PCE module providing the path computation capabilities (referred in this document as path computation engine) needs to be updated.

- Existing constraints, objective functions and their relaxations may be aggregated and otherwise associated together, thus, producing new more complex ones that do not require change of code even on PCEs supporting them.
- Different elements such as conditions, actions, variables, etc. may be re-used by multiple constraints, diversities, and optimizations.
- PCCs and PCEs need to handle other (that is, not request specific) policies. Path computation related policies of all types can be placed within the same policy repositories, can be managed by the same policy management tools and can be interpreted using the same mechanisms. Also policies need to be supported by PCCs and PCEs independently from peculiarities of a specific PCC-PCE communication protocol. Thus, introducing a new (request specific) type of policies describing constraints and other elements of a path computation request will be a natural and relatively inexpensive addition to the policy-enabled path computation architecture.

3.2. Representative Policy Scenarios

This section provides example scenarios of how policy may be applied using the PCE policy framework within the PCE architecture context. Actual networks may use one of the scenarios discussed, some combination of the presented scenarios or other scenarios (not discussed). This section should not be viewed as limiting other applications of policy within the PCE architecture.

3.2.1. Scenario: Policy Configured Paths

A very simple usage scenario for PCE policy would be to use PCE to centrally administer configured paths. Configured paths are composed of strict and loose hops in the form of EROs (see [RFC3209]), and are used by one or more LSPs. Typically such paths are configured at the LSP ingress. In the context of policy-enabled path computation, an alternate approach is possible.

In particular, service specific policies can be installed that will provide configured path(s) for a specific service request. The

request may be identified based on service parameters such as end-points, requested QoS or even a token that identifies the end-user. The configured path(s) would then be used as input to PCE computation process, which would return explicit routes by expanding of all specified loose hops.

The described policies may be applied at either PCC or PCE, see Figure 5. In the PCC case, the configured path would be processed at the PCC and then passed to the PCE along with the PCE request, probably in the form of (inclusion) constraints. When applied at the PCE, the configured path would be used locally. Both cases require some method to configure and manage policies. In the PCC case, the real benefit would come when there is an automated policy distribution mechanism.

Policy-configured paths may also be used in multiple PCE environments (see Figures 7 and 8). For example, consider the case when there is limited TE visibility and independent PCEs are used to determine path(s) within each area of the TE visibility. In such a case, it may not be possible (or desirable) to configure entire explicit path(s) on a single PCE. However, it is possible to configure explicit path(s) for each area of the TE visibility and each responsible PCE. One by one, the PCEs would then map an incoming signaling request to appropriate configured path(s). Note that to make such scenario work it would likely be necessary to start and finish the configured paths on TE domain boundary nodes. Clearly, consistent PC Policy Repositories are also critical in this example.

3.2.2. Scenario: Provider Selection Policy

A potentially more interesting usage scenario is applying PC policies in multi-domain multi-provider networks. There are numerous interesting policy applications in such networks. A rudimentary example is simple access control, that is, deciding which clients are permitted to request inter-domain path computation.

A more complicated example is applying policy to determine which domain or provider network will be used to support a particular PCE request. Consider the topology presented in Figure 1. In this example there are multiple transit networks available to provide a path from a source domain to a destination domain. Furthermore, each transit network may have one or more options for reaching a particular domain. Each domain will need to select which of the multiple available paths will be used to satisfy a particular PCE request.

Clearly, TE reachability, availability and optimality are the basic criteria for path selection, however, policies can provide an

important added consideration in the decision process. For example, transit network A may be more expensive and provide lower delay or loss than transit network C. Likewise, a transit network may wish to treat PCE requests from its own customers differently than requests from other providers. In both cases, computation based on traffic engineering databases will result in multiple transit networks that provide reachability, and policies can be used to govern which PCE requests get better service.

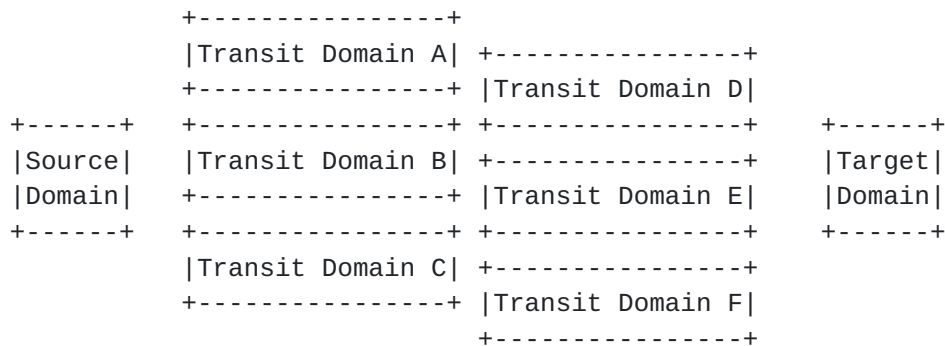


Figure 1: Multi-domain network with multiple transit options

There are multiple options for differentiating which PCE requests use a particular transit domain and get a particular (better or worse) level of service. For example, a PCE in the source domain may use user and request specific policies to determine the level of service to provide. A PCE in the source domain may also use domain specific policies to choose which transit domains are acceptable. A PCE in a transit domain may use request specific policies to determine if a request is from a direct customer or another provider, and then use domain specific policies to identify how the request should be processed.

3.2.3. Scenario: Policy Based Constraints

Another usage scenario is the use of policy to provide new constraints in a PCE request. Consider an LSR with a policy enabled PCC, as shown in Figure 3, which receives a service request via signaling, including over a NNI or UNI reference point, or receives a configuration request over a management interface to establish a service. In either case the path(s) needed to support the service are not explicitly specified in the message/request, and hence path computation is needed.

In this case, the PCC may apply user or service specific policies to decide how the path selection process should be constrained, that is, which constraints, diversities, optimizations and relaxations should

be applied in order for the service LSP(s) to have a likelihood to be successfully established and provide necessary QoS and resilience against network failures. When deciding on the set of constraints the PCC uses as an input all information it knows about the user and service, such as the contents of the received message, port ID over which message was received, associated VPN ID, signaling/reference point type, request time, etc. Once the constraints and other parameters of the required path computation are determined, the PCC generates a path computation request which includes the request-specific policies that describe the determined set of constraints, optimizations, and other parameters that indicate how the request is to be considered in the path computation process.

The PCC may also apply server specific policies for each of the known (i.e., discovered or configured) PCEs in order to select which PCE to use. The PCC may also use server specific policies to form the request to match the PCE's capabilities so that the request will not be rejected and has a higher likelihood of being satisfied in an efficient way. An example of a request modification as the result of a server specific policy is removing a constraint not supported by the PCE. Once the policy processing is completed at the PCC, and the path computation request resulting from the original service request is updated by the policy processing, the request is sent to the PCE.

The PCE that receives the request validates and otherwise processes the request, applying the policies found in the request as well as any policies that are available at the PCE, e.g., client and domain specific policies. As a result of the policy processing, the PCE may decide to reject the request. It also may decide to respond with one or several pre-computed paths if user or client specific policies instruct the PCE to do so. If the PCE decides to satisfy the request by performing a path computation, it determines if it needs the cooperation of other PCEs and defines parameters for path computations to be performed locally and remotely. After that, the PCE instructs a co-located path computation engine to perform the local path computation(s) and, if necessary, sends path computation requests to one or more other PCEs. It then waits for the responses from the local path computation engine and, when used, the remote PCE. It then combines the resulting paths and sends the result back to the requesting PCC. The response may indicate policies describing the resulting paths, their characteristics (summary cost, expected end-to-end delay, etc.) as well as additional information related to the request, e.g., which of constraints were honored, which were dismissed and which were relaxed and in what way.

The PCC processes the response and instructs the LSR to encode the received path(s) into the outgoing signaling message(s).

4. Requirements

The following requirements must be addressed by mechanisms and protocols that enable policy-based control over path computation requests and decisions:

- (G)MPLS path computation-specific
The mechanisms must meet the policy-based control requirements specific to the problem of path computation using RSVP-TE as the signaling protocol on MPLS and GMPLS LSRs.
- Support for non (G)MPLS PCCs
The mechanisms must be sufficiently generic to support non-(G)MPLS (LSR) clients such as an NMS, or network planner, etc.
- Support for many policies
The mechanisms must include support for many policies and policy configurations. In general, the determination and configuration of viable policies are the responsibility of the service provider.
- Provision for monitoring and accounting information
The mechanisms must include support for monitoring policy state, and provide access information. In particular, mechanisms must provide usage and access information that may be used for accounting purposes.
- Fault tolerance and recovery
The mechanisms must include provisions for fault tolerance and recovery from failure cases such as failure of PCC/PCE PDPs, disruption in communication that separate a PCC/PCE PDP from its associated PCC/PCE PEPs.
- Support for policy-ignorant nodes
The mechanisms should not be mandatory for every node in a network. Policy based path computation control may be enforced at a subset of nodes, for example, on boundary nodes within an administrative domain. These policy-capable nodes will function as trusted nodes from the point of view of the policy-ignorant nodes in that administrative domain. Alternatively, policy may be applied solely on PCEs with all PCCs being policy-ignorant nodes.
- Scalability
One of the important requirements for the mechanisms is scalability. The mechanisms must scale at least to the same extent that RSVP-TE signaling scales in terms of accommodating multiple LSPs and network nodes in the path of an LSP. There are

several sensitive areas in terms of scalability of policy based path computation control. First, not every policy aware node in an infrastructure should be expected to contact a remote PDP. This would cause potentially long delays in verifying requests. Additionally, the policy control architecture must scale at least as well as RSVP-TE protocol based on factors such as the size of RSVP-TE messages, the time required for the network to service an RSVP-TE request, local processing time required per node, and local memory consumed per node. These scaling considerations are of particular importance during re-routing of a set of LSPs.

- Security and denial of service considerations
The policy control architecture, protocols and mechanisms must be secure as far as the following aspects are concerned:
 - o First, the mechanisms proposed must minimize theft and denial of service threats.
 - o Second, it must be ensured that the entities (such as PEPs and PDPs) involved in policy control can verify each other's identity and establish necessary trust before communicating.
- Inter-AS and inter-area requirements
There are several inter-AS policy related requirements discussed in [RFC4216] and [INTERAS-PCEP], and inter-area policy related requirements discussed in [PCECP-IA-REQ]. These requirements must be addressed by policy-enabled PCE mechanisms and protocols.

It should be noted that this document only outlines the communication elements and mechanisms needed to allow a wide variety of possible policies to be applied for path computation control. It does not include any discussion of any specific policy behavior. Nor does it define or require use of specific policies.

5. Path Computation Policy Information Model (PCPIM)

The Policy Core Information Model (PCIM) introduced in [RFC3060] and expanded in [RFC3460] presents the object-oriented information model for representing general policy information.

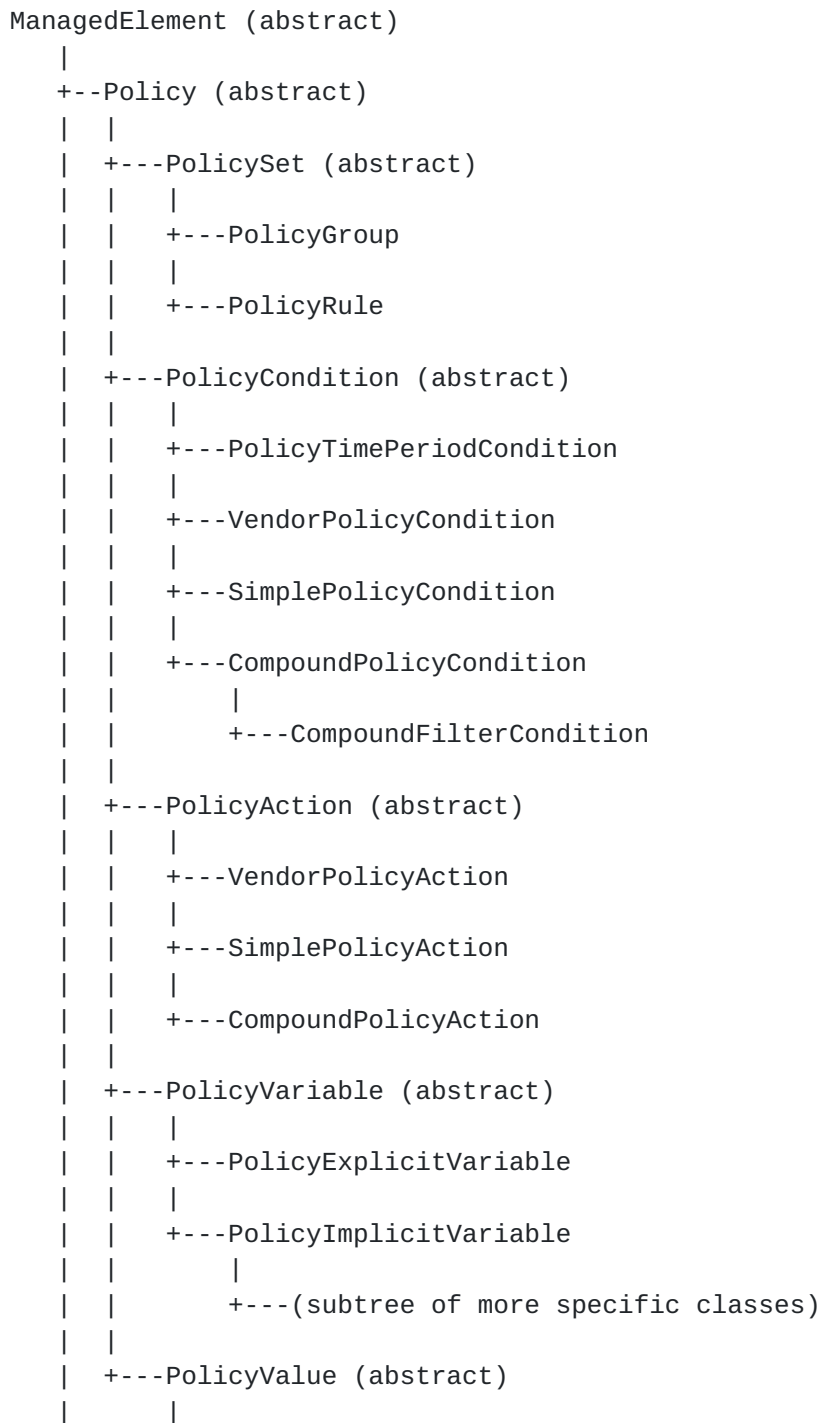
This model defines two hierarchies of object classes:

- structural classes representing policy information and control of policies
- association classes that indicate how instances of the structural

classes are related to each other.

These classes can be mapped to various concrete implementations, for example, to a directory that uses LDAPv3 as its access protocol.

Figure 2 shows an abstract from the class inheritance hierarchy for PCIM.




```
|      +---(subtree of more specific classes)
|
```

Figure 2: PCIM Class Inheritance

The policy classes and associations defined in PCIM are sufficiently generic to allow them to represent policies related to anything.

Policy models for application-specific areas such as the Path Computation Service may extend the PCIM in several ways. The preferred way is to use the PolicyGroup, PolicyRule, and PolicyTimePeriodCondition classes directly as a foundation for representing and communicating policy information. Then, specific subclasses derived from PolicyCondition and PolicyAction can capture application-specific definitions of conditions and actions of policies. Two subclasses, VendorPolicyCondition and VendorPolicyAction, are also defined to provide a standard extension mechanism for vendor-specific extensions to the Policy Core Information Model.

Policy Quality of Service Information Model [[RFC3644](#)] further extends the PCIM to represent QoS policy information for large-scale policy domains. New classes introduced in this document describing QoS and RSVP related variables, conditions and actions can be used as a foundation for the PCPIM.

Detailed description of the PCPIM will be provided in a separate documents.

6. Policy-Enabled Path Computation Framework Components

The following components are defined as part of the framework to support policy-enabled path computation:

- PC Policy Repository
A database from which PC policies are available in the form of instances of PCPIM classes. PC Policies are configured and managed by PC Policy Management Tools.
- PCE Policy Decision Point (PCE-PDP)
A logical entity capable of retrieving relevant path computation policies from one or more Policy Repositories and delivering the information to associated PCE-PEP(s);
- PCE Policy Enforcement Point (PCE-PEP)
A logical entity capable of issuing device specific Path Computation Engine configuration requests for the purpose of

enforcing the policies

- PCC Policy Decision Point (PCC-PDP)
A logical entity capable of retrieving relevant path computation policies from one or more Policy Repositories and delivering the information to associated PCC-PEP(s);
- PCC Policy Enforcement Point (PCC-PEP)
A logical entity capable of issuing device specific Path Computation Service User configuration requests for the purpose of enforcing the policies.

From the policy perspective a PCC is logically decomposed into two parts: PCC-PDP and PCC-PEP. When present, a PCC-PEP is co-located with a Path Computation Service User entity that requires remote path computation (for example, the GMPLS control plane of an LSR). The PCC-PEP and PCC-PDP may be physically co-located (as per [RFC2748]) or separated. In the later case they talk to each other via such protocols as COPS and/or COPS-PR [RFC3084].

Likewise, a PCE is logically decomposed into two parts: PCE-PEP and PCE-PDP. When present, PCE-PEP is co-located with a Path Computation Engine entity that actually provides the Path Computation Service (that is, runs path computation algorithms). PCE-PEP and PCE-PDP may be physically co-located or separated. In the later case they communicate using such protocols as COPS and/or COPS-PR.

PCC-PDP/PCE-PDP may be co-located with, or separated from, an associated PC Policy Repository. In the latter case, the PDPs use some access protocol (for example, LDAPv3 or SNMP). The task of PDPs is to retrieve policies from the repository(ies) and convey them to respective PEPs either in unsolicited way or upon the PEPs requests.

A PCC-PEP may receive policy information not only from PCC-PDPs(s) but also from PCE-PEP(s) via PCC-PCE communication and/or PCE discovery protocols. Likewise, a PCE-PEP may receive policy information not only from PCE-PDPs(s) but also from PCC-PEP(s) via PCC-PCE communication protocol.

Any given policy can be interpreted (that is, translated into a sequence of concrete device specific configuration requests) either on a PDP or on the associated PEP or partly on the PDP and partly on the PEP.

Generally speaking, the task of the PCC-PEP is to select the PCE and build path computation requests applying service specific policies provided by the PCC-PDP. The task of the PCE-PEP is to control path computations by applying request-specific policies found in the

requests as well as client-specific and domain-specific policies supplied by the PCE-PDP.

7. Policy Component Configurations

7.1. PCC-PCE Configurations

The PCE policy architecture supports policy being applied at a PCC and at a PCE. While the architecture supports policy being applied at both, there is no requirement for policy to always be applied at both, or even at either. The use of policy in a network, on PCCs and on PCEs, is a specific network design choice. Some networks may choose to apply policy only at PCCs (Figure 3), some at PCEs (Figure 4), and others at both PCCs and PCEs (Figure 5). Regardless of where policy is applied it must be applied in a consistent fashion in order to achieve the intended results.

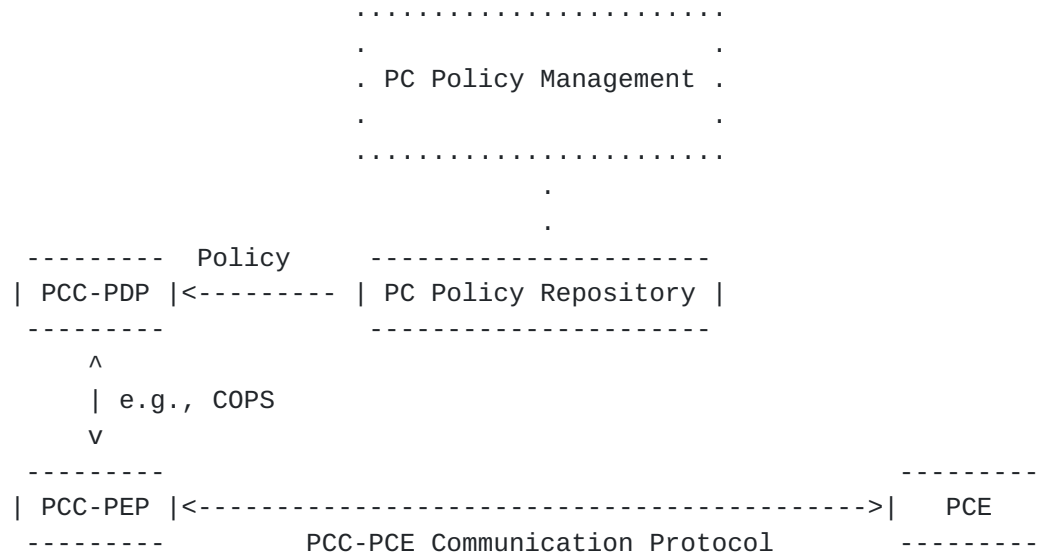


Figure 3: Policies are applied on PCC only

Along with supporting flexibility in where policy may be applied, the PCE architecture is also flexible in terms of where specific types of policies may be applied. Also the PCE architecture allows for the application of only a subset of policy types. [PCE-ARCH] defines several PC policy types. Each of these may be applied at either a PCC or a PCE or both. Clearly when policy is only applied at PCCs or at PCEs, all PC policy types used in the network must be applied at those locations.

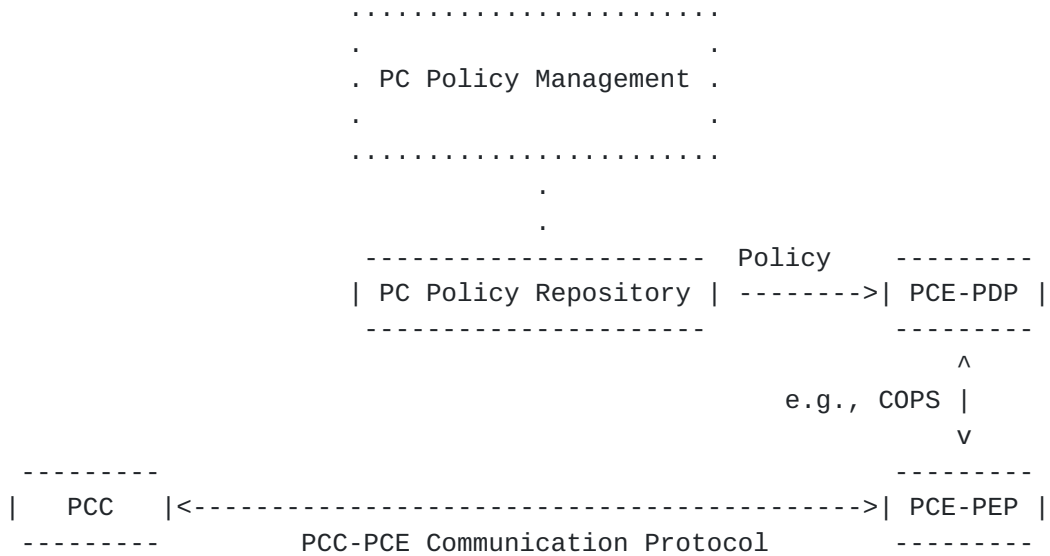


Figure 4: Policies are applied on PCE only

In the case where policy is only applied at a PCE, it is expected that the PCC will pass to the PCE all information about the service that it can gather in the path computation request (most likely in the form of PCPIM policy variables). The PCE is expected to understand this information and apply appropriate policies while defining the actual parameters of the path computation to be performed. Note that in this scenario the PCC cannot apply server-specific or any other policies, and PCE selection is static.

When applying policy at both PCC and PCE, it is necessary to select which types of policies are applied at each. In such configurations, it is likely that the application of policy types will be distributed across PCC and PCE rather than applying all of them at both. For example, user specific and server specific policies may be applied at a PCC, request and client specific policies may be applied at a PCE, while domain specific policies may be applied at both the PCC and PCE.

In the case when policy is only applied at a PCC, the PCC must apply all the types of required policies, for example user, service, server and domain-specific policies. The PCC uses the policies to construct a path computation request that appropriately represents the applied policies. The request will necessarily be limited to the set of "basic" (that is, non-policy capable) constraints explicitly defined by the PCC-PCE communication protocol in use.

7.2. Policy Repositories

Within the policy-enabled path computation framework policy repositories may be used in a single or multiple PC policy repository configuration:

o) Single PC Policy Repository

In this configuration there is a single PC Policy Repository shared between PCCs and PCEs.

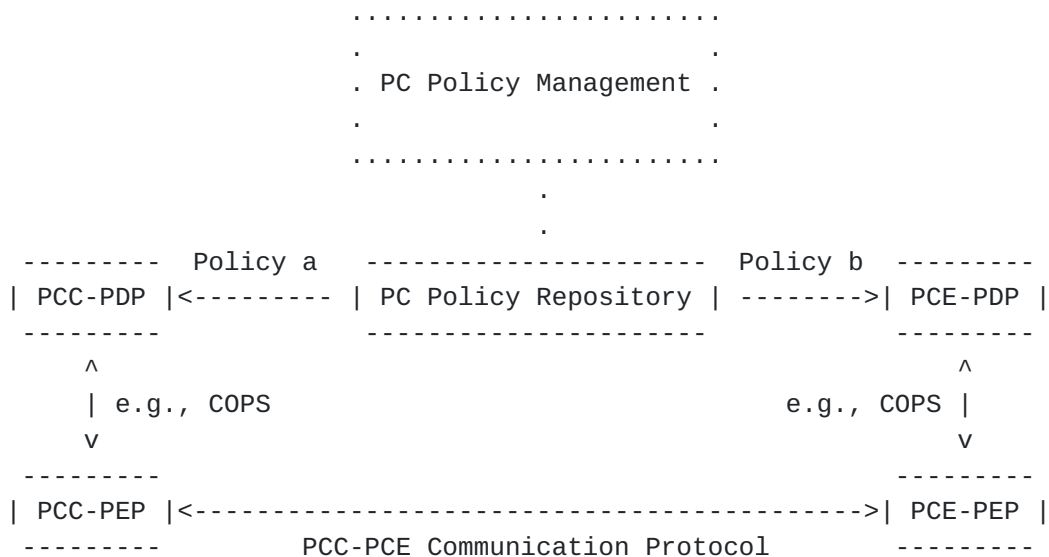


Figure 5: Single PCC/PCE Policy Repository

o) Multiple PC Policy Repositories

The repositories in this case may be fully or partially synchronized by some discovery/ synchronization management protocol or may be completely independent. Note that the situation when PC Policy Repository A exactly matches PC Policy Repository B, results in the single PC Policy Repository configuration case.

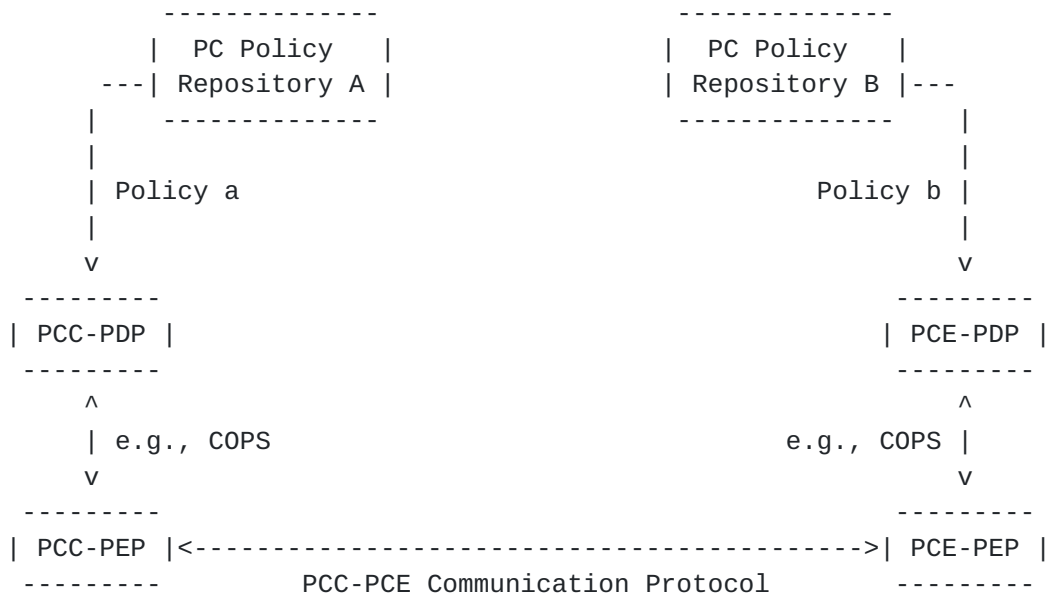


Figure 6: Multiple PCE/PCC Policy Repositories

7.3. Cooperating PCE Configurations

The previous section shows the relationship between PCCs and PCEs. A parallel relationship exists between cooperating PCEs, and, in fact, this relationship can be viewed as the same as the relationship between PCCs and PCEs. The one notable difference is that there will be cases where having a shared PC Policy Repository will not be desirable, for example, when the PCEs are managed by different entities. Note that in this case it still remains necessary for the policies to be consistent across the domains in order to identify usable paths. The other notable difference is that a PCE, while processing a path computation request, may need to apply requestor-specific (that is, client-specific) policies in order to modify the request before sending it to other cooperating PCE(s). This relationship is particularly important as the PCE Architecture allows for configuration where all PCCs are not policy-enabled.

The following are example configurations. These examples do not represent an exhaustive list and other configurations are expected.

o) Single Policy Repository

In this configuration there is a single PC Policy repository shared between PCEs. This configuration is likely to be useful within a single administrative domain where multiple PCEs are provided for redundancy or load distribution purposes.

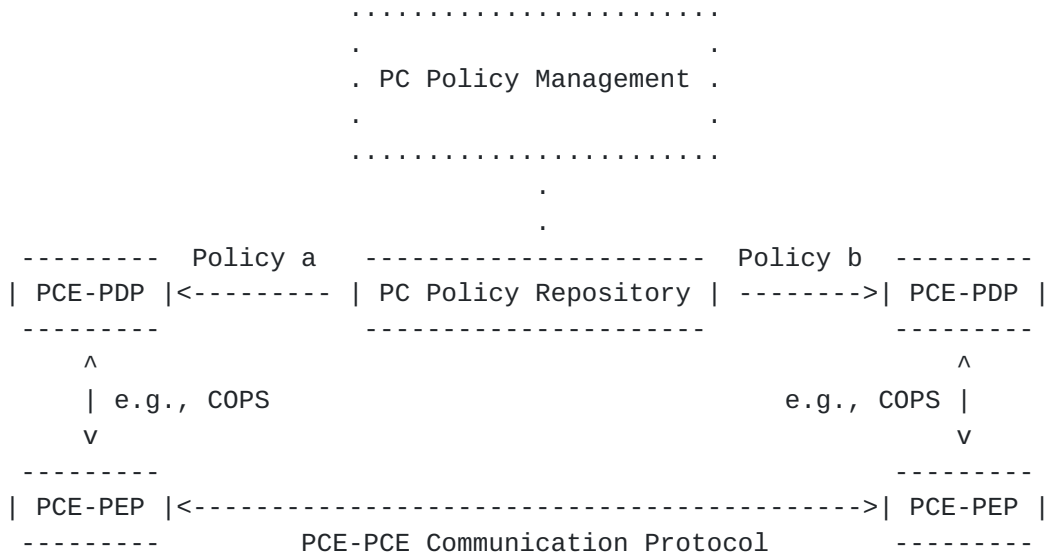


Figure 7: Single PCC Policy Repository

o) Multiple Policy Repositories

The repositories in this case may be fully or partially synchronized by some discovery/synchronization management protocol(s) or may be completely independent. In the multi-domain case it is expected that the repositories will be distinct, providing, however, consistent policies.

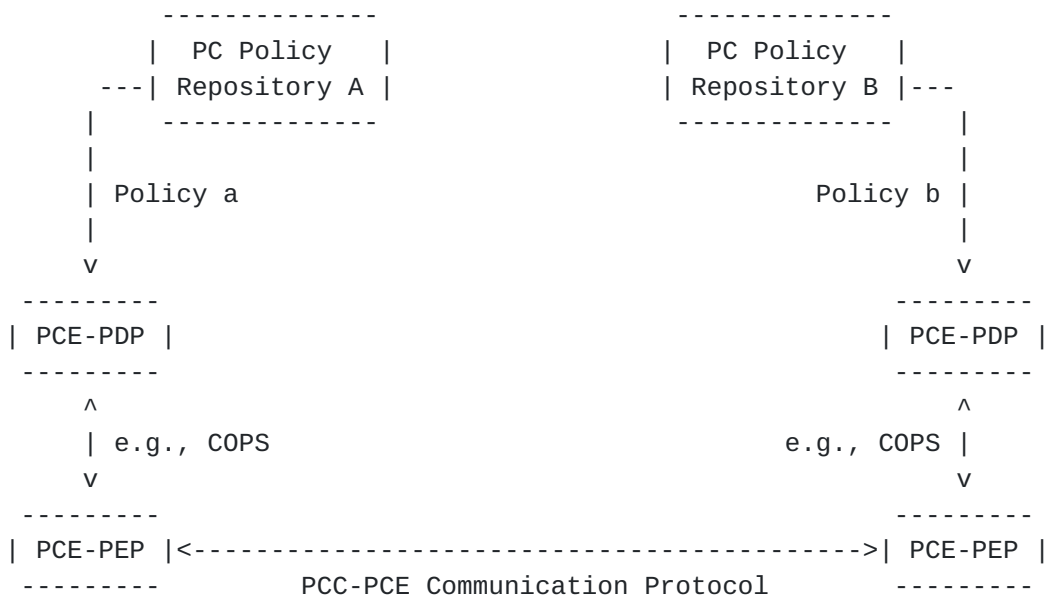


Figure 8: Multiple PCC Policy Repositories

7.4. Policy Configuration Management

The management of path computation policy information used by PCCs and PCEs is largely out of scope of the described framework. The framework assumes that such information is installed, removed and otherwise managed using typical policy management techniques. Policy Repositories may be populated and managed via static configuration, standard and proprietary policy management tools, or even dynamically via policy management/discovery protocols and applications.

8. Inter-Component Communication

8.1. Policy Communication

Flexibility in the application of policy types is imperative from the architecture perspective. However, this commodity implies added complexity on the part of the PCE related communication protocols.

One added complexity is that PCE communication protocols must carry certain information to support various policy types that may be applied. For example, in the case where policy is only applied at a PCE, a PCC-PCE request must carry sufficient information for the PCE to apply service or user specific policies. This does imply that a PCC must have sufficient understanding of what policies can be applied at the PCE. Such information may be obtained via local configuration, static coding or even via a PCE discovery mechanism. The PCC must also have sufficient understanding to properly encode the required information for each policy type.

Another added complexity is that PCE communication protocols must also be able to carry information that may result from a policy decision. For example, user or service specific policy applied at a PCC may result in policy related information that must be carried along with the request for use by a PCE. This complexity is particularly important as it may be used to introduce new path computation parameters (e.g. constraints, objection functions, etc.) without modification of the core PCC and PCE. This communication will likely simply require the PCE communication protocols to support opaque policy related information elements.

A final added complexity is that PCE communication protocols must also be able to support updated or unsolicited responses from a PCE. For example, changes in PCE policy may force a change to a previously provided path. Such updated or unsolicited responses may contain information that the PCC must act on, and may contain policy information that must be provided to a PCC.

PCC-PEP and PCE-PEP or a pair of PCE-PEPs communicate via a request-response type PCC-PCE Communication Protocol. This document makes no assumptions as to what exact protocol is used to support this communication. This document does assume that the semantics of a path computation request are sufficiently abstract and general, and support both PCE-PCC and PCE-PCE communication.

A path computation request at minimum should include:

- o One or several source addresses;
- o One or several destination addresses;
- o Computation type (P2P, P2MP, MP2P, etc.);
- o Number of required paths;
- o Zero or more policy descriptors in the following format:
 - <policy name>,
 - <policy variable1 name>, <param11>, <param12>, ..., <param1N>
 - <policy variable2 name>, <param21>, <param22>, ..., <param2N>
 - ...
 - <policy variableM name>, <paramM1>, <paramM2>, ..., <paramMN>

A successful path computation response, at minimum, should include the list of computed paths and may include policies (in the form of policy descriptors as in path computation request, see above) for use in evaluating and otherwise applying the computed paths.

PCC-PCE Communication Protocol provides transport for policy information and should not understand nor make any assumptions about the semantics of policies specified in path computation requests and responses.

Note: This document explicitly allows for (but does not require) the PCC to decide that all necessary constraints, objective functions, etc. pertinent to the computation of paths for the service in question are to be determined by the PCE performing the computation. In this case the PCC will use a set of policies (more precisely, PCPIM policy variables) describing the service specific information. These policies may be placed within the path computation request and delivered to the PCE via a PCC-PCE communication protocol. The PCE (more precisely, PCE-PEP) is expected to understand this information and use it to determine the constraints and optimization functions applying local policies (that is, policies locally configured or provided by the associated PCE-PDP(s)).

8.2. PCE Discovery Policy Considerations

Dynamic PCE discovery allows for PCCs and PCEs to automatically discover a set of PCEs (including information required for the PCE selection). It also allows for PCCs and PCEs to dynamically detect new PCEs or any modification of PCEs information. Policy can be applied in two ways in this context:

1. Restricting the scope of information distribution for the mandatory set of information (in particular the PCE presence and location).
2. Restricting the type and nature of the optional information distributed by the discovery protocol. The latter is also subject to policy since the PCE architecture allows for distributing this information using either PCE discovery protocol(s) or PCC-PCE communication protocol(s). One important policy decision in this context is the nature of the information to be distributed, especially, when this information is not strictly speaking a "discovery" information, rather, the PCE state changes. Client-specific and domain-specific policies may be applied when deciding whether this information should be distributed and to which clients of the path computation service (that is, which PCCs and/or PCEs)

Another place where policy applies is at the administrative boundaries. In multi-domain networks multiple PCEs will communicate with each other and across administrative boundaries. In such cases, domain-specific policies would be applied to 1) filter the information exchanged between peering PCEs during the discovery process (to the bare minimum in most cases if at all allowed by the security policy) and 2) limit the content of information being passed in path computation request and responses.

9. Path Computation Sequence of Events

This section presents representative scenarios; other scenarios are also possible.

9.1. Policy-enabled PCC, Policy-enabled PCE

When a GMPLS LSR receives a Setup (RSVP Path) message from an upstream LSR, the LSR may decide to use a remote Path Computation Entity. The following sequence of events occurs in this case:

- A PCC-PEP co-located with the LSR applies the service specific policies to select a PCE for the service path computation as well as to build the path computation request (that is, to select a list of policies, their variables, conditions and actions expressing constraints, diversities, objective functions and relaxation strategies appropriate for the service path computation). The policies may be:

- a) Statically configured on the PCC-PEP;
- b) Communicated to the PCC-PEP by a remote or local PCC-PDP via protocol such as COPS either pro-actively (most of the cases) or upon an explicit request by the PCC-PEP in case when some specifics of the new service have not been covered yet by the policies so far known to the PCC-PEP)

The input for the decision process on the PCC-PEP is the information found in the signaling message as well as any other service specific information such as port ID over which the message was received, associated VPN ID, the reference point type (UNI, E-NNI, etc.) and so forth. After the path computation request is built it is sent directly to the PCE-PEP using the PCC-PCE Communication Protocol.

- PCE-PEP validates and otherwise processes the request applying the policies found in the request as well as client and domain specific polices. The latter, again, may be either statically configured on the PCE-PEP or provided by the associated local or remote PCE-PDP via a protocol such as COPS. The outcome of the decision process is the following information:
 - a) Whether the request should be satisfied, rejected or dismissed.
 - b) The sets of sources and destinations for which paths should be locally computed.
 - c) The set of constraints, diversities, optimization functions and relaxations to be considered in each of locally performed path computation.
 - d) The address of the next-in-chain PCE.

- e) The path computation request to be sent to the next-in-chain PCE.

The PCE-PEP instructs a co-located path computation engine to perform the local path computation(s) and, if necessary, sends the path computation request to the next-in-chain PCE using a PCC-PCE Communication Protocol. Then it waits for the responses from the local path computation engine and the remote PCE, combines the resulting paths and sends them back to the PCC-PEP using the PCC-PCE Communication Protocol. The response contains the resulting paths as well as policies describing some additional information (for example, which of constraints were honored, which were dismissed and which were relaxed and in what way)

- PCC-PEP instructs the signaling sub-system of the GMPLS LSR to encode the received path(s) into the outgoing Setup message(s).

9.2. Policy-ignorant PCC, Policy-enabled PCE

This case parallels the previous example, but the user and service specific policies should be applied at the PCE as the PCC is policy ignorant. Again, when a GMPLS LSR has received a Setup (RSVP Path) message from an upstream LSR, the LSR may decide to use a non co-located Path Computation Entity. The following sequence of events occurs in this case:

- The PCC constructs a PCE request using information found in the signaling/provisioning message as well as any other service specific information such as port ID over which the message was received, associated VPN ID, the reference point type (UNI, E-NNI, etc.) and so forth. This information is encoded in the request in the form of policy variables. After the request is built it is sent directly to the PCE-PEP using a PCC-PCE Communication Protocol.
- PCE-PEP validates and otherwise processes the request interpreting the policy variables found in the request and applying user, service- and also client- and domain- specific policies to build the actual path computation request. The policies, again, may be either statically configured on the PCE-PEP or provided by the associated local or remote PCE-PDP via a protocol such as COPS. The outcome of the decision process is the following information:
 - a) Whether the request should be satisfied, rejected or dismissed.

- b) The sets of sources and destinations for which paths should be locally computed.
- c) The set of constraints, diversities, optimization functions and relaxations to be considered in each of locally performed path computation.
- d) The address of the next-in-chain PCE.
- e) The path computation request to be sent to the next-in-chain PCE.

The PCE-PEP instructs a co-located path computation engine to perform the local path computation(s) and, if necessary, sends the path computation request to the next-in-chain PCE using the PCC-PCE Communication Protocol. Then it waits for the responses from the local path computation engine and the remote PCE, combines the resulting paths and sends them back to the PCC-PEP using the PCC-PCE Communication Protocol. The response contains the resulting paths as well as policies describing some additional information (for example, which of constraints were honored, which were dismissed and which were relaxed and in what way)

- PCC-PEP instructs the signaling sub-system of the GMPLS LSR to encode the received path(s) into the outgoing Setup message(s).

10. Introduction of New Constraints

An important aspect of the policy-enable path computation framework discussed above is the ability to introduce new constraints with minimal impact. In particular, only those components and mechanisms that will use a new constraint need to change in order to support a new constraints. Importantly, those components and mechanisms that will not use a new constraint must not require change for a new constraint to be utilized. For example, the PCE communication protocols must not require any changes to support a new constraint. Likewise, PCC and PCEs that will not process a new constraint must not require any modification.

Consider the case where a PCE has been upgraded with software supporting optical physical impairment constraint, such as Polarization Mode Dispersion (PMD), that previously was not supported in the domain. In this case, one or more new policies will be installed in the PC Policy Repository (associated with the PCE) defining the constraint (rules that determine application criteria, set of policy variables, conditions, actions, etc.) and its relaxation strategy(ies). The new policies will be also propagated

into other PC Policy Repositories within the domain via discovery and synchronization protocols or via local configuration. PCE-PDPs and PCC-PDPs will then retrieve the corresponding policies from the repository(ies). From then on PCC-PDPs will instruct associated PCC-PEPs to add the new policy information into path computation requests for services with certain parameters (for example, for services provisioned in the OCh layer).

It is important to note that policy-enabled path computation model naturally solves the PCE capability discovery issues. Suppose a PCE working in a single PC Policy Repository configuration starts to support a new constraint. Once a corresponding policy installed in the repository, it automatically becomes available for all repository users, that is, PCCs. In the multi-repository case some policy synchronization must be provided, however, this problem is one of the management plane which is solved already.

11. Security Considerations

This document adds to the policy security considerations mentioned in [[PCE-ARCH](#)]. In particular it is now necessary to consider the security of policy information maintained in PC Policy Repositories and policy related transactions. The most notable issues, some of which are also listed in [[PCE-ARCH](#)], are:

- Unauthorized access to the PC Policy Repositories;
- Interception of policy information when it is retrieved from the repositories and/or transported from PDPs to PEPs;
- Interception of policy related information in path computation requests and responses;
 - o Impersonation of user and client identities;
 - o Falsification of policy information and/or PCE capabilities;
 - o Denial of service attacks on policy related communication mechanisms.

As with [[PCE-ARCH](#)], it is expected that PCE solutions will address these issues in detail.

12. Acknowledgements

We would like to thank Bela Berde for fruitful discussions on PBM and Policy-driven path computation.

13. IANA Considerations

None.

14. References

14.1. Normative References

- [PCE-ARCH] Farrel, A., Vasseur, JP., Ash, J., "Path Computation Element (PCE) Architecture", April 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2753] R. Yavatkar, D. Pendarakis, R. Guerin, A Framework for Policy Based Admission Control, [RFC 2753](#), January 2000.
- [RFC2748] D. Durham, et al., The COPS (Common Open Policy Service) protocol, [RFC 2748](#), IETF, January 2000.
- [RFC3060] B. Moore, et al., Policy Information Model Version1 Specification, [RFC 3060](#), February 2001.
- [RFC3209] Awduche, D., et al., "Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC3460] Moore, B. Ed., "Policy Core Information Model (PCIM) Extensions", [RFC 3460](#), January 2003.
- [RFC3473] Berger, L., et al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [RFC3644] Y. Snir, et al., Policy Quality of Service (QoS) Information Model, [RFC 3644](#), November 2003.
- [RFC4216] Zhang, R., Vasseur, J-P., Eds., "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", [RFC4216](#), November 2005.

- [INTERAS-PCEP] Bitar, N., Zhang, R., Kumaki, K., Eds., "Inter-AS Requirements for the Path Computation Element Communication Protocol (PCECP)", February 2006.
- [PCECP-IA-REQ] Le Roux, J-L., Ed., "PCE Communication Protocol (PCECP) Specific Requirements for Inter-Area (G)MPLS Traffic Engineering", February 2006.

14.2. Informative References

- [DMTF] Common Information Model (CIM) Schema, version 2.x. Distributed Management Task Force, Inc. The components of the CIM v2.x schema are available via links on the following DMTF web page:
http://www.dmtf.org/standards/standard_cim.php.
- [RFC3031] Rosen, E., Viswanathan, V. Callon, R., "Multiprotocol Label Switching Architecture", [RFC 3031](#), January 2001.
- [RFC3084] Chan, K., Seligson, J., Durham, D., Gai, S., McCloghrie, K., Herzog, S., Reichmeyer, F., Yavatkar, R. and A. Smith, "COPS Usage for Policy Provisioning (COPS-PR)", [RFC 3084](#), February 2001.
- [RFC3198] Westerinen, A., et al., "Terminology for Policy-Based Management", [RFC 3198](#), November 2001.
- [RFC3630] Katz, D., Kompella, K., Yeung., D., "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.

15. Authors' Addresses

Igor Bryskin
Movaz Networks, Inc.
7926 Jones Branch Drive
Suite 615
McLean, VA 22102
Email: ibryskin@movaz.com

Dimitri Papadimitriou (Alcatel)
Fr. Wellesplein 1,
B-2018 Antwerpen, Belgium
Phone: +32 3 240-8491
Email: dimitri.papadimitriou@alcatel.be

Lou Berger
LabN Consulting, LLC
Phone: +1 301 468 9228
Email: lberger@labn.net

16. Full Copyright Statement

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

17. Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Generated on: Thu Jun 22 19:34:07 EDT 2006