

Network Working Group
Internet Draft
Document: [draft-jacquetnet-ip-fwd-pib-01.txt](#)
Category: Experimental
Expires December 2003

M. Boucadair
C. Jacquenet
France Telecom
June 2003

An IP Forwarding Policy Information Base
<[draft-jacquetnet-ip-fwd-pib-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#) [1].

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 draft specifies a set of Policy Rule Classes (PRC) for the enforcement of an IP forwarding policy by network devices. Instances of such classes reside in a virtual information store, which is called the IP Forwarding Policy Information Base (PIB). The corresponding IP forwarding policy provisioning data are intended for use by a COPS-PR IP TE Client-Type, and they complement the PRC classes that have been defined in the Framework PIB.

Table of Contents

1.	Introduction.....	2
2.	Conventions used in this document.....	3
3.	Changes since the Previous Version.....	3
4.	PIB Overview.....	3
5.	The IP Forwarding Policy Information Base.....	4
6.	Security Considerations.....	9
7.	References.....	9
8.	Acknowledgments.....	10

[1.](#) Introduction

The deployment of value-added IP services over the Internet has become one of the most competing challenges for service providers, as well as a complex technical issue.

Within the context of network resource provisioning and allocation, the Common Open Policy Service protocol (COPS, [2]) and its usage for the support of Policy Provisioning ([3]) is one of the most promising candidate protocols that should help service providers in dynamically enforcing IP routing and traffic engineering policies.

An IP routing/TE policy consists in appropriately provisioning and allocating/de-allocating the switching and the transmission resources of an IP network (i.e. the routers and the links that connect these routers, respectively), according to e.g. rate, one-way delay, inter-packet delay variation, etc.) that have been possibly negotiated between the customers and the service providers, and according to (a set of) routing metrics, which can also reflect the network conditions.

Thus, the enforcement of IP routing/TE policies yields the need for an introduction of a high level of automation for the dynamic provisioning of the configuration data that will be taken into account by the routers to select the appropriate IP routes.

Within the context of this document, the actual enforcement of an IP forwarding policy is primarily based upon the activation of both intra- and inter-domain dynamic routing protocols that will be activated by the routers to select, install, maintain and possibly withdraw IP routes.

Such routes have been selected so that they comply as much as possible with the aforementioned QoS requirements and/or specific routing constraints, possibly depending on the type of traffic that will be conveyed along these routes.

It is therefore necessary to provide the route selection processes with the information that will depict the routing policies that are to be enforced within a domain and, whenever appropriate, the

aforementioned constraints and metrics, given the dynamic routing protocols actually support traffic engineering capabilities for the calculation and the selection of such routes.

Some of these capabilities are currently being specified in [4] and [5] for the OSPF (Open Shortest Path First) and the IS-IS (Intermediate System to Intermediate System routing protocol, [6]) interior routing protocols respectively, while there is a comparable effort for the BGP4 (Border Gateway Protocol, version 4) protocol, as described in [7], for example.

To provide the route selection processes with the aforementioned information, one possibility is to use the COPS-PR protocol, together with a collection of policy provisioning data that will be stored in a virtual information store, called a Policy Information Base.

This draft describes a collection of Policy Rule Classes that will be stored and dynamically maintained in an IP forwarding PIB. The "rule" and "role" concepts, which have been defined in [8], are adopted by this document to distribute the IP routing policy provisioning data over the COPS-PR protocol.

The corresponding IP forwarding policy provisioning data are intended for use by a COPS-PR IP TE Client-Type ([9]), and they complement the PRC classes that have been defined in the Framework PIB ([10]).

This document is organized as follows:

- [Section 4](#) provides an overview of the organization of the IP forwarding PIB,
- [Section 5](#) provides a description of the PRC classes of the IP forwarding PIB, according to the semantics of the Structure of Policy Provisioning Information (SPPI, [11]).

[2.](#) 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](#) [12].

[3.](#) Changes since the Previous Version

- Some elaboration has been provided as far as the interaction with other relevant PIB specification efforts is concerned,
- The references section has been updated.

[4.](#) PIB Overview

The dynamic enforcement of an IP forwarding policy relies upon the activation of intra- and inter-domain routing protocols that will have the ability to take into account configuration information for the computation and the selection of routes, which will comply as much as possible with the constraints and requirements that MAY have been contractually defined between customers and service providers.

This document specifies an IP forwarding PIB that mainly aims at storing and maintaining the information related to the IP routes that have been installed in the routers' Forwarding Information Bases, so that service providers maintain and update the adequate knowledge of the network's resources availability, from an IP routing perspective.

Jacquenet et al. Experimental - Expires Dec. 2003

[Page 3]

Internet Draft

An IP Forwarding PIB

June 2003

As such, this PIB has been designed so that it SHOULD be gracefully complemented by PIB modules that will reflect the IGP- and BGP- inferred routing policies to be enforced, in terms of cost metrics' values to be assigned and updated whenever needed.

Also, the accounting PIB module which is described in [13] aims at providing the most accurate feedback (to service providers) on how efficient the enforcement of a given IP forwarding policy (as specified in this document) actually is.

The choice of this PIB organization is basically twofold:

- Make the PIB implementation simple,
- Provide the appropriate granularity of policy provisioning data that will be manipulated according to the requirements and technical choices of service providers.

Therefore, the IP forwarding PIB is currently organized into the following provisioning classes:

1. The Forwarding Classes (ipFwdClasses): the information contained in these classes is meant to provide a detailed

description of the IP routes as they have been selected by the routers of a given domain,

2. The Statistics Classes (ipFwdStatsClasses): the information contained in these classes is meant to provide statistics on the use of the IP routes currently depicted in the IP forwarding PIB.

5. The IP Forwarding Policy Information Base

```
IP-FWD-PIB PIB-DEFINITIONS ::= BEGIN
```

```
IMPORTS
```

```
    Unsigned32, Integer32, MODULE-IDENTITY,
    MODULE-COMPLIANCE, OBJECT-TYPE, OBJECT-GROUP
        FROM COPS-PR-SPPI
    InstanceId, ReferenceId, Prid, TagId
        FROM COPS-PR-SPPI-TC
    InetAddress, InetAddressType
        FROM INET-ADDRESS-MIB
    Count, TEXTUAL-CONVENTION
        FROM ACCT-FR-PIB-TC
    TruthValue, TEXTUAL-CONVENTION
        FROM SNMPv2-TC
    RoleCombination, PrcIdentifier
        FROM FRAMEWORK-ROLE-PIB
    SnmpAdminString
        FROM SNMP-FRAMEWORK-MIB;
```

Jacquenet et al. Experimental - Expires Dec. 2003

[Page 4]

Internet Draft

An IP Forwarding PIB

June 2003

```
ipFwdPib      MODULE-IDENTITY
```

```
    SUBJECT-CATEGORIES { tbd }      -- IP TE client-type to be
                                         -- assigned by IANA

    LAST-UPDATED      "200301220900Z"
    ORGANIZATION      "France Telecom"
    CONTACT-INFO      "
                        Mohamed Boucadair
                        France Telecom R & D
                        42, rue des Coutures
                        BP 6243
                        14066 CAEN CEDEX 04
                        France
```

Phone: +33 2 31 75 92 31
E-Mail: mohamed.boucadair@francetelecom.com"

DESCRIPTION

"The PIB module containing a set of policy rule classes that describe the IP routes that have been computed by means of routing/TE policy enforcement, as well as route traffic statistics."

REVISION "200306251000Z"

DESCRIPTION

"Initial version."

::= { pib tbd } -- tbd to be assigned by IANA

ipFwdClasses OBJECT IDENTIFIER ::= { ipFwdPib 1 }
ipFwdStatsClasses OBJECT IDENTIFIER ::= { ipFwdPib 2 }

--
-- Forwarding classes. The information contained in these classes
-- is meant to provide a detailed description of the available IP
-- routes. One table has been specified so far, but there is room
-- for depicting different kinds of routes, like MPLS (MultiProtocol
-- Label Switching, ([14]) LSP (Label switched Paths) paths.
--
--
--

--
-- The ipFwdTable
--

ipFwdTable	OBJECT-TYPE
SYNTAX	SEQUENCE OF ipRouteEntry
PIB-ACCESS	notify
STATUS	current
DESCRIPTION	

"This table describes the IP routes that are installed in the forwarding tables of the routers."

::= { ipFwdClasses 1 }

ipRouteEntry OBJECT-TYPE

SYNTAX	ipRouteEntry
STATUS	current
DESCRIPTION	"A particular route to a particular destination."

PIB-INDEX	{ ipRoutePrid }
UNIQUENESS	{ ipRouteDest, ipRouteMask, ipRoutePhbId, ipRouteNextHopAddress ipRouteNextHopMask ipRouteIfIndex }

::= { ipFwdTable 1 }

ipRouteEntry ::= SEQUENCE {	
ipRoutePrid	InstanceId,
ipRouteDestAddrType	InetAddressType,
ipRouteDest	InetAddress,
ipRouteMask	Unsigned32,
ipRouteNextHopAddrType	InetAddressType,
ipRouteNextHopAddress	InetAddress,
ipRouteNextHopMask	Unsigned32,
ipRoutePhbId	Integer32,
ipRouteOrigin	Integer32,
ipRouteIfIndex	Unsigned32
}	

ipRoutePrid	OBJECT-TYPE
-------------	-------------

SYNTAX	InstanceId
STATUS	current
DESCRIPTION	"An integer index that uniquely identifies this route entry among all the route entries."

::= { ipRouteEntry 1 }

ipRouteDestAddrType	OBJECT-TYPE
---------------------	-------------

SYNTAX	InetAddressType
STATUS	current
DESCRIPTION	"The address type enumeration value ([15]) used to specify the type of a route's destination IP address."

::= { ipRouteEntry 2 }

ipRouteDest OBJECT-TYPE

SYNTAX InetAddress

STATUS current

DESCRIPTION

"The IP address to match against the packet's
destination address."

::= { ipRouteEntry 3 }

ipRouteMask OBJECT-TYPE

SYNTAX Unsigned32 (0..128)

STATUS current

DESCRIPTION

"Indicates the length of a mask for the matching of the
destination IP address. Masks are constructed by
setting bits in sequence from the most-significant bit
downwards for ipRouteMask bits length. All other bits
in the mask, up to the number needed to fill the length
of the address ipRouteDest are cleared to zero. A zero
bit in the mask then means that the corresponding bit
in the address always matches."

::= { ipRouteEntry 4 }

ipRouteNextHopAddrType OBJECT-TYPE

SYNTAX InetAddressType

STATUS current

DESCRIPTION

"The address type enumeration value used to specify the
type of the next hop's IP address."

::= { ipRouteEntry 5 }

ipRouteNextHopAddress OBJECT-TYPE

SYNTAX InetAddress

STATUS current

DESCRIPTION

"On remote routes, the address of the next router en
route; Otherwise, 0.0.0.0."

::= { ipRouteEntry 6 }

ipRouteNextHopMask OBJECT-TYPE

SYNTAX Unsigned32 (0..128)

Jacquenet et al. Experimental - Expires Dec. 2003

[Page 7]

Internet Draft

An IP Forwarding PIB

June 2003

STATUS current

DESCRIPTION

"Indicates the length of a mask for the matching of the next hop's IP address. Masks are constructed by setting bits in sequence from the most-significant bit downwards for ipRouteNextHopMask bits length. All other bits in the mask, up to the number needed to fill the length of the address ipRouteNextHop are cleared to zero. A zero bit in the mask then means that the corresponding bit in the address always matches."

::= { ipRouteEntry 7 }

ipRoutePhbId OBJECT-TYPE

SYNTAX Integer32 (-1 | 0..63)

STATUS current

DESCRIPTION

"The binary encoding that uniquely identifies a Per Hop Behaviour (PHB, [16]) or a set of PHBs associated to the DiffServ Code Point (DSCP) marking of the IP datagrams that will be conveyed along this route. A value of -1 indicates that a specific PHB ID value has not been defined, and thus, all PHB ID values are considered a match."

::= { ipRouteEntry 8 }

ipRouteOrigin OBJECT-TYPE

SYNTAX INTEGER {
 OSPF (0)
 IS-IS (1)
 BGP (2)
 STATIC (3)
 OTHER (4)
}

STATUS current

DESCRIPTION

"The value indicates the origin of the route. Either the route has been computed by OSPF, by IS-IS, announced by BGP4, is static, or else."

::= { ipRouteEntry 9 }

ipRouteIfIndex OBJECT-TYPE

SYNTAX Unsigned32 (0..65535)

STATUS current

DESCRIPTION

"The ifIndex value that identifies the local interface through which the next hop of this route is accessible."

::= { ipRouteEntry 10 }

--

-- Route statistics classes. The information contained
-- in the yet-to-be defined tables aim at reporting statistics about
-- COPS control traffic, route traffic and potential errors. The
-- next version of the draft will provide a first table that will be
-- based upon the use of the "count" clause.

--

--

END

[6.](#) Security Considerations

The traffic engineering policy provisioning data as they are described in this PIB will be used for configuring the appropriate network elements that will be involved in the dynamic enforcement of the corresponding routing and traffic engineering policies, by means of a COPS-PR communication that will convey this information.

The function of dynamically provisioning network elements with such configuration information implies that only an authorized COPS-PR communication takes place.

From this perspective, this draft does not introduce any additional security issues other than those that have been identified in the COPS-PR specification, and it is therefore recommended that the IPSec ([17]) protocol suite be used to secure the above-mentioned authorized communication.

7. References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [2] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja R., Sastry A., "The COPS (Common Open Policy Service) Protocol", [RFC 2748](#), Proposed Standard, January 2000.
- [3] Ho Chan, K., Durham, D., Gai, S., Herzog, S., McLoghrie, K., Reichmeyer, F., Seligson, J., Smith, A., Yavatkar, R., "COPS Usage for Policy Provisioning (COPS-PR)", [RFC 3084](#), March 2001.
- [4] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF", [draft-katz-yeung-ospf-traffic-10.txt](#), Work in Progress, June 2003.

- [5] Smit, H., Li, T., "IS-IS Extensions for Traffic Engineering", [draft-ietf-isis-traffic-04.txt](#), Work in Progress, December 2002.
- [6] ISO/IEC 10589, "Intermediate System to Intermediate System, Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", June 1992.
- [7] Jacquenet, C., "Providing Quality of Service Indication by the BGP-4 Protocol: the QOS_NLRI Attribute", [draft-jacquenet-qos-nrli-05.txt](#), Work in Progress, June 2003.
- [8] Moore, B. et al., "Policy Core Information Model -- Version 1 Specification", [RFC 3060](#), February 2001.
- [9] Jacquenet, C., "A COPS client-type for IP traffic engineering", [draft-jacquenet-ip-te-cops-04.txt](#), Work in Progress, January 2003.
- [10] Sahita, R., et al., "Framework Policy Information Base", [RFC 3318](#), March 2003.
- [11] McLoghrie, K., et al., "Structure of Policy Provisioning Information (SPPI)", [RFC 3159](#), August 2001.

- [12] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997
- [13] Boucadair, M., "An IP TE PIB for Accounting purposes", [draft-boucadair-ip-te-acct-pib-02.txt](#), Work in Progress, June 2003.
- [14] Rosen, E., et al., "Multiprotocol Label Switching Architecture", [RFC 3031](#), January 2001.
- [15] Daniele, M., Haberman, B., Routhier, S., Schoenwaelder, J., "Textual Conventions for Internet Network Addresses", [RFC 3291](#), May 2002.
- [16] Black, D., Brim, S., Carpenter, B., Le Faucheur, F., "Per Hop Behaviour Identification Codes", [RFC 3140](#), June 2001.
- [17] Kent, S., Atkinson, R., "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.

8. Acknowledgments

Part of this work is funded by the European Commission, within the context of the MESCAL (Management of End-to-End Quality of Service Across the Internet At Large, <http://www.mescal.org>) project, which is itself part of the IST (Information Society Technologies) research program.

The authors would also like to thank all the partners of the MESCAL project for the fruitful discussions that have been conducted so far within the context of the traffic engineering specification effort of the project.

9. Authors' Addresses

Mohamed Boucadair
France Telecom R & D

Jacquenet et al. Experimental – Expires Dec. 2003

[Page 10]

Internet Draft

An IP Forwarding PIB

June 2003

DMI/SIR
42, rue des Coutures
BP 6243
14066 Caen Cedex 4
France
Phone: +33 2 31 75 92 31
Email: mohamed.boucadair@francetelecom.com

Christian Jacquenet
France Telecom Long Distance
3, avenue François Château

CS 36901
35069 Rennes CEDEX
France
Phone: +33 2 99 87 63 31
Email: christian.jacquet@francetelecom.com

10. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist its implementation may be prepared, coed, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

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

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