Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 21, 2013

M. Boucadair
C. Jacquenet
France Telecom
N. Wang
Centre for Communication System Research
September 17, 2012

IP/MPLS Connectivity Provisioning Profile
draft-boucadair-connectivity-provisioning-profile-00

Abstract

This document describes the Connectivity Provisioning Profile (CPP) and proposes a CPP Template to capture IP connectivity requirements to be met in the context of delivery of services (e.g. Voice over IP or IP TV) which are to be plugged upon an IP/MPLS infrastructure.

The CPP defines the set of IP transfer parameters to be supported by the underlying IP/MPLS transport network together with a reachability scope and bandwidth/capacity needs. Appropriate performance metrics such as one-way delay, one-way delay variation are used to characterize an IP transfer service. Both global and restricted reachability scopes can be captured in the CPP.

Table of Contents

1. Introduction

This document describes the Connectivity Provisioning Profile (CPP) and proposes a CPP Template to capture IP/MPLS connectivity requirements to be met in the context of delivery of services (e.g., Voice over IP, IP TV, VPN services) which are to be plugged upon an IP/MPLS infrastructure.

The CPP defines the set of IP/MPLS transfer guarantees to be offered by the underlying IP/MPLS transport network together with a reachability scope and capacity needs. Appropriate performance metrics such as one-way delay or one-way delay variation are used to characterize the IP transfer service. Guarantees related to availability and resiliency are also included in the CPP.

The CPP assumes service differentiation at the network layer can be enforced by tweaking various parameters which belong to distinct dimensions (e.g, forwarding, routing, traffic access management, traffic classification, etc.).

The CPP can be used in both the vertical model (i.e., the service and network infrastructures are managed by the same administrative entity) or the functional separation model (i.e., where distinct administrative entities mange the service and the network infrastructures). In the following sections, no assumption is made about the deployment model (vertical or separation).

The CPP also aims at rationalizing the connectivity needs of above-deployed services and then to handle in a generic fashion all these requirements. Service-specific IP provisioning rules may lead to increase the required IP transfer classes to be (pre)-engineered in the operational network. As such, the use of the CPP allows to engineer generic and limited number of classes and then map individual CPP to these classes. Instantiating each CPP into a distinct class of service should be avoided. Therefore, application-agnostic IP provisioning practices should be recommended since the requirements captured in the CPP can be used to identify which network class of service is to be used to meet those requiremenst/ guarantees.

CPP may also be used as a "hint" or a guideline for the network dimensioning and planning departments to ensure that appropriate resources (e.g., network cards, routers, upgrade link capacity, etc.) have been provisioned. Otherwise, (underlying) IP/MPLS networks would not be able to meet the targets expressed in all CPP requests (see Figure 1).

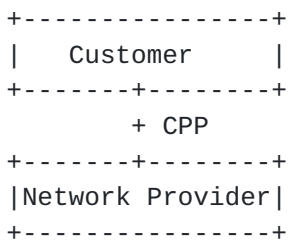

Figure 1: CPP: Interactions

The customer shown in Figure 1 may be a service provider (e.g., IP telephony service provider) which requires invoking resources provided by an underlying network provider or an enterprise which wants to interconnect its sites using VPN services offered by a network provider.

The definition of a clear interface between the service and the network layers has various advantages, such as rationalizing the engineering of network infrastructures. The CPP interface aims at exposing and characterizing the IP transfer requirements to be met between the Customer Nodes (e.g., Media Gateway, Session Border Controller, etc.) when invoking IP transfer capabilities. These requirements include: reachability scope (e.g., limited scope, Internet-wide), bandwidth requirements, QoS parameters (e.g., one-way delay [RFC2679], loss [RFC2680] or one-way delay variation [RFC3393]), protection and high availability guidelines (e.g., sub-50ms/sub-100ms/second restoration). These requirements will then be translated into IP/MPLS-related technical clauses (e.g., need for recovery means, definition of the class of services, need for control plane protection, etc.) and in a further stage be addressed by the activation of adequate network features and technology-specific actions (e.g., MPLS-TE [RFC3346], RSVP [RFC2205], OSPF or IS-IS configuration, etc.).

Customer Nodes belong to a service infrastructure or an enterprise site (see Figure 2, Figure 3 and Figure 4). The connectivity between these Customer Nodes is implemented owing to the IP transfer capability implemented through a collaboration of a set of IP resources. IP transfer capabilities are considered by the above services as black boxes. Appropriate notifications and reports would be communicated (through dedicated means) to Customer Nodes to assess the level of the experienced IP transfer service. These notifications may also be used to assess the efficiency of the various policies enforced in the networking infrastructure to accommodate the requirements detailed in the CPP.

Having this CPP abstraction makes a clear distinction between the

connectivity provisioning requirements and the associated technology-specific rules that have been (or are to be) enforced in operational nodes, and which are meant to accommodate such requirements.

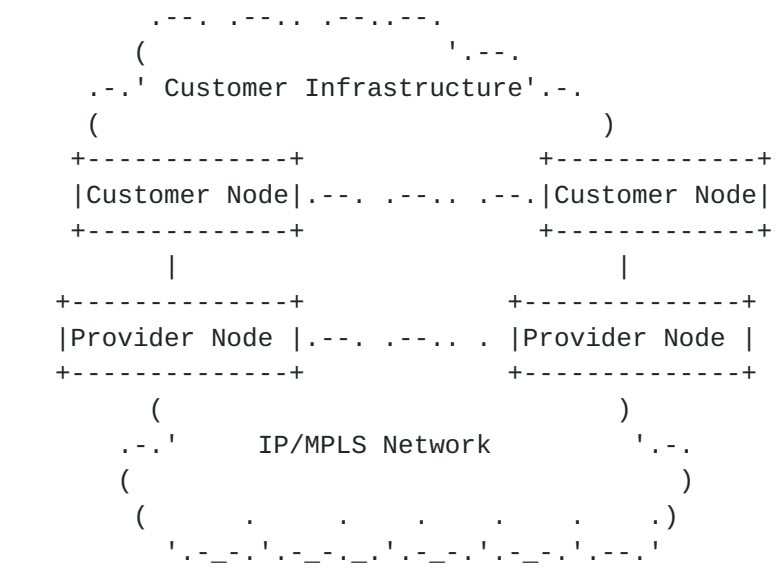

Figure 2: Reference Architecture

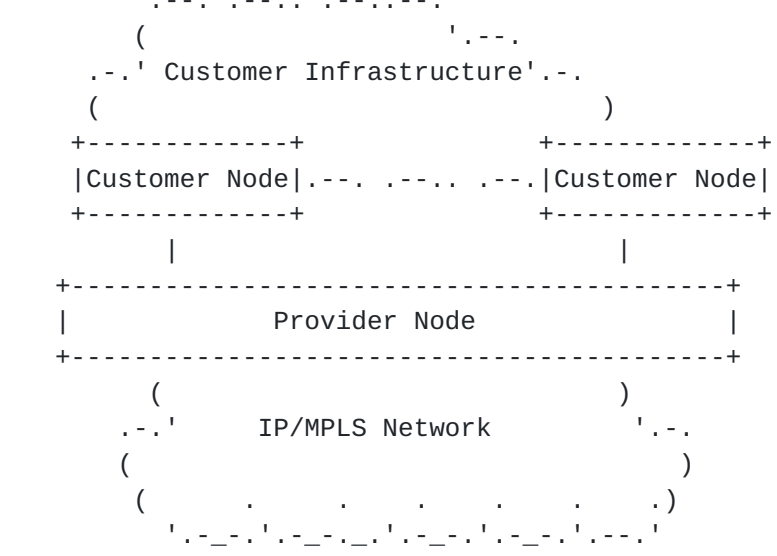

Figure 3: Reference Architecture (2)

As shown in Figure 4, the customer infrastructure can be connected over IP/MPLS networks managed by distinct network providers.

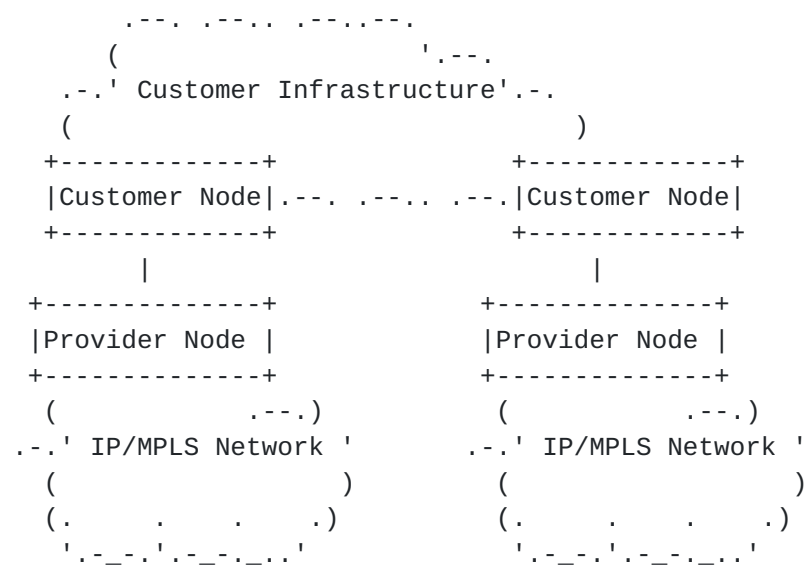

Figure 4: Reference Architecture (3)

1.1. Scope

This document details the clauses of the CPP. Protocols used to negotiate and to enforce a CPP are out of scope.

In addition to CPP clauses, other clauses may be included in an agreement between a customer and a provider (e.g., contact point, escalation procedure, incidents management, billing, etc.). It is out of scope of this document to detail all those additional clauses.

Examples of how to translate CPP clauses into technology-specific policies are provided for illustration purposes. It is out of scope of this document to provide an exhaustive list of the technical means to meet the objective included in a CPP.

2. Connectivity Provisioning Profile (CPP)

A CPP can be seen as an inventory of connectivity provisioning requirements with regard to IP transfer service. This section lists the clauses of the CPP. Figure 5 provides an overview of the CPP template. CPP clauses are elaborated in the following sub-sections.

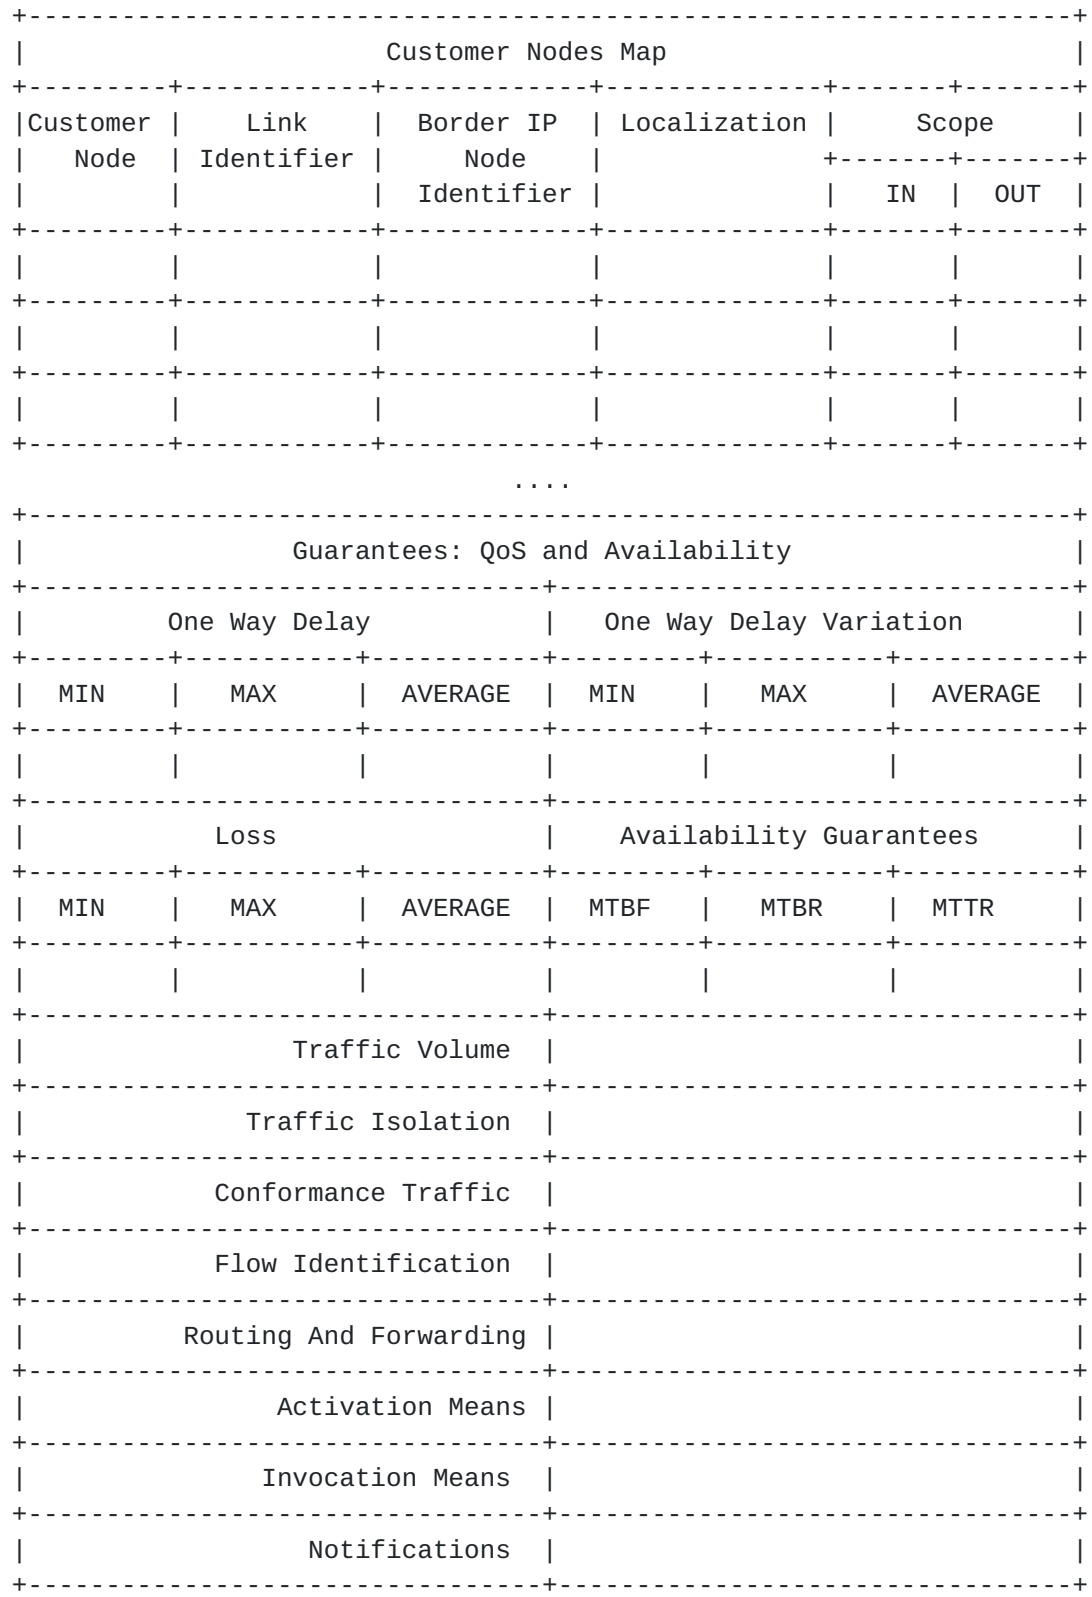

Figure 5: CPP Template

2.1. Customer Nodes or Service Nodes

A CPP must include the list of Customer Nodes (e.g., CEs) to be connected to the underlying IP transport network.

These nodes should be unambiguously identified (e.g., using a unique Service_identifier). For each Customer Node, a border link or a node part of the connectivity domain which is connected to the Customer Node should be identified.

Based on the location of the Customer Node (e.g., CE), appropriate operations to retrieve the corresponding border link or "Provider Node" (e.g., PE) should be undertaken. This operation can be manual or automated.

A "service site" would be located behind a given Customer Node. An identifier of the site may also be pertinent to be captured in the CPP for the provisioning of managed VPN [RFC4026] for instance (e.g., Site_identifier).

A Customer Node may be connected to several Provider Nodes and multiple Customer Nodes may be connected to the same Provider Node (see Figure 3).

2.2. Scope

The Scope specifies the connection between involved Customer Nodes. It is defined as an unidirectional parameter. Both directions should be described in the CPP.

The reachability scope may be defined as allowed destination IP prefixes that can be reached from the customer site.

Both IPv4 and IPv6 scopes may be distinguished.

A "Scope" delimits a topological (or geographical) network portion over which the performance and availability guarantees are not valid.

A scope may be defined by an "Ingress" and "Egress" points. Several types may be envisaged. Examples are listed below:

(1) "1:1" Pipe model. Only point to point communications are allowed.

(2) "1:N" Hose model. Only communications destined to a set of destinations are allowed.

(3) "1:any" Unspecified hose model. All outbound communications destined to whatever reachable destinations are to be offered.

A Scope indicating external "Ingress" or "Egress" nodes (i.e., not

connected to the Provider Network or Customer Network) may also be accepted provided that these nodes are unambiguously identified (e.g., IPv6 prefix).

2.3. QoS Guarantees

QoS guarantees denote a set of IP transfer performance metrics which characterize the quality of the IP transfer treatment to be experienced (when crossing an IP transport infrastructure) by a flow issued or destined to a (set of) "Customer Node(s)".

IP performance metrics can be expressed as qualitative or quantitative parameters. When quantitative metrics are used, maximum or average numerical values are provided together with a validity interval which should be indicated in the measurement method.

Several performance metrics have been defined such as:
- Traffic Loss [RFC2680]
- One way delay [RFC2679]
- One way delay variation [RFC3393]

The value of these parameters may be specific to a given path or a given scope (e.g., between two "Customer Nodes"). Concretely, IP performance metric value indicated in a CPP should reflect the measurement between a set of "Customer Node" or between a "Customer Node" and Provider Nodes attached to the IP network.

Meta-QoS class concept can be used when qualitative metrics are used [RFC5160].

2.4. Availability Guarantees

This clause specifies the percentage of the time when the agreed IP performance guarantees must be met. The guarantees cover both QoS deterioration (i.e., IP transfer service is available but it is below the agreed performance bounds), physical failures or service unavailability in general. In order to meet the availability guarantees, several engineering practices may be enforced in the border link such as allowing for multi-homed "Customer Nodes".

The following mechanisms are provided as illustration examples to show that several technical choices may be enforced to meet the service availability needs:
- When an IGP instance is running between the "Customer Node" and the "Provider Node", activate a dedicated protocol, such as BFD (Bi-directional Forwarding Detection [RFC5881][RFC5883]), to control IGP availability and then to ensure sub-second IGP adjacency failure detection.

- Use of LSP ping capability to detect LSP availability (check if the LSP is in place or not) [RFC4379].
- Pre-install backup LSPs for fast-rerouting when MPLS is used between "Customer Nodes" [RFC4090].
- Enable VRRP [RFC5798].
- Enable IP Fast Reroute features (e.g., [RFC5286]).

2.5. Traffic Volume

This clause characterizes the required capacity to be provided by the underlying IP transport network. This capacity is bound to a defined "Scope" (See Section 2.2) and IP transfer performance guarantees (see (Section 2.3)and (Section 2.4)).

Traffic volume may be expressed per border link and for both directions (i.e., incoming and outgoing). It is up to the administrative entity, which manages the IP transport network, to appropriately dimension its network [RFC5136] to meet the capacity requirements expressed in all negotiated CPPs.

2.6. Conformance Traffic

When capacity information (see Section 2.5) is included in the CPP, out-of-profile traffic would be handled separately.

Shaping/policing filters may be applied so as to assess wither the traffic is within the capacity profile or out of profile.

Out-of-profile traffic may be discared or under-classed (e.g., using the Lower than Best Effort PDB [RFC3662]).

Conditions on the injected packets MTU may also be indicated in the CPP.

2.7. Traffic Isolation

This clause indicates if the traffic issued by/destined to "Customer Nodes" should be isolated when crossing the IP transport network.

This clause is then translated into IP engineering policies such as activating dedicated tunnels using IPsec or establish BGP/MPLS VPN facilities [RFC4364], or a combination thereof. Activated means should be consistent with those used to meet the availability and performance guarantees.

When a "M:N" Scope is defined, optimization should be encouraged and not systematically assume a fully meshed tunnel topology.

2.8. Flow Identification

To identify the flows that need to be handled in the context of a given CPP, flow identifiers should be indicated in the CPP. This identifier is used for traffic classification purposes.

A flow identifier may be composed of the following parameters:

- source IP address,
- source port number,
- destination IP address,
- destination port number,
- ToS or DSCP field,
- tail-end tunnel endpoint, or
- a combination thereof.

Distinct treatments may be implemented for elastic and non elastic traffic (e.g., see the "Constraints on traffic" clause defined in [RFC5160]).

Flow classification rules may be specific to a given link or a given rule may be applied for all border links. This should be clearly captured in the CPP. For incoming traffic, some practices such as re-marking the DSCP field should be indicated in CPP. Re- marking action is under the responsibility of IP nodes, but this should be inferred by some constraints such as maintaining the service transparency (e.g., VPN services).

2.9. Routing & Forwarding

When outsourced routing actions are required, dedicated routes may be installed so as to convey the traffic to its (service) destination and avoiding some nodes (or ASes).

A requirement to dedicate a logical topology may also be envisaged owing to the activation of [RFC4915] or [RFC5120] .

This practice should be indicated in the CPP, otherwise routing actions belong to the underlying IP routing capabilities. Forwarding behavior (e.g., Per Domain Behaviour) may also be specified in a CPP. Nevertheless, it is optional. If indicated, consistency with the IP performance bounds defined in the CPP should be carefully ensured.

In the context of VoIP deployments for instance, a routing policy would be to avoid satellite links since this may degrade the offered service.

2.10. Activation Means

This clause indicates the required action(s) to be undertaken to activate access to the IP connectivity service.

Examples of these actions would be the activation of an IGP instance, the establishment of a BGP [RFC4271] or MB-BGP session [RFC4760], etc.

2.11. Invocation Means

Two types are defined:

Implicit: This clause when indicated means that no explicit means to invoke the connectivity service is required. Access to connectivity service is conditioned by the requested network capacity.

Explicit: This clause indicates the need for an explicit means to access the connectivity service. Examples of explicit invocation means include the use of RSVP [RFC2205] or RSVP-TE [RFC3209]. Appropriate access control procedures [RFC6601] would have to be enforced to check if the capacity actually used is not above the agreed threshold.

2.12. Notifications

For operation purposes (e.g., supervision) and service fulfillment needs, added value service platforms need to be notified about critical events which may impact the delivery of the service.

The notification procedure should be indicated in the CPP. This procedure may specify the type of information to be sent, the interval, data model, etc.

This may be enforced using SNMP, Syslog notifications, or a phone call!

3. IANA Considerations

This document does not require any action from IANA.

4. Security Considerations

This document does not define an architecture nor specify a protocol.

5. Acknowledgements

Some of the items listed above are results of discussions with P. Georgatsos, E. Mykoniati and D. Griffin. Special thanks to them.

6. References

6.1. Normative References

[RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF",

RFC 4915, June 2007.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

[RFC5136] Chimento, P. and J. Ishac, "Defining Network Capacity", RFC 5136, February 2008.

[RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, September 2008.

[RFC5798] Nadas, S., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, March 2010.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

[RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.

6.2. Informative References

[RFC3346] Boyle, J., Gill, V., Hannan, A., Cooper, D., Awduche, D., Christian, B., and W. Lai, "Applicability Statement for Traffic Engineering with MPLS", RFC 3346, August 2002.

[RFC3662] Bless, R., Nichols, K., and K. Wehrle, "A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services", RFC 3662, December 2003.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC5160] Levis, P. and M. Boucadair, "Considerations of Provider-to-Provider Agreements for Internet-Scale Quality of Service (QoS)", RFC 5160, March 2008.

[RFC6601] Ash, G. and D. McDysan, "Generic Connection Admission Control (GCAC) Algorithm Specification for IP/MPLS Networks", RFC 6601, April 2012.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
France Telecom
Rennes, 35000
France

Email: christian.jacquenet@orange.com

Ning Wang
Centre for Communication System Research
University of Surrey
Guildford,
UK

Email: n.wang@surrey.ac.uk