

Network Working Group
Internet Draft
Intended status: Informational
Expires: February 2017

A. Bashandy, Ed.
C. Filsfils
Cisco Systems
P. Mohapatra
Sproute Networks
August 1, 2016

BGP Prefix Independent Convergence
[draft-ietf-rtgwg-bgp-pic-02.txt](#)

Abstract

In the network comprising thousands of iBGP peers exchanging millions of routes, many routes are reachable via more than one next-hop. Given the large scaling targets, it is desirable to restore traffic after failure in a time period that does not depend on the number of BGP prefixes. In this document we proposed an architecture by which traffic can be re-routed to ECMP or pre-calculated backup paths in a timeframe that does not depend on the number of BGP prefixes. The objective is achieved through organizing the forwarding data structures in a hierarchical manner and sharing forwarding elements among the maximum possible number of routes. The proposed technique achieves prefix independent convergence while ensuring incremental deployment, complete automation, and zero management and provisioning effort. It is noteworthy to mention that the benefits of BGP-PIC are hinged on the existence of more than one path whether as ECMP or primary-backup.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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/1id-abstracts.txt>

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

This Internet-Draft will expire on December 1, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction.....	3
1.1.	Conventions used in this document.....	4
1.2.	Terminology.....	4
2.	Overview.....	5
3.	Constructing the Shared Hierarchical Forwarding Chain.....	7
3.1.	Example 1: Primary-Backup Path Scenario.....	8
3.2.	Example 2: Platforms with Limited Levels of Hierarchy.....	9
4.	Forwarding Behavior.....	13
5.	Forwarding Chain Adjustment at a Failure.....	15
5.1.	BGP-PIC core.....	16
5.2.	BGP-PIC edge.....	17
5.2.1.	Adjusting forwarding Chain in egress node failure...17	
5.2.2.	Adjusting Forwarding Chain on PE-CE link Failure....17	
5.3.	Handling Failures for Flattened Forwarding Chains.....	18
6.	Properties.....	19
6.1.	Coverage.....	19
6.1.1.	A remote failure on the path to a BGP next-hop.....	19

6.1.2.	A local failure on the path to a BGP next-hop.....	19
6.1.3.	A remote iBGP next-hop fails.....	20
6.1.4.	A local eBGP next-hop fails.....	20
6.2.	Performance.....	20
6.2.1.	Perspective.....	20
6.3.	Automated.....	21
6.4.	Incremental Deployment.....	22
7.	Dependency.....	22
7.1.	Hierarchical Hardware FIB.....	22
7.2.	Availability of more than one primary or secondary BGP next-hops.....	22
7.3.	Pre-Computation of a secondary BGP next-hop.....	23
8.	Security Considerations.....	23
9.	IANA Considerations.....	23
10.	Conclusions.....	23
11.	Acknowledgments.....	25
12.	References.....	23
12.1.	Normative References.....	23
12.2.	Informative References.....	24

1. Introduction

As a path vector protocol, BGP is inherently slow due to the serial nature of reachability propagation. BGP speakers exchange reachability information about prefixes[2][3] and, for labeled address families, namely AFI/SAFI 1/4, 2/4, 1/128, and 2/128, an edge router assigns local labels to prefixes and associates the local label with each advertised prefix such as L3VPN [8], 6PE [9], and Softwire [7] using BGP label unicast technique[4]. A BGP speaker then applies the path selection steps to choose the best path. In modern networks, it is not uncommon to have a prefix reachable via multiple edge routers. In addition to proprietary techniques, multiple techniques have been proposed to allow for BGP to advertise more than one path for a given prefix [6][11][12], whether in the form of equal cost multipath or primary-backup. Another more common and widely deployed scenario is L3VPN with multi-homed VPN sites with unique Route Distinguisher.

This document proposes a hierarchical and shared forwarding chain organization that allows traffic to be restored to pre-calculated alternative equal cost primary path or backup path in a time period that does not depend on the number of BGP prefixes. The technique relies on internal router behavior that is completely transparent to the operator and can be incrementally deployed and enabled with zero operator intervention.

1.1. 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](#) [1].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC-2119](#) significance.

1.2. Terminology

This section defines the terms used in this document. For ease of use, we will use terms similar to those used by L3VPN [8]

- o BGP prefix: It is a prefix P/m (of any AFI/SAFI) that a BGP speaker has a path for.
- o IGP prefix: It is a prefix P/m (of any AFI/SAFI) that is learnt via an Interior Gateway Protocol, such as OSPF and ISIS, has a path for. The prefix may be learnt directly through the IGP or redistributed from other protocol(s)
- o CE: It is an external router through which an egress PE can reach a prefix P/m.
- o Ingress PE, "iPE": It is a BGP speaker that learns about a prefix through a IBGP peer and chooses an egress PE as the next-hop for the prefix..
- o Path: It is the next-hop in a sequence of unique connected nodes starting from the current node and ending with the destination node or network identified by the prefix.
- o Recursive path: It is a path consisting only of the IP address of the next-hop without the outgoing interface. Subsequent lookups are needed to determine the outgoing interface.
- o Non-recursive path: It is a path consisting of the IP address of the next-hop and one outgoing interface
- o Primary path: It is a recursive or non-recursive path that can be used all the time. A prefix can have more than one primary path
- o Backup path: It is a recursive or non-recursive path that can be used only after some or all primary paths become unreachable

- o Leaf: A leaf is container data structure for a prefix or local label. Alternatively, it is the data structure that contains prefix specific information.
- o IP leaf: Is the leaf corresponding to an IPv4 or IPv6 prefix
- o Label leaf. It is the leaf corresponding to a locally allocated label such as the VPN label on an egress PE [8].
- o Pathlist: It is an array of paths used by one or more prefix to forward traffic to destination(s) covered by a IP prefix. Each path in the pathlist carries its "path-index" that identifies its position in the array of paths. A pathlist may contain a mix of primary and backup paths
- o OutLabel-List: Each labeled prefix is associated with an OutLabel-List. The OutLabel-List is an array of one or more outgoing labels and/or label actions where each label or label action has 1-to-1 correspondence to a path in the pathlist. Label actions are: push the label, pop the label, or swap the incoming label with the label in the Outlabel-Array entry. The prefix may be an IGP or BGP prefix
- o Adjacency: It is the layer 2 encapsulation leading to the layer 3 directly connected next-hop
- o Dependency: An object X is said to be a dependent or Child of object Y if Object Y cannot be deleted unless object X is no longer a dependent/child of object Y
- o Route: It is a prefix with one or more paths associated with it. Hence the minimum set of objects needed to construct a route is a leaf and a pathlist.

2. Overview

The idea of BGP-PIC is based on two pillars

- o A shared hierarchical Forwarding Chain
- o A forwarding plane that supports multiple levels of indirection

To illustrate the two pillars above, we will use an example of a simple multihomed L3VPN [8] prefix in a BGP-free core running LDP [5] or segment routing over MPLS forwarding plane [14].

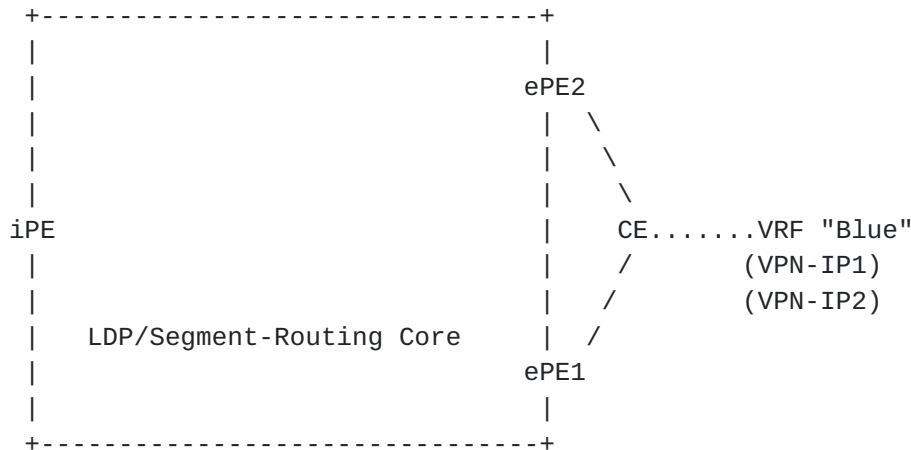


Figure 1 VPN prefix reachable via multiple PEs

Referring to Figure 1, suppose the iPE (the ingress PE) receives NLRI for the VPN prefixes VPN-IP1 and VPN-IP2 from two egress PEs, ePE1 and ePE2 with next-hop BGP-NH1 and BGP-NH2, respectively. Assume that ePE1 advertise the VPN labels VPN-L11 and VPN-L12 while ePE2 advertise the VPN labels VPN-L21 and VPN-L22 for VPN-IP1 and VPN-IP2, respectively. Suppose that BGP-NH1 and BGP-NH2 are resolved via the IGP prefixes IGP-IP1 and IGP-P2, where each happen to have 2 ECMP paths with IGP-NH1 and IGP-NH2 reachable via the interfaces I1 and I2, respectively. Suppose that local labels (whether LDP[5] or segment routing [14]) on the downstream LSRs for IGP-IP1 are IGP-L11 and IGP-L12 while for IGP-P2 are IGP-L21 and IGP-L22.

Based on the information about NLRI and the resolving IGP prefixes, a hierarchical forwarding chain can be constructed as shown in Figure 2.

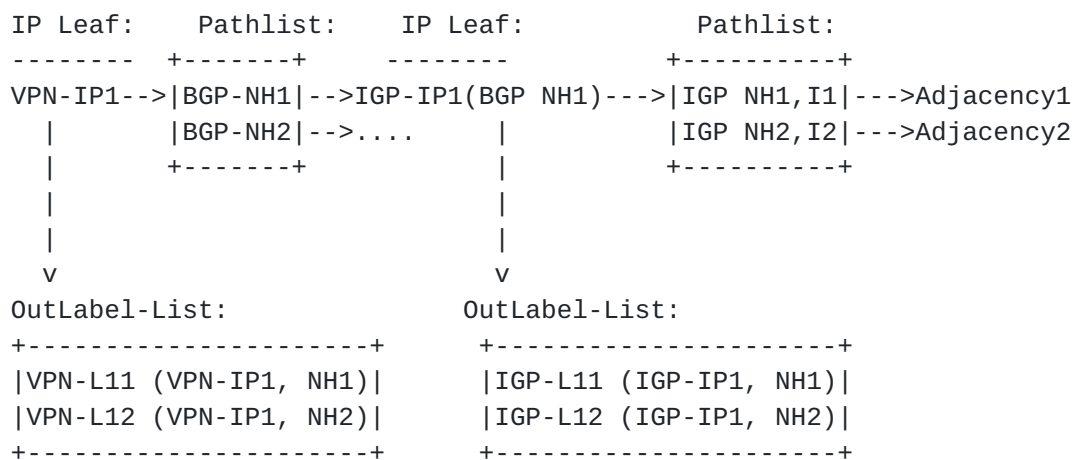


Figure 2 Shared Hierarchical Forwarding Chain at iPE

The forwarding chain depicted in Figure 2 illustrates the first

pillar, which is sharing and hierarchy. We can see that the BGP

Bashandy

Expires February 1, 2017

[Page 6]

pathlist consisting of BGP-NH1 and BGP-NH2 is shared by all NLRIs reachable via ePE1 and ePE2. As such, it is possible to make changes to the pathlist without having to make changes to the NLRIs. For example, if BGP-NH2 becomes unreachable, there is no need to modify any of the possibly large number of NLRIs. Instead only the shared pathlist needs to be modified. Likewise, due to the hierarchical structure of the forwarding chain, it is possible to make modifications to the IGP routes without having to make any changes to the BGP NLRIs. For example, if the interface "I2" goes down, only the shared IGP pathlist needs to be updated, but none of the IGP prefixes sharing the IGP pathlist nor the BGP NLRIs using the IGP prefixes for resolution need to be modified.

Figure 2 can also be used to illustrate the second BGP-PIC pillar. Having a deep forwarding chain such as the one illustrated in Figure 2 requires a forwarding plane that is capable of accessing multiple levels of indirection in order to calculate the outgoing interface(s) and next-hops(s). While a deeper forwarding chain minimizes the re-convergence time on topology change, there will always exist platforms with limited capabilities and hence imposing a limit on the depth of the forwarding chain. The example in [Section 3.2](#) illustrates how to gracefully trade off convergence speed with the number of hierarchical levels to support platforms with different capabilities.

3. Constructing the Shared Hierarchical Forwarding Chain

Constructing the forwarding chain is an application of the two pillars described in [Section 2](#).

The whole process starts when BGP downloads a prefix to FIB. The prefix contains one or more outgoing paths. For certain labeled prefixes, such as VPN [\[8\]](#) prefixes, each path may be associated with an outgoing label and the prefix itself may be assigned a local label. The list of outgoing paths defines a pathlist. If such pathlist does not already exist, then FIB creates a new pathlist, otherwise the existing pathlist is used. The BGP prefix is added as a dependent of the pathlist.

The previous step constructs the upper part of the hierarchical forwarding chain. The forwarding chain is completed by resolving the paths of the pathlist. A BGP path usually consists of a next-hop. The next-hop is resolved by finding a matching IGP prefix.

The end result is a hierarchical shared forwarding chain where the BGP pathlist is shared by all BGP prefixes that use the same list of paths and the IGP prefix is shared by all pathlists that have a path resolving via that IGP prefix. It is noteworthy to mention that the forwarding chain is constructed without any operator intervention at

all.

Bashandy

Expires February 1, 2017

[Page 7]

The remainder of this section illustrates two examples. The first example illustrates the applicability of BGP-PIC in a primary-backup path deployment. The second example illustrates how BGP-PIC can be applied in cases where the forwarding plane supports limited number of indirections.

3.1. Example 1: Primary-Backup Path Scenario

Consider the egress PE ePE1 in the case of the multi-homed VPN prefixes in the BGP-free core depicted in Figure 1. Suppose ePE1 determines that the primary path is the external path but the backup path is the iBGP path to the other PE ePE2 with next-hop BGP-NH2. ePE2 constructs the forwarding chain depicted in Figure 3. We are only showing a single VPN prefix for simplicity. But all prefixes that are multihomed to ePE1 and ePE2 share the BGP pathlist.

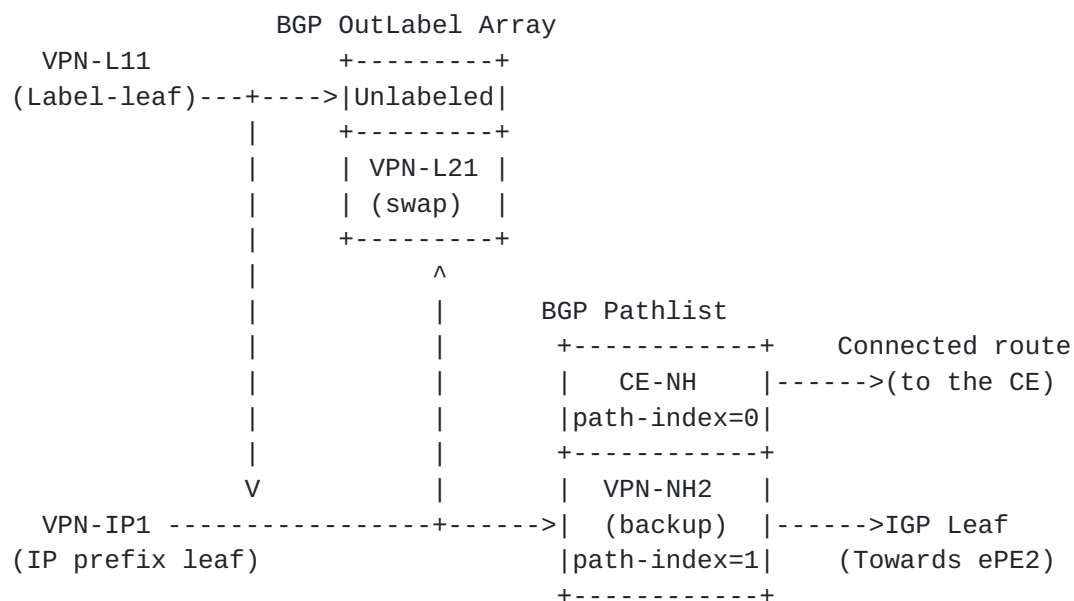


Figure 3 : VPN Prefix Forwarding Chain with eiBGP paths on egress PE

The example depicted in Figure 3 differs from the example in Figure 2 in two main aspects. First, as long as the primary path towards the CE (external path) is useable, it will be the only path used for forwarding while the OutLabel-List contains both the unlabeled label (primary path) and the VPN label (backup path) advertised by the backup path ePE2. The second aspect is presence of the label leaf corresponding to the VPN prefix. This label leaf is used to match VPN traffic arriving from the core. Note that the label leaf shares the OutLabel-List and the pathlist with the IP prefix.

3.2. Example 2: Platforms with Limited Levels of Hierarchy

This example uses a case of inter-AS option C [8] where there are 3 levels of hierarchy. Figure 4 illustrates the sample topology. To force 3 levels of hierarchy, the ASBRs on the ingress domain (domain 1) advertise the core routers of the egress domain (domain 2) to the ingress PE (iPE) via BGP-LU [4] instead of redistributing them into the IGP of domain 1. The end result is that the ingress PE (iPE) has 2 levels of recursion for the VPN prefix VPN-IP1 and VPN2-P2.

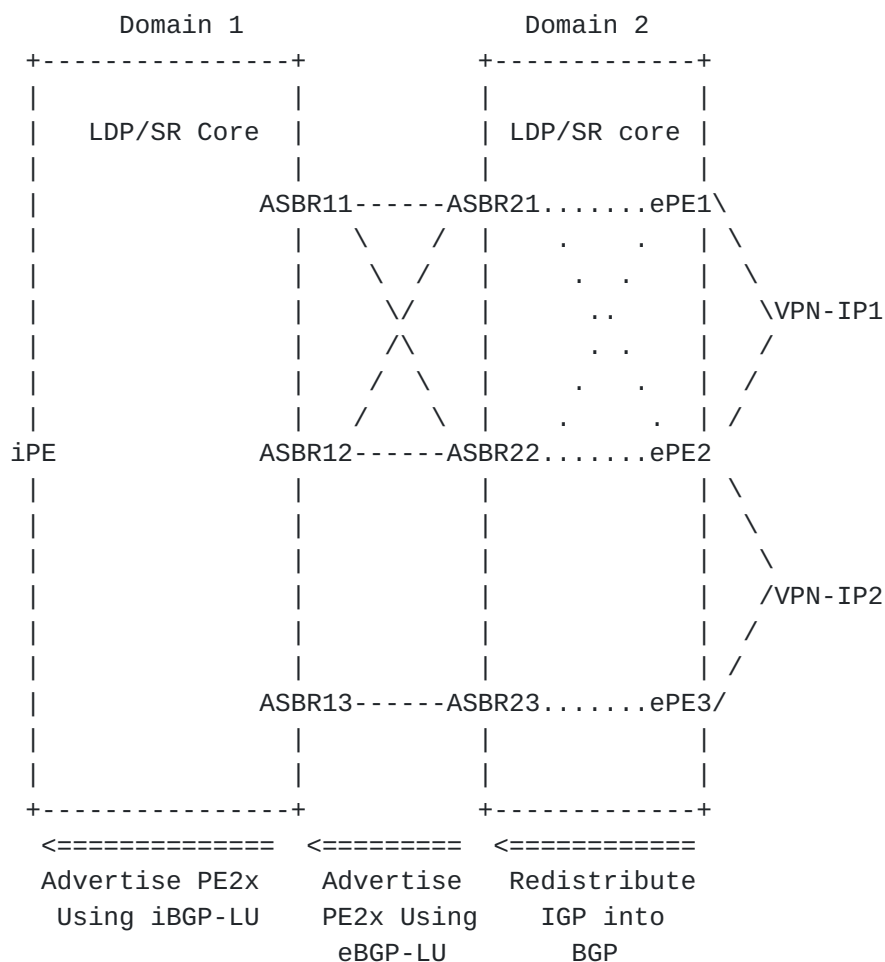


Figure 4 Sample 3-level hierarchy topology

We will make the following assumptions about connectivity

- o In "domain 2", both ASBR21 and ASBR22 can reach both ePE1 and ePE2 using the same distance
- o In "domain 2", only ASBR23 can reach ePE3

- o In "domain 1", iPE (the ingress PE) can reach ASBR11, ASBR12, and ASBR13 via IGP using the same distance.

We will make the following assumptions about the labels

- o The VPN labels advertised by ePE1 and ePE2 for prefix VPN-IP1 are VPN-L11 and VPN-L21, respectively
- o The VPN labels advertised by ePE2 and ePE3 for prefix VPN-IP2 are VPN-L22 and VPN-L32, respectively
- o The labels advertised by ASBR11 to iPE using BGP-LU [4] for the egress PEs ePE1 and ePE2 are LASBR11(ePE1) and LASBR11(ePE2), respectively.
- o The labels advertised by ASBR12 to iPE using BGP-LU [4] for the egress PEs ePE1 and ePE2 are LASBR12(ePE1) and LASBR12(ePE2), respectively
- o The label advertised by ASBR11 to iPE using BGP-LU [4] for the egress PE ePE3 is LASBR13(ePE3)
- o The IGP labels advertised by the next hops directly connected to iPE towards ASBR11, ASBR12, and ASBR13 in the core of domain 1 are IGP-L11, IGP-L12, and IGP-L13, respectively.

The diagram in Figure 5 illustrates the forwarding chain in iPE assuming that the forwarding hardware in iPE supports 3 levels of hierarchy. The leaves corresponding to the ASBRs on domain 1 (ASBR11, ASBR12, and ASBR13) are at the bottom of the hierarchy. There are few important points:

- o Because the hardware supports the required depth of hierarchy, the sizes of a pathlist equal the size of the label list associated with the leaves using this pathlist
- o The index inside the pathlist entry indicates the label that will be picked from the Outlabel-List if that path is chosen by the forwarding engine hashing function.

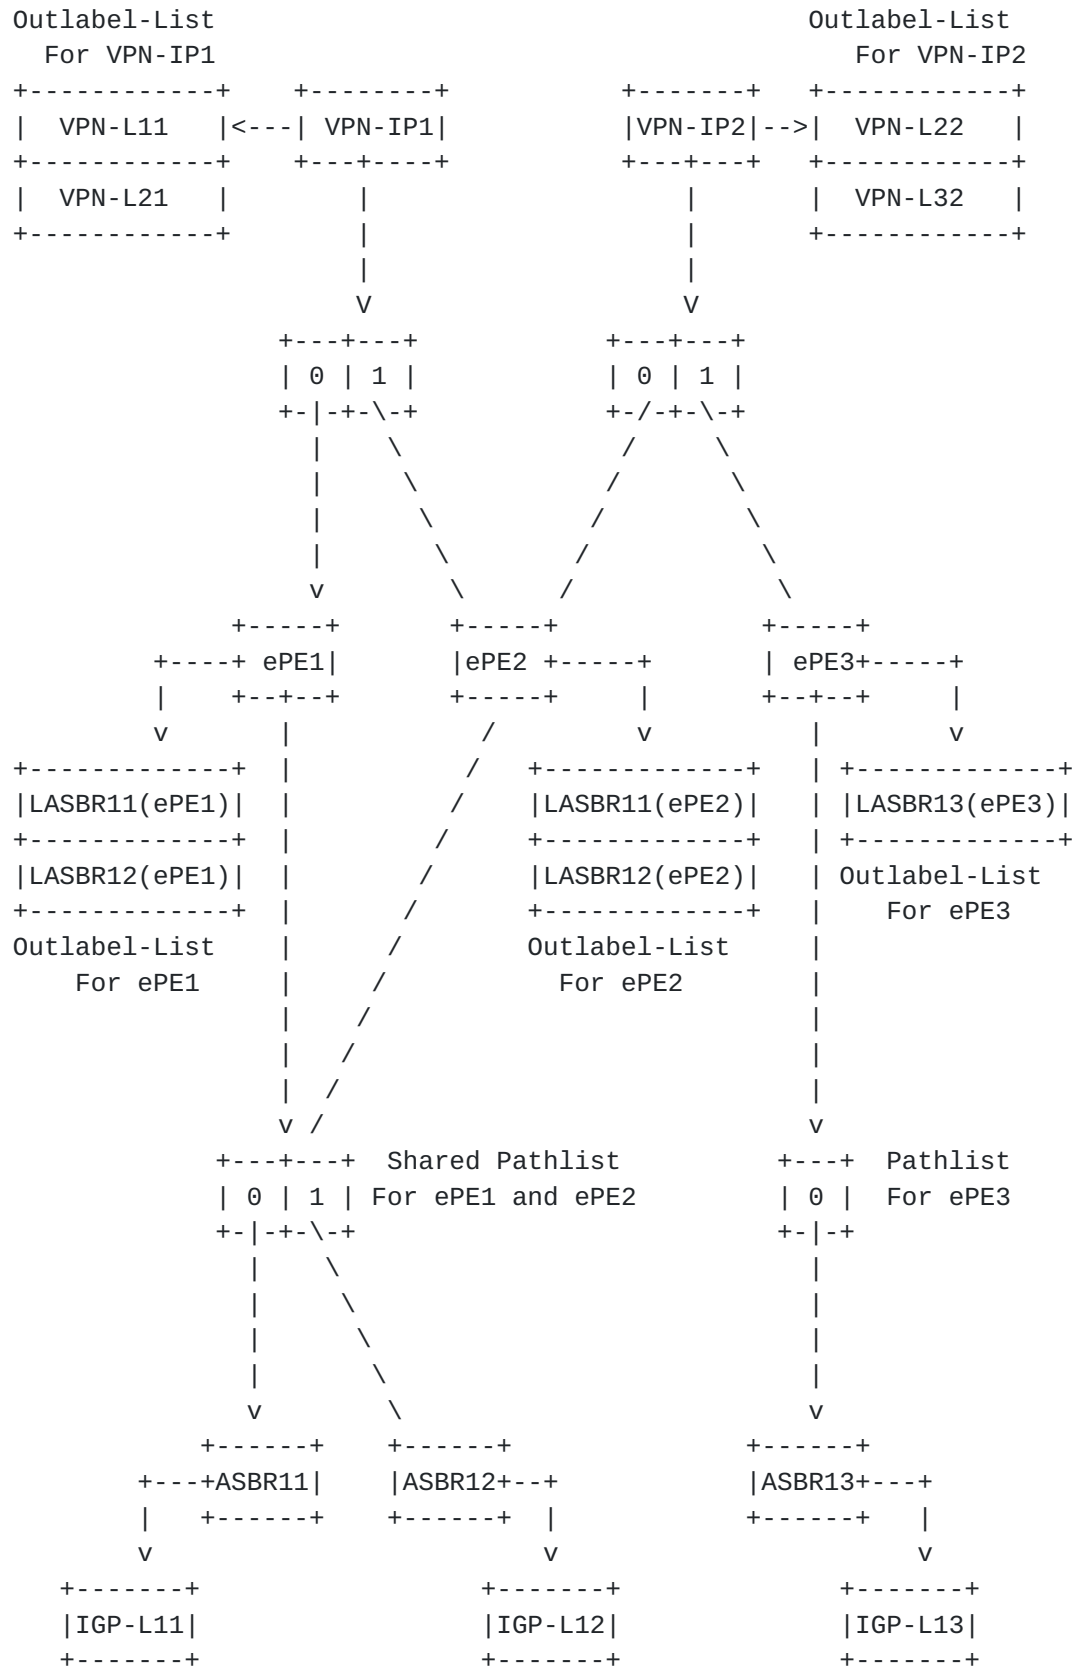


Figure 5 : Forwarding Chain for hardware supporting 3 Levels

Now suppose the hardware on iPE (the ingress PE) supports 2 levels of hierarchy only. In that case, the 3-levels forwarding chain in Figure 5 needs to be "flattened" into 2 levels only.




Figure 6 : Flattening 3 levels to 2 levels of Hierarchy on iPE

Figure 6 represents one way to "flatten" a 3 levels hierarchy into two levels. There are few important points:

- o The flattened pathlists have label lists associated with them. The size of the label list associated with the flattened pathlist equals the size of the pathlist. Hence it is possible that an implementation includes these label lists in the flattened pathlist itself
- o Because of "flattening", the size of a flattened pathlist may not be equal to the size of the label lists of leaves using the flattened pathlist.
- o The indices inside a flattened pathlist still indicate the label index in the Outlabel-Lists of the leaves using that pathlist. Because the size of the flattened pathlist may be different from the size of the label lists of the leaves, the indices may be repeated
- o Let's take a look at the flattened pathlist used by the prefix "VPN-IP2". The pathlist associated with the prefix "VPN-IP2" has three entries.
 - o The first and second entry have index "0". This is because both entries correspond to ePE2. Hence when hashing performed by the forwarding engine results in using first or the second entry in the pathlist, the forwarding engine will pick the correct VPN label "VPN-L22", which is the label advertised by ePE2 for the prefix "VPN-IP2"
 - o The third entry has the index "1". This is because the third entry corresponds to ePE3. Hence when the hashing is performed by the forwarding engine results in using the third entry in the flattened pathlist, the forwarding engine will pick the correct VPN label "VPN-L32", which is the label advertised by "ePE3" for the prefix "VPN-IP2"

4. Forwarding Behavior

This section explains how the forwarding plane uses the hierarchical shared forwarding chain to forward a packet.

When a packet arrives at a router, it matches a leaf. A labeled packet matches a label leaf while an IP packet matches an IP prefix leaf. The forwarding engines walks the forwarding chain starting from the leaf until the walk terminates on an adjacency. Thus when a packet arrives, the chain is walked as follows:

1. Lookup the leaf based on the destination address or the label at the top of the packet
2. Retrieve the parent pathlist of the leaf

3. Pick the outgoing path from the list of resolved paths in the pathlist. The method by which the outgoing path is picked is beyond the scope of this document (i.e. flow-preserving hash exploiting entropy within the MPLS stack and IP header). Let the "path-index" of the outgoing path be "i".
4. If the prefix is labeled, use the "path-index" "i" to retrieve the ith label "Li" stored the ith entry in the OutLabel-List and apply the label action of the label on the packet (e.g. for VPN label on the ingress PE, the label action is "push").
5. Move to the parent of the chosen path "i"
6. If the chosen path "i" is recursive, move to its parent prefix and go to step 2
7. If the chosen path "i" is non-recursive move to its parent adjacency
8. Encapsulate the packet in the L2 string specified by the adjacency and send the packet out.

Let's apply the above forwarding steps to the forwarding chain depicted in Figure 2 in [Section 2](#). Suppose a packet arrives at ingress PE iPE from an external neighbor. Assume the packet matches the VPN prefix VPN-IP1. While walking the forwarding chain, the forwarding engine applies a hashing algorithm to choose the path and the hashing at the BGP level yields path 0 while the hashing at the IGP level yields path 1. In that case, the packet will be sent out of interface I2 with the label stack "IGP-L12,VPN-L11".

Now let's try and apply the above steps to the flattened forwarding chain illustrated in Figure 6.

- o Suppose a packet arrives at "iPE" and matches the VPN prefix "VPN-IP2"
- o The forwarding engine walks to the parent of the "VPN_P2", which is the flattened pathlist and applies a hashing algorithm to pick a path
- o Suppose the hashing by the forwarding engine picks the second entry in the flattened pathlist associated with the leaf "VPN-IP2".
- o Because the second entry has the index "0", the label "VPN-L22" is pushed on the packet

- o At the same time, the forwarding engine picks the second label from the Outlabel-Array associated with the flattened pathlist. Hence the next label that is pushed is "LASBR12(ePE2)"
- o The forwarding engine now moves to the parent of the flattened pathlist corresponding to the second entry. The parent is the IGP label leaf corresponding to "ASBR12"
- o So the packet is forwarded towards the ASBR "ASBR12" and the IGP label at the top will be "L12"

Based on the above steps, a packet arriving at iPE and destined to the prefix VPN-L22 reaches its destination as follows

- o iPE sends the packet along the shortest path towards ASBR12 with the following label stack starting from the top: {L12, LASBR12(ePE2), VPN-L22}.
- o The penultimate hop of ASBR12 pops the top label "L12". Hence the packet arrives at ASBR12 with the label stack {LASBR12(ePE2), VPN-L22} where "LASBR12(ePE2)" is the top label.
- o ASBR12 swaps "LASBR12(ePE2)" with the label "LASBR22(ePE2)", which is the label advertised by ASBR22 for the ePE2 (the egress PE).
- o ASBR22 receives the packet with "LASBR22(ePE2)" at the top.
- o Hence ASBR22 swaps "LASBR22(ePE2)" with the IGP label for ePE2 advertised by the next-hop towards ePE2 in domain 2, and sends the packet along the shortest path towards ePE2.
- o The penultimate hop of ePE2 pops the top label. Hence ePE2 receives the packet with the top label VPN-L22 at the top
- o ePE2 pops "VPN-L22" and sends the packet as a pure IP packet towards the destination VPN-IP2.

5. Forwarding Chain Adjustment at a Failure

The hierarchical and shared structure of the forwarding chain explained in [Section 2](#) allows modifying a small number of forwarding chain objects to re-route traffic to a pre-calculated equal-cost or backup path without the need to modify the possibly very large number of BGP prefixes. In this section, we go over various core and edge failure scenarios to illustrate how FIB manager can utilize the forwarding chain structure to achieve BGP prefix independent convergence.

5.1. BGP-PIC core

This section describes the adjustments to the forwarding chain when a core link or node fails but the BGP next-hop remains reachable.

There are two case: remote link failure and attached link failure. Node failures are treated as link failures.

When a remote link or node fails, IGP on the ingress PE receives advertisement indicating a topology change so IGP re-converges to either find a new next-hop and/or outgoing interface or remove the path completely from the IGP prefix used to resolve BGP next-hops. IGP and/or LDP download the modified IGP leaves with modified outgoing labels for labeled core.

When a local link fails, FIB manager detects the failure almost immediately. The FIB manager marks the impacted path(s) as unusable so that only useable paths are used to forward packets. Hence only IGP pathlists with paths using the failed local link need to be modified. All other pathlists are not impacted. Note that in this particular case there is actually no need even to backwalk to IGP leaves to adjust the OutLabel-Lists because FIB can rely on the path-index stored in the useable paths in the pathlist to pick the right label.

It is noteworthy to mention that because FIB manager modifies the forwarding chain starting from the IGP leaves only, BGP pathlists and leaves are not modified. Hence traffic restoration occurs within the time frame of IGP convergence, and, for local link failure, assuming a backup path has been precomputed, within the timeframe of local detection (e.g. 50ms). Examples of solutions that pre-computing backup paths are IP FRR [[16](#)] remote LFA [[17](#)], Ti-LFA [[15](#)] and MRT [[18](#)] or eBGP path having a backup path [[10](#)].

Let's apply the procedure to the forwarding chain depicted in Figure 2. Suppose a remote link failure occurs and impacts the first ECMP IGP path to the remote BGP next-hop. Upon IGP convergence, the IGP pathlist used by the BGP next-hop is updated to reflect the new topology (one path instead of two). As soon as the IGP convergence is effective for the BGP next-hop entry, the new forwarding state is immediately available to all dependent BGP prefixes. The same behavior would occur if the failure was local such as an interface going down. As soon as the IGP convergence is complete for the BGP next-hop IGP route, all its BGP depending routes benefit from the new path. In fact, upon local failure, if LFA protection is enabled for the IGP route to the BGP next-hop and a backup path was pre-computed and installed in the pathlist, upon the local interface failure, the LFA backup path is immediately activated (sub-50msec) and thus protection benefits all the depending BGP traffic through

the hierarchical forwarding dependency between the routes.

Bashandy

Expires February 1, 2017

[Page 16]

5.2. BGP-PIC edge

This section describes the adjustments to the forwarding chains as a result of edge node or edge link failure.

5.2.1. Adjusting forwarding Chain in egress node failure

When an edge node fails, IGP on neighboring core nodes send route updates indicating that the edge node is no longer reachable. IGP running on the iBGP peers instructs FIB to remove the IP and label leaves corresponding to the failed edge node from FIB. So FIB manager performs the following steps:

- o FIB manager deletes the IGP leaf corresponding to the failed edge node
- o FIB manager backwalks to all dependent BGP pathlists and marks that path using the deleted IGP leaf as unresolved
- o Note that there is no need to modify BGP leaves because each path in the pathlist carries its path index and hence the correct outgoing label will be picked. Consider for example the forwarding chain depicted in Figure 2. If the 1st BGP path becomes unresolved, then the forwarding engine will only use the second path for forwarding. Yet the pathindex of that single resolved path will still be 1 and hence the label VPN-L12 will be pushed.

5.2.2. Adjusting Forwarding Chain on PE-CE link Failure

Suppose the link between an edge router and its external peer fails. There are two scenarios (1) the edge node attached to the failed link performs next-hop self and (2) the edge node attached to the failure advertises the IP address of the failed link as the next-hop attribute to its iBGP peers.

In the first case, the rest of iBGP peers will remain unaware of the link failure and will continue to forward traffic to the edge node until the edge node attached to the failed link withdraws the BGP prefixes. If the destination prefixes are multi-homed to another iBGP peer, say ePE2, then FIB manager on the edge router detecting the link failure applies the following steps:

- o FIB manager backwalks to the BGP pathlists marks the path through the failed link to the external peer as unresolved
- o Hence traffic will be forwarded used the backup path towards ePE2

- o For labeled traffic
 - o The Outlabel-List attached to the BGP leaf already contains an entry corresponding to the backup path.
 - o The label entry in OutLabel-List corresponding to the internal path to backup egress PE has swap action to the label advertised by backup egress PE
 - o For an arriving label packet (e.g. VPN), the top label is swapped with the label advertised by backup egress PE and the packet is sent towards that backup egress PE
- o For unlabeled traffic, packets are simply redirected towards backup egress PE.

In the second case where the edge router uses the IP address of the failed link as the BGP next-hop, the edge router will still perform the previous steps. But, unlike the case of next-hop self, IGP on failed edge node informs the rest of the iBGP peers that IP address of the failed link is no longer reachable. Hence the FIB manager on iBGP peers will delete the IGP leaf corresponding to the IP prefix of the failed link. The behavior of the iBGP peers will be identical to the case of edge node failure outlined in [Section 5.2.1](#).

It is noteworthy to mention that because the edge link failure is local to the edge router, sub-50 msec convergence can be achieved as described in [\[10\]](#).

Let's try to apply the case of next-hop self to the forwarding chain depicted in Figure 3. After failure of the link between ePE1 and CE, the forwarding engine will route traffic arriving from the core towards VPN-NH2 with path-index=1. A packet arriving from the core will contain the label VPN-L11 at top. The label VPN-L11 is swapped with the label VPN-L21 and the packet is forwarded towards ePE2.

5.3. Handling Failures for Flattened Forwarding Chains

As explained in the Example in [Section 3.2](#) if the number of hierarchy levels of a platform cannot support the native number of hierarchy levels of a recursive forwarding chain, the instantiated forwarding chain is constructed by flattening two or more levels. Hence a 3 levels chain in Figure 5 is flattened into the 2 levels chain in Figure 6.

While reducing the benefits of BGP-PIC, flattening one hierarchy into a shallower hierarchy does not always result in a complete loss of the benefits of the BGP-PIC. To illustrate this fact suppose ASBR12 is no longer reachable in domain 1. If the platform supports the full hierarchy depth, the forwarding chain is the one depicted

in Figure 5 and hence the FIB manager needs to backwalk one level to the pathlist shared by "ePE1" and "ePE2" and adjust it. If the platform supports 2 levels of hierarchy, then a useable forwarding chain is the one depicted in Figure 6. In that case, if ASBR12 is no longer reachable, the FIB manager has to backwalk to the two flattened pathlists and update both of them.

The main observation is that the loss of convergence speed due to the loss of hierarchy depth depends on the structure of the forwarding chain itself. To illustrate this fact, let's take two extremes. Suppose the forwarding objects in level $i+1$ depend on the forwarding objects in level i . If every object on level $i+1$ depends on a separate object in level i , then flattening level i into level $i+1$ will not result in loss of convergence speed. Now let's take the other extreme. Suppose " n " objects in level $i+1$ depend on 1 object in level i . Now suppose FIB flattens level i into level $i+1$. If a topology change results in modifying the single object in level i , then FIB has to backwalk and modify " n " objects in the flattened level, thereby losing all the benefit of BGP-PIC. Experience shows that flattening forwarding chains usually results in moderate loss of BGP-PIC benefits. Further analysis is needed to corroborate and quantify this statement.

6. Properties

6.1. Coverage

All the possible failures, except CE node failure, are covered, whether they impact a local or remote IGP path or a local or remote BGP next-hop as described in [Section 5](#). This section provides details for each failure and now the hierarchical and shared FIB structure proposed in this document allows recovery that does not depend on number of BGP prefixes.

6.1.1. A remote failure on the path to a BGP next-hop

Upon IGP convergence, the IGP leaf for the BGP next-hop is updated upon IGP convergence and all the BGP depending routes leverage the new IGP forwarding state immediately.

This BGP resiliency property only depends on IGP convergence and is independent of the number of BGP prefixes impacted.

6.1.2. A local failure on the path to a BGP next-hop

Upon LFA protection, the IGP leaf for the BGP next-hop is updated to use the precomputed LFA backup path and all the BGP depending routes leverage this LFA protection.

This BGP resiliency property only depends on LFA protection and is independent of the number of BGP prefixes impacted.

6.1.3. A remote iBGP next-hop fails

Upon IGP convergence, the IGP leaf for the BGP next-hop is deleted and all the depending BGP Path-Lists are updated to either use the remaining ECMP BGP best-paths or if none remains available to activate precomputed backups.

This BGP resiliency property only depends on IGP convergence and is independent of the number of BGP prefixes impacted.

6.1.4. A local eBGP next-hop fails

Upon local link failure detection, the adjacency to the BGP next-hop is deleted and all the depending BGP pathlists are updated to either use the remaining ECMP BGP best-paths or if none remains available to activate precomputed backups.

This BGP resiliency property only depends on local link failure detection and is independent of the number of BGP prefixes impacted.

6.2. Performance

When the failure is local (a local IGP next-hop failure or a local eBGP next-hop failure), a pre-computed and pre-installed backup is activated by a local-protection mechanism that does not depend on the number of BGP destinations impacted by the failure. Sub-50msec is thus possible even if millions of BGP routes are impacted.

When the failure is remote (a remote IGP failure not impacting the BGP next-hop or a remote BGP next-hop failure), an alternate path is activated upon IGP convergence. All the impacted BGP destinations benefit from a working alternate path as soon as the IGP convergence occurs for their impacted BGP next-hop even if millions of BGP routes are impacted.

6.2.1. Perspective

The following table puts the BGP PIC benefits in perspective assuming

- o 1M impacted BGP prefixes
- o IGP convergence ~ 500 msec
- o local protection ~ 50msec
- o FIB Update per BGP destination ~ 100usec conservative,

~ 10usec optimistic

- o BGP Convergence per BGP destination ~ 200usec conservative,

~ 100usec optimistic

	Without PIC	With PIC
Local IGP Failure	10 to 100sec	50msec
Local BGP Failure	100 to 200sec	50msec
Remote IGP Failure	10 to 100sec	500msec
Local BGP Failure	100 to 200sec	500msec

Upon local IGP next-hop failure or remote IGP next-hop failure, the existing primary BGP next-hop is intact and usable hence the resiliency only depends on the ability of the FIB mechanism to reflect the new path to the BGP next-hop to the depending BGP destinations. Without BGP PIC, a conservative back-of-the-envelope estimation for this FIB update is 100usec per BGP destination. An optimistic estimation is 10usec per entry.

Upon local BGP next-hop failure or remote BGP next-hop failure, without the BGP PIC mechanism, a new BGP Best-Path needs to be recomputed and new updates need to be sent to peers. This depends on BGP processing time that will be shared between best-path computation, RIB update and peer update. A conservative back-of-the-envelope estimation for this is 200usec per BGP destination. An optimistic estimation is 100usec per entry.

6.3. Automated

The BGP PIC solution does not require any operator involvement. The process is entirely automated as part of the FIB implementation.

The salient points enabling this automation are:

- o Extension of the BGP Best Path to compute more than one primary ([11] and [12]) or backup BGP next-hop ([6] and [13]).
- o Sharing of BGP Path-list across BGP destinations with same primary and backup BGP next-hop
- o Hierarchical indirection and dependency between BGP pathlist and IGP pathlist

6.4. Incremental Deployment

As soon as one router supports BGP PIC solution, it benefits from all its benefits without any requirement for other routers to support BGP PIC.

[7. Dependency](#)

This section describes the required functionality in the forwarding and control planes to support BGP-PIC described in this document

7.1. Hierarchical Hardware FIB

BGP PIC requires a hierarchical hardware FIB support: for each BGP forwarded packet, a BGP leaf is looked up, then a BGP Pathlist is consulted, then an IGP Pathlist, then an Adjacency.

An alternative method consists in "flattening" the dependencies when programming the BGP destinations into HW FIB resulting in potentially eliminating both the BGP Path-List and IGP Path-List consultation. Such an approach decreases the number of memory lookup's per forwarding operation at the expense of HW FIB memory increase (flattening means less sharing hence duplication), loss of ECMP properties (flattening means less pathlist entropy) and loss of BGP PIC properties.

7.2. Availability of more than one primary or secondary BGP next-hops

When the primary BGP next-hop fails, BGP PIC depends on the availability of a pre-computed and pre-installed secondary BGP next-hop in the BGP Pathlist.

The existence of a secondary next-hop is clear for the following reason: a service caring for network availability will require two disjoint network connections hence two BGP next-hops.

The BGP distribution of the secondary next-hop is available thanks to the following BGP mechanisms: Add-Path [[11](#)], BGP Best-External [[6](#)], diverse path [[12](#)], and the frequent use in VPN deployments of different VPN RD's per PE. It is noteworthy to mention that the availability of another BGP path does not mean that all failure scenarios can be covered by simply forwarding traffic to the available secondary path. The discussion of how to cover various failure scenarios is beyond the scope of this document

7.3. Pre-Computation of a secondary BGP next-hop

[13] describes how a secondary BGP next-hop can be precomputed on a per BGP destination basis.

8. Security Considerations

The behavior described in this document is internal functionality to a router that result in significant improvement to convergence time as well as reduction in CPU and memory used by FIB while not showing change in basic routing and forwarding functionality. As such no additional security risk is introduced by using the mechanisms proposed in this document.

9. IANA Considerations

No requirements for IANA

10. Conclusions

This document proposes a hierarchical and shared forwarding chain structure that allows achieving BGP prefix independent convergence, and in the case of locally detected failures, sub-50 msec convergence. A router can construct the forwarding chains in a completely transparent manner with zero operator intervention thereby supporting smooth and incremental deployment.

11. References

11.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006
- [3] Bates, T., Chandra, R., Katz, D., and Rekhter Y., "Multiprotocol Extensions for BGP", [RFC 4760](#), January 2007
- [4] Y. Rekhter and E. Rosen, " Carrying Label Information in BGP-4", [RFC 3107](#), May 2001
- [5] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007

11.2. Informative References

- [6] Marques, P., Fernando, R., Chen, E., Mohapatra, P., Gredler, H., "Advertisement of the best external route in BGP", [draft-ietf-idr-best-external-05.txt](#), January 2012.
- [7] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Softwire Mesh Framework", [RFC 5565](#), June 2009.
- [8] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [9] De Clercq, J. , Ooms, D., Prevost, S., Le Faucheur, F., "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", [RFC 4798](#), February 2007
- [10] O. Bonaventure, C. Filsfils, and P. Francois. "Achieving sub-50 milliseconds recovery upon bgp peering link failures, " IEEE/ACM Transactions on Networking, 15(5):1123-1135, 2007
- [11] D. Walton, A. Retana, E. Chen, J. Scudder, "Advertisement of Multiple Paths in BGP", [draft-ietf-idr-add-paths-12.txt](#), November 2015
- [12] R. Raszuk, R. Fernando, K. Patel, D. McPherson, K. Kumaki, "Distribution of diverse BGP paths", [RFC 6774](#), November 2012
- [13] P. Mohapatra, R. Fernando, C. Filsfils, and R. Raszuk, "Fast Connectivity Restoration Using BGP Add-path", [draft-pmohapat-idr-fast-conn-restore-03](#), Jan 2013
- [14] C. Filsfils, S. Previdi, A. Bashandy, B. Decraene, S. Litkowski, M. Horneffer, R. Shakir, J. Tansura, E. Crabbe "Segment Routing with MPLS data plane", [draft-ietf-spring-segment-routing-mpls-02](#) (work in progress), October 2015
- [15] C. Filsfils, S. Previdi, A. Bashandy, B. Decraene, " Topology Independent Fast Reroute using Segment Routing", [draft-francois-spring-segment-routing-ti-lfa-02](#) (work in progress), August 2015
- [16] M. Shand and S. Bryant, "IP Fast Reroute Framework", [RFC 5714](#), January 2010
- [17] S. Bryant, C. Filsfils, S. Previdi, M. Shand, N So, " Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", [RFC 7490](#) April 2015

- [18] A. Atlas, C. Bowers, G. Enyedi, " An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees", [draft-ietf-rtgwg-mrt-frr-architecture-10](#) (work in progress), February 2016

12. Acknowledgments

Special thanks to Neeraj Malhotra, Yuri Tsier for the valuable help

Special thanks to Bruno Decraene for the valuable comments

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Ahmed Bashandy
Cisco Systems
170 West Tasman Dr, San Jose, CA 95134, USA
Email: bashandy@cisco.com

Clarence Filsfils
Cisco Systems
Brussels, Belgium
Email: cfilsfil@cisco.com

Prodosh Mohapatra
Sproute Networks
Email: mpradosh@yahoo.com

