

# UNDERSTANDING AND DEPLOYING LOOP-FREE ALTERNATE FEATURE

Theoretical Overview and Operational Examples  
of Loop-Free Alternates in the Junos OS

Although Juniper Networks has attempted to provide accurate information in this guide, Juniper Networks does not warrant or guarantee the accuracy of the information provided herein. Third party product descriptions and related technical details provided in this document are for information purposes only and such products are not supported by Juniper Networks. All information provided in this guide is provided "as is", with all faults, and without warranty of any kind, either expressed or implied or statutory. Juniper Networks and its suppliers hereby disclaim all warranties related to this guide and the information contained herein, whether expressed or implied of statutory including, without limitation, those of merchantability, fitness for a particular purpose and noninfringement, or arising from a course of dealing, usage, or trade practice.

## Table of Contents

Introduction .....	3
Scope .....	3
Design Considerations .....	3
Defining the Problem .....	4
The Solution: Loop-free Alternates .....	5
Operational Theory .....	5
Loop-free Alternates According to RFC5286 .....	5
Scaling Consideration .....	6
Fate Sharing .....	7
Loop-Free Alternate Highlights .....	7
Extending Coverage .....	8
RSVP-TE Tunnels .....	9
LDP Tunneling .....	10
Implementation .....	10
Physical and Logical Topology .....	10
Base Configuration Before Enabling Loop-Free Alternate .....	11
Loop-Free Alternate Configuration .....	11
Enabling the IGP for Loop-Free Alternate .....	11
Enabling the PFE to Store Multiple Next Hops for a Given Prefix .....	12
Verification .....	12
Verification Summary .....	12
The Expected Result for All Remaining Egress Routers .....	13
Backup Coverage .....	13
Checking Loop-Free Alternate in Details .....	13
Checking the Coverage for Cyprus as Seen by the Loop-Free Alternate-Enabled Node Germany .....	14
Identifying the Transport Label Used by Spain to Reach the Egress FEC Cyprus .....	16
Checking the Label Operation of the LSR Germany When There is No Link Failure .....	17
Checking the Label Operation of the LSR Germany in Case of a Link Failure .....	18
Summary .....	19
About Juniper Networks .....	20

## List of Figures

Figure 1: Local convergence breakdown .....	4
Figure 2: Sample topology .....	6
Figure 3: Sample network .....	8
Figure 4: Lab simulation and implementation topology .....	10

## Introduction

As IP networks are increasingly used for transporting mission-critical and sensitive applications such as VoIP and pseudowires, high availability and fast recovery times are becoming very critical considerations in designing networks. In networks configured with RSVP-TE, MPLS fast reroute offers a solution that delivers sub-50 ms recovery times on a par with traditional SONET architectures by providing pre-calculated, pre-signaled, and preinstalled backup paths. Until recently, fast reroute (and therefore RSVP-TE) was the only option for network operators seeking sub-50 ms recovery.

Increasingly, network operators have been seeking a method to provide similarly fast recovery in LDP-based networks. In traditional implementations, this is difficult because restoration in an LDP network is based on interior gateway protocol (IGP) convergence time, which, in turn, is a function of the total number of links and routing nodes in a network. Since convergence time increases whenever the network grows, recovery time could be in the hundreds of milliseconds in a large-scale network. While this recovery time is acceptable for some types of traffic, it is not acceptable for real-time applications like video and VoIP.

Juniper addresses this challenge with the Loop-free Alternates feature in Juniper Networks® Junos® operating system. By adding a preinstalled backup next hop into the forwarding plane (similar to frame reroute), Loop-Free Alternate provides consistent sub-50 ms IGP convergence times, independent of the size of the network and without adding operational burden. Loop-Free Alternate can be implemented on a per-router basis and does not introduce any new protocols to the network, ensuring that it is also suitable for multivendor environments.

Loop-Free Alternate is an attractive feature for service providers and enterprises that require rapid and reliable failover protection in OSPF and IS-IS based networks.

## Scope

This document provides an overview of Juniper's Loop-free Alternates feature, a solution that delivers fast restoration and convergence for OSPF and IS-IS based networks. This paper provides a theoretical overview of how Loop-Free Alternate works and then presents several operational examples that illustrate implementation scenarios.

## Design Considerations

Today, many sensitive, real-time applications such as voice and video require IP networks capable of recovering from temporary fault conditions so quickly that the fault and associated recovery are imperceptible to the end user, and occur under the industry-standard 50 milliseconds. With capabilities such as MPLS frame reroute, traffic engineered IP networks are capable of delivering this level of restoration speed. However, in IGP-based networks, recovery traditionally requires network-wide convergence to occur first—and this may stretch recovery times past the 50 ms standard.

To better understand the factors that contribute to restoration times, it is helpful to consider the two steps that must occur:

- Detecting the failure
- Reacting to the failure

Depending on the type of failure (node vs. link), time to detect the failure can vary but is usually very quick—for example, 3-5 ms when loss-of-light is detected on the phy-level. Reacting to failures can take much longer because it is traditionally a control plane-intensive process that requires routes to be recalculated and distributed across the network. Until now, the only possible exception, as noted above, was if a standby path were calculated and installed in the forwarding table in advance using RSVP-TE and frame reroute. This would allow the router adjacent to the failure to switch to a predetermined path without involving the control plane, dramatically improving recovery times.

With the introduction of the Loop-free Alternates feature in Junos OS, Juniper is delivering automated frame reroute-like capabilities and restoration times in networks that use LDP as opposed to RSVP-TE. Loop-Free Alternate is a software feature available in Junos OS 9.5, and is supported across all Juniper Networks M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers.

### Defining the Problem

When a link or node failure occurs, routers adjacent to the error detect the failure, which in turn triggers a predefined set of actions (Figure 1). These actions can be triggered via loss-of-light or an Operation, Administration, and Maintenance (OAM) mechanism such as Bidirectional Forwarding Detection (BFD), link-fault-management (802.3ah), or connectivity-fault-management (802.1ag).

Failure detection first occurs in the Packet Forwarding Engine (PFE). The PFE signals the information up to the routing engine operating system. On failure detection, the event is pushed to the routing protocol process (routing daemon). At this stage, the only thing known is that there is a failed link. There is no information available as to whether the neighboring node is down, or if it is just the link facing toward the neighboring link that is down.

On notification of the failure, the routing protocol process performs two tasks:

- Calculates a shortest-path towards the destination to choose a next hop that circumvents the failed link. Most commonly a link-state protocol such as OSPF or IS-IS is used as the IGP, and both of these protocols use the Dijkstra algorithm to calculate a loop-free topology. This is known as the shortest-path-first (SPF) calculation.
- In addition to the SPF calculation, the neighbors need to be informed of the link failure so that they can update their routing tables with the new loop-free topology as well.

The SPF calculation is a very slow process. However, with the latest Routing Engines, a full SPF calculation for large networks can often be completed in the range of a few milliseconds (1-5 ms). In reality, performing the SPF calculation is a very small part of the overall convergence time. Rather, it's the whole of the process and communication between the various elements that consumes the most time.

After the routing protocol process calculates all new next hops for all given prefixes, this new next hop information is pushed down to the PFE. Once the PFE adjacent to the failure is updated, traffic is rerouted and the appropriate data path in a given node is maintained. However, other nodes in the network may not have received the updates regarding the link failure, or they may still be processing the updates. Only when all nodes in the network have received the IGP updates regarding the failure and updated their routing tables can global convergence be considered complete.

Figure 1 summarizes the process, starting with the event propagation in the PFE, and the steps that must be followed as information travels upwards to the Routing Engine OS and from there towards the routing protocol process (routing daemon). Junos OS uses an event-driven mechanism, which reacts immediately on a received event to deliver fast convergence. However, due to the nature of the convergence process, total convergence time may be hundreds of milliseconds.

As the graphic below shows, though the individual steps are relatively quick, global convergence can take time.

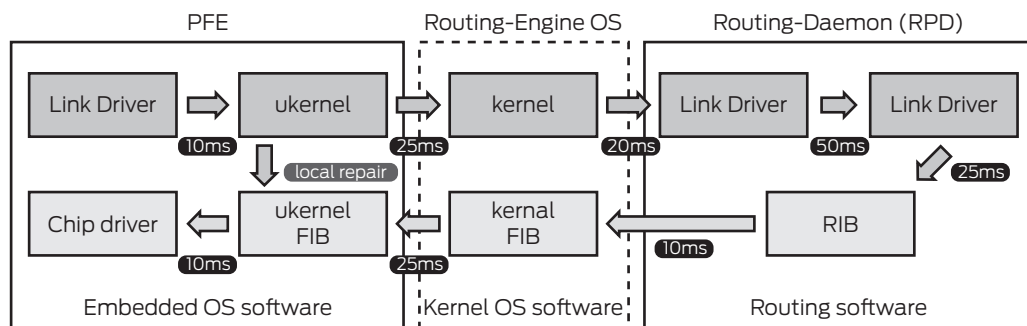


Figure 1: Local convergence breakdown

## The Solution: Loop-free Alternates

The goal of Loop-Free Alternate is to deliver loop-free rerouting capabilities without involving the Routing Engine's OS or the routing daemon for route calculation. This is also known as performing a local repair, as opposed to the typical procedure in IGP networks which, as described above, requires a global repair. In a local repair, the routing protocol software pre-installs an alternate next hop which gets locally activated in case the PFE detects a link break. The loop-free local repair path is active until the global repair route changes kick in. Both local repair and global repair are complementary technologies. The main advantage of the local repair path is that it is blazing fast—it can send traffic to potential backups as fast as 20 milliseconds after a network failure. Moreover, it can do that virtually irrespective of the affected number of prefixes.

Loop-Free Alternate implements a loop-free local repair by changing the event flow described above. Rerouting begins much earlier, as soon as the failure is detected. By pre-installing an alternate backup next hop on the forwarding table, Loop-Free Alternate eliminates the need to wait for the routing daemon to be notified and perform an SPF calculation, as well as the need to wait for the convergence of other routers in the network (global convergence). The delay introduced by sending messages from PFE to the Routing Engine OS and even to the routing daemon is avoided by Loop-Free Alternate. This eliminates the time to update the forwarding table, hence greatly reducing recovery and convergence time.

Loop-Free Alternate enables the node adjacent to the failure for a quick bypass of the failed link/node. This in turn allows other non-converged nodes to still use the Loop-Free Alternate-enabled node as a transit path without any traffic being “blackholed.”

## Operational Theory

This section provides theoretical background for the Loop-Free Alternate feature in Junos OS.

### Loop-free Alternates According to RFC5286

Common IGPs (like OSPF and IS-IS) only calculate the best path, or alternatively a set of equal-cost paths between a given source/destination router pair. Most implementations also support equal-cost multipath (ECMP) routes. There is no reason why an ECMP route could not be used as a mechanism for connectivity restoration when one of the paths becomes unavailable. Expanding on that model, less than equal-cost routes (LECMP) can even be used for the purpose of providing connectivity restoration. As long as the LECMP route does not cause a forwarding loop for a given destination, it can be used.

To avoid forwarding loops, a router needs to execute additional calculations to verify that an LECMP route does not generate a forwarding loop, and may therefore be elected as a viable backup route. An LECMP route that does not cause forwarding loops is called a Loop-free Alternate. An Loop-Free Alternate-enabled node performs a calculation in advance to determine a backup path that takes into consideration all adjacent neighbors and checks if they (the neighbors) can provide loop-free forwarding in case the primary next hop fails.

IETF's RFC5286 (Basic Specification for IP Fast Reroute: Loop-Free Alternates) offers a method for calculating Loop-Free Alternates based on the calculation of route inequalities. Juniper's Loop-Free Alternate implementation for OSPF/IS-IS networks does not rely on these inequalities; but rather introduces the tracklist concept. Using the tracklist (described in the next section) provides a better method of calculating loop-free routers using less computational overhead. This results in much better scaling in regard to IGP adjacencies for the Loop-Free Alternate-enabled router.


Here is a short example for link protection applying inequality #1:

$$\text{Distance\_opt}(N, D) < \text{Distance\_opt}(N, S) + \text{Distance\_opt}(S, D)$$

Distance_opt	is used to indicate the shortest distance from X and Y.
S	is used to indicate the calculating router.
N	is a neighbor of S.
D	is the destination under consideration.

The metric as calculated by SPF from the neighbor N towards the destination D must be smaller than the summary metric of (source S to neighbor N+ neighbor N to destination D).

Figure 2 provides a sample topology to help illustrate this calculation.



**Figure 2: Sample topology**

The primary path (as calculated via SPF) from R1 to destination R3 is R1->R2->R3 which results in a metric of  $\text{Distance\_opt}(S, D)=2$ .

The Loop-Free Alternate enabled routing node R1 checks its available neighbors, which are R2 and R4. It is obvious that R2 can't be used as backup next hop here, since it is the primary path and protection is needed in case of failure in the R1->R2 link.

However, R4 is a candidate to be used as a backup next hop. Applying the inequality calculation, we can see that

inequality #1 is applicable:

$$\begin{aligned}
 &\text{Distance\_opt}(N, D) < ( \text{Distance\_opt}(N, S) + \text{Distance\_opt}(S, D) ) \\
 &= \text{Distance\_opt}(R4, R3) < ( \text{Distance\_opt}(R4, R1) + \text{Distance\_opt}(R1, R3) ) \\
 &= 2 < (2+1) \\
 &= \text{true}
 \end{aligned}$$

Therefore, router R4 can truly be called a loop-free alternate next hop for source node R1 for the destination node R3. Router R4 can even use equal-cost load balancing or ECMP over router R2 or R5. In both cases, packets with a destination of router R3 can be loop-free forwarded. R2 is valid, because router R1 is only configured for link protection.

All routers enabled for link state protocols share a common link state database. Thus, they have the identical knowledge of the network topology. R1 is able to calculate the shortest-path tree for itself as well as R4, considering R4 is on the top of the shortest-path tree. As a result of having R1 running shortest-path tree with R4 as root of the tree, R1 is able to determine the distance from neighbor R4 to the destination D ( $\text{Distance\_opt}(N, D)$ ).

Having this knowledge, node R1 can solve the inequality. As the "true" result for inequality #1, R4 becomes the loop-free alternate for destination node R3. In case the link between R1 and R2 goes down, R1 can immediately reroute to R4, even if R4 is not even aware of the link failure between R1 and R2.

The shortest-path tree as seen by node R4 is resulting in an ECMP condition, either R4->R5->R3 or R4->R2->R3. In both paths, the rerouted packets are not traversing the failed link R1->R2.

## Scaling Consideration

### Running Multiple OSPF/IS-IS Calculations

As mentioned above, the Loop-Free Alternate-enabled routing node has to run shortest-path tree calculations for each of its neighbors, which can be a significant processing burden if Loop-Free Alternate is enabled on a larger node with many adjacent nodes. This is one of the powerful strengths of Juniper's Loop-Free Alternate implementation and the Junos OS control plane. By using a separate process to maintain adjacency processing protocol-hellos, Junos OS delivers fast convergence and provides robustness in shaky environments without the tuning of parameters such as shortest-path tree delay. Additionally, Juniper's implementation of Loop-Free Alternate relies on innovative new techniques like the tracklist to further reduce the processing requirements of shortest-path tree calculations.

Even if the routing-daemon consumes 100% of the CPU, a separate daemon can preempt the routing protocol process to send out the protocol-hellos. After having the hello send out and maintain the adjacency, this separate daemon is giving back the CPU cycles to the system which enables the routing protocol process to continue its work. As a summary of this scaling-review, Loop-Free Alternate does run nicely on Junos OS in large carrier-grade networks.

## The Tracklist

One of the limitations of the traditional implementations of Loop-Free Alternate, as outlined in RFC5286, is scaling. This limitation occurs because in a large network, the number of CPU cycles needed to calculate the backup next hop can be significant. In the Loop-Free Alternates section above, a simple example is used to show how inequalities are calculated, but in a real-world network, the calculation is much more complex. While the inequalities deliver a loop-free alternate, they require extensive control plane resources to calculate, especially in large-scale implementations with multiple hops and a large mesh.

Juniper's Loop-Free Alternate implementation for OSPF and IS-IS uses a much less CPU-intensive method, the tracklist, instead of relying on inequalities. The tracklist acts like a flight recorder of sorts and is capable of "remembering" each node towards the destination D that a packet would traverse when using neighbor N as primary next hop. Said simply, the tracklist ensures that the backup path puts the packet nearer to its destination than the rerouting node. This shorter distance, and the fact that Loop-Free Alternate avoids the failed node/link, results in loop-free forwarding in case of single failures. When calculating the Loop-Free Alternate route, the Juniper router can refer back to the tracklist, so that if it detects duplicate IDs, either system IDs or its own, (here router R1 in topology 1) on a possible backup path to destination D, that path is ruled out immediately since it would not be loop-free.

The result is that the tracklist provides the same loop-free backup next hop as would RFC5286, however, it allows Loop-Free Alternate to operate efficiently even in nodes with many adjacencies delivering dramatically reduced CPU resources. In short, Junos operating system's unique Loop-Free Alternate implementation allows service providers and enterprises to scale Loop-Free Alternate in even the largest network implementations.

## Fate Sharing

The UI provides extensive insight into the Loop-Free Alternate operation. The UI indicates when and why alternative next hops can't be used as Loop-Free Alternate. Fate sharing is one of these instances, as it makes the link or node unusable as Loop-Free Alternate.

Fate sharing is a condition where a link or a node is being "shared" for forwarding by the primary next hop during normal operations, and by the backup next hop in case of failure. Loop-Free Alternate protects either a link failure or a node failure. If, for example, link A will be protected by Loop-Free Alternate, then it makes no sense to use link A in a backup scenario. For obvious reasons, fate sharing is something that has to be avoided during the election process for an Loop-Free Alternate next hop. Fate sharing gets detected and avoided by Juniper's Loop-Free Alternate implementation as the default.

### Link Fate Sharing

In the topology example used above (Figure 2), such a fate-sharing link can be seen with Loop-Free Alternate link protection configured on router R1. When router R1 checks if router R2 serves as a loop-free next hop for destination R3, link fate sharing for the link R1->R2 applies. The link R1->R2 is used to reach the primary next hop R2. When using R2 as backup next hop, the shortest path to R2 is link R1->R2, which means a fate-sharing link. This means that only R4 can be used as a loop-free alternate next hop for destination R3.

### Node Fate Sharing

Node fate sharing applies the same basic principle at the node level. For example, in topology with node protection configured on router R1, such node fate sharing is seen as well for router R2. Router R2 is the primary next hop for destination router R3. When trying to make the router R2 the Loop-Free Alternate for destination R3, node fate sharing applies to router R2.

## Loop-Free Alternate Highlights

In summary, Loop-Free Alternate addresses the issue of restoration and fast convergence in OSPF and IS-IS networks by enabling the PFE to implement a local repair immediately following a failure. This provides sub-50 ms restoration on an automated basis, with minimal operational impact. Below we summarize some of the key benefits and characteristics of Loop-Free Alternate:


- Loop-Free Alternate does not require the addition of any new protocols. Loop-Free Alternate is entirely self-contained and does not rely on any "helper node" to work properly. Rollout can be done node by node.
- Loop-Free Alternate installs a backup next hop in the forwarding plane in advance. The backup next hop is elected by running multiple SPF calculations, with different neighbors as root of the tree. Juniper's unique tracklist functionality ensures high scaling when evaluating all eligible next hops

- Upon link failure, the backup next hop can be immediately selected, as this next hop is guaranteed to provide loop-free forwarding for a given destination.
- Loop-Free Alternate can start rerouting without awaiting results from the control plane route calculations. As a result of this, sub-50 ms failover time is achievable in an Loop-Free Alternate-enabled network. Loop-Free Alternate is implemented to operate with MPLS VPNs using LDP as a label distribution protocol. Loop-Free Alternate provides improved convergence for MPLS VPNs as well.
- The scalable and robust Junos OS routing protocol process implementation allows routers to run multiple SPF calculations without negatively impacting network stability.
- Loop-Free Alternate is very easy to turn on, and does not need to be configured on all routers in a network. It can be implemented in selected nodes only. As the Loop-Free Alternate Configuring section shows, the keyword "link-protection" under [edit protocols isis interface] is enough to enable Loop-Free Alternate for the given router.
- The Junos OS implementation of Loop-Free Alternate is highly scalable and is designed to operate in large, carrier-grade networks.

### Extending Coverage

Based on network topology, it may not be possible for Loop-Free Alternate to automatically provide 100% coverage for all routes. In some cases, network topology may make it impossible to select a loop-free alternate for each destination, even if multiple paths exist. A corner case example is used below to illustrate how this situation could arise, and examine workarounds using the rich feature set of Junos OS.

In the example below (Figure 3), router R1 tries to elect a loop-free alternate for destination 10/8 in case its primary next hop R2 fails. R1 has another available path over router R4 towards destination R3.



**Figure 3: Sample network**

Assuming the metrics for all routes are equal to one (1) except for R4->R5, which has a metric of 10, node R1 is configured for Loop-Free Alternate link node protection. In this situation, it is impossible to provide a loop-free alternate for Loop-Free Alternate-enabled node R1 with destination 10/8 if the primary link R1->R2 fails.

The reason is explained below.

If Loop-Free Alternate selects R4 as an alternate for destination 10/8, R1 will initiate a local repair and start rerouting packets before even R4 receives any update about the failed link R1>R2, (the global repair). However, because of the high metric between R4 and R5, R4 sees the shortest path to R3 as going back through R1, so R4 will send all packets with destination 10/8 back to R1. This forms a routing loop, which will perpetuate until a global repair is complete and node R4 converges. Once global convergence is completed, R4 is aware that the link R1->R2 is down and the shortest path to reach R3 is now over R5 instead of rerouting back to R1.

Using Junos OS, there are a number of ways to solve the problem described above and allow node R1 to find a loop-free alternate next hop for prefix 10/8:

- Add additional links: By placing additional link(s) between R1 and R5 (e.g., with a metric of 5), coverage for node R3 is achieved as well. The primary path still stays over node R2 and the loop-free alternate next hop will be R5. In case the primary link R1-R2 goes down, R1 can safely reroute traffic for destination prefix 10/8 to node R5 without looping traffic.
- Provision tunnels: Instead of using physical cabling, operators can provision a virtual tunnel to make node R5 appear as a "direct neighbor" of R1. The most likely option in this case would be to provision a new label-switched path (LSP) with RSVP-TE. While adding physical links can add costs, provisioning tunnels is essentially a no-cost solution to provide full Loop-Free Alternate coverage.

When a tunnel is provisioned, the Loop-Free Alternate-enabled router PFE sees the tunnel endpoint as a valid next hop. The only difference here is that the endpoint of the tunnel does not necessarily need to be a physically adjacent router. The tunnel endpoint becomes an Loop-Free Alternate candidate and is used during the Loop-Free Alternate election process.

In Junos OS, a RSVP-TE tunnel has to be configured with the “backup” keyword under [edit protocols mpls label-switched-path name] to use the tunnel endpoint in the Loop-Free Alternate election process. A configured tunnel endpoint matching the router’s system-id is considered a candidate next hop for Loop-Free Alternate similar to the direct IGP adjacency case.

### RSVP-TE Tunnels

As seen in the previous section, a tunnel endpoint is considered a candidate next hop for Loop-Free Alternate when being configured, so source routing must be a property of the tunnel. This means that the Loop-Free Alternate-enabled router (“the ingress”) decides which path to take to a given tunnel endpoint. And the tunnel must be able to ignore the path an IGP might choose to reach the tunnel endpoint.

The idea is to overrule node R4, which would reroute back to router R1 based on the IGP metric. Instead of using IGP, one can use a source routed RSVP-TE tunnel towards node R5 through R4, bypassing R4’s routing decision. RSVP is a powerful tool when it comes to traffic engineering and makes it easy to provide a tunnel that does not follow the shortest path as calculated by the IGP.

The RSVP-TE tunnel starts at router R1, traverses router R4, and ends at router R5 (R1-over-R4-to-R5). Router R5 is the next hop after R4. There is no need to construct a tunnel towards the egress router R3 as is done in RSVP-TE fast reroute. For Loop-Free Alternate to work, it is enough if this one tunnel is targeted to the next hop R5 to allow loop-free forwarding...

If R1 needs to rely on the IGP, the tunnel path towards node R5 would look like R1->R2->R3->R5. However, to protect against a failure of link R1->R2, the link R1->R2 must be avoided in the path calculation (see fate sharing section). Using RSVP’s source route capability, the tunnel can be forced to traverse node R4 instead of following the shortest path.

For IP and MPLS VPN traffic, the RSVP-TE tunnel provides coverage towards node R3 for ingress R1 in case of an R1->R2 link failure in the following way:

1. An RSVP-TE tunnel R1-over-R4-to-R5 is preinstalled and ready to use. The “backup” keyword allows the Loop-Free Alternate election process to use the tunnel endpoint as candidate next hop. R1 considers the router R5 as “neighbor” and checks if R5 provides coverage for destination 10/8.
2. The link between R1->R2 fails.
3. Upon detecting this, the PFE of router R1 places all traffic local to R3 into the RSVP-TE signaled tunnel-lsp. The tunnel-lsp follows the path R1-over-R4-to-R5 and thus successfully bypasses the broken link R1->R2. An `rsvp-transport-label` is pushed onto the packet towards node R4. In case of IPv4, this results in a MPLS-tagged frame with a single transport label. In case of an MPLS VPN, this results in a 3 label stack.
4. For IP traffic, R4 receives an MPLS-tagged frame. As R4 is the penultimate router towards the tunnel endpoint, this mpls-tag is being removed and a native IP frame is forwarded to R5. In case the frame belongs to an MPLS VPN, R4 receives a 3 label MPLS stack with the top RSVP transport label and determines the egress interface towards R5. R4 performs penultimate-hop-popping and removes the top RSVP-TE label before sending the packet (with a 2 MPLS label stack) to R5.
5. In case the rerouted traffic is IP, R5 just does a destination-ip lookup and forwards the packet. In case the rerouted frame belongs to an MPLS VPN, R5 is itself a penultimate router towards router R3 and pops the LDP transport label, resulting in a single MPLS tagged frame just containing the vpn-label.

Finally the tunnel endpoint R5 is able to forward IP packets and MPLS VPN frames in a loop-free manner to R3.

As a result, the R1->R2 link failure recovers in just sub 50 ms without the need for full mesh RSVP-TE tunnels.

## LDP Tunneling

Bullet 3 and especially 5 from the previous partner paragraph need some further explanation about the LDP transport label which is being swapped by the rerouting node R1. R1 needs to be aware which transport label R5 is assigning for forward error correction (FEC) R3 (see step 5). Only neighbors to R5 have this knowledge.

A targeted LDP session from router R1 to R5 provides this knowledge to R1. With a targeted LDP session, R1 becomes LDP “neighbor” with R5. Being neighbors, R5 provides information to R1 about the transport label to reach R3 FECs.

In Junos OS, the keyword “ldp-tunneling” is setting up the targeted LDP session.

```

protocols {
  mpls {
    label-switched-path TUNNEL-R1-over-R4-to-R5 {
      backup;
      to 192.168.53.101;
      ldp-tunneling;
      primary OVER-R4;
    }
    path OVER-R4 {
      10.102.103.1 loose;
    }
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
}

```

```

cgraf@Germany# run show isis backup label-switched-path
Backup MPLS LSPs:
TUNNEL-R1-over-R4-to-R5, Egress: 192.168.53.101, Status: up, Last change:
00:00:01
  TE-metric: 19, Metric: 0, Refcount: 2

```

## Implementation

### Physical and Logical Topology

This chapter addresses the configuration tasks and the commands which operators can use to manage Loop-Free Alternate. We provide a sample MPLS Layer 2 VPN topology, and provide a step-by-step guide to enabling Loop-Free Alternate link protection.

In Figure 4, we assume that the transit node in Germany is configured for Loop-Free Alternate. The actual implementation of Loop-Free Alternate in Junos OS release 9.5 and above supports IP, L2VPN, virtual private LAN service (VPLS), and L3VPN on either ingress, egress, or transit node. There is also a focus on LDP interoperability, as the rerouting node needs to handle the two label stack of MPLS VPNs properly.

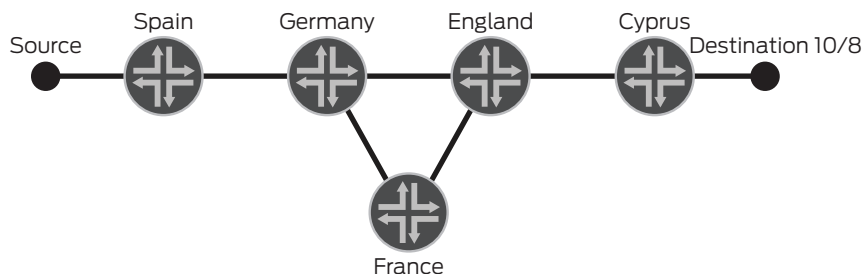


Figure 4: Lab simulation and implementation topology

## Base Configuration Before Enabling Loop-Free Alternate

- Ingress provider edge (PE) router is Spain, egress PE router is Cyprus. Germany, France, and England act as P-nodes or label-switching routers (LSRs).
- Loopback addresses:
  - 192.168.53.110 Cyprus
  - 192.168.53.103 Germany
  - 192.168.53.106 Spain
- An L2VPN is configured between Spain and Cyprus. In this example, we are only interested in the route from Spain to Cyprus.
- Germany is the only Loop-Free Alternate-enabled router in this example. This highlights another important aspect of Loop-Free Alternate. Namely that Loop-Free Alternate is a per-hop-behavior, so Germany's neighbors do not need to support Loop-Free Alternate at all to benefit from the fast reroute capabilities.
- For MPLS VPN implementations, it is very important that Loop-Free Alternate is integrated into LDP, as node Spain sends a double tagged MPLS frame to egress PE router Cyprus, traversing the Loop-Free Alternate enabled node Germany. Upon rerouting, Germany needs to push/swap labels accordingly.
- In this case, the targeted LDP session is not required, as Germany is already a neighbor to the Loop-Free Alternate France.
- The BGP/L2VPN configuration is not listed here. Only the Loop-Free Alternate-related configuration is listed.

## Loop-Free Alternate Configuration

Loop-Free Alternate configuration needs a minimum of two configuration statements:

- Enabling the IGP for Loop-Free Alternate
- Enabling the PFE to store multiple next hops for a given prefix

### Enabling the IGP for Loop-Free Alternate

Enabling IS-IS for Loop-Free Alternate just needs a single new configuration statement. The options are no-eligible-backup, node-link-protection, or link-protection:

```

protocols {
  isis {
    interface all {
      link-protection; # This command configure Loop-Free Alternate for IGP
at given interface
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```

## Enabling the PFE to Store Multiple Next Hops for a Given Prefix

Within Junos OS, the handling of prefixes is always done via policies. To allow the PFE to store more than a single next hop for a given prefix, a load-balancing policy must be applied.

```

routing-options {
  forwarding-table {
    export PLB;
  }
}

policy-options {
  policy-statement PLB {
    then {
      load-balance per-packet;
      accept;
    }
  }
}

```

### Verification

#### Verification Summary

Loop-Free Alternate installs a backup next hop per prefix if the egress router can be protected. An existing “coverage” means that the egress routing node can be successfully protected against the configured failure (e.g., link or node protection). All user interface output and operational commands are based on the fact that egress routers need to be protected. Being listed as “no-coverage” means that the Loop-Free Alternate-enabled router (either ingress or transit) cannot provide an Loop-Free Alternate next hop.

The “show isis backup” command covers all IS-IS related operational commands for Loop-Free Alternate. The command has the “no-coverage” knob to list uncovered nodes. Any node listed with “no-coverage” will not be protected through Loop-Free Alternate. From Germany’s point of view, Spain is not covered. According to the topology, this makes sense, as there is no redundant path from Germany (the root) to node Spain.

#### No Coverage

The “show isis backup spf results no-coverage” command lists all egress routers for which Loop-Free Alternate coverage can’t be provided. As only node Spain is listed here, all other egress nodes like England, Spain, Cyprus, and France are protected via a loop-free alternate.

```

cgraf@Germany> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
Spain.00
Primary next hop: fe-0/3/0.0, Spain, SNPA: 0:17:cb:42:42:81
Root: Spain, Root Metric: 10, Metric: 0
  Not eligible, Reason: Primary next hop link fate sharing
Root: France, Root Metric: 10, Metric: 20
  track-item: Spain.00-00
  track-item: Germany.00-00
  Not eligible, Reason: Path loops
Root: England, Root Metric: 10, Metric: 20
  track-item: Spain.00-00
  track-item: Germany.00-00
  Not eligible, Reason: Path loops
1 nodes

```

The output of this command is reviewed in detail in the section titled, Checking the Coverage for Cyprus as Seen by the Loop-Free Alternate-Enabled Node Germany. This command is executed on the Loop-Free Alternate-enabled router Germany. Loop-Free Alternate checks all neighbors of Germany, so node Spain, France, and England are listed. All are marked “not eligible.” For Spain, there is no protection available.

### The Expected Result for All Remaining Egress Routers

The operational command “show isis backup spf results” provides much useful information. This command displays which egress nodes can be protected via a loop-free alternate next hop.

From Germany's point of view, this command shows the results for each egress router in the network and whether those nodes can be protected against the configured failure condition (e.g., link protection):

1. Protecting egress node England: Instead of sending packets over so-0/2/2 (the primary path to England), R1 can send packets to France without running into any loop for destination England. It is expected to see France as an eligible next hop to protect the link to England.
2. Protecting egress node France: Instead of sending packets over so-0/2/3 (the primary path to France), R1 can send packets to England without running into any loop for destination France. England is expected as an eligible next hop to protect the link to France.
3. Protecting egress node Cyprus: Instead of sending packets over so-0/2/2 (the primary path to Cyprus), R1 can send packets over France without running into any loop for destination Cyprus. It is expected to see France as an eligible next hop to protect the link to Cyprus.
4. As mentioned before, there is no protection for Spain possible, as Spain is not reachable via a redundant path.

### Backup Coverage

Each Loop-Free Alternate-enabled node has an understanding of the number of destinations it is protecting. The “show isis backup coverage” command quickly reveals this information:

```
cgraf@Germany> show isis backup coverage
Backup Coverage:
Topology      Level  Node   IPv4   IPv6   CLNS
IPV4 Unicast  1      0.00%  85.71% 0.00%  0.00%
IPV4 Unicast  2      75.00%  85.71% 0.00%  0.00%
```

### Checking Loop-Free Alternate in Details

This section provides detailed verification of the Loop-Free Alternate implementation on the example network topology shown in Figure 4.

1. Check the coverage for Cyprus as seen by the Loop-Free Alternate-enabled node Germany
2. Check L2VPN on the ingress PE router Spain
3. Identify the transport label used by Spain to reach the egress FEC Cyprus
4. Check if the Loop-Free Alternate node Germany has coverage for Cyprus
5. Check the label operation of the LSR Germany when there is no link failure
6. Check the label operation of the LSR Germany in case of a link failure

As previously mentioned, the quick “show isis backup spf results no-coverage” would have been sufficient in most cases, but the steps listed in this section are intended to provide more insight.

## Checking the Coverage for Cyprus as Seen by the Loop-Free Alternate-Enabled Node Germany

As the egress PE router is Cyprus, below we illustrate a command to check Loop-Free Alternate coverage for Cyprus. Checking for the destination router (not neighbor nodes) is important to remember while verifying Loop-Free Alternate operation on a given ingress node.

```
cgraf@Germany# run show isis backup spf results
IS-IS level 1 SPF results:
  0 nodes

IS-IS level 2 SPF results:
Cyprus.00
  Primary next hop: so-0/2/2.0, England
    Root: England, Root Metric: 10, Metric: 10
    Not eligible, Reason: Primary next hop link fate sharing
    Root: France, Root Metric: 10, Metric: 20
    track-item: England.00-00
    Eligible, Backup next hop: so-0/2/3.0, France
    Root: Spain, Root Metric: 10, Metric: 30
    track-item: England.00-00
    track-item: Germany.00-00
    Not eligible, Reason: Interface is already covered
France.00
  Primary next hop: so-0/2/3.0, France
    Root: England, Root Metric: 10, Metric: 10
    track-item: France.00-00
    Eligible, Backup next hop: so-0/2/2.0, England
    Root: France, Root Metric: 10, Metric: 0
    Not eligible, Reason: Interface is already covered
    Root: Spain, Root Metric: 10, Metric: 20
    track-item: France.00-00
    track-item: Germany.00-00
    Not eligible, Reason: Interface is already covered
England.00
  Primary next hop: so-0/2/2.0, England
    Root: England, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next hop link fate sharing
    Root: France, Root Metric: 10, Metric: 10
    track-item: England.00-00
    Eligible, Backup next hop: so-0/2/3.0, France
    Root: Spain, Root Metric: 10, Metric: 20
    track-item: England.00-00
    track-item: Germany.00-00
    Not eligible, Reason: Interface is already covered
Spain.00
  Primary next hop: fe-0/3/0.0, Spain, SNPA: 0:17:cb:42:42:81
    Root: England, Root Metric: 10, Metric: 20
    track-item: Spain.00-00
    track-item: Germany.00-00
    Not eligible, Reason: Path loops
    Root: Spain, Root Metric: 10, Metric: 0
    Not eligible, Reason: Primary next hop link fate sharing
    Root: France, Root Metric: 10, Metric: 20
    track-item: Spain.00-00
    track-item: Germany.00-00
    Not eligible, Reason: Path loops
  4 nodes

[1] Cyprus.00
  Primary next hop: [2] so-0/2/2.0, England
  [3] Root: England, [4] Root Metric: 10, [5] Metric: 10
  [6] track-item: England.00-00
  [8] Not eligible, Reason: Primary next hop [7] link fate sharing
```

1. Cyprus is the egress node where Loop-Free Alternate checks coverage in case the primary link (over so-0/2/2.0) that is being used towards Cyprus fails.
2. Actual path taken to Cyprus is over so-0/2/2.0 with England as primary next hop.
3. Each neighbor (here Root: England) local to Germany is checked to see if it can serve as a loop-free candidate next hop.
4. Root Metric: 10 (the Root Metric) is the metric from the candidate next-hop England towards the root Germany.
5. Metric 10 is the distance from candidate next-hop England towards the destination node Cyprus.00.
6. Track-item, as discussed earlier, is like a flight recorder. All nodes on the primary path towards the egress router Cyprus are being listed here. The result of "Not eligible" can also be seen, with the primary path traversing England and the candidate next hop being England itself. The candidate next-hop England is "Not eligible," as of link fate sharing
7. Link fate sharing is occurring on neighbor England. The primary path towards Cyprus is over node England and England cannot be made the backup next hop. While traffic flows through a backup next hop, the rerouted packets will try to use the link towards England as well.
8. Node England is not eligible. Being listed as "Not eligible" simply means that this neighbor (England) can't serve as the loop-free alternate for destination Cyprus here

The Loop-Free Alternate election process examines each neighbor to find an eligible next hop for egress router Cyprus. As part of this process, neighbor router France is checked. The ui-output shows, "Eligible, Backup next hop: so-0/2/3.0, France." This clearly identifies that the Loop-Free Alternate process found that router France can serve as a loop-free alternate next hop for egress node Cyprus. If node Spain forwards IP or MPLS VPN packets to Cyprus, and if the primary link Germany->England fails, then Germany activates the loop-free alternate to next hop France, which is capable of forwarding the packets in a loop-free manner to Cyprus.

```

Root: France, Root Metric: 10, Metric: 20
  track-item: England.00-00
    Eligible, Backup next hop: so-0/2/3.0, France  << Being Eligible is means
that France can be                                used as loop-free alternate next
hop here
  Root: Spain, Root Metric: 10, Metric: 30
    track-item: England.00-00
    track-item: Germany.00-00
    Not eligible, Reason: Interface is already cove

```

As a result of "show isis backup spf results," we see that Germany is able to use the loop-free alternate France for egress node Cyprus in case the primary link towards England fails.

This matches the expectation described above.

The next steps provide insight into the integration of Loop-Free Alternate into LDP.

```

Checking L2VPN on the Ingress PE Router Spain
cgraf@Spain# run show l2vpn connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch    WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down  NP -- interface hardware not present
CM -- control-word mismatch    -> -- only outbound connection is up
CN -- circuit not provisioned  <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: l2vpn
Local site: Spain (1)
  connection-site      Type  St      Time last up      # Up trans
  2                    rmt   Up      Mar 12 16:11:51 2009      1
  Local interface: ge-0/0/0.0, Status: Up, Encapsulation: ETHERNET
  Remote PE: 192.168.53.110, Negotiated control-word: Yes (Null)
  Incoming label: 800001, Outgoing label: 800000

```

The above “show l2vpn connections” command provides insight about the VPN label that is being used for the given L2VPN connection. Spain pushes VPN label “800000” towards Cyprus for a given L2VPN.

### Identifying the Transport Label Used by Spain to Reach the Egress FEC Cyprus

The next task is to identify the transport label that Spain uses to reach the egress FEC Cyprus. Knowing the label that Cyprus uses towards Germany, it is important to check what action the Loop-Free Alternate-enabled router Germany takes upon receiving it. There are two valid approaches. One is to check the MPLS routing table via “show route table mpls.0,” and the second option is to check table inet.3, where all available FECs should be seen. In this example, node Spain uses transport label “299824” to reach the egress FEC Cyprus (the endpoint of the L2VPN).

```

cgraf@Spain> show route table mpls.0

mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 15:50:33, metric 1
                  Receive
1                *[MPLS/0] 15:50:33, metric 1
                  Receive
2                *[MPLS/0] 15:50:33, metric 1
                  Receive
100064           *[LDP/9] 13:00:30, metric 1
                  > to 10.103.106.1 via ge-0/0/1.0, Pop
100064(S=0)     *[LDP/9] 13:00:30, metric 1

```

```

> to 10.103.106.1 via ge-0/0/1.0, Pop
100128      *[LDP/9] 01:21:42, metric 1
> to 10.103.106.1 via ge-0/0/1.0, Swap 299792
100144      *[LDP/9] 01:21:42, metric 1
> to 10.103.106.1 via ge-0/0/1.0, Swap 299808
100160      *[LDP/9] 01:21:42, metric 1
> to 10.103.106.1 via ge-0/0/1.0, Swap 299824
800001      *[L2VPN/7] 12:59:51
> via ge-0/0/0.0, Pop      Offset: 4
ge-0/0/0.0  *[L2VPN/7] 12:59:51, metric2 1
> to 10.103.106.1 via ge-0/0/1.0, Push 800000, Push
299824(top) Offset: -4

```

The customer edge (CE)-facing interface is ge-0/0/0.0 for the L2VPN. Two labels are being pushed over the core-facing interface ge-0/0/1.0 towards the next hop 10.103.106.1 (Germany). The two label stack is the transport label "299824" and the vpn-label "800000."

The transport label 299824 is used to reach FEC Cyprus (lo0.0 192.168.53.110) and can be seen in the inet.3 as well:

```

cgraf@Spain> show route table inet.3 | find 192.168.53.110

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.53.110/32  *[LDP/9] 01:49:27, metric 1
> to 10.103.106.1 via ge-0/0/1.0, Push 299824

```

#### Checking the Label Operation of the LSR Germany When There is No Link Failure

All label swap operations are available via the routing table mpls.0. Checking the routing table mpls.0 on router Germany for label "299824" reveals the forwarding path from Spain for given L2VPN to Cyprus as egress PE router.

```

cgraf@Germany> show route table mpls.0 detail | find ^299824
299824 (1 entry, 1 announced)
  *LDP      Preference: 9
           Next hop type: Router, Next hop index: 262149
           Next hop reference count: 2
           Next hop: via so-0/2/2.0 weight 0x1, selected
           Label operation: Swap 299888
           Next hop: via so-0/2/3.0 weight 0x4000
           Label operation: Swap 299888
           State: <Active Int>
           Local AS: 100
           Age: 1:23:09      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.53.110/32

```

Two very important items are highlighted here:

- Primary forwarding action taken: Germany swaps the label 299824 with 299888 over the primary path so-0/2/2.0 to England. The keyword “selected” identifies this next hop over so-0/2/2.0 as active.
- Loop-Free Alternate action taken: There is a second next hop via so-0/2/3.0 listed (to France). This next hop has a weight of 0x4000 and was installed because of Loop-Free Alternate. This next hop is only being used if the primary next hop is down. When using the loop-free alternate France, Germany performs the same swap operation of label “299888.” This is essentially a coincidence that Germany swaps to the same label. Within the LDP database, each neighbor (France 192.168.53.102 and England 192.168.53.101) advertise the same transport label to reach the egress FEC Cyprus (192.168.53.110).

In this example, node England and France are both advertising label 299888 by coincidence towards Germany for FEC Cyprus (lo0.0 Cyprus: 192.168.53.110):

```
cgraf@Germany> show ldp database
...
Input label database, 192.168.53.103:0--192.168.53.101:0
  Label      Prefix
    3        192.168.53.101/32
 299840     192.168.53.102/32
 299904     192.168.53.103/32
 299920     192.168.53.106/32
 299888     192.168.53.110/32
  ....
Input label database, 192.168.53.103:0--192.168.53.102:0
  Label      Prefix
 299840     192.168.53.101/32
    3        192.168.53.102/32
 299904     192.168.53.103/32
 299920     192.168.53.106/32
 299888     192.168.53.110/32
```

### Checking the Label Operation of the LSR Germany in Case of a Link Failure

The “show route table mpls.0 detail | find ^299824” discussed earlier would have been sufficient to show what action is taken in the case of rerouting to a loop-free alternate next hop.

However, another very convenient command is the “show route detail” on the Loop-Free Alternate-enabled node Germany of node Cyprus’ loopback 192.168.53.110 to verify the detailed operation of Loop-Free Alternate for a given prefix. In the case of an L2VPN, it is enough to check the FEC Cyprus, since it is the egress PE router for the given L2VPN. Loop-Free Alternate is so incredibly fast because it is pre-installing a backup next hop in the PFE (Loop-Free Alternate uses a different weight (0x4000) for the backup next hop). Those multiple backup next hops are highlighted in the “show route detail” command.

The important line to watch is the weighted next hop with a weight of 0x4000:

Next hop: 10.102.103.1 via so-0/2/3.0 weight 0x4000

```
cgraf@Germany# run show route detail 192.168.53.110

inet.0: 22 destinations, 22 routes (21 active, 0 holddown, 1 hidden)
192.168.53.110/32 (1 entry, 1 announced)
  *IS-IS Preference: 18
    Level: 2
    Next hop type: Router, Next hop index: 262143
    Next hop reference count: 6
    Next hop: 10.101.103.1 via so-0/2/2.0 weight 0x1, selected
    Next hop: 10.102.103.1 via so-0/2/3.0 weight 0x4000
    State: <Active Int>
    Local AS: 100
```

```

Age: 20:39      Metric: 20
Task: IS-IS
Announcement bits (2): 0-KRT 2-LDP
AS path: I

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

192.168.53.110/32 (1 entry, 1 announced)
State: <FlashAll>
*LDP Preference: 9
Next hop type: Router
Next hop reference count: 2
Next hop: via so-0/2/2.0 weight 0x1, selected
Label operation: Push 299888
Next hop: via so-0/2/3.0 weight 0x4000
Label operation: Push 299888
State: <Active Int>
Local AS: 100
Age: 20:38      Metric: 1
Task: LDP
Announcement bits (1): 1-Resolve tree 1
AS path: I

```

In inet.0 (see above), the loop-free alternate 10.102.103.1 via so-0/2/3/0 (to France) is used for IPv4 packets if the egress router is Cyprus (.192.168.53.110).

The routing table inet.3 may be confusing, as a single label push operation is listed. The push is only used if Germany is the ingress node. In case of being transit (as in this example), please see the section titled, “Checking the Label Operation of the LSR Germany When There is No Link Failure.”

Alternatively, to see what LSR Germany is doing upon receiving a transport label of “299824,” one can check the forwarding table as well.

```

cgraf@Germany> show route forwarding-table detail | find ^299824
299824          user      0                ulst 262149      2
                                                Swap 299888      581          1 so-
0/2/2.0
                                                Swap 299888      586          1 so-
0/2/3.0

```

## Summary

Loop-free Alternatives is a new feature that helps service providers and large enterprises meet the availability demands of sensitive, real-time applications with minimal operational impact. Loop-Free Alternate adds fast reroute capabilities to IP/LDP-based networks, providing SONET-like sub-50 ms restoration times and ensuring minimal packet loss or service disruption in the event of a node or link failure. Because Loop-Free Alternate does not require the introduction of any new protocols to the network, and can be done on a node-by-node basis, it is very simple to implement and can be optimized for many different business requirements and network architectures.

Because Juniper’s Loop-Free Alternate implementation is based on the powerful Junos OS, it is highly scalable, stable, and easy to configure. Juniper’s unique and innovative Loop-Free Alternate implementation is suitable for even the world’s largest networks, providing network operators with another option for ensuring that their networks operate with the highest levels of performance and reliability.

## About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at [www.juniper.net](http://www.juniper.net).

---

### Corporate and Sales Headquarters

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089 USA  
Phone: 888.JUNIPER (888.586.4737)  
or 408.745.2000  
Fax: 408.745.2100  
[www.juniper.net](http://www.juniper.net)

### APAC Headquarters

Juniper Networks (Hong Kong)  
26/F, Cityplaza One  
1111 King's Road  
Taikoo Shing, Hong Kong  
Phone: 852.2332.3636  
Fax: 852.2574.7803

### EMEA Headquarters

Juniper Networks Ireland  
Airside Business Park  
Swords, County Dublin, Ireland  
Phone: 35.31.8903.600  
EMEA Sales: 00800.4586.4737  
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2009 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.