Internet Engineering Task Force (IETF)
Request for Comments: 6571
Category: Informational
ISSN: 2070-1721

C. Filsfils, Ed.
Cisco Systems
P. Francois, Ed.
Institute IMDEA Networks
M. Shand

B. Decraene
France Telecom
J. Uttaro
AT&T
N. Leymann
M. Horneffer
Deutsche Telekom

June 2012

Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks

Abstract

In this document, we analyze the applicability of the Loop-Free Alternate (LFA) method of providing IP fast reroute in both the core and access parts of Service Provider networks. We consider both the link and node failure cases, and provide guidance on the applicability of LFAs to different network topologies, with special emphasis on the access parts of the network.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6571.

Table of Contents

1. Introduction

In this document, we analyze the applicability of the Loop-Free Alternate (LFA) [RFC5714] [RFC5286] method of providing IP fast reroute (IPFRR) in both the core and access parts of Service Provider (SP) networks. We consider both the link and node failure cases, and provide guidance on the applicability of LFAs to different network topologies, with special emphasis on the access parts of the network.

We first introduce the terminology used in this document in Section 2. In Section 3, we describe typical access network designs, and we analyze them for LFA applicability. In Section 4, we describe a simulation framework for the study of LFA applicability in SP core networks, and present results based on various SP networks. We then emphasize the independence between protection schemes used in the core and at the access level of the network. Finally, we discuss the key benefits of the LFA method, which stem from its simplicity, and we draw some conclusions.

2. Terminology

We use IS-IS [RFC1195] [IS-IS] as a reference. It is assumed that normal routing (i.e., when traffic is not being fast-rerouted around a failure) occurs along the shortest path. The analysis is equally applicable to OSPF [RFC2328] [RFC5340].

A per-prefix LFA for a destination D at a node S is a pre-computed backup IGP next hop for that destination. This backup IGP next hop can be link-protecting or node-protecting. In this document, we assume that all links to be protected with LFAs are point-to-point.

Link-protecting: A neighbor N is a link-protecting per-prefix LFA for S's route to D if equation eq1 is satisfied. This is in line with the definition of an LFA in [RFC5714].

eq1: ND < NS + SD

where XY refers to the IGP distance from X to Y

Equation eq1

Node-protecting: A neighbor N is a node-protecting LFA for S's route to D with initial IGP next hop F if N is a link-protecting LFA for D and equation eq2 is satisfied. This is in line with the definition of a Loop-Free Node-Protecting Alternate (also known as a node-protecting LFA) in [RFC5714].

eq2: ND < NF + FD

Equation eq2

De facto node-protecting LFA: This is a link-protecting LFA that turns out to be node-protecting. This occurs in cases illustrated by the following examples:

- The LFA candidate that is picked by S actually satisfies Equation eq2, but S did not verify that property. The show command issued by the operator would not indicate this LFA as "node-protecting", while in practice (de facto), it is.

- A cascading effect of multiple LFAs can also provide de facto node protection. Equation eq2 is not satisfied, but the combined activation of LFAs by some other neighbors of the failing node F provides (de facto) node protection. In other words, it puts the data plane in a state such that packets forwarded by S ultimately

reach a neighbor of F that has a node-protecting LFA. Note that in this case, S cannot indicate the node-protecting behavior of the repair without running additional computations.

Per-link LFA: A per-link LFA for the link SF is one pre-computed backup IGP next hop for all of the destinations reached through SF. This is a neighbor of the repairing node that is a per-prefix LFA for all of the destinations that the repairing node reaches through SF. Note that such a per-link LFA exists if S has a per-prefix LFA for destination F.

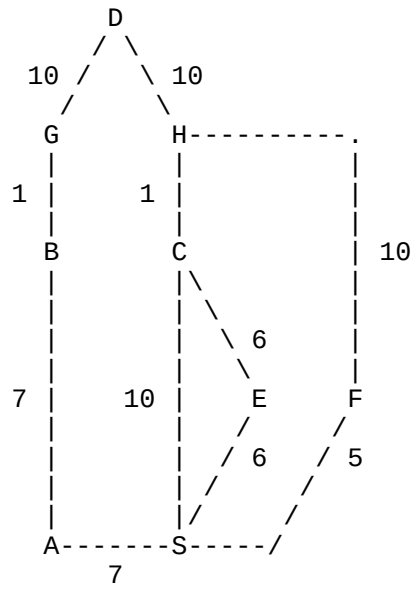

Figure 1: Example 1

In Figure 1, considering the protection of link SC, we can see that A, E, and F are per-prefix LFAs for destination D, as none of them use S to reach D.

For destination D, A and F are node-protecting LFAs, as they do not reach D through node C, while E is not node-protecting for S, as it reaches D through C.

If S does not compute and select node-protecting LFAs, there is a chance that S picks the non-node-protecting LFA E, although A and F were node-protecting LFAs. If S enforces the selection of node-protecting LFAs, then in the case of the single failure of link SC,

S will first activate its LFA, deviate traffic addressed to D along S-A-B-G-D and/or S-F-H-D, and then converge to its post-convergence optimal path S-E-C-H-D.

A reaches C via S; thus, A is not a per-link LFA for link SC. E reaches C through link EC; thus, E is a per-link LFA for link SC. This per-link LFA does not provide de facto node protection. Upon failure of node C, S would fast-reroute D-destined packets to its per-link LFA (= E). E would itself detect the failure of EC; hence, it would activate its own per-link LFA (= S). Traffic addressed to D would be trapped in a loop; hence, there is no de facto node protection behavior.

If there were a link between E and F that E would pick as its LFA for destination D, then E would provide de facto node protection for S, as upon the activation of its LFA, S would deviate traffic addressed to D towards E. In turn, E deviates that traffic to F, which does not reach D through C.

F is a per-link LFA for link SC, as F reaches C via H. This per-link LFA is de facto node-protecting for destination D, as F reaches D via F-H-D.

Micro-Loop (uLoop): the occurrence of a transient forwarding loop during a routing transition (as defined in [RFC5715]).

In Figure 1, the loss of link SE cannot create any uLoop, because of the following:

1. The link is only used to reach destination E.

2. S is the sole node changing its path to E upon link SE failure.

3. S’s shortest path to E after the failure goes via C.

4. C’s best path to E (before and after link SC failure) is via CE.

On the other hand, upon failure of link AB, a micro-loop may form for traffic destined to B. Indeed, if A updates its Forwarding Information Base (FIB) before S, A will reroute B-destined traffic towards S, while S is still forwarding this traffic to A.

3. Access Network

The access part of the network often represents the majority of the nodes and links. It is organized in several tens or more of regions interconnected by the core network. Very often, the core acts as an IS-IS level-2 domain (OSPF area 0), while each access region is

confined in an IS-IS level-1 domain (OSPF non-0 area). Very often, the network topology within each access region is derived from a unique template common across the whole access network. Within an access region itself, the network is made of several aggregation regions, each following the same interconnection topologies.

For these reasons, in the next sections, we base the analysis of the LFA applicability in a single access region, with the following assumptions:

- Two routers (C1 and C2) provide connectivity between the access region and the rest of the network. If a link connects these two routers in the region area, then it has a symmetric IGP metric c.

- We analyze a single aggregation region within the access region. Two aggregation routers (A1 and A2) interconnect the aggregation region to the two routers C1 and C2 for the analyzed access region. If a link connects A1 to A2, then it has a symmetric IGP metric a. If a link connects a router A to a router C, then for the sake of generality we will call d the metric for the directed link CA and u the metric for the directed link AC.

- We analyze two edge routers, E1 and E2, in the access region. Each is dual-homed directly either to C1 and C2 (Section 3.1) or to A1 and A2 (Sections 3.2, 3.3, and 3.4). The directed link metric between Cx/Ax and Ey is d and u in the opposite direction.

- We assume a multi-level IGP domain. The analyzed access region forms a level-1 (L1) domain. The core is the level-2 (L2) domain. We assume that the link between C1 and C2, if it exists, is configured as L1L2. We assume that the loopbacks of the C routers are part of the L2 topology. L1 routers learn about them as propagated routes (L2=>L1 with the Down bit set). We remind the reader that if an L1L2 router learns about X/x as an L1 path P1, an L2 path P2, and an L1L2 path P12, then it will prefer path P1. If path P1 is lost, then it will prefer path P2.

- We assume that all of the C, A, and E routers may be connected to customers; hence, we analyze LFA coverage for the loopbacks of each type of node.

- We assume that no useful traffic is directed to router-to-router subnets; hence, we do not analyze LFA applicability for such subnets.

- A prefix P models an important IGP destination that is not present in the local access region. The IGP metric from C1 to P is x, and the metric from C2 to P is x + e.

- We analyze LFA coverage against all link and node failures within the access region.
- WxYz refers to the link from Wx to Yz.
- We assume that c < d + u and a < d + u (a commonly agreed-upon design rule).
- In the square access design (Section 3.3), we assume that c < a (a commonly agreed-upon design rule).
- We analyze the most frequent topologies found in an access region.
- We first analyze per-prefix LFA applicability and then per-link.
- The topologies are symmetric with respect to a vertical axis; hence, we only detail the logic for the link and node failures of the left half of the topology.

3.1. Triangle

We describe the LFA applicability for the failures of C1E1, E1, and C1 (Figure 2).

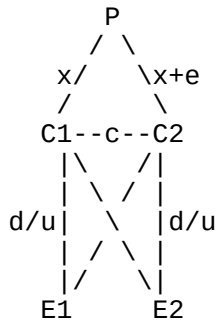

Figure 2: Triangle

3.1.1. E1C1 Failure

3.1.1.1. Per-Prefix LFA

Three destinations are impacted by this link failure: C1, E2, and P.

The LFA for destination C1 is C2, because eq1: c < d + u. Node protection for route C1 is not applicable. (If C1 goes down, traffic destined to C1 is lost anyway.)

The LFA to E2 is via C2, because eq1: d < d + u + d. It is node-protecting, because eq2: d < c + d.

The LFA to P is via C2, because c < d + u. It is node-protecting if eq2: x + e < x + c, i.e., if e < c. This relationship between e and c is an important aspect of the analysis, which is discussed in detail in Sections 3.5 and 3.6.

Conclusion: All important intra-PoP (Point of Presence) routes with primary interface E1C1 benefit from LFA link and node protection. All important inter-PoP routes with primary interface E1C1 benefit from LFA link protection, and also from node protection if e < c.

3.1.1.2. Per-Link LFA

We have a per-prefix LFA to C1; hence, we have a per-link LFA for link E1C1. All impacted destinations are protected against link failure. In the case of C1 node failure, the traffic to C1 is lost (by definition), the traffic to E2 is de facto protected against node failure, and the traffic to P is de facto protected when e < c.

3.1.2. C1E1 Failure

3.1.2.1. Per-Prefix LFA

C1 only has one primary route via C1E1: the route to E1 (because c < d + u).

C1’s LFA to E1 is via C2, because eq1: d < c + d.

Node protection upon E1’s failure is not applicable, as the only impacted traffic is sinked at E1 and hence is lost anyway.

Conclusion: All important routes with primary interface C1E1 benefit from LFA link protection. Node protection is not applicable.

3.1.2.2. Per-Link LFA

We have a per-prefix LFA to E1; hence, we have a per-link LFA for link C1E1. De facto node protection is not applicable.

3.1.3. uLoop

The IGP convergence cannot create any uLoop. See Section 3.7.

3.1.4. Conclusion

All important intra-PoP routes benefit from LFA link and node protection or de facto node protection. All important inter-PoP routes benefit from LFA link protection. De facto node protection is ensured if e < c. (This is particularly the case for dual-plane core or two-tiered IGP metric design; see Sections 3.5 and 3.6.)

The IGP convergence does not cause any uLoop.

Per-link LFAs and per-prefix LFAs provide the same protection benefits.

3.2. Full Mesh

We describe the LFA applicability for the failures of C1A1, A1E1, E1, A1, and C1 (Figure 3).

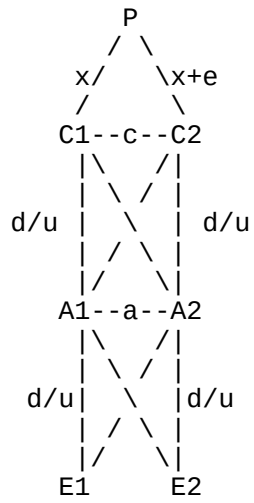

Figure 3: Full Mesh

3.2.1. E1A1 Failure

3.2.1.1. Per-Prefix LFA

Four destinations are impacted by this link failure: A1, C1, E2, and P.

The LFA for A1 is A2: eq1: a < d + u. Node protection for route A1 is not applicable. (If A1 goes down, traffic to A1 is lost anyway.)

The LFA for C1 is A2: eq1: u < d + u + u. Node protection for route C1 is guaranteed: eq2: u < a + u.

The LFA to E2 is via A2: eq1: d < d + u + d. Node protection is guaranteed: eq2: d < a + d.

The LFA to P is via A2: eq1: u + x < d + u + u + x. Node protection is guaranteed: eq2: u + x < a + u + x.

Conclusion: All important intra-PoP and inter-PoP routes with primary interface E1A1 benefit from LFA link and node protection.

3.2.1.2. Per-Link LFA

We have a per-prefix LFA to A1; hence, we have a per-link LFA for link E1A1. All impacted destinations are protected against link failure. De facto node protection is provided for all destinations (except to A1, which is not applicable).

3.2.2. A1E1 Failure

3.2.2.1. Per-Prefix LFA

A1 only has one primary route via A1E1: the route to E1 (because a < d + u).

A1’s LFA to E1 is via A2: eq1: d < a + d.

Node protection upon E1’s failure is not applicable, as the only impacted traffic is sinked at E1 and hence is lost anyway.

Conclusion: All important routes with primary interface A1E1 benefit from LFA link protection. Node protection is not applicable.

3.2.2.2. Per-Link LFA

We have a per-prefix LFA to E1; hence, we have a per-link LFA for link C1E1. De facto node protection is not applicable.

3.2.3. A1C1 Failure

3.2.3.1. Per-Prefix LFA

Two destinations are impacted by this link failure: C1 and P.

The LFA for C1 is C2, because eq1: c < d + u. Node protection for route C1 is not applicable. (If C1 goes down, traffic to C1 is lost anyway.)

The LFA for P is via C2, because c < d + u. It is de facto protected against node failure if eq2: x + e < x + c.

Conclusion: All important intra-PoP routes with primary interface A1C1 benefit from LFA link protection. (Node protection is not applicable.) All important inter-PoP routes with primary interface E1C1 benefit from LFA link protection (and from de facto node protection if e < c).

3.2.3.2. Per-Link LFA

We have a per-prefix LFA to C1; hence, we have a per-link LFA for link A1C1. All impacted destinations are protected against link failure. In the case of C1 node failure, the traffic to C1 is lost (by definition), and the traffic to P is de facto node protected if e < c.

3.2.4. C1A1 Failure

3.2.4.1. Per-Prefix LFA

C1 has three routes via C1A1: A1, E1, and E2. E2 behaves like E1 and hence is not analyzed further.

C1’s LFA to A1 is via C2, because eq1: d < c + d. Node protection upon A1’s failure is not applicable, as the traffic to A1 is lost anyway.

C1’s LFA to E1 is via A2: eq1: d < u + d + d. Node protection upon A1’s failure is guaranteed, because eq2: d < a + d.

Conclusion: All important routes with primary interface C1A1 benefit from LFA link protection. Node protection is guaranteed where applicable.

3.2.4.2. Per-Link LFA

We have a per-prefix LFA to A1; hence, we have a per-link LFA for link C1E1. De facto node protection is available.

3.2.5. uLoop

The IGP convergence cannot create any uLoop. See Section 3.7.

3.2.6. Conclusion

All important intra-PoP routes benefit from LFA link and node protection.

All important inter-PoP routes benefit from LFA link protection. They benefit from node protection upon failure of A nodes. They benefit from node protections upon failure of C nodes if e < c. (This is particularly the case for dual-plane core or two-tiered IGP metric design; see Sections 3.5 and 3.6.)

The IGP convergence does not cause any uLoop.

Per-link LFAs and per-prefix LFAs provide the same protection benefits.

3.3. Square

We describe the LFA applicability for the failures of C1A1, A1E1, E1, A1, and C1 (Figure 4).

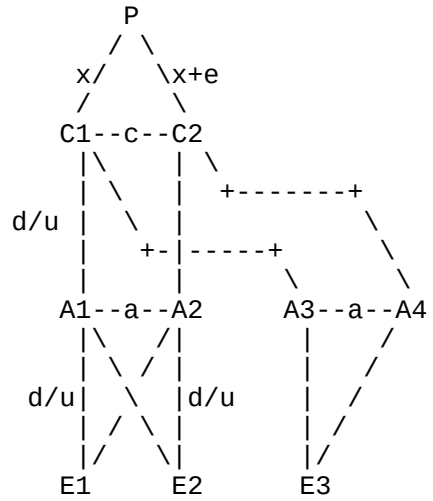

Figure 4: Square

3.3.1. E1A1 Failure

3.3.1.1. Per-Prefix LFA

E1 has six routes via E1A1: A1, C1, P, E2, A3, and E3.

E1's LFA route to A1 is via A2, because eq1: a < d + u. Node protection for traffic to A1 upon A1 node failure is not applicable.

E1's LFA route to A3 is via A2, because eq1: u + c + d < d + u + u + d. This LFA is guaranteed to be node-protecting, because eq2: u + c + d < a + u + d.

E1's LFA route to C1 is via A2, because eq1: u + c < d + u + u. This LFA is guaranteed to be node-protecting, because eq2: u + c < a + u.

E1's primary route to E2 is via ECMP(E1A1, E1A2) (Equal-Cost Multi-Path). The LFA for the first ECMP path (via A1) is the second ECMP path (via A2). This LFA is guaranteed to be node-protecting, because eq2: d < a + d.

E1's primary route to E3 is via ECMP(E1A1, E1A2). The LFA for the first ECMP path (via A1) is the second ECMP path (via A2). This LFA is guaranteed to be node-protecting, because eq2: u + d + d < a + u + d + d.

If e = 0: E1's primary route to P is via ECMP(E1A1, E1A2). The LFA for the first ECMP path (via A1) is the second ECMP path (via A2). This LFA is guaranteed to be node-protecting, because eq2: u + x + 0 < a + u + x.

If e <> 0: E1's primary route to P is via E1A1. Its LFA is via A2, because eq1: u + c + x < d + u + u + x. This LFA is guaranteed to be node-protecting, because eq2: u + c + x < a + u + x.

Conclusion: All important intra-PoP and inter-PoP routes with primary interface E1A1 benefit from LFA link protection and node protection.

3.3.1.2. Per-Link LFA

We have a per-prefix LFA for A1; hence, we have a per-link LFA for link E1A1. All important intra-PoP and inter-PoP routes with primary interface E1A1 benefit from LFA per-link protection and de facto node protection.

3.3.2. A1E1 Failure

3.3.2.1. Per-Prefix LFA

A1 only has one primary route via A1E1: the route to E1.

A1's LFA for route E1 is the path via A2, because eq1: d < a + d. Node protection is not applicable.

Conclusion: All important routes with primary interface A1E1 benefit from LFA link protection. Node protection is not applicable.

3.3.2.2. Per-Link LFA

All important routes with primary interface A1E1 benefit from LFA link protection. De facto node protection is not applicable.

3.3.3. A1C1 Failure

3.3.3.1. Per-Prefix LFA

Four destinations are impacted when A1C1 fails: C1, A3, E3, and P.

A1's LFA to C1 is via A2, because eq1: u + c < a + u. Node protection is not applicable for traffic to C1 when C1 fails.

A1's LFA to A3 is via A2, because eq1: u + c + d < a + u + d. It is de facto node-protecting, as a < u + c + d (as we assumed a < u + d). Indeed, for destination A3, A2 forwards traffic to C2, and C2 has a node-protecting LFA -- A4 -- for the failure of link C2C1, as a < u + c + d. Hence, the cascading application of LFAs by A1 and C2 during the failure of C1 provides de facto node protection.

A1's LFA to E3 is via A2, because eq1: u + d + d < a + u + d + d. It is node-protecting, because eq2: u + d + d < u + c + d + d.

A1's primary route to P is via C1 (even if e = 0, u + x < u + c + x). The LFA is via A2, because eq1: u + c + x < a + u + x (case where c <= e) and eq1: u + x + e < a + u + x (case where c >= e). This LFA is node-protecting (from the viewpoint of A1 computing eq2) if eq2: u + x + e < u + c + x. This inequality is true if e < c.

Conclusion: All important intra-PoP routes with primary interface A1C1 benefit from LFA link protection and node protection. Note that A3 benefits from de facto node protection. All important inter-PoP routes with primary interface A1C1 benefit from LFA link protection. They also benefit from node protection if e < c.

3.3.3.2. Per-Link LFA

All important intra-PoP routes with primary interface A1C1 benefit from LFA link protection and de facto node protection. All important inter-PoP routes with primary interface A1C1 benefit from LFA link protection. They also benefit from de facto node protection if e < c.

3.3.4. C1A1 Failure

3.3.4.1. Per-Prefix LFA

Three destinations are impacted by C1A1 link failure: A1, E1, and E2. E2's analysis is the same as E1 and hence is omitted.

C1 has no LFA for A1. Indeed, its neighbors (C2 and A3) have a shortest path to A1 via C1. This is due to the assumption (c < a).

C1's LFA for E1 is via C2, because eq1: d + d < c + d + d. It provides node protection, because eq2: d + d < d + a + d.

Conclusion: All important intra-PoP routes with primary interface A1C1, except A1, benefit from LFA link protection and node protection.

3.3.4.2. Per-Link LFA

C1 does not have a per-prefix LFA for destination A1; hence, there is no per-link LFA for link C1A1.

3.3.4.3. Assumptions on the Values of c and a

The commonly agreed-upon design rule (c < a) is especially beneficial for a deployment using per-link LFA: it provides a per-link LFA for the most important direction (A1C1). Indeed, there are many more destinations reachable over A1C1 than over C1A1. As the IGP convergence duration is proportional to the number of routes to update, there is a better benefit in leveraging LFA FRR for link A1C1 than for link C1A1.

Note as well that the consequence of this assumption is much more important for per-link LFA than for per-prefix LFA.

For per-prefix LFAs, in the case of link C1A1 failure, we do have a per-prefix LFA for E1, E2, and any node subtended below A1 and A2. Typically, most of the traffic traversing link C1A1 is directed to these E nodes; hence, the lack of per-prefix LFAs for the destination A1 might be insignificant. This is a good example of the coverage benefit of per-prefix LFAs over per-link LFAs.

In the remainder of this section, we analyze the consequence of not having c < a.

It definitely has a negative impact upon per-link LFAs.

With c > a, C1A1 has a per-link LFA, while A1C1 has no per-link LFA. The number of destinations impacted by A1C1 failure is much larger than the direction C1A1; hence, the protection is provided for the wrong direction.

For per-prefix LFAs, the availability of an LFA depends on the topology and needs to be assessed individually for each per-prefix LFA. Some backbone topologies will lead to very good protection coverage, while some others might provide very poor coverage.

More specifically, upon A1C1 failure, the coverage of a remote destination P depends on whether e < a. In such a case, A2 is de facto node-protecting per-prefix LFA for P.

Such a study likely requires a planning tool, as each remote destination P would have a different e value (exception: all of the edge devices of other aggregation pairs within the same region, as for these e = 0 by definition, e.g., E3.)

Finally, note that c = a is the worst choice. In this case, C1 has no per-prefix LFA for A1 (and vice versa); hence, there is no per-link LFA for C1A1 and A1C1.

3.3.5. Conclusion

All important intra-PoP routes benefit from LFA link and node protection with one exception: C1 has no per-prefix LFA to A1.

All important inter-PoP routes benefit from LFA link protection. They benefit from node protection if e < c.

Per-link LFA provides the same protection coverage as per-prefix LFA, with two exceptions: first, C1A1 has no per-link LFA at all. Second, when per-prefix LFA provides node protection (eq2 is satisfied), per-link LFA provides effective de facto node protection.

3.3.6. A Square Might Become a Full Mesh

If the vertical links of the square are made of parallel links (at the IP topology or below), then one should consider splitting these "vertical links" into "vertical and crossed links". The topology becomes "full mesh". One should also ensure that the two resulting sets of links (vertical and crossed) do not share any Shared Risk Link Group (SRLG).

A typical scenario in which this is prevented would be when the A1C1 bandwidth may be within a building while the A1C2 is between buildings. Hence, while from a router-port viewpoint the operation is cost-neutral, from a cost-of-bandwidth viewpoint it is not.

3.3.7. A Full Mesh Might Be More Economical Than a Square

In a full mesh, the vertical and crossed links play the dominant role, as they support most of the primary and backup paths. The capacity of the horizontal links can be dimensioned on the basis of traffic destined to a single C node or a single A node, and to a single E node.

3.4. Extended U

For the Extended U topology, we define the following terminology:

C1L1: the node "C1" as seen in topology L1.

C1L2: the node "C1" as seen in topology L2.

C1LO: the loopback of C1. This loopback is in L2.

C2LO: the loopback of C2. This loopback is in L2.

We remind the reader that C1 and C2 are L1L2 routers and that their loopbacks are in L2 only.

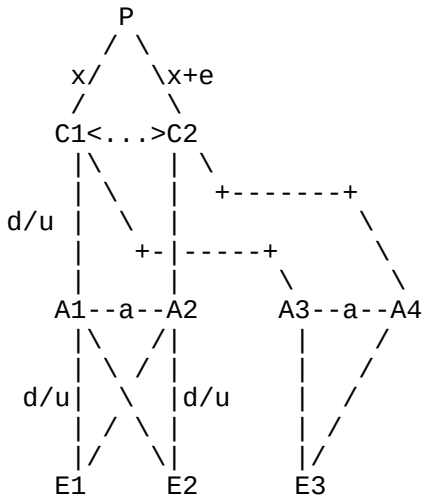

Figure 5: Extended U

There is no L1 link between C1 and C2. There might be an L2 link between C1 and C2. This is not relevant, as this is not seen from the viewpoint of the L1 topology, which is the focus of our analysis.

It is guaranteed that there is a path from C1LO to C2LO within the L2 topology (except if the L2 topology partitions, which is very unlikely and hence not analyzed here). We call "c" its path cost. Once again, we assume that c < a.

We exploit this property to create a tunnel T between C1LO and C2LO. Once again, as the source and destination addresses are the loopbacks of C1 and C2 and these loopbacks are in L2 only, it is guaranteed that the tunnel does not transit via the L1 domain.

IS-IS does not run over the tunnel; hence, the tunnel is not used for any primary paths within the L1 or L2 topology.

Within level-1, we configure C1 (C2) with a level-1 LFA extended neighbor "C2 via tunnel T" ("C1 via tunnel T").

A router supporting such an extension learns that it has one additional potential neighbor in topology level-1 when checking for LFAs.

The L1 topology learns about C1LO as an L2=>L1 route with the Down bit set, propagated by C1L1 and C2L1. The metric advertised by C2L1 is bigger than the metric advertised by C1L1 by "c".

The L1 topology learns about P as an L2=>L1 route with the Down bit set, propagated by C1L1 and C2L1. The metric advertised by C2L1 is bigger than the metric advertised by C1L1 by "e". This implies that e <= c.

3.4.1. E1A1 Failure

3.4.1.1. Per-Prefix LFA

Five destinations are impacted by E1A1 link failure: A1, C1LO, E2, E3, and P.

The LFA for A1 is via A2, because eq1: a < d + u. Node protection for traffic to A1 upon A1 node failure is not applicable.

The LFA for E2 is via A2, because eq1: d < d + u + d. Node protection is guaranteed, because eq2: d < a + d.

The LFA for E3 is via A2, because eq1: u + d + d < d + u + d + d. Node protection is guaranteed, because eq2: u + d + d < a + u + d + d.

The LFA for C1LO is via A2, because eq1: u + c < d + u + u. Node protection is guaranteed, because eq2: u + c < a + u.

If e = 0: E1’s primary route to P is via ECMP(E1A1, E1A2). The LFA for the first ECMP path (via A1) is the second ECMP path (via A2). Node protection is possible, because eq2: u + x < a + u + x.

If e <> 0: E1's primary route to P is via E1A1. Its LFA is via A2, because eq1: a + c + x < d + u + u + x. Node protection is guaranteed, because eq2: u + x + e < a + u + x <=> e < a. This is true, because e <= c and c < a.

Conclusion: Same as that for the square topology.

3.4.1.2. Per-Link LFA

Same as the square topology.

3.4.2. A1E1 Failure

3.4.2.1. Per-Prefix LFA

Same as the square topology.

3.4.2.2. Per-Link LFA

Same as the square topology.

3.4.3. A1C1 Failure

3.4.3.1. Per-Prefix LFA

Three destinations are impacted when A1C1 fails: C1, E3, and P.

A1's LFA to C1LO is via A2, because eq1: u + c < a + u. Node protection is not applicable for traffic to C1 when C1 fails.

A1's LFA to E3 is via A2, because eq1: u + d + d < d + u + u + d + d. Node protection is guaranteed, because eq2: u + d + d < a + u + d + d.

A1's primary route to P is via C1 (even if e = 0, u + x < a + u + x). The LFA is via A2, because eq1: u + x + e < a + u + x <=> e < a (which is true; see above). Node protection is guaranteed, because eq2: u + x + e < a + u + x.

Conclusion: Same as that for the square topology.

3.4.3.2. Per-Link LFA

Same as the square topology.

3.4.4. C1A1 Failure

3.4.4.1. Per-Prefix LFA

Three destinations are impacted by C1A1 link failure: A1, E1, and E2. E2's analysis is the same as E1 and hence is omitted.

C1L1 has an LFA for A1 via the extended neighbor C2L1 reachable via tunnel T. Indeed, eq1 is true: d + a < d + a + u + d. From the viewpoint of C1L1, C2L1's path to C1L1 is C2L1-A2-A1-C1L1. Remember that the tunnel is not seen by IS-IS for computing primary paths! Node protection is not applicable for traffic to A1 when A1 fails.

C1L1's LFA for E1 is via extended neighbor C2L1 (over tunnel T), because eq1: d + d < d + a + u + d + d. Node protection is guaranteed, because eq2: d + d < d + a + d.

3.4.4.2. Per-Link LFA

C1 has a per-prefix LFA for destination A1; hence, there is a per-link LFA for the link C1A1. Node resistance is applicable for traffic to E1 (and E2).

3.4.5. Conclusion

The Extended U topology is as good as the square topology.

It does not require any crossed links between the A and C nodes within an aggregation region. It does not need an L1 link between the C routers in an access region. Note that a link between the C routers might exist in the L2 topology.

3.5. Dual-Plane Core and Its Impact on the Access LFA Analysis

A dual-plane core is defined as follows:

- Each access region k is connected to the core by two C routers (C(1,k) and C(2,k)).
- C(1,k) is part of plane-1 of the dual-plane core.
- C(2,k) is part of plane-2 of the dual-plane core.
- C(1,k) has a link to C(2, l) iff k = l.
- {C(1,k) has a link to C(1, l)} iff {C(2,k) has a link to C(2, l)}.

In a dual-plane core design, e = 0; hence, the LFA node-protection coverage is improved in all of the analyzed topologies.

3.6. Two-Tiered IGP Metric Allocation

A two-tiered IGP metric allocation scheme is defined as follows:

- All of the link metrics used in the L2 domain are part of range R1.
- All of the link metrics used in an L1 domain are part of range R2.
- Range R1 << range R2 such that the difference e = C2P - C1P is smaller than any link metric within an access region.

Assuming such an IGP metric allocation, the following properties are guaranteed: c < a, e < c, and e < a.

3.7. uLoop Analysis

In this section, we analyze a case where the routing transition following the failure of a link may have some uLoop potential for one destination. Then, we show that all of the other cases do not have uLoop potential.

In the square design, upon the failure of link C1A1, traffic addressed to A1 can undergo a transient forwarding loop as C1 reroutes traffic to C2, which initially reaches A1 through C1, as c < a. This loop will actually occur when C1 updates its FIB for destination A1 before C2.

It can be shown that all of the other routing transitions following a link failure in the analyzed topologies do not have uLoop potential. Indeed, in each case, for all destinations affected by the failure, the rerouting nodes deviate their traffic directly to adjacent nodes whose paths towards these destinations do not change. As a consequence, all of these routing transitions cannot undergo transient forwarding loops.

For example, in the square topology, the failure of directed link A1C1 does not lead to any uLoop. The destinations reached over that directed link are C1 and P. A1’s and E1’s shortest paths to these destinations after the convergence go via A2. A2’s path to C1 and P is not using A1C1 before the failure; hence, no uLoop may occur.

3.8. Summary

In this section, we summarize the applicability of LFAs detailed in the previous sections. For link protection, we use "Full" to refer to the applicability of LFAs for each destination, reached via any link of the topology. For node protection, we use "Yes" to refer to the fact that node protection is achieved for a given node.

1. Intra-Area Destinations

   Link Protection

   + Triangle: Full
   + Full Mesh: Full
   + Square: Full, except C1 has no LFA for dest A1
   + Extended U: Full

   Node Protection

   + Triangle: Yes

   + Full Mesh: Yes
   + Square: Yes
   + Extended U: Yes

2. Inter-Area Destinations

   Link Protection

   + Triangle: Full
   + Full Mesh: Full
   + Square: Full
   + Extended U: Full

   Node Protection

   + Triangle: Yes, if e < c
   + Full Mesh: Yes for A failure, if e < c for C failure
   + Square: Yes for A failure, if e < c for C failure
   + Extended U: Yes, if e <= c and c < a

3. uLoops

   * Triangle: None
   * Full Mesh: None
   * Square: None, except traffic to A1 when C1A1 fails
   * Extended U: None, if a > e

4. Per-Link LFA vs. Per-Prefix LFA

* Triangle: Same
* Full Mesh: Same
* Square: Same, except C1A1 has no per-link LFA. In practice, this means that per-prefix LFAs will be used. (Hence, C1 has no LFA for dest = E1 and dest = A1.)
* Extended U: Same

4. Core Network

In the backbone, the optimization of the network design to achieve the maximum LFA protection is less straightforward than in the case of the access/aggregation network.

The main optimization objectives for backbone topology design are cost, latency, and bandwidth, constrained by the availability of fiber. Optimizing the design for local IP restoration is more likely to be considered as a non-primary objective. For example, the way the fiber is laid out and the resulting cost to change it lead to ring topologies in some backbone networks.

Also, the capacity-planning process is already complex in the backbone. The process needs to make sure that the traffic matrix (demand) is supported by the underlying network (capacity) under all possible variations of the underlying network (what-if scenario related to one-SRLG failure). Classically, "supported" means that no congestion is experienced and that the demands are routed along the appropriate latency paths. Selecting the LFA method as a deterministic FRR solution for the backbone would require enhancement of the capacity-planning process to add a third constraint: Each variation of the underlying network should lead to sufficient LFA coverage. (We detail this aspect in Section 7.)

On the other hand, the access network is based on many replications of a small number of well-known (well-engineered) topologies. The LFA coverage is deterministic and is independent of additions/insertions of a new edge device, a new aggregation sub-region, or a new access region.

In practice, we believe that there are three profiles for the backbone applicability of the LFA method:

In the first profile, the designer plans all of the network resilience on IGP convergence. In such a case, the LFA method is a free bonus. If an LFA is available, then the loss of connectivity is likely reduced by a factor of 10 (50 msec vs. 500 msec); otherwise, the loss of connectivity depends on IGP

convergence, which is the initial target anyway. The LFA method should be very successful here, as it provides a significant improvement without any additional cost.

In the second profile, the designer seeks a very high and deterministic FRR coverage, and he either does not want or cannot engineer the topology. The LFA method should not be considered in this case. MPLS Traffic Engineering (TE) FRR would perform much better in this environment. Explicit routing ensures that a backup path exists, whatever the underlying topology.

In the third profile, the designer seeks a very high and deterministic FRR coverage, and he does engineer the topology. The LFA method is appealing in this scenario, as it can provide a very simple way to obtain protection. Furthermore, in practice, the requirement for FRR coverage might be limited to a certain part of the network (e.g., a given sub-topology) and/or is likely limited to a subset of the demands within the traffic matrix. In such a case, if the relevant part of the network natively provides a high degree of LFA protection for demands of interest, it might actually be straightforward to improve the topology and achieve the level of protection required for the sub-topology and the demands that matter. Once again, the practical problem needs to be considered (which sub-topology, and which real demands need 50 msec), as it is often simpler than the theoretical generic one.

For the reasons explained previously, the backbone applicability should be analyzed on a case-by-case basis, and it is difficult to derive generic rules.

In order to help the reader to assess the LFA applicability in his own case, we provide some simulation results based on 11 real backbone topologies in the next section.

4.1. Simulation Framework

In order to perform an analysis of LFA applicability in the core, we usually receive the complete IS-IS/OSPF linkstate database taken on a core router. We parse it to obtain the topology. During this process, we eliminate all nodes connected to the topology with a single link and all prefixes except a single "node address" per router. We compute the availability of per-prefix LFAs to all of these node addresses, which we hereafter call "destinations". We treat each link in each direction.

For each (directed) link, we compute whether we have a per-prefix LFA to the next hop. If so, we have a per-link LFA for the link.

The per-link-LFA coverage for a topology T is the fraction of the number of links with a per-link LFA divided by the total number of links.

For each link, we compute the number of destinations whose primary path involves the analyzed link. For each such destination, we compute whether a per-prefix LFA exists.

The per-prefix LFA coverage for a topology T is the following fraction:

(the sum across all links of the number of destinations with a primary path over the link and a per-prefix LFA)

divided by

(the sum across all links of the number of destinations with a primary path over the link)

4.2. Data Set

Our data set is based on 11 SP core topologies with different geographical scopes: worldwide, national, and regional. The number of nodes ranges from 600 to 16. The average link-to-node ratio is 2.3, with a minimum of 1.2 and maximum of 6.

4.3. Simulation Results

Topology	Per-Link LFA	Per-Prefix LFA
T1	45%	76%
T2	49%	98%
T3	88%	99%
T4	68%	84%
T5	75%	94%
T6	87%	98%
T7	16%	67%
T8	87%	99%
T9	67%	79%
T10	98%	99%
T11	59%	77%
Average	67%	89%
Median	68%	94%

Table 1: Core LFA Coverages

In Table 1, we observe a wide variation in terms of LFA coverage across topologies: from 67% to 99% for the per-prefix LFA coverage, and from 16% to 98% for the per-link LFA coverage. Several topologies have been optimized for LFAs (T3, 6, 8, and 10). This illustrates the need for case-by-case analysis when considering LFAs for core networks.

It should be noted that, contrary to the access/aggregation topologies, per-prefix LFA outperforms per-link LFA in the backbone.

5. Core and Access Protection Schemes Are Independent

Specifically, a design might use LFA FRR in the access and MPLS TE FRR in the core.

The LFA method provides great benefits for the access network, due to its excellent access coverage and its simplicity.

MPLS TE FRR’s topology independence might prove beneficial in the core when the LFA FRR coverage is judged too small and/or the designer feels unable to optimize the topology to improve the LFA coverage.

6. Simplicity and Other LFA Benefits

The LFA solution provides significant benefits that mainly stem from its simplicity.

Behavior of LFAs is an automated process that makes fast restoration an intrinsic part of the IGP, with no additional configuration burden in the IGP or any other protocol.

Thanks to this integration, the use of multiple areas in the IGP does not make fast restoration more complex to achieve than in a single area IGP design.

There is no requirement for network-wide upgrade, as LFAs do not require any protocol change and hence can be deployed router by router.

With LFAs, the backup paths are pre-computed and installed in the data plane in advance of the failure. Assuming a fast enough FIB update time compared to the total number of (important) destinations, a "<50-msec repair" requirement becomes achievable. With a prefix-independent implementation, LFAs have a fixed repair time, as the repair time depends on the failure detection time and the time required to activate the behavior of an LFA, which does not scale with the number of destinations to be fast-rerouted.

Link and node protection are provided together and without any operational differences. (As a comparison, MPLS TE FRR link and node protections require different types of backup tunnels and different grades of operational complexity.)

Also, compared to MPLS TE FRR, an important simplicity aspect of the LFA solution is that it does not require the introduction of yet another virtual layer of topology. Maintaining a virtual topology of explicit MPLS TE tunnels clearly increases the complexity of the network. MPLS TE tunnels would have to be represented in a network management system in order to be monitored and managed. In large networks, this may significantly contribute to the number of network entities polled by the network management system and monitored by operational staff. An LFA, on the other hand, only has to be monitored for its operational status once per router, and it needs to be considered in the network-planning process. If the latter is done based on offline simulations for failure cases anyway, the incremental cost of supporting LFAs for a defined set of demands may be relatively low.

The per-prefix mode of LFAs allows for simpler and more efficient capacity planning. As the backup path of each destination is optimized individually, the load to be fast-rerouted can be spread on a set of shortest repair paths (as opposed to a single backup tunnel). This leads to a simpler and more efficient capacity-planning process that takes congestion during protection into account.

7. Capacity Planning with LFA in Mind

We briefly describe the functionality a designer should expect from a capacity-planning tool that supports LFAs, and the related capacity-planning process.

7.1. Coverage Estimation - Default Topology

Per-Link LFA Coverage Estimation: The tool would color each unidirectional link in, depending on whether or not per-link LFAs are available.

Per-Prefix LFA Coverage Estimation: The tool would color each unidirectional link with a colored gradient, based on the percent of destinations that have a per-prefix LFA.

In addition to the visual GUI reporting, the tool should provide detailed tables that list, on a per-interface basis, the percentage of LFAs, the number of prefixes with LFAs, the number of prefixes without LFAs, and a list of those prefixes without LFAs.

Furthermore, the tool should list and provide percentages for the traffic matrix demands with less than 100% source-to-destination LFA coverage, as well as average coverage (number of links on which a demand has an LFA/number of links traversed by this demand) for every demand (using a threshold).

The user should be able to alter the color scheme to show whether these LFAs are guaranteed node-protecting or de facto node-protecting, or only link-protecting.

This functionality provides the same level of information as we described in Sections 4.1 to 4.3.

7.2. Coverage Estimation in Relation to Traffic

Instead of reporting the coverage as a ratio of the number of destinations with a backup, one might prefer a ratio of the amount of traffic on a link that benefits from protection.

This is likely much more relevant, as not all destinations are equal, and it is much more important to have an LFA for a destination attracting lots of traffic rather than an unpopular destination.

7.3. Coverage Verification for a Given Set of Demands

Depending on the requirements on the network, it might be more relevant to verify the complete LFA coverage of a given sub-topology, or a given set of demands, rather than to calculate the relative coverage of the overall traffic. This is most likely true for the third engineering profile described in Section 4.

In that case, the tool should be able to separately report the LFA coverage on a given set of demands and highlight each part of the network that does not support 100% coverage for any of those demands.

7.4. Modeling - What-If Scenarios - Coverage Impact

The tool should be able to compute the coverage for all of the possible topologies that result from a set of expected failures (i.e., one-SRLG failure).

Filtering the key information from the huge amount of generated data should be a key property of the tool.

For example, the user could set a threshold (at least 80% per-prefix LFA coverage in all one-SRLG what-if scenarios), and the tool would report only the cases where this condition is not met, hopefully with some assistance on how to remedy the problem (IGP metric optimization).

As an application example, a designer who is not able to ensure that c < a could leverage such a tool to assess the per-prefix LFA coverage for square aggregation topologies grafted to the backbone of his network. The tool would analyze the per-prefix LFA availability for each remote destination and would help optimize the backbone topology to increase the LFA protection coverage for failures within the square aggregation topologies.

7.5. Modeling - What-If Scenarios - Load Impact

The tool should be able to compute the link load for all routing states that result from a set of expected failures (i.e., one-SRLG failure).

The routing states that should be supported are 1) network-wide converged state before the failure, 2) state in which all of the LFAs protecting the failure are active, and 3) network-wide converged state after the failure.

Filtering the key information from the huge amount of generated data should be a key property of the tool.

For example, the user could set a threshold (at most 100% link load in all one-SRLG what-if scenarios), and the tool would report only the cases where this condition is violated, hopefully with some assistance on how to remedy the problem (IGP metric optimization).

The tool should be able to do this for the aggregate load, and on a per-class-of-service basis as well.

> Note: In cases where the traffic matrix is unknown, an intermediate solution consists of identifying the destinations that would attract traffic (i.e., Provider Edge (PE) routers), and those that would not (i.e., Provider (P) routers). One could achieve this by creating a traffic matrix with equal demands between the sources/destinations that would attract traffic (PE to PE). This will be more relevant than considering all demands between all prefixes (e.g., when there is no customer traffic from P to P).

7.6. Discussion on Metric Recommendations

While LFA FRR has many benefits (Section 6), LFA FRR's applicability depends on topology.

The purpose of this document is to show how to introduce a level of control over this topology parameter.

On the one hand, we wanted to show that by adopting a small set of IGP metric constraints and a repetition of well-behaved patterns, the designer could deterministically guarantee maximum link and node protection for the vast majority of the network (the access/ aggregation). By doing so, he would obtain an extremely simple resiliency solution.

On the other hand, we also wanted to show that it might not be so bad to not apply (all of) these constraints.

Indeed, we explained in Section 3.3.4.3 that the per-prefix LFA coverage in a square where c >= a might still be very good, depending on the backbone topology.

We showed in Section 4.3 that the median per-prefix LFA coverage for 11 SP backbone topologies still provides 94% coverage. (Most of these topologies were built without any idea of LFA!)

Furthermore, we showed that any topology may be analyzed with an LFA-aware capacity-planning tool. This would readily assess the coverage of per-prefix LFAs and would assist the designer in fine-tuning it to obtain the level of protection he seeks.

While this document highlights LFA applicability and benefits for SP networks, it also notes that LFAs are not meant to replace MPLS TE FRR.

With a very LFA-unfriendly topology, a designer seeking guaranteed <50-msec protection might be better off leveraging the explicit-routed backup capability of MPLS TE FRR to provide 100% protection while ensuring no congestion along the backup paths during protection.

But when LFAs provide 100% link and node protection without any uLoop, then clearly the LFA method seems a technology to consider to drastically simplify the operation of a large-scale network.

8. Security Considerations

The security considerations applicable to LFAs are described in [RFC5286]. This document does not introduce any new security considerations.

9. Conclusions

The LFA method is an important protection alternative for IP/MPLS networks.

Its simplicity benefit is significant, in terms of automation and integration with the default IGP behavior and the absence of any requirement for network-wide upgrade. The technology does not require any protocol change and hence can be deployed router by router.

At first sight, these significant simplicity benefits are negated by the topological dependency of its applicability.

The purpose of this document is to highlight that very frequent access and aggregation topologies benefit from excellent link and node LFA coverage.

A second objective consists of describing the three different profiles of LFA applicability for the IP/MPLS core networks and illustrating them with simulation results based on real SP core topologies.

10. Acknowledgments

We would like to thank Alvaro Retana and especially Stewart Bryant for their valuable comments on this work.

11. References

11.1. Normative References

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

11.2. Informative References

[RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, January 2010.

[RFC5715] Shand, M. and S. Bryant, "A Framework for Loop-Free Convergence", RFC 5715, January 2010.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.

[IS-IS] ISO/IEC 10589:2002, Second Edition, "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 2002.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

Authors' Addresses

Clarence Filsfils (editor)
Cisco Systems
Brussels 1000
BE

EMail: cf@cisco.com

Pierre Francois (editor)
Institute IMDEA Networks
Avda. del Mar Mediterraneo, 22
Leganese 28918
ES

EMail: pierre.francois@imdea.org

Mike Shand

EMail: imc.shand@googlemail.com

Bruno Decraene
France Telecom
38-40 rue du General Leclerc
92794 Issy Moulineaux cedex 9
FR

EMail: bruno.decraene@orange.com

James Uttaro
AT&T
200 S. Laurel Avenue
Middletown, NJ 07748
US

EMail: uttaro@att.com

Nicolai Leymann
Deutsche Telekom
Winterfeldtstrasse 21
10781, Berlin
DE

EMail: N.Leymann@telekom.de

Martin Horneffer
Deutsche Telekom
Hammer Str. 216-226
48153, Muenster
DE

EMail: Martin.Horneffer@telekom.de