



Loop Free Alternates for Unicast SPBM

David Allan, János Farkas, Panagiotis Saltsidis,
András Császár, Jeff Tantsura

May 15, 2013




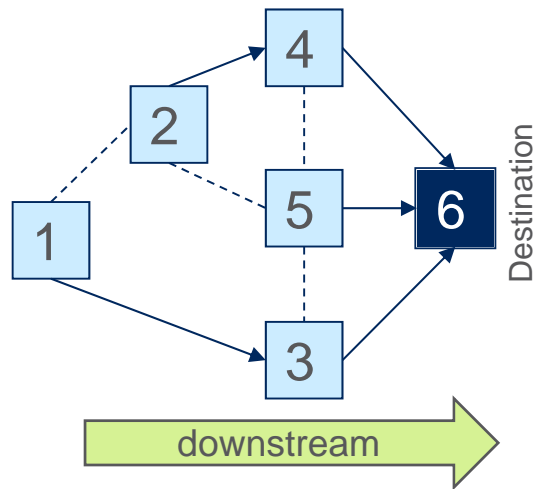
Introduction

- › Leveraging Loop Free Alternate (LFA) paths by means of local repair, i.e. Fast Re-Route (FRR), provides very fast failover
- › Fault handling in a few milliseconds
- › Very simple
- › Effective
- › Easy add-on to distributed control protocols
- › LFA FRR was introduced to packet networks by RSTP
 - See the Alternate Port
- › **We should leverage LFA for SPB too!**

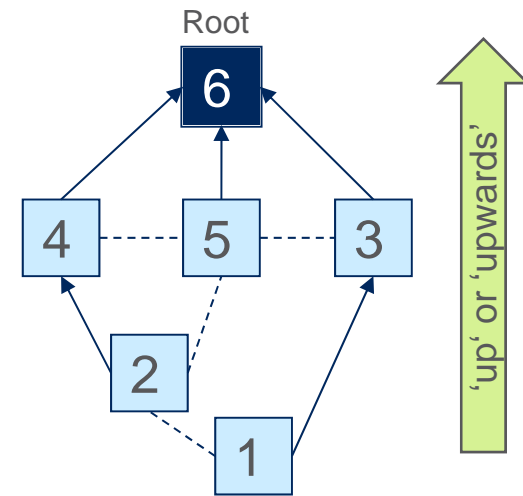
Background




a destination rooted Shortest Path Tree (SPT)



the same SPT as a spanning tree



PLR: Point of Local Repair





Loop-free Forwarding

- › Conditions for loop-free forwarding are specified by M. Seaman, “Link state agreement”
<http://www.ieee802.org/1/files/public/docs2010/aq-seaman-link-state-0910-v4.pdf>
 - Unicast forwarding: Equation 1 in Section 3.1
 - › Forwarding is loop-free even during topology changes if the forwarding node is
 - a) not closer to the destination than its downstream neighbour
 - b) not farther from the destination than any of its upstream neighbours
- › Criteria given by RFC 5286 <http://tools.ietf.org/html/rfc5286>
 - Loop-Free Criterion
 - › A **Neighbour N** can provide a loop-free alternate (LFA) to a **Destination D** with respect to a **Calculating Node S** if and only if
$$Distance(\mathbf{N}, \mathbf{D}) < Distance(\mathbf{N}, \mathbf{S}) + Distance(\mathbf{S}, \mathbf{D})$$
 - Downstream Path Criterion
 - › $Distance(\mathbf{N}, \mathbf{D}) < Distance(\mathbf{S}, \mathbf{D})$
- › **Downstream alternate paths are always loop-free!**




RSTP/MSTP

○ Root, Forwarding


● Designated, Forwarding

|| Alternate, Discarding



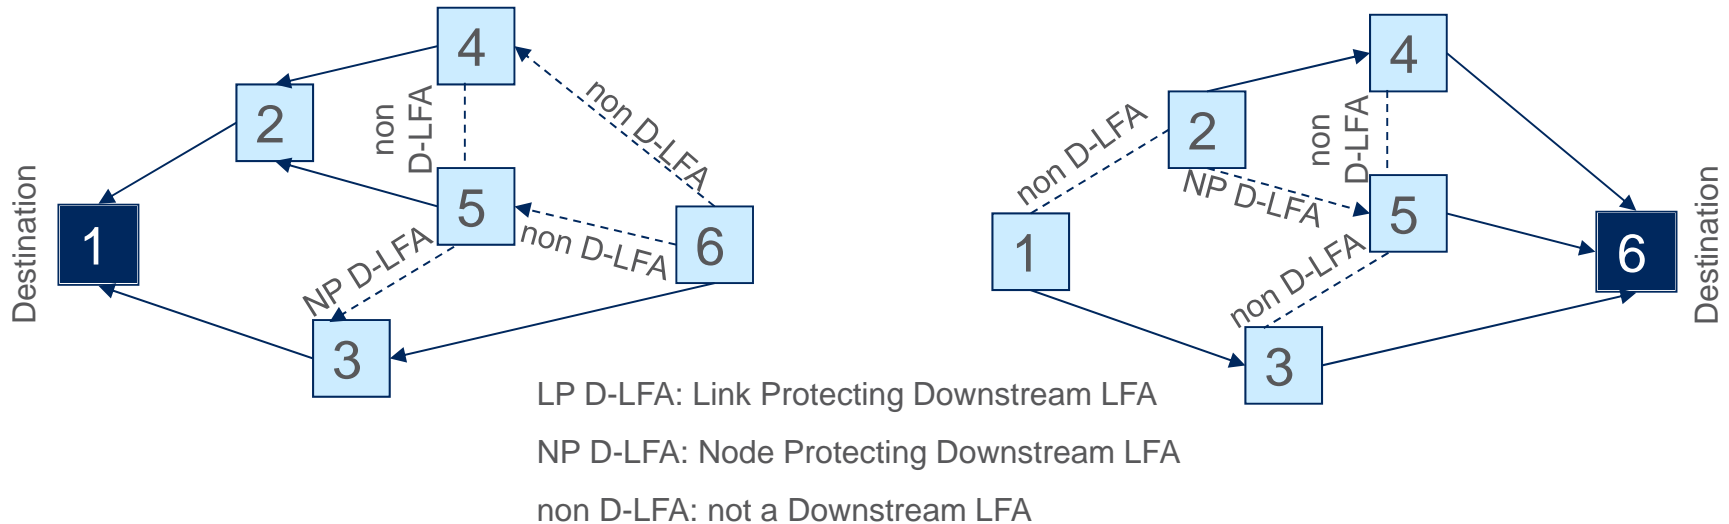
- › Forwarding is along a (non-directed) SPT rooted at the Root Bridge
- › RSTP implements downstream LFA by means of Alternate Ports
 - If connectivity is lost at a Root Port, then the Alternate Discarding port becomes Root Forwarding (e.g. 2-5 link goes down, then the formerly Alternate Discarding port of 5 towards 4 becomes Root Forwarding)
 - It is just a local repair! (the only action is at bridge 5 in the example)
- › Loop prevention
 - Alternate Ports do not cause loop
 - Proposal-Agreement handshake during restoration
- › Symmetry is kept as there is only one active path between any node pair

SPBV



- › SPTs can be implemented by the per tree Port State and Port Role variables (13.17)
 - Alternate Port based LFA is just there
- › Loop prevention
 - Alternate Ports do not cause loop
 - Agreement Protocol during restoration
- › Symmetry is maintained if the same ECT-Algorithm is used for the selection of the Alternate Ports as for the computation of the SPT

IP Fast Re-Route (FRR) – Downstream LFAs





- › IP forwarding is based on destination rooted directed trees
 - The Shortest Path First (SPF) next hops span a destination rooted SPT
- › Loops are avoided by means of only using loop-free alternate paths
- › Symmetry is not required
 - Even the cost of the same link may differ in the two directions

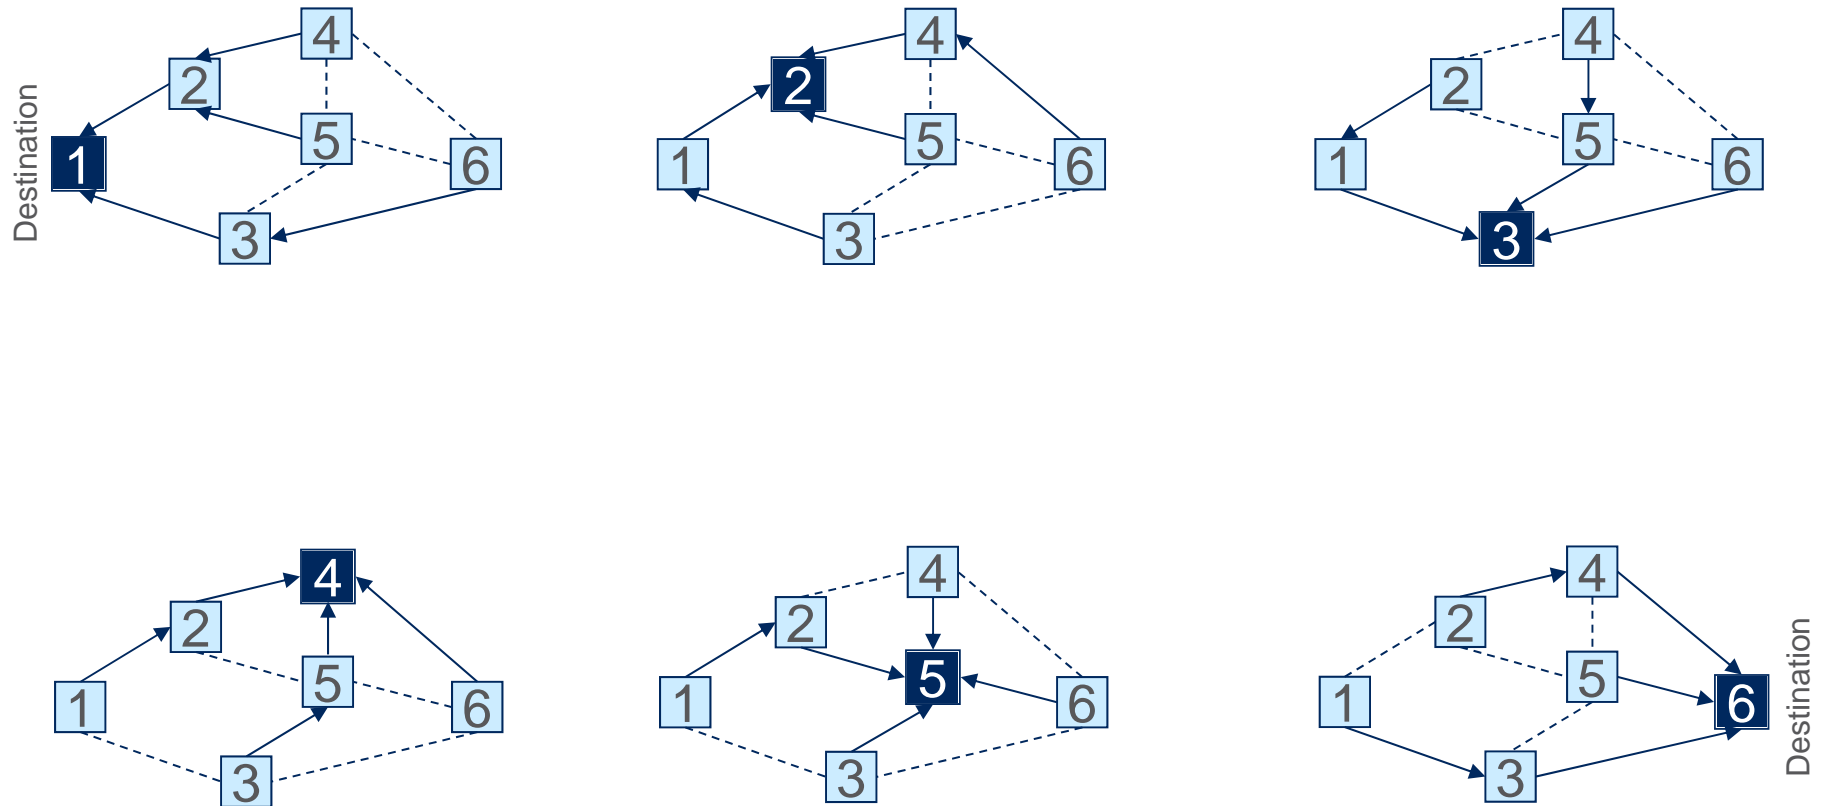
SPBM Unicast Forwarding




- › Destination rooted SPTs (as in case of IP)
- › Unicast loops are mitigated by ingress checking (also referred to as RPFC: Reverse Path Forwarding Check)
- › Symmetry during basic forwarding is ensured by ISIS-SPB




SPBM Unicast Example – Destination rooted SPTs




SPBM Unicast Example – Downstream LFA for Link 2-4 Failure




- nothing can be done
(4 has no downstream LFA to 1)




- nothing can be done
(2 has no downstream LFA to 2)




- nothing to do



- 2 has to install state
- 5 has to admit 1→4, 2→4 frames




- nothing to do




- 2 has to install state
- 5 has to admit 2→6 frames


SPBM Unicast Example – Downstream LFA for Link 1-2 Failure




- nothing can be done
(2 has no downstream LFA to 1)




- nothing can be done
(1 has no downstream LFA to 2)




- 2 has to install state
- 5 has to admit 2→3 frames



- 1 has to install state
- 3 has to admit 1→4 frames




- 1 has to install state
- 3 has to admit 1→5 frames




- nothing to do


SPBM Unicast Example – Downstream LFA for Node 2 Failure




- 4 and 5 have to install state
- 5 has to admit 4→1 frames
- 3 has to admit 4→1, 5→1 frames




- nothing to do




- 1 has to install state



- 1 has to install state
- 3 has to admit 1→4 frames



- 1 has to install state
- 3 has to admit 1→5 frames



- nothing to do



ERICSSON

Proposal: Add LFA to SPBM for Unicast Traffic

What Needs to Be Done for SPBM Unicast?

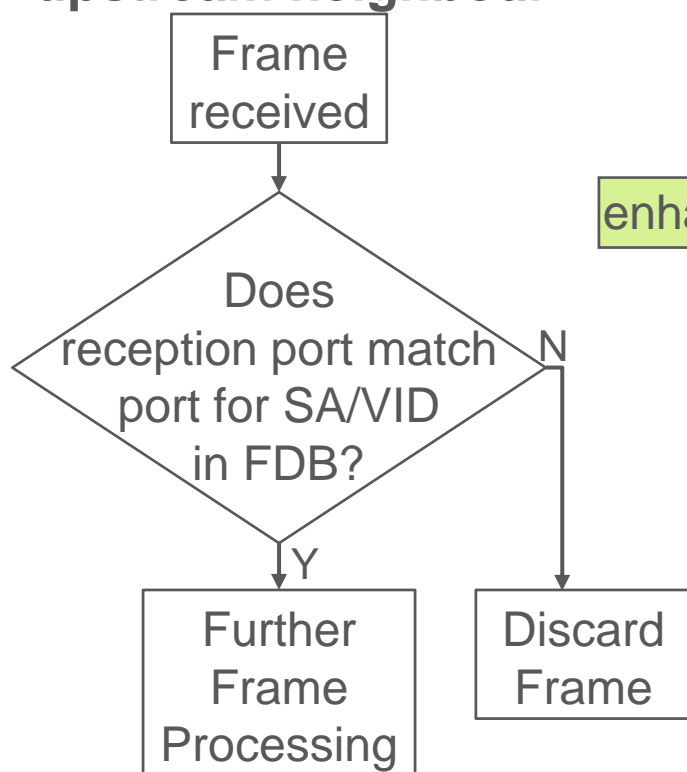


- › Enhance ingress checking = extend, relax it
 - Frame reception from all upstream neighbours has to be allowed
 - It remains loop-free as the loop-free conditions are still met
- › Ensure safe updates after a topology change
 - Change detected
 - › Update ingress checking to block unsafe reception
 - › Remove unsafe LFAs
 - Re-convergence is over
 - › Install states for primary paths
 - › Install states for new LFA path
 - › Update ingress checking to allow reception from upstream neighbours

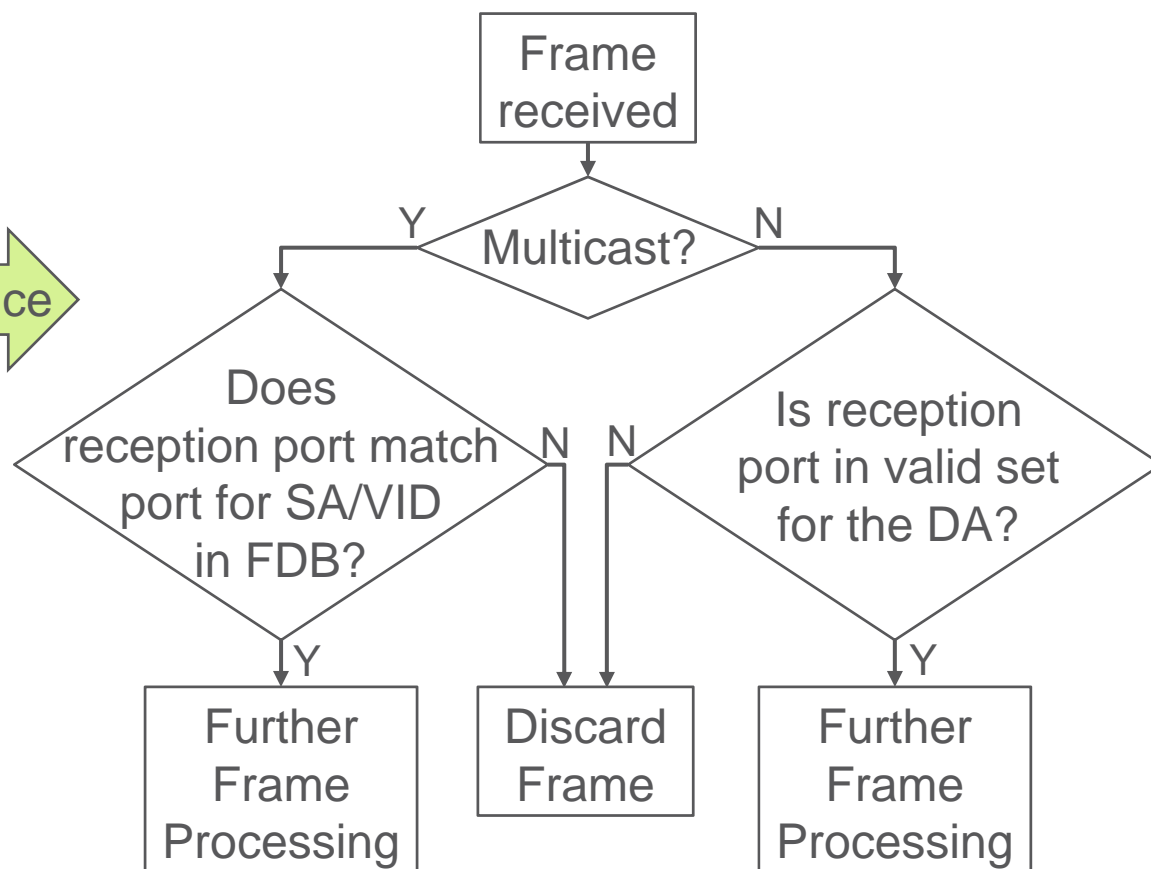


Ingress Checking

Reception from single upstream neighbour



Unicast frame admitted from any upstream neighbour



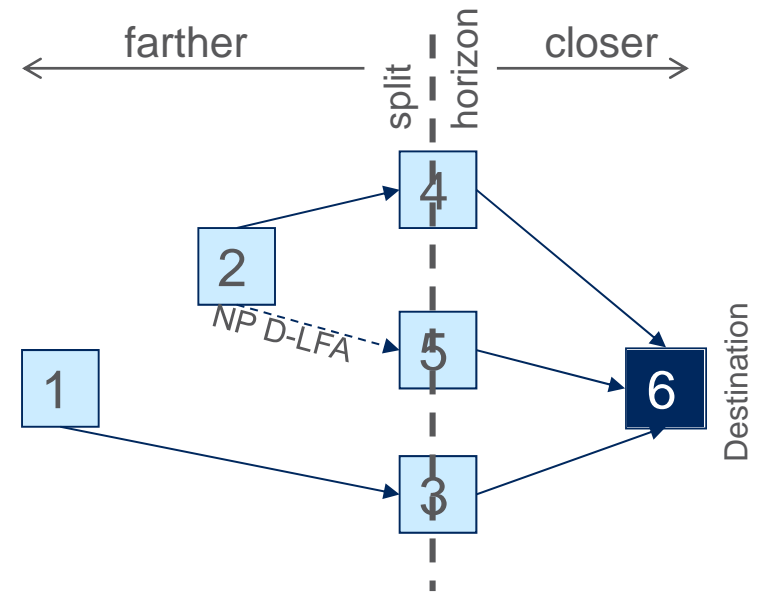
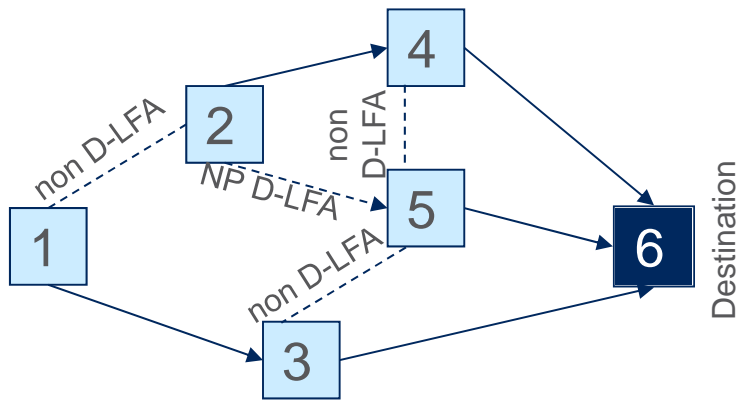
Port for SA/VID in FDB is the port providing the shortest path to SA

The valid set of ports for a DA is comprised of all upstream neighbours with respect to the given DA

Enhanced Ingress Check Provides Split Horizon



- › Enhanced ingress checking still provides loop mitigation
- › It is in fact a split horizon



Only frames coming from farther, i.e. received from upstream neighbors are admitted 'through' the horizon by enhanced ingress checking

LFA Updates After A Topology Change



› After topology change

1. Detecting mismatch in the Agreement Digest
2. Update ingress checking: remove ports that have become invalid
3. Update LFA paths: remove paths that have become non-downstream

› After re-convergence

1. Multicast states have been updated
 - › Installing multicast states is the last step, hence everything else has been done including All Pairs Shortest Path, which gives LFA paths too
2. Update LFA paths: add new LFA paths
3. Updated ingress checking: add ports to new upstream neighbors to the valid set

Summary



- › We can easily leverage LFA for SPB
- › Ingress checking (RPFC) has to be enhanced for SPBM unicast