IETF 116 – 6lo

Reliability Considerations of Path-Aware Semantic Addressing

draft-li-6lo-pasa-reliability

IETF 116 – Yokohama

Since IETF 115

Reliability Considerations of Path-Aware Semantic Addressing

draft-li-6lo-pasa-reliability-01

Status | IESG evaluation record | IESG writeups | Email expansions | History

Versions:

00 01

New Revision: draft-li-6lo-pasa-reliability-01.txt
March 2023

Main Content

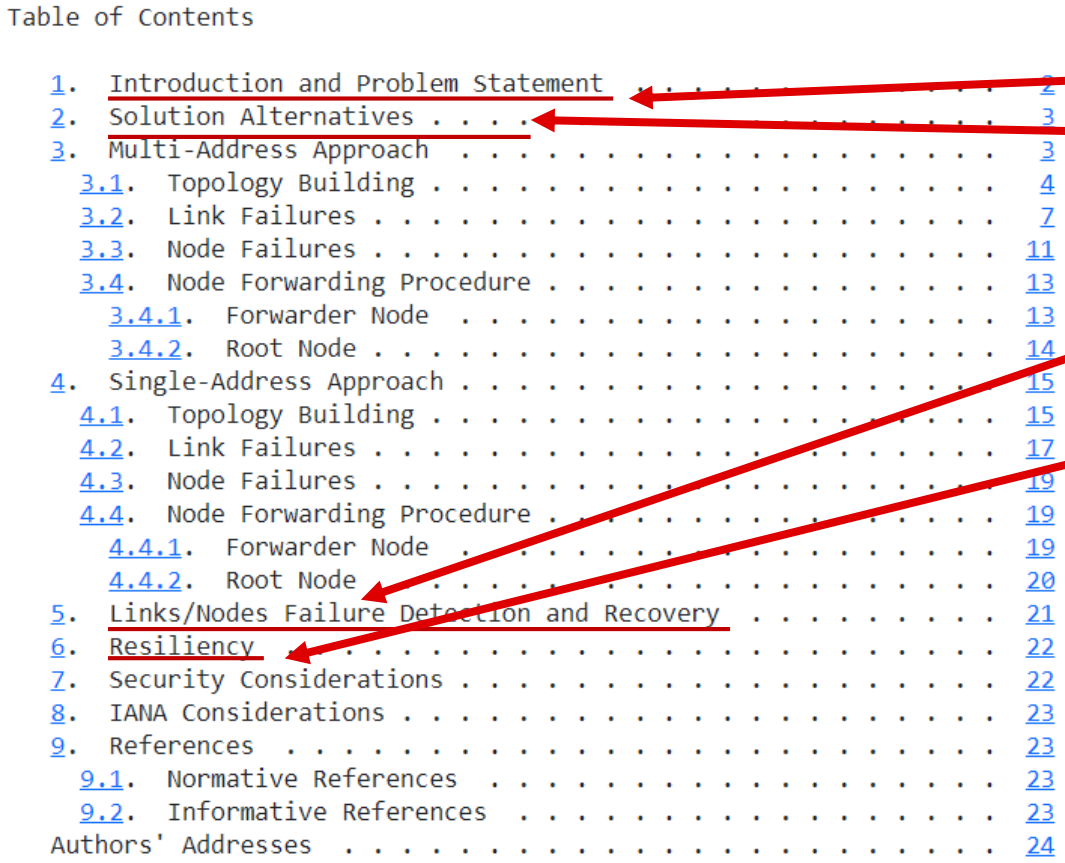

Table of Contents

Main content/structure not changed:

- A general introduction to the problem
- Possible solution alternatives
  - See rest of the presentation
- Some considerations about failure detection and recovery
- Some considerations about resiliency

- But quite some:
  - Rewriting for clarification
  - Fix examples bugs
  - Align text with main specs

Solutions Alternatives

- Pre-requisite
  - Presence of redundant links
  - Some not actively used
  - Active links forming a tree
  - All nodes have a secondary parent
    - Except root node
    - Alternative parent is connected through a non actively used link

- Multi-Address:
  - using multiple addresses per node, one for each alternative parent (logically creating multiple topologies)

- Single-Address:
  - using one single address per node, store the addresses of alternative parents/children

Multi-address: Primary Topology

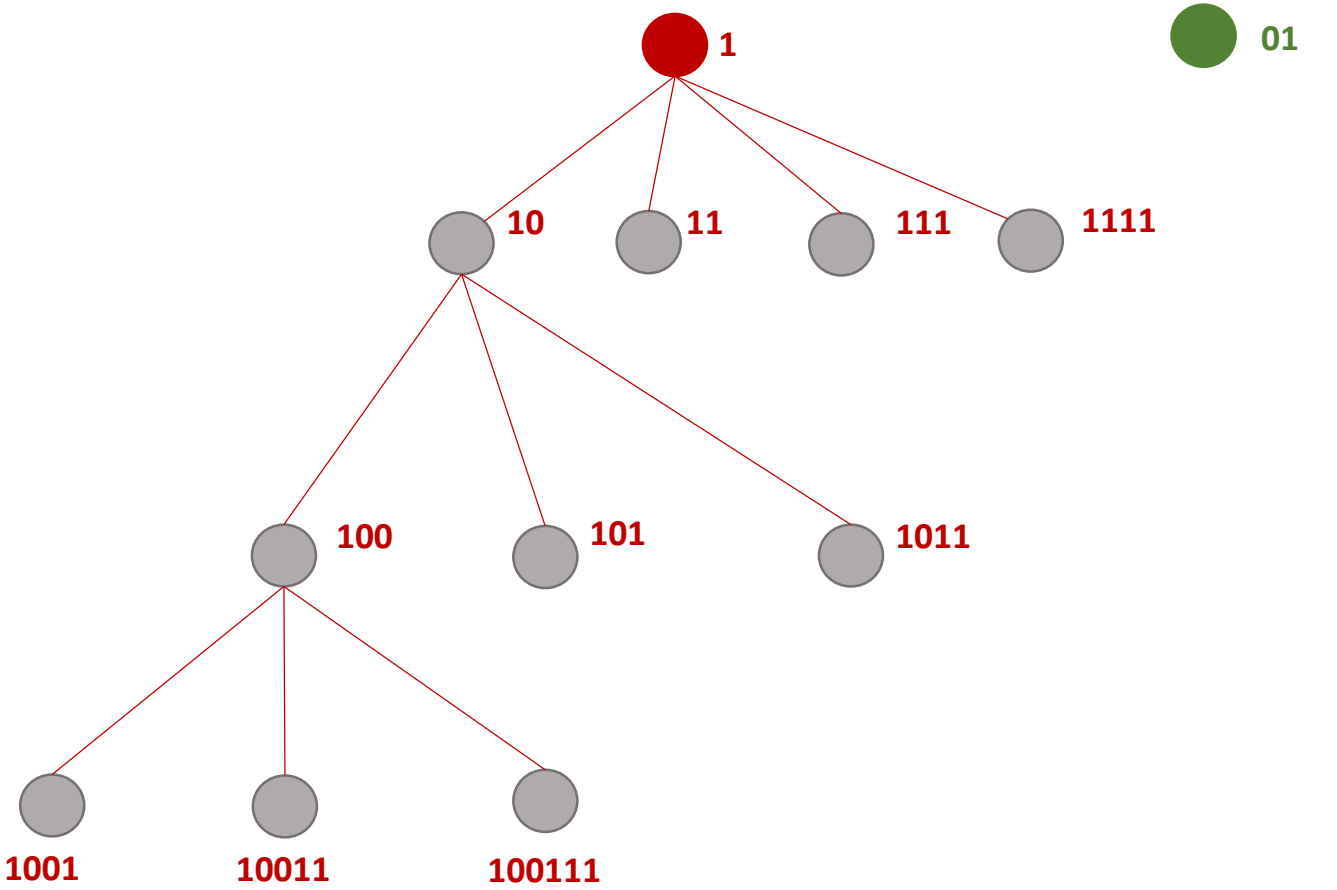

xxx1: PASA Host
xxx0: PASA Router

Multi-address: Secondary Topology

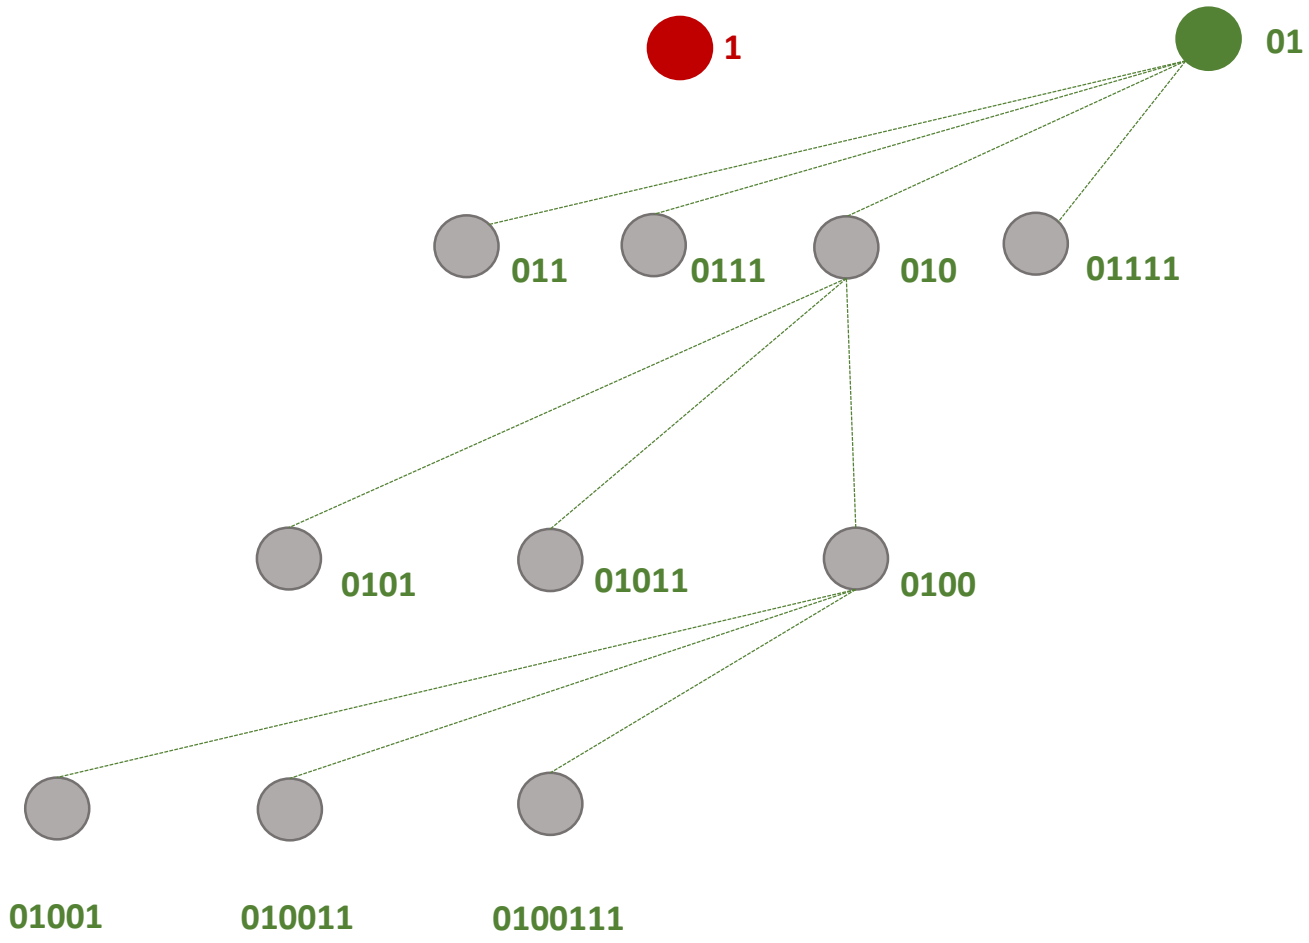

xxx1: PASA Host
xxx0: PASA Router

Multi-address: Overall Topology

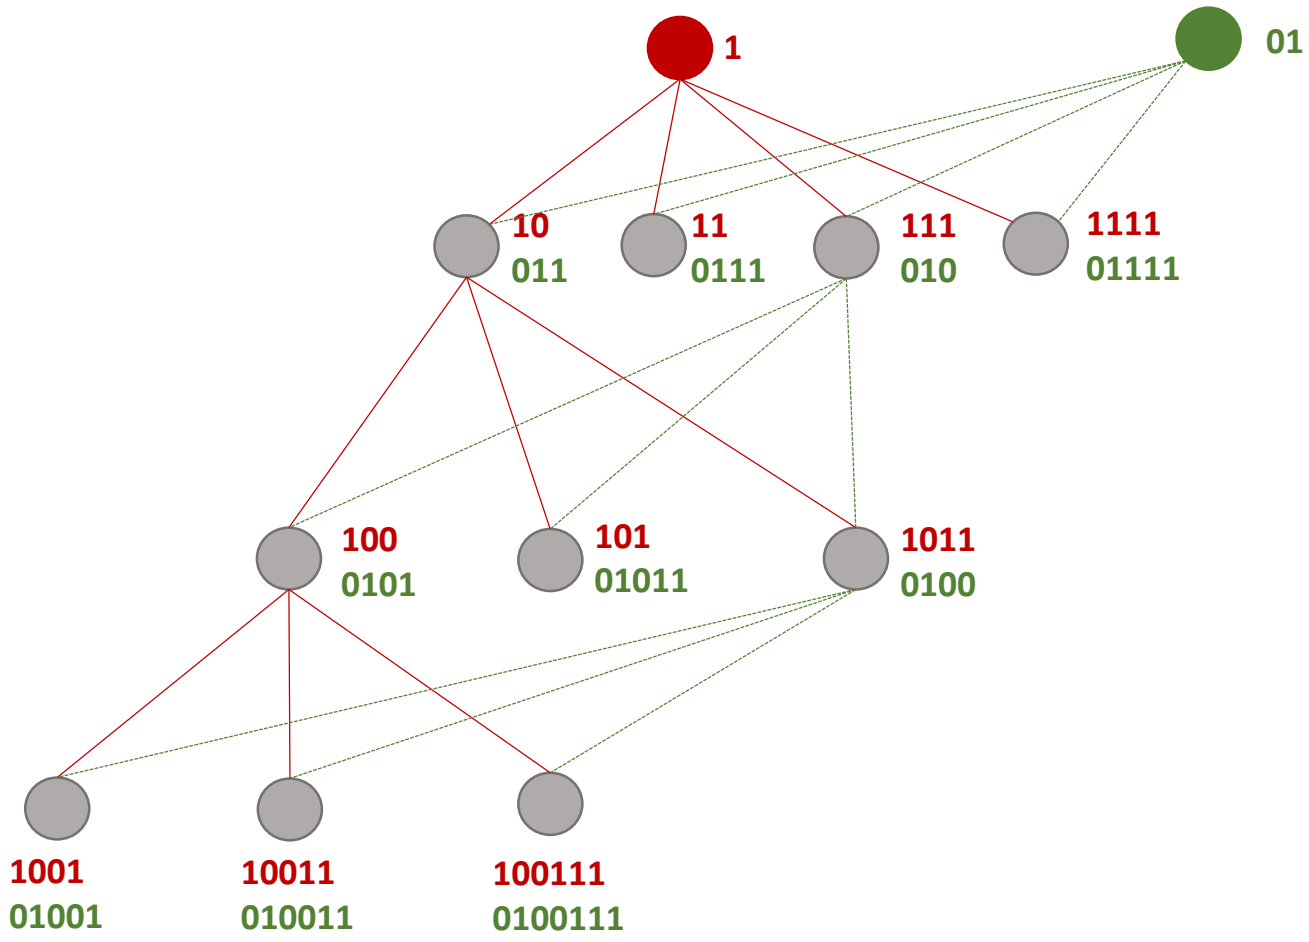

xxx1: PASA Host
xxx0: PASA Router

- Each node in the PASA only has two addresses built using the same AF but different roots (1 and 01)
- Each node except the root has at least 2 parents (primary and secondary)
- Forwarding done following only primary address (and topology)
- Each node maintains a redirect table if something has to be tunneled via secondary topology

Multi-address: Link Failure

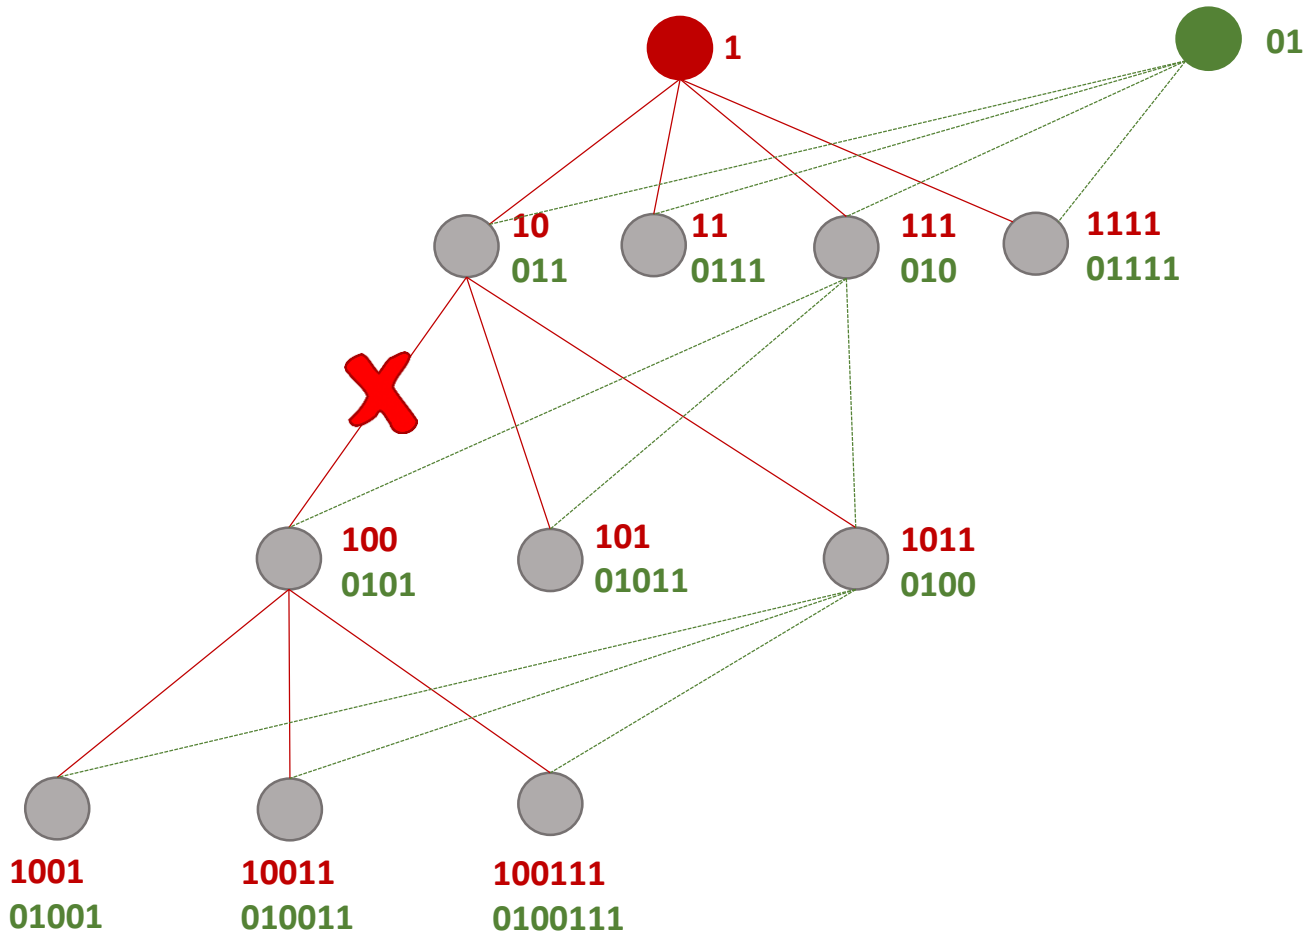

xxx1: PASA Host
xxx0: PASA Router

- Each node in the PASA only has two addresses built using the same AF but different roots (1 and 01)
- Each node except the root has at least 2 parents (primary and secondary)
- Forwarding done following only primary address (and topology)
- Each node maintains a redirect table if something has to be tunneled via secondary topology

Multi-address: Link Failure

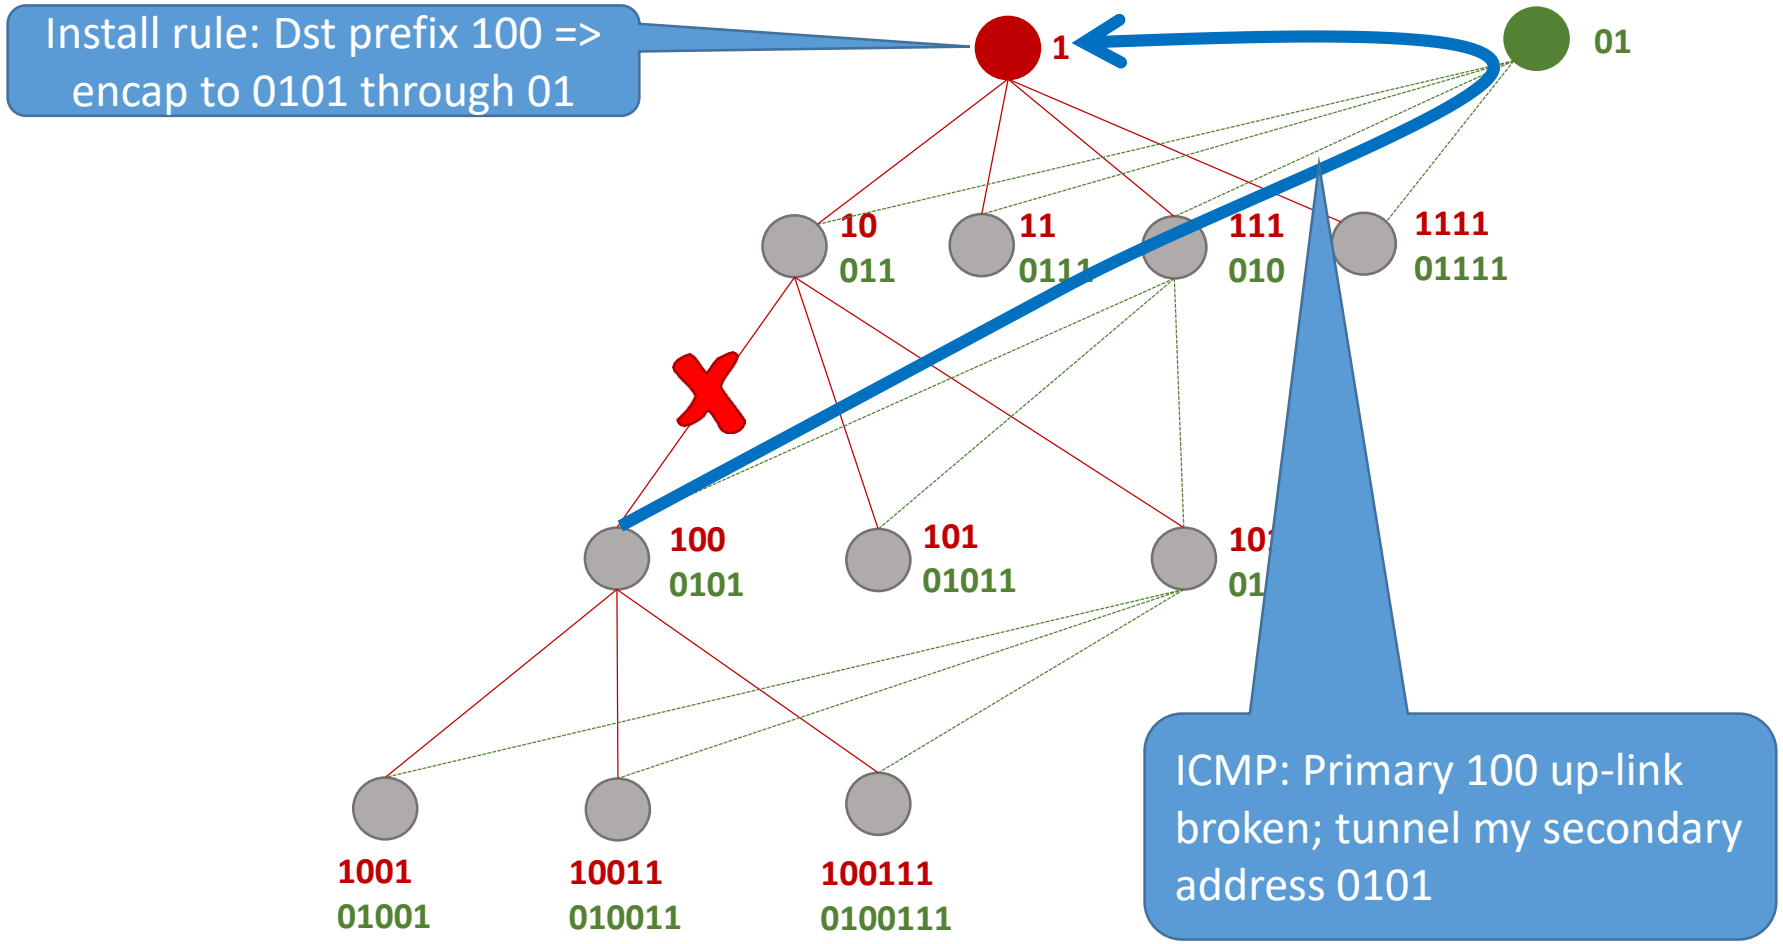

- Downstream node (w.r.t. failure) 100 to notify root to send through secondary topology
  - Either via encapsulation
  - Or routing state created by ICMP message

xxx1: PASA Host
xxx0: PASA Router

Multi-address: Link Failure

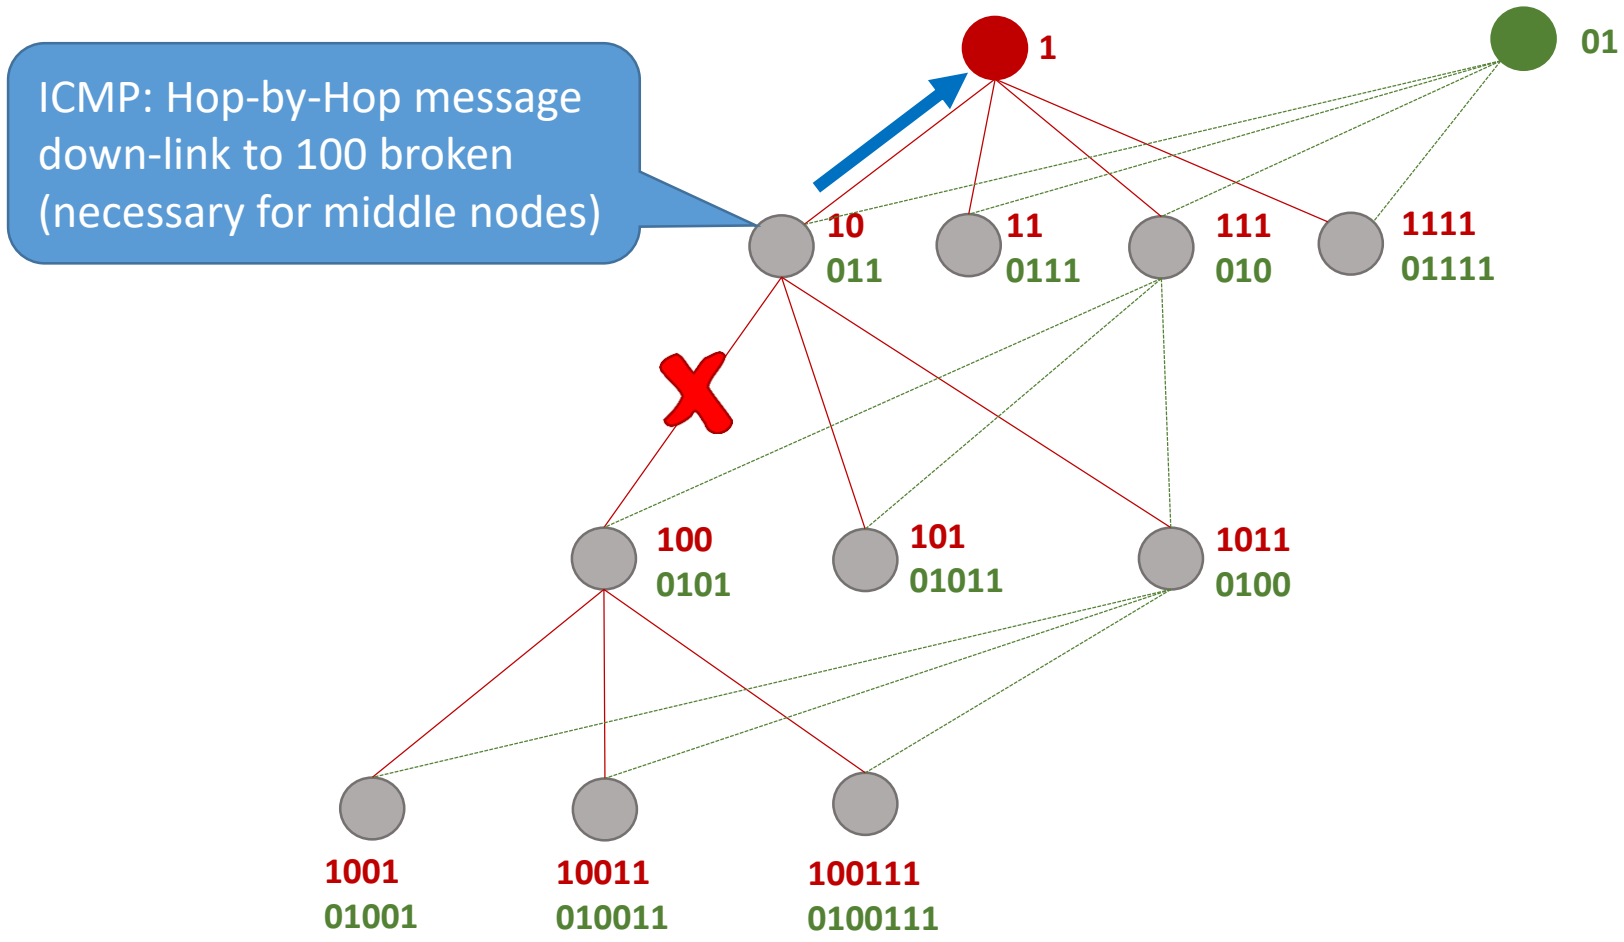

- Downstream node (w.r.t. failure) 100 to notify root to send through secondary topology
  - Either via encapsulation
  - Or routing state created by ICMP message
- Upstream node (w.r.t. failure) 10 notify broken link so that all nodes between the root and itself will install a redirect rule to send to the root

xxx1: PASA Host
xxx0: PASA Router

Multi-address: Link Failure

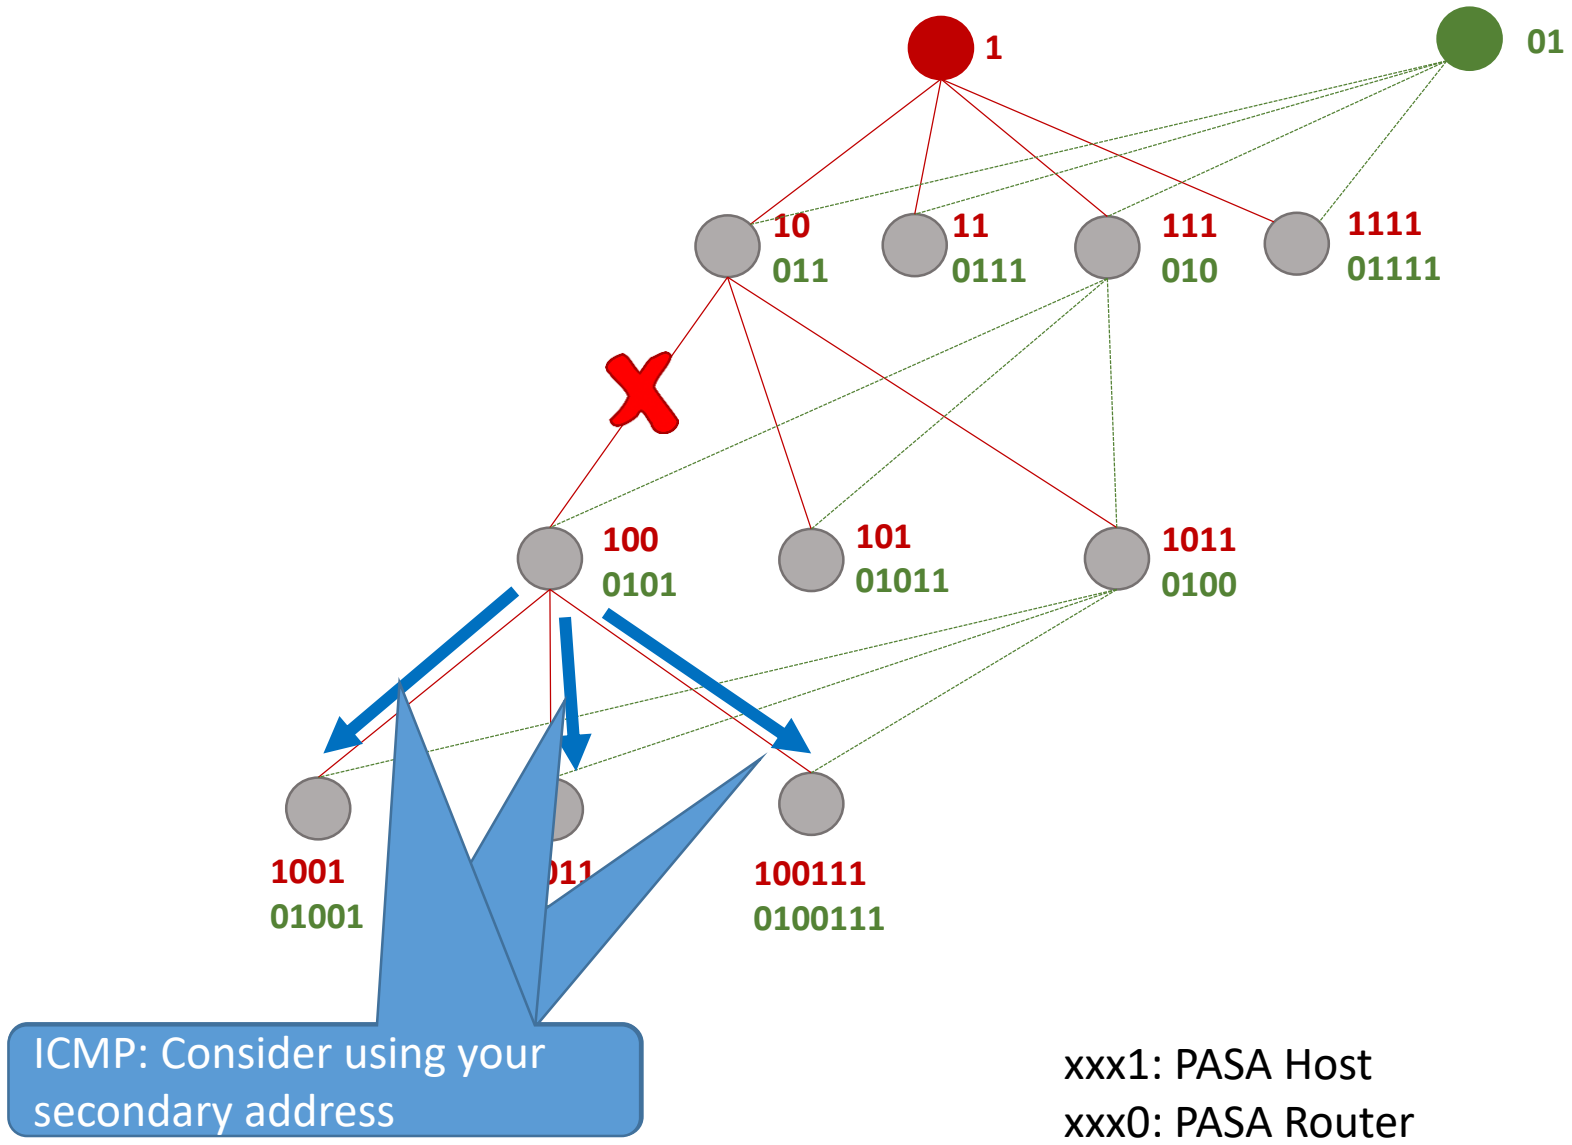

- Downstream node (w.r.t. failure) 100 to notify root to send through secondary topology
  - Either via encapsulation
  - Or routing state created by ICMP message
- Upstream node (w.r.t. failure) 10 notify broken link so that all nodes between the root and itself will install a redirect rule to send to the root
- Optimization: Downstream node to notify sub-tree to use secondary topology

Multi-address: Communication in case of Link Failure

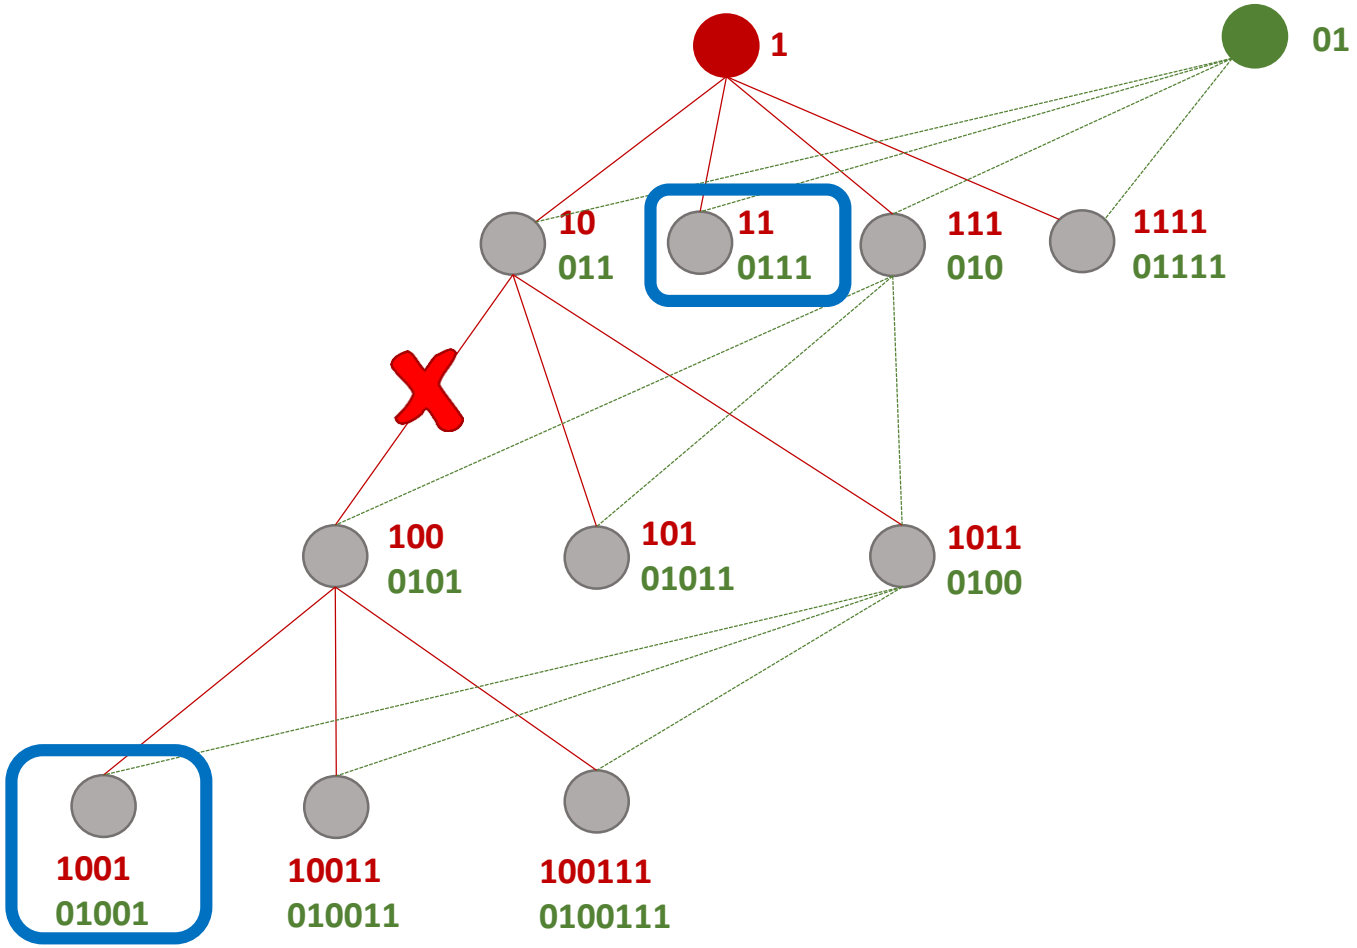

From 11 to 1001:

1. Packet is transmitted from 11 to 1 (on the primary tree).
2. Because of the redirect entry, 1 encapsulates packet toward 0101 and transmits it to 01 (root secondary).
3. 01 will use PASA stateless forwarding to transmit the packet to 010 (on the secondary tree).
4. 010 will use PASA stateless forwarding to transmit the packet to 0101 (on the secondary tree).
5. 0101 will decapsulate, note the destination is on the primary tree, use the PASA stateless forwarding to transmit the packet to 1001 (on the primary tree).

From 1001 to 11:

1. Packet is transmitted from 1001 to 100 (on the primary tree).
2. Because 100 knows the upstream link is broken it encapsulates the packet with source 0101 and destination 01 (root primary tree) then transmits the packet to 010 (on the secondary tree).
3. 010 will use PASA stateless forwarding to transmit the packet to 01 (on the secondary tree).
4. 01 will decapsulate and see that packet is destined a node in the primary tree and transmits it to 1.
5. 1 will use the PASA stateless forwarding to transmit the packet to 11 (on the primary tree).

Multi-Address: Forwarding Operation

PASA Router

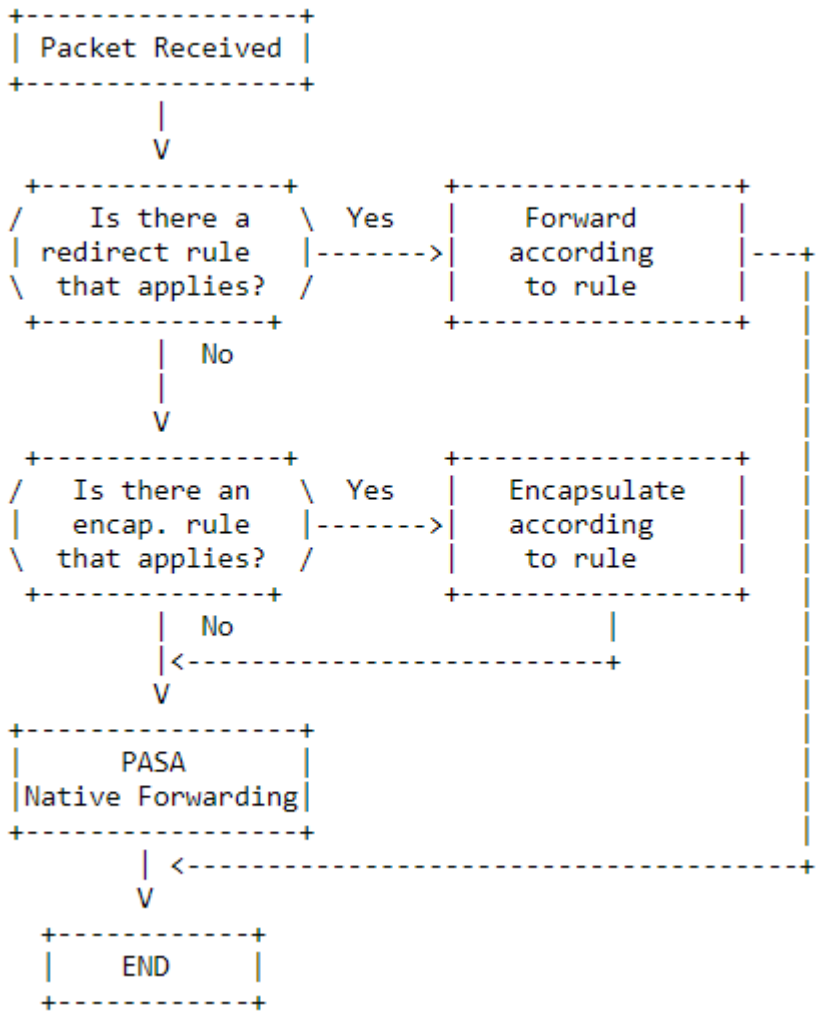

PASA Root

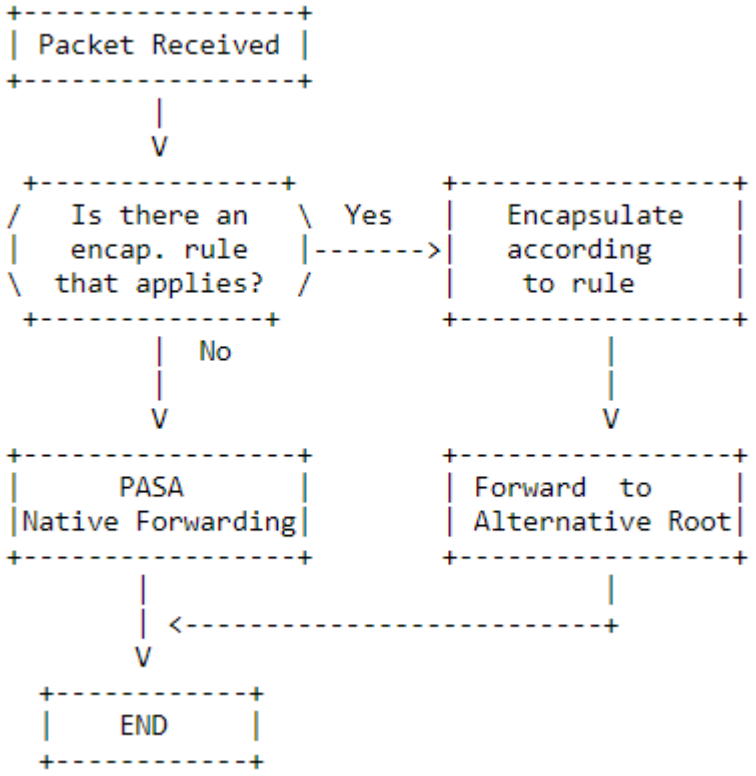

Single Address: Topology Building

- **Single PASA domain + Locally store alternative parents/children**

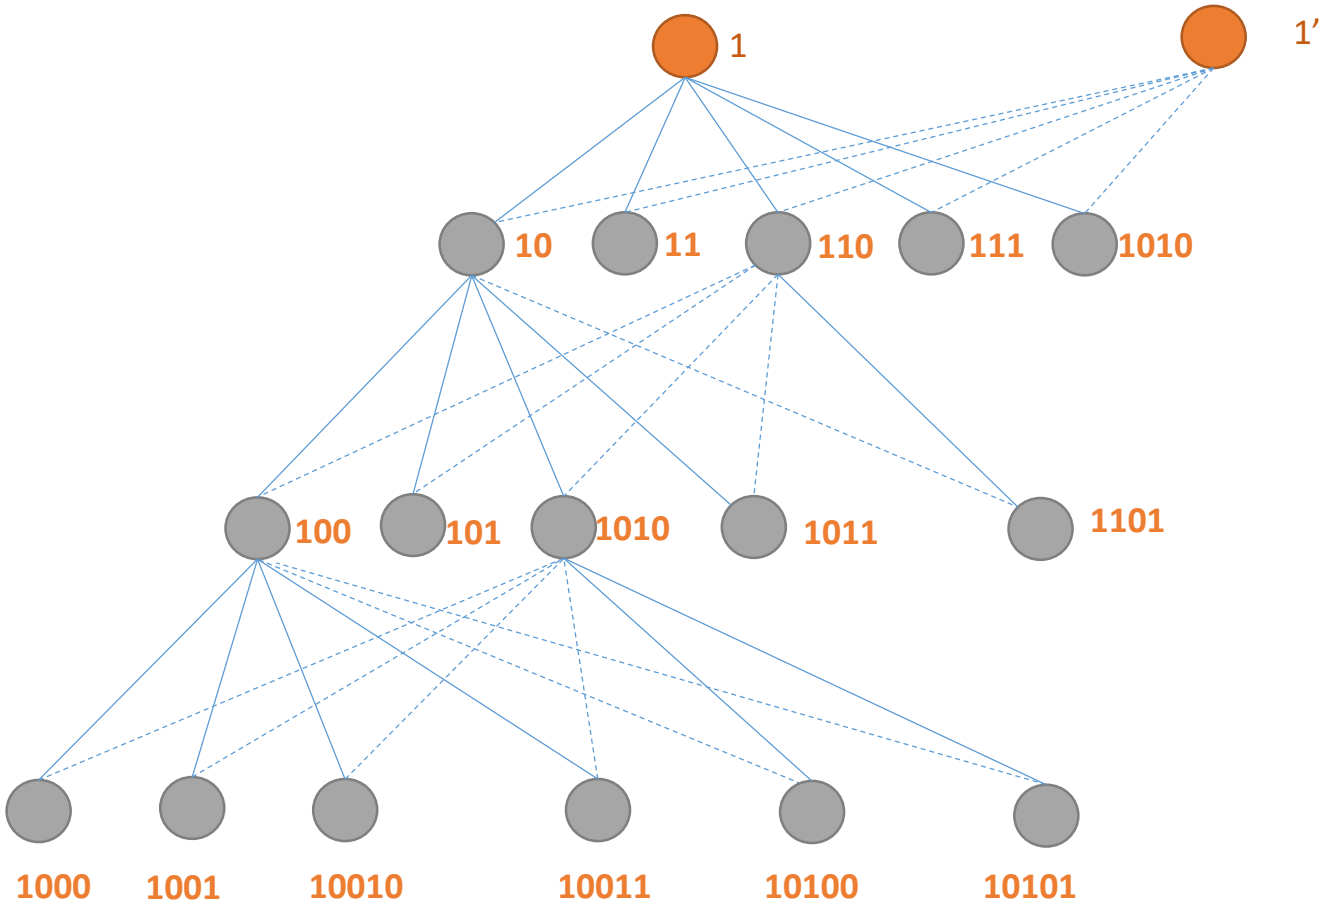

xxx1: PASA Host
xxx0: PASA Router

Single Address: Topology Building

- **Locally store alternative parents/children**

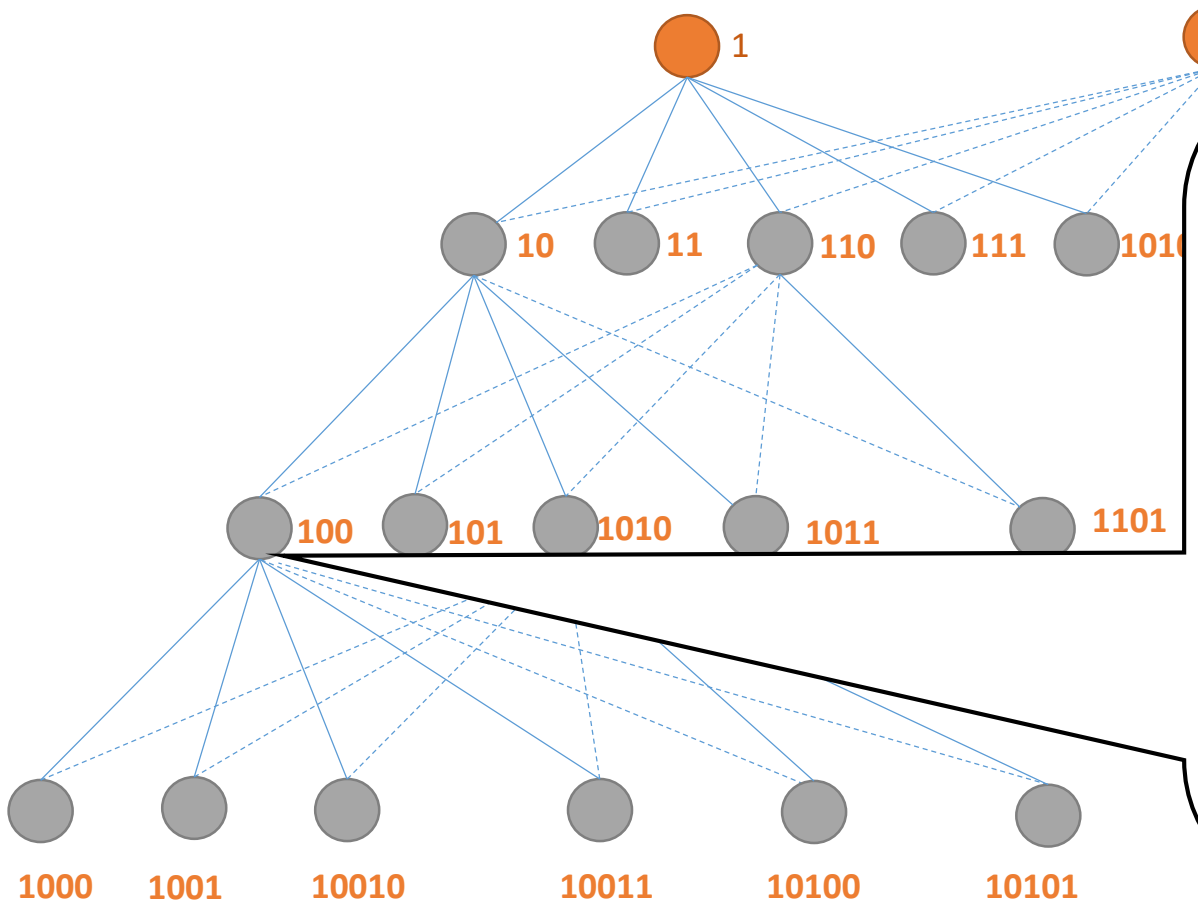

Destination	Flags	
100	I	I = Current Node
10	PP	PP = Primary Parent
1000	PFC	PFC = Primary Forwarder Child
10010	PFC	
1001	PLC	PLC = Primary Leaf Child
10011	PLC	
110	AP	AP = Alternative Parent
10100	AFC	AFC = Alternative Forwarder Child whose alternative parent is the current node
10101	ALC	ALC = Alternative Leaf Child whose alternative parent is the current node

xxx1: PASA Host
xxx0: PASA Router

Single Address: Link Failure

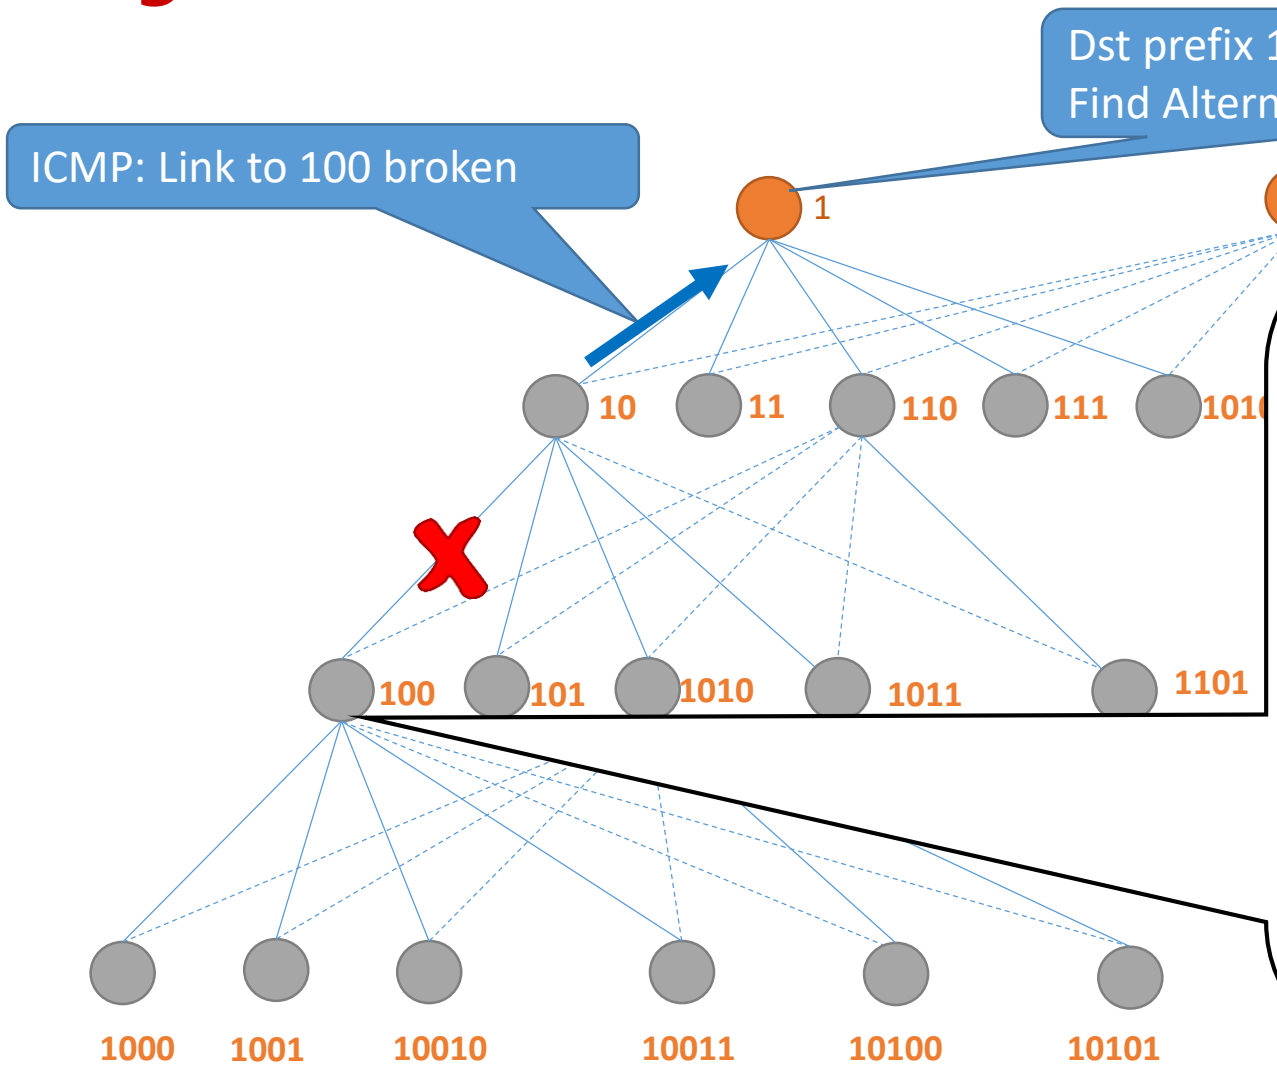

Destination	Flags	
100	I	I = Current Node
~~10~~	~~PP~~	PP = Primary Parent
1000	PFC	PFC = Primary Forwarder Child
10010	PFC	
1001	PLC	PLC = Primary Leaf Child
10011	PLC	
110	AP	AP = Alternative Parent
10100	AFC	AFC = Alternative Forwarder Child whose alternative parent is the current node
10101	ALC	ALC = Alternative Leaf Child whose alternative parent is the current node

xxx1: PASA Host
xxx0: PASA Router

Single Address: Communication in case of Link Failure

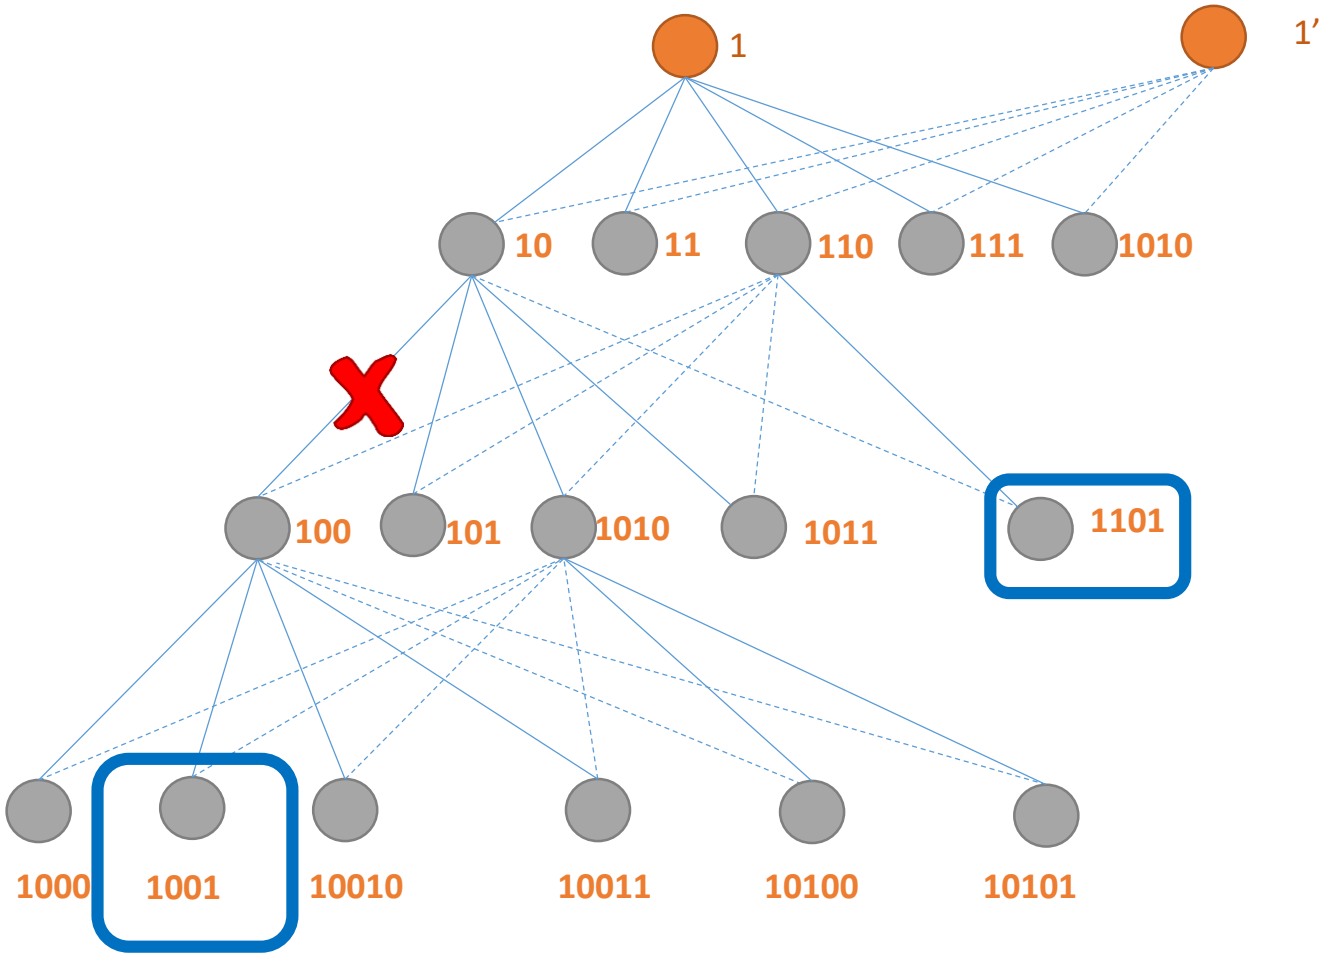

From 1101 to 1001:

1. The packet is transmitted from node 1101 to the root 1, using PASA stateless forwarding.
2. Root 1 computing an alternative path 1 -> 110 -> 100 -> 1001.
   - PASA forwarding does not allow to go first through node 110 and then node 100, the root 1 will encapsulate the addresses of node 110 and node 100 in an extension header so to perform segmented routing.
3. Once the packet reaches 100, the segment routing extension is dropped, and the packet is sent to its destination 1001 by using PASA native forwarding.

From 1001 to 1101:

1. The packet is transmitted from node 1001 to node 100 using PASA native forwarding.
2. Because of the failed link, node 100 sends the packet to an alternative parent node.
3. PASA native forwarding is used then. In this example, the resulting path is: 100 -> 110 -> 1101.
   - If the alternative parent is in the same sub-tree like the destination, the packet is forwarded downward to the right child, otherwise it is sent upward to the its own parent. This goes on recursively until the packet reaches the root in the worst case, where it is then sent downward to the correct forwarder child, until it reaches the destination

xxx1: PASA Host
xxx0: PASA Router

Single-Address: Forwarding Operation

PASA Router

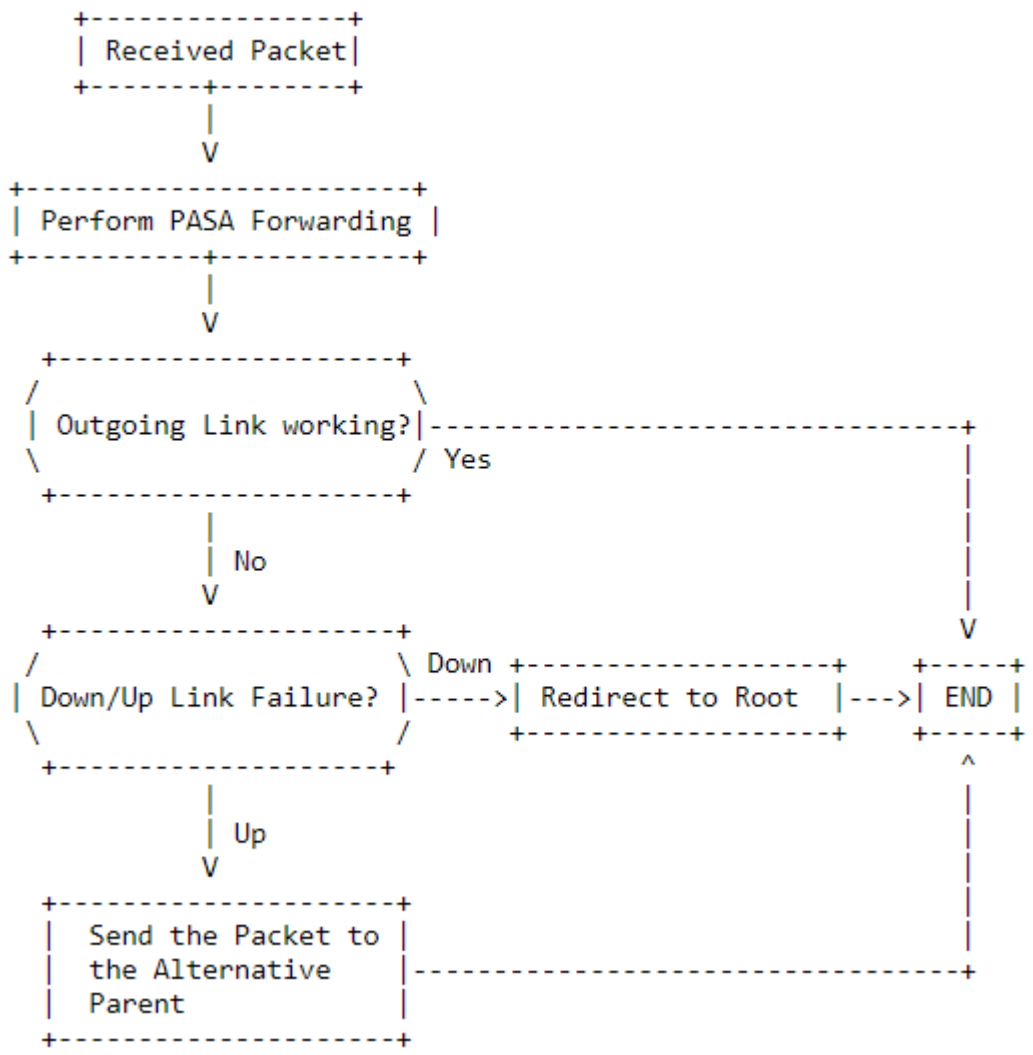

PASA Root

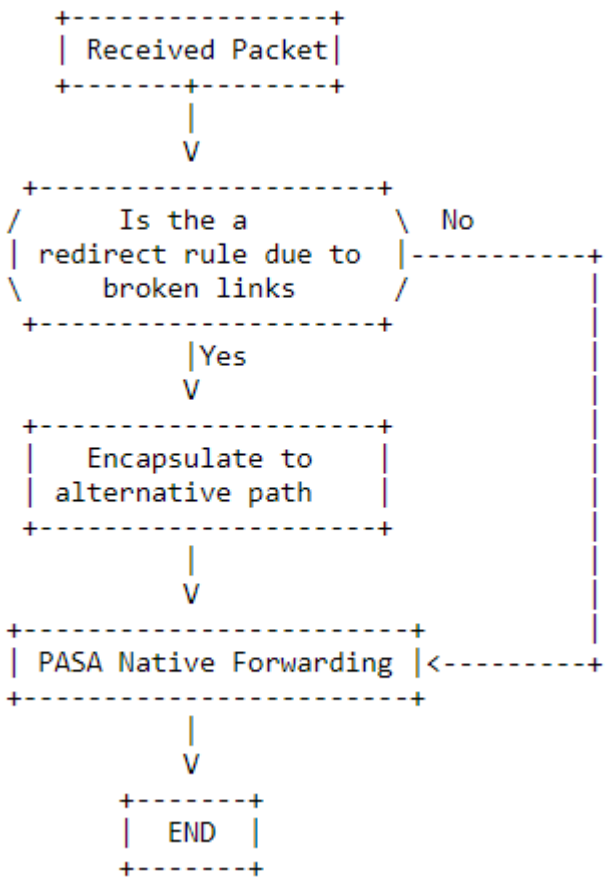

Multiple Addresses vs Single Address

Approach	Root State	Forwarder State	Robustness
Multiple Addresses	Low (redirect rule)	Low (redirect rule)	Lower for multiple failures (limited knowledge)
Single Address	High (topology)	Low (neighborhood)	Higher for multiple failures (root to find feasible path)

Next Steps

- Any question/comment welcome!
  - Update the draft based on comments and remarks
  - Any new contributor welcome
- Better align examples
- Refine wording
  - Clarification
  - NSA relic
- Security Considerations to be added
- See you @ IETF 117

THANKS!