Scalable PE-CE BGP FRR Using Repair Label

draft-bashandy-idr-bgp-repair-label-03

Authors :

Ahmed Bashandy, Cisco Systems
Burjiz Pithawala, Cisco Systems
Jakob Heitz, Ericsson

Presenter :

Ahmed Bashandy

Cisco Systems

IETF83, Jul/2012

Paris, France

What do we want to do

- BGP free core
  - Packets are ***tunneled*** between edge routers
- On loss of PE-CE link or CE node, we want to
  - Restore traffic through a pre-calculated PE
    - That PE can get to same prefix
    - Do not wait for control plane convergence, and
    - avoid *loops* when restoring traffic

Problem: If CE fails

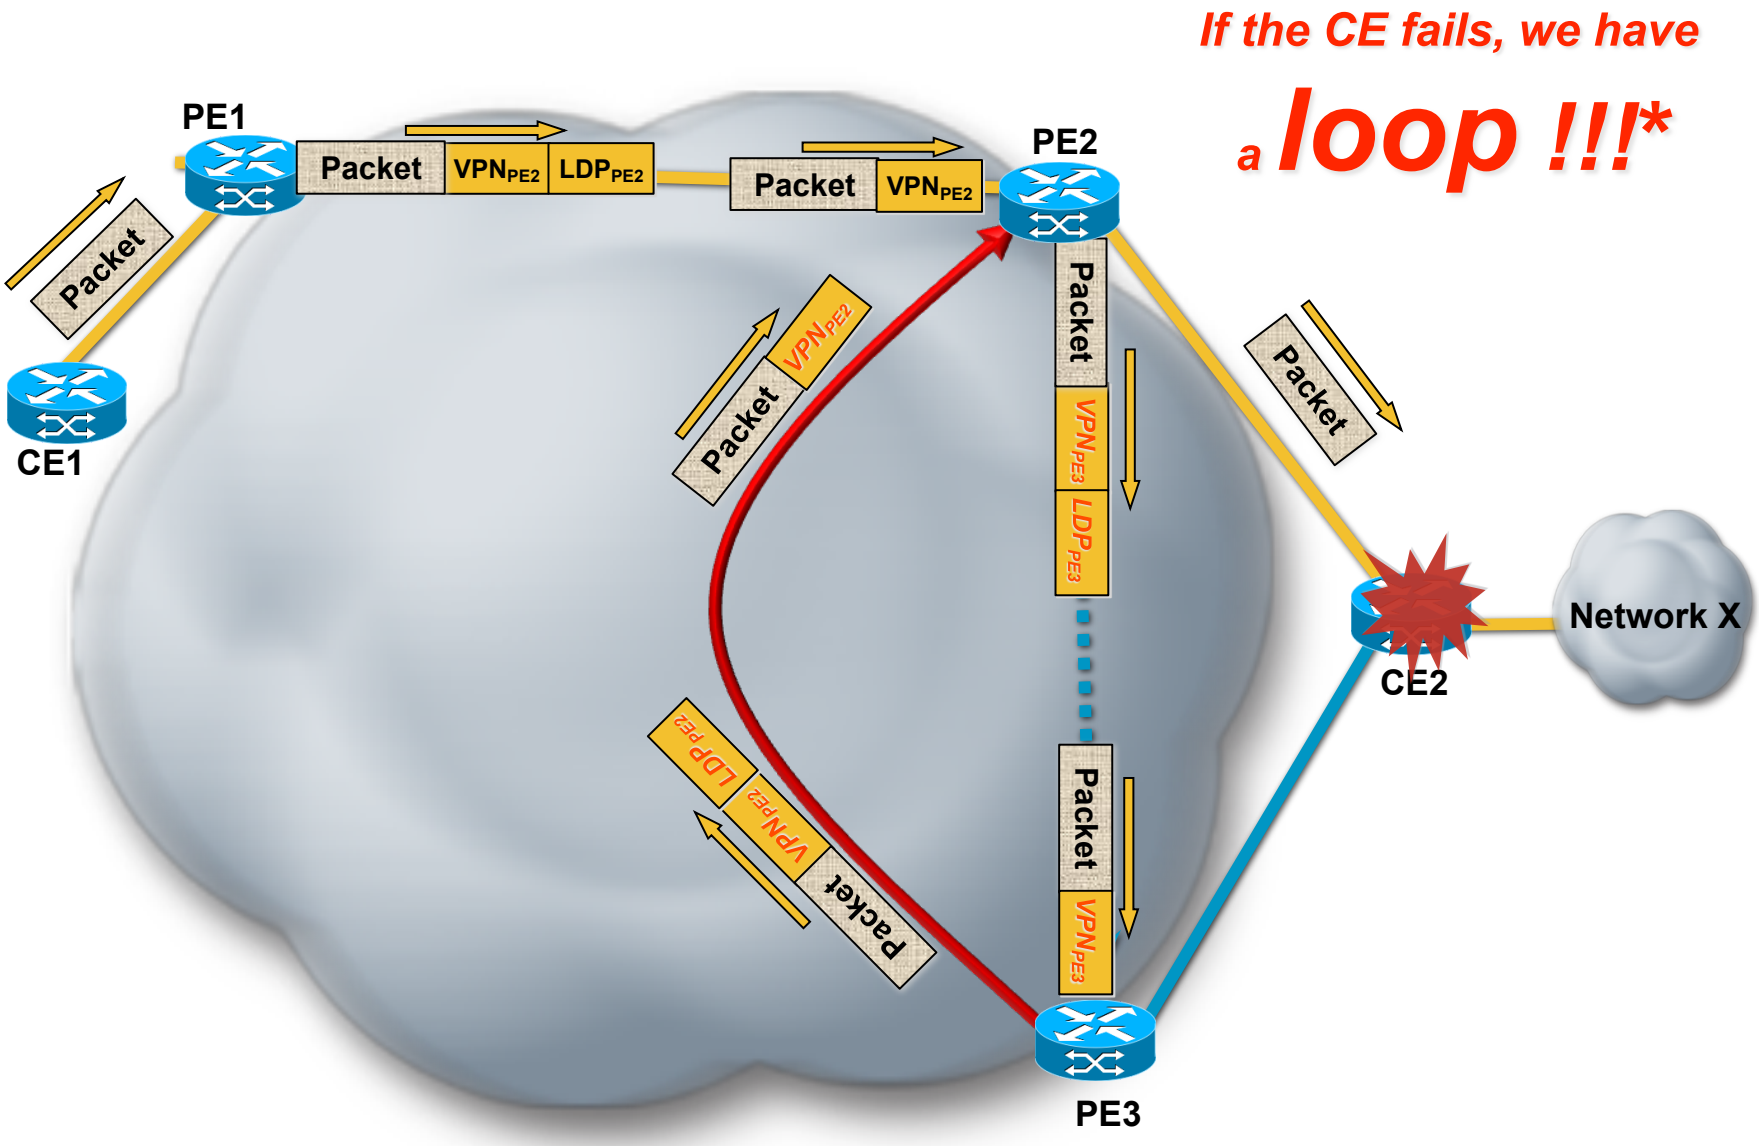

Proposed Solution

1. A PE having an external path advertises a "***repair***" label
   - Optional non-transitive attribute

2. When repairing a failed NH, the repairing router uses the ***repair label*** instead of the primary label advertised by the repair PE

3. The repair PE ***never*** repairs repaired traffic:
   Deliver the packet to the external next hop
   OR
   Drop it

Proposed Solution

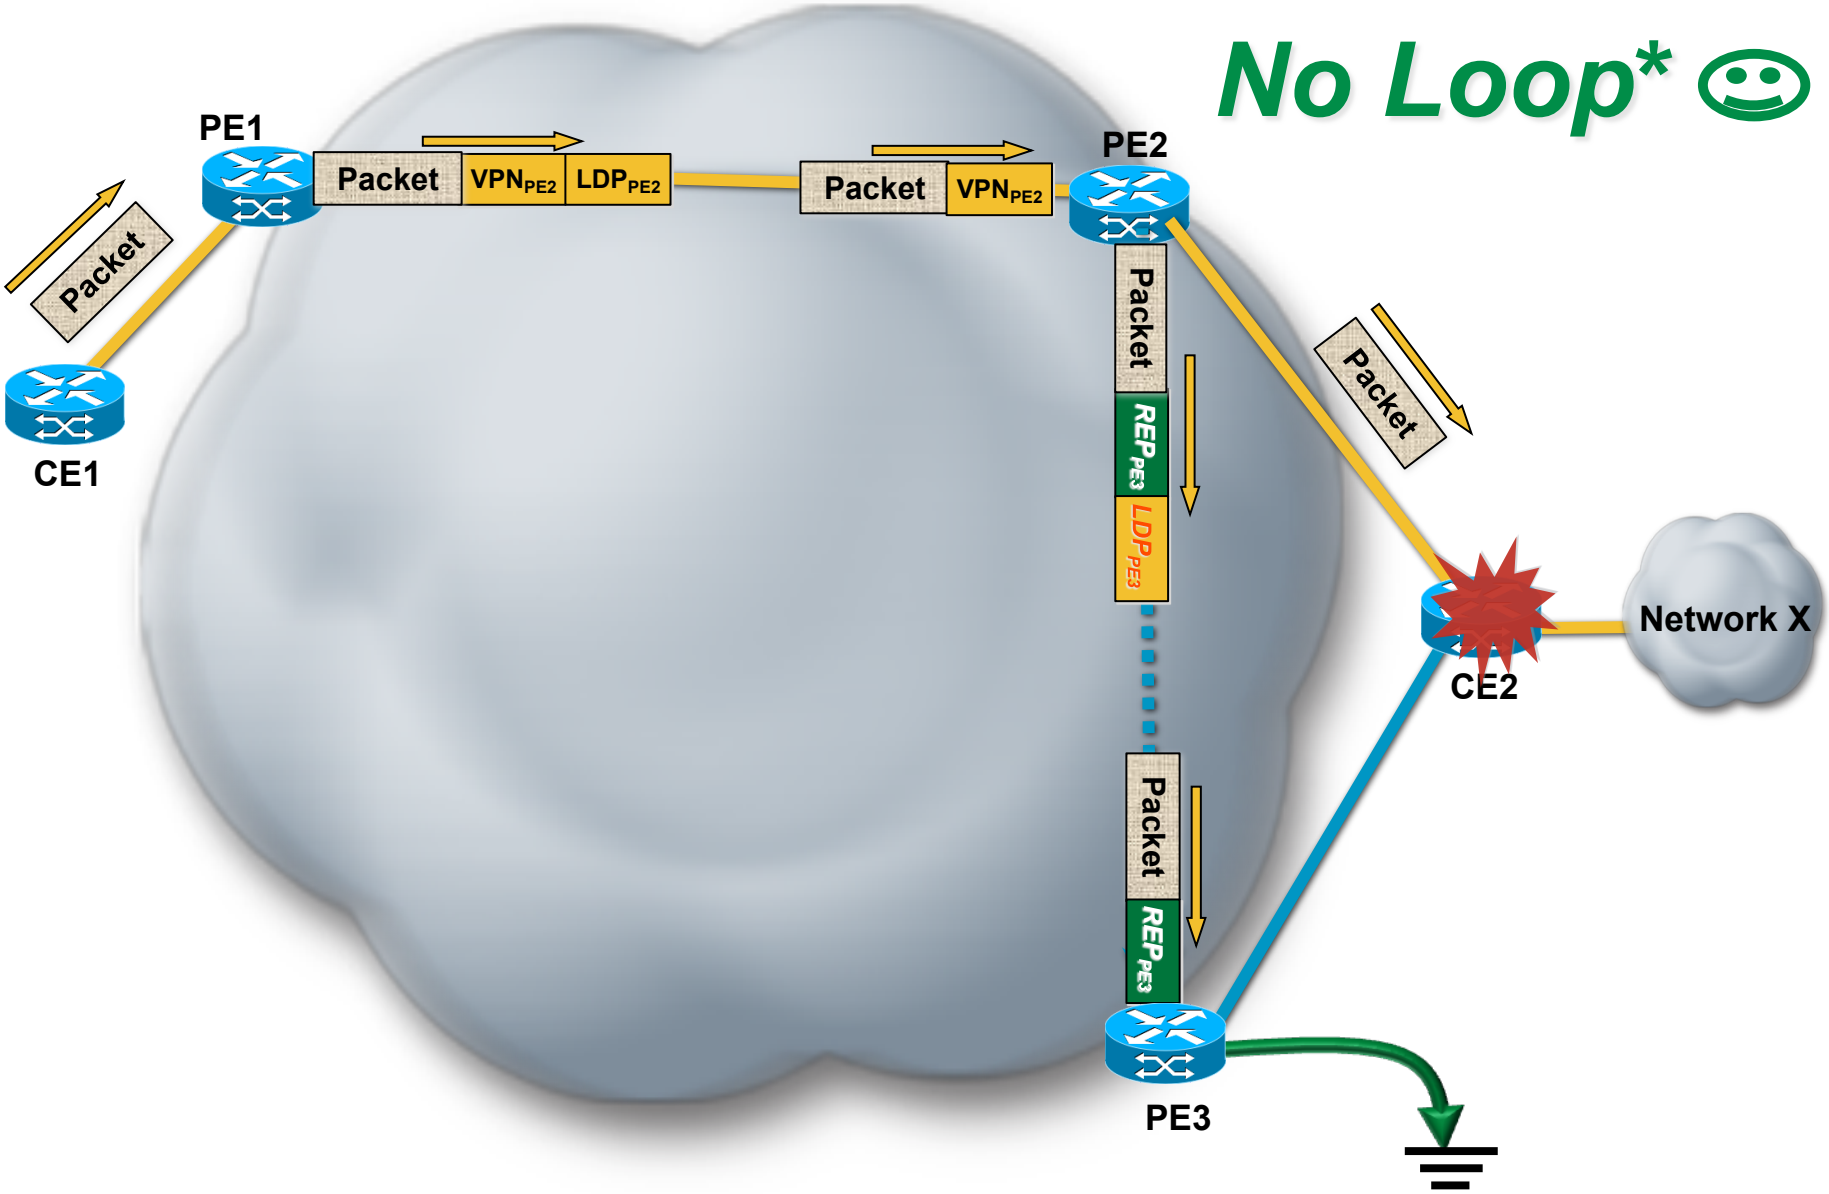

Changes from Previous Version

- Re-worded so that the repair label MUST point to an external path
- Removed the handling of a packet arriving from the core to a router that does not have a local label for it because this behavior is beyond the scope of the draft

Q & A