An analysis of the AMS-IX outage on 22/23 – Nov 2023

Stavros Konstantaras
Sr. Network Engineer – AMS-IX NOC

amsix

Before we start

- Vendor names will be mentioned
- We do NOT blame anyone
- Knowledge sharing is the power of our community
- However, have completed 6+1 migrations with great success

The AMS-IX platform

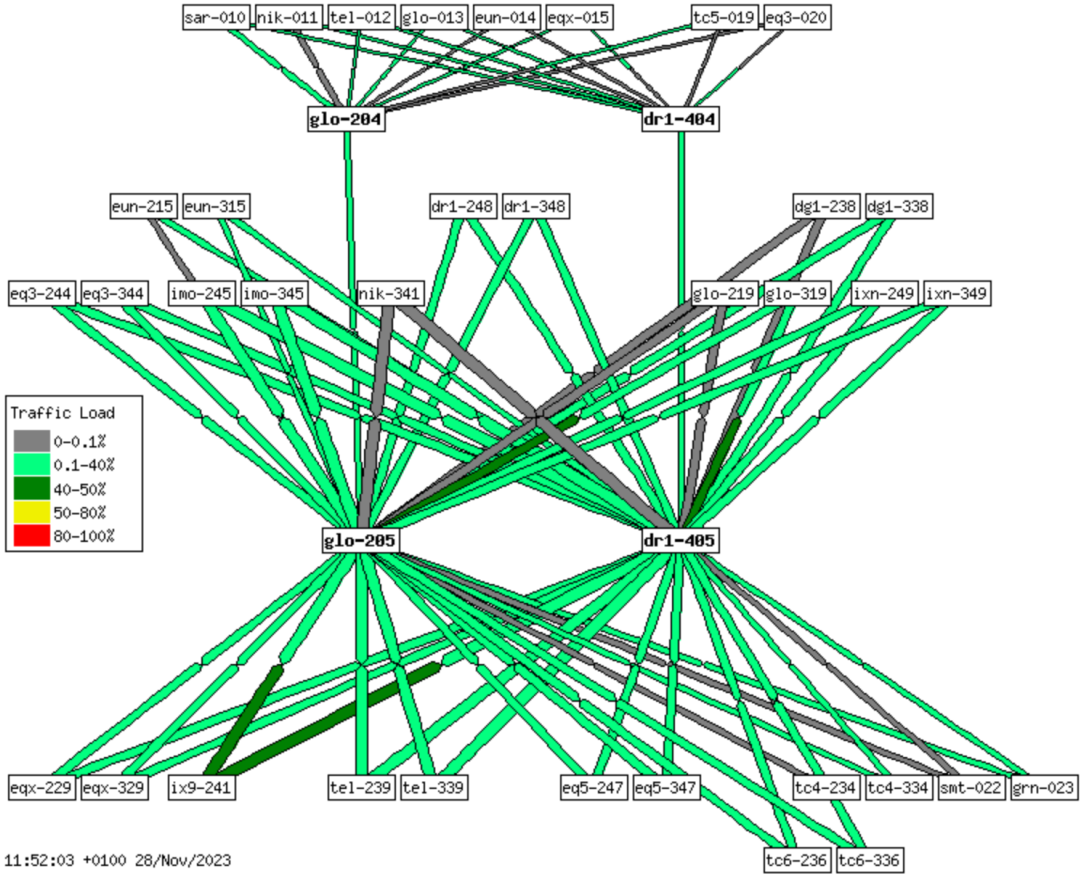

- Spine-Leaf network with dual asterisk topology
- Every colocation has 2 PEs with a PXC in front as a demarc point
- 3 generations of equipment
  - Brocade MLXe-16/32
  - Extreme SLX9850
  - Juniper MX10k8
- Protocols used:
  - OSPF
  - MPLS/VPLS
  - LDP
  - RSVP-TE
  - LACP

Wednesday afternoon (1/2)

- Unexpected automated swaps (19:08 CET) of customers at Science Park campus
  - stub-nik-341 → stub-ix9-241
- But there is more meat on that. Customers reported:
  - Loss of traffic
  - Unstable connections
  - Unreachable remote peers
- Could not point to a single PE router or Colocation
- Logs where overwhelmed with RSVP messages
- A small clue: the unstable connections where LAGs

Wednesday afternoon (2/2)

- All NOC hands on the table
  - Engaged Extreme TAC immediately in call
  - Identified several RSVP sessions flapping
  - Disconnected a suspicious customer but things didn’t improve
- We recognized a known interop issue between Extreme SLXs and Juniper MX10k8s
  - We placed an RSVP policer to cut the excess amount of RSVP messages
  - Results in much slower convergence
- We managed to stabilize the network
- Concluded work around 1am

Thursday morning

- Started investigation around 9am
  - Fine tuned the RSVP rate-limiters
- At 9:38am the issue came back
  - Rolled back everything but didn’t work
- At 10:22 isolated the Juniper core router
  - Issue wasn’t fixed
- At 10:52 identified the suspicious customer was connected
  - We disconnected the customer immediately
  - Networks started calming down and converge!!!

Our analysis

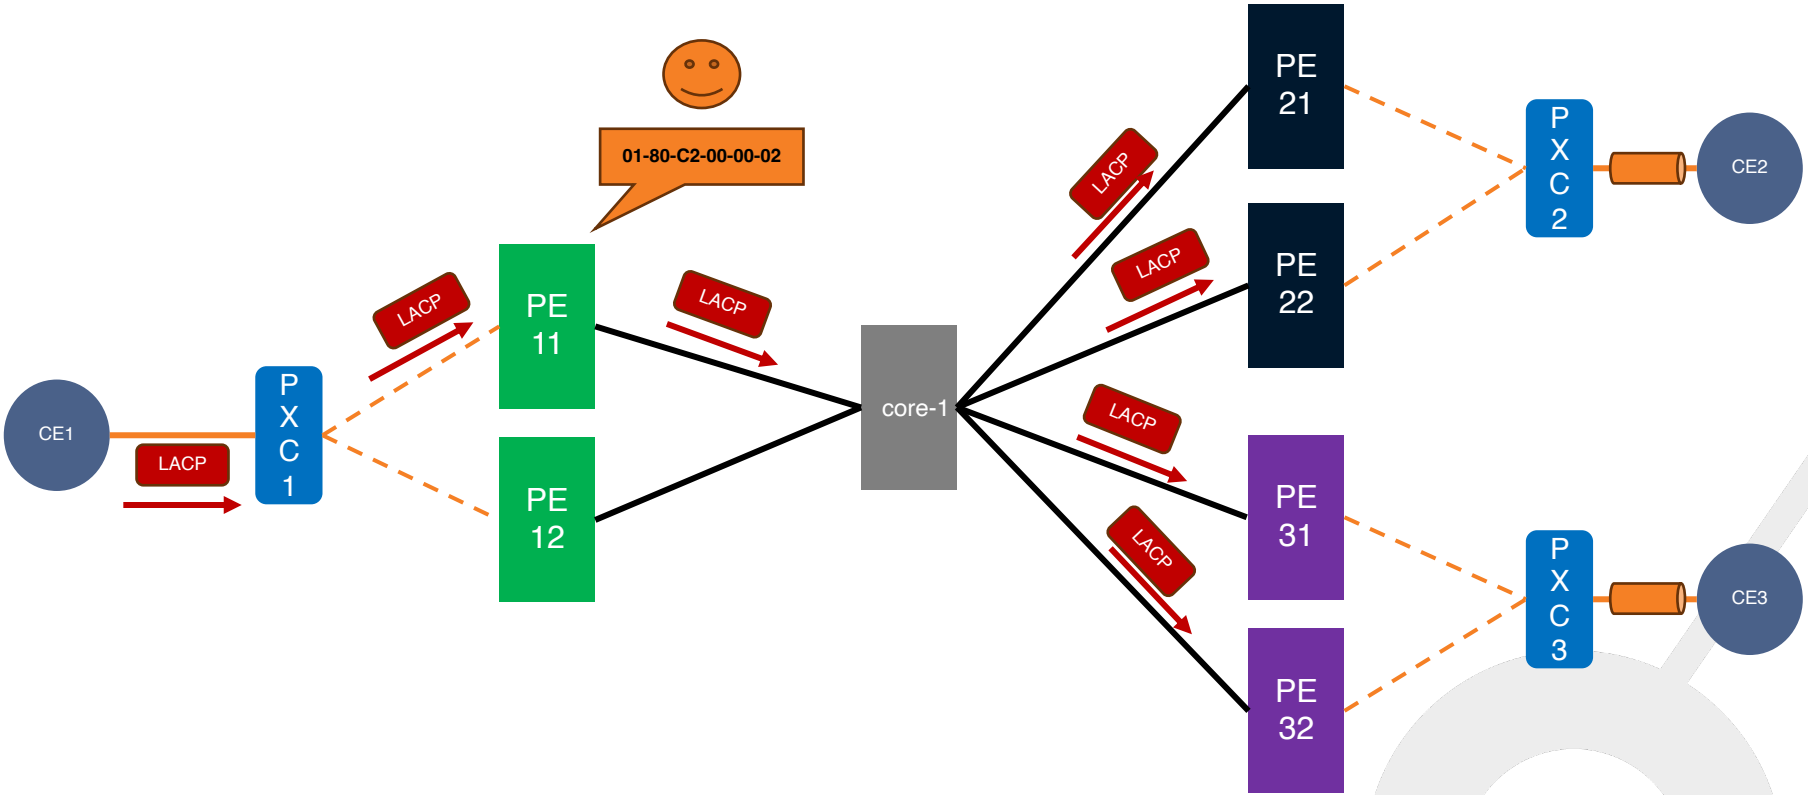

Our analysis

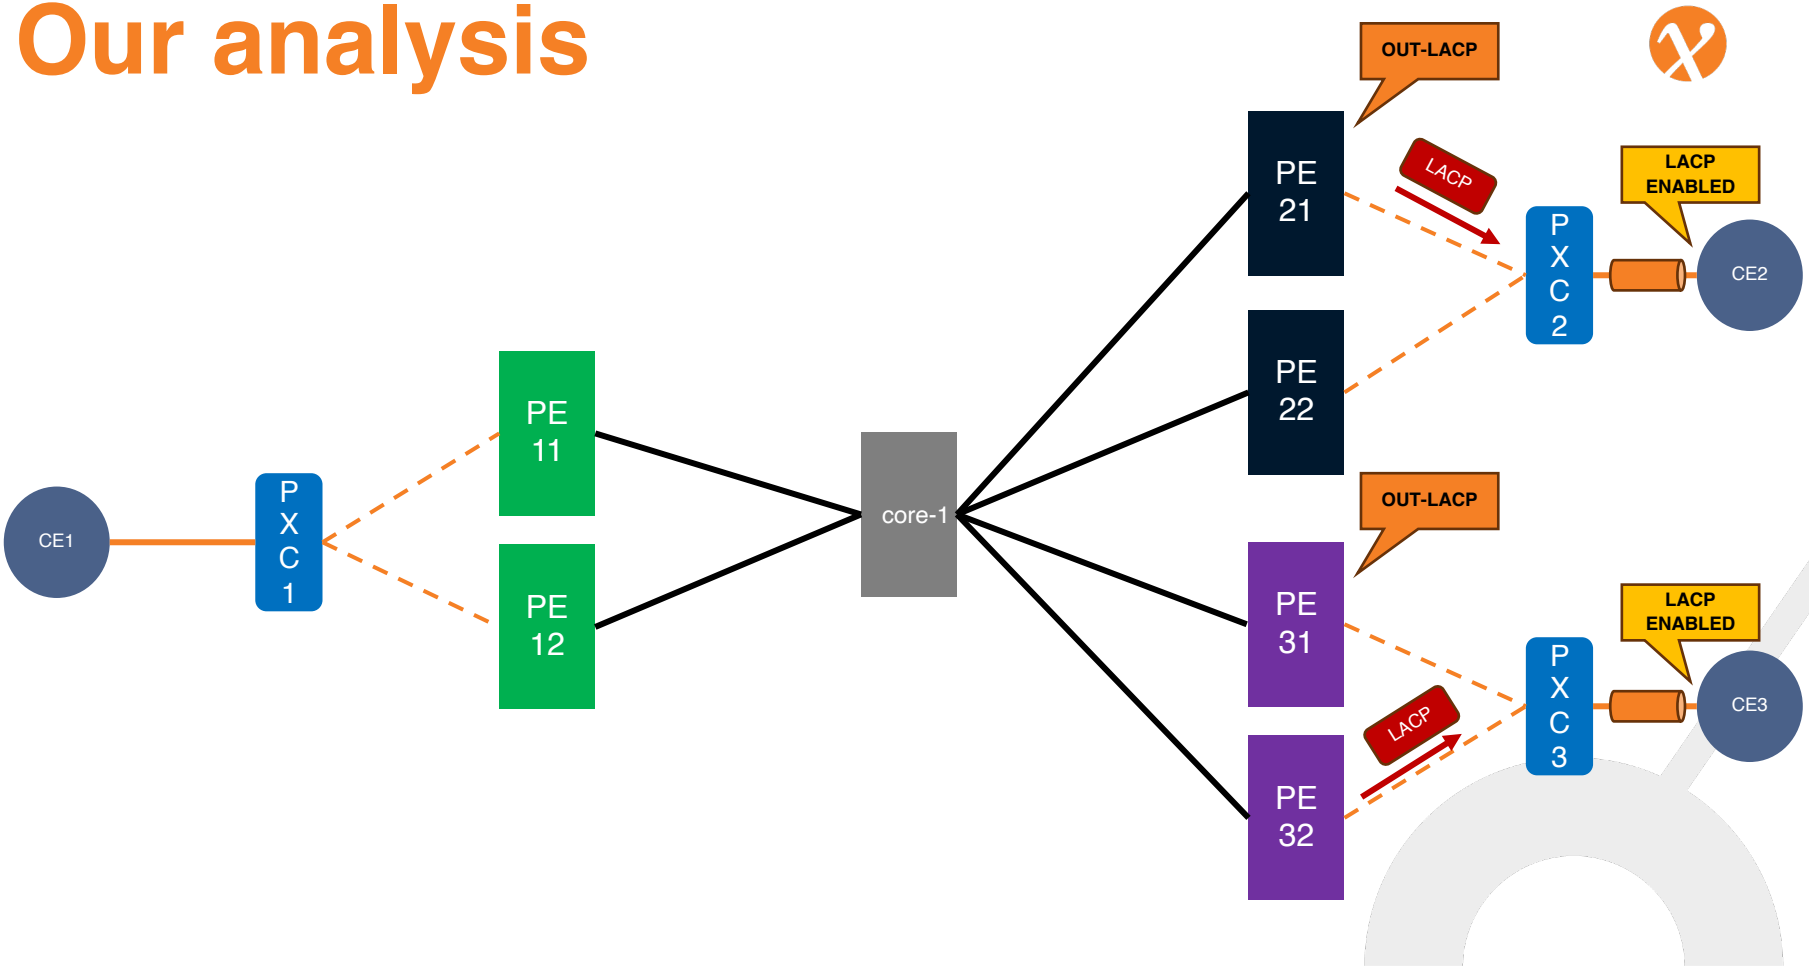

Our analysis

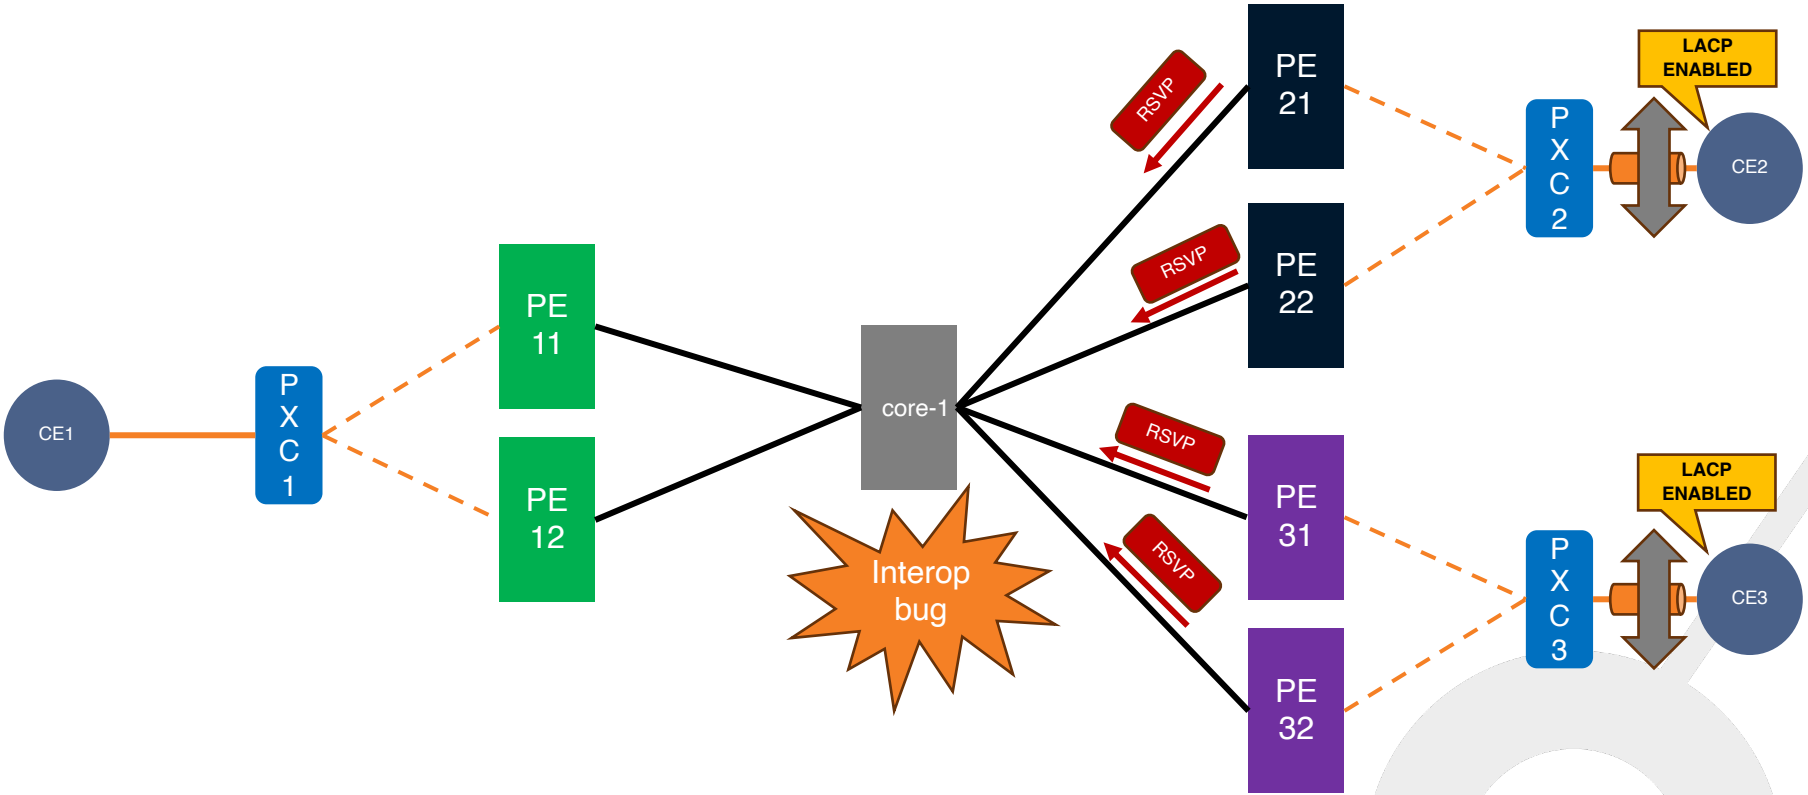

The interop issue

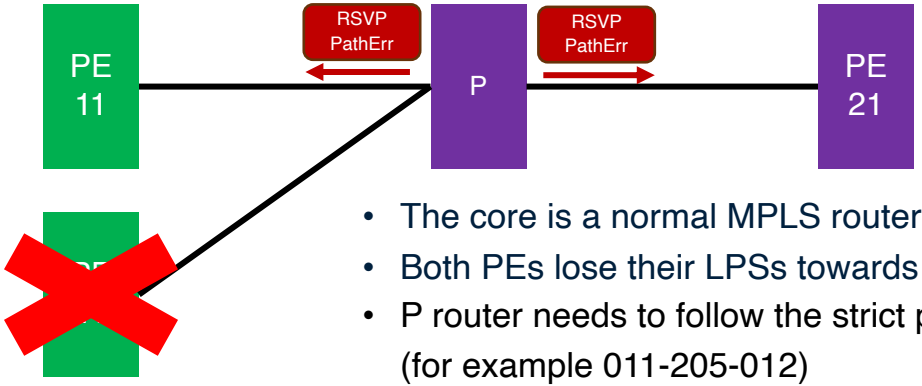

- The core is a normal MPLS router (RFC3209 compatible)
- Both PEs lose their LPSs towards PE12 → try to re-establish their RSVP session
- P router needs to follow the strict path(s) defined from AMS-IX configuration (for example 011-205-012)
- But PE 012 is offline:
  - The P router generates PathError messages with the "**Path State Removed**" flag
  - PE11 follows the flag and removes RSVP state
  - PE11 does not respect the ***retry-timer*** (default 30 seconds)
  - PE11 creates new RSVP state
  - PE11 sends new RSVP PATH messages to core router
  - (… endless loop) → RSVP Storm occurs

The interop issue

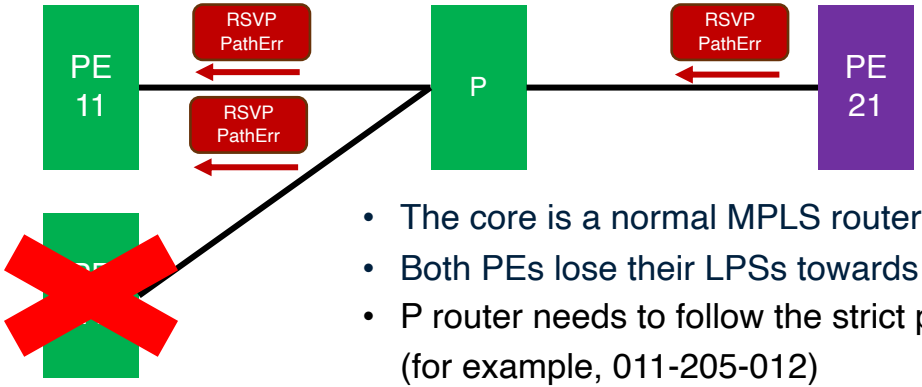

- The core is a normal MPLS router (RFC3209 compatible)
- Both PEs lose their LPSs towards PE12 → try to re-establish their RSVP session
- P router needs to follow the strict path(s) defined from AMS-IX configuration (for example, 011-205-012)
- But OSPF topology is unstable:
  - Loopbacks are reachable via PE routers and not via P routers
  - The PE routers return PathError messages (wrong destination loopback) with the "**Path State Removed**" flag
  - PE11 follows the flag and removes RSVP state without respect to the ***retry-timer***
  - PE11 creates new RSVP state and sends new RSVP PATH messages to core router
  - (… endless loop) → RSVP Storm occurs

In short

- LACP packets leaked into the platform
  - IEEE 802.3 Annex 57A.5 describes handling of “slow protocols”
  - Those packets shall be dropped unless allowed explicitly
  - We never experienced such behavior with our previous vendor

- OUT-LACP MAC ACLs didn’t work as expected
  - Both for SLX-OS (globally) and JunOS (locally)

- Interop issue magnified the bad situation ☹

Measures taken

- All JunOS firewall filters for NON-LACP enabled customers updated
- AMS-IX Provisioning updated for handling future customer ports
- LAB setup was updated to work on the review of the OUT-LACP ACLs
- Working on new filters to protect PEs from leaked traffic on VPLS level (JunOS only)
- Working with both vendors to provide them data for their cases
  - Still unclear to us if LACP handling is a bug or a default behavior

Thank you

stavros.konstantaras@ams-ix.net