Are my LSPs up?
Are they up now?

Ina Minei
ina@juniper.net

MPLS 2004

Juniper® NETWORKS

Are my LSPs up?

Silent failures

- Silent errors are not reported by the control plane.
  - Requires "polling" the LSP periodically – ping, lsping.
  - How often is good enough?
    - Tradeoff between quick detection and resources spent on detection.
    - Scaling improvements using BFD
  - Where does it make sense? When the failure causes violation of the SLAs

Categories of failures

- Protocol dependent – correct operation of the label distribution protocols yielding a failure scenario
- Application dependent – failure in the PE-PE tunnels not propagated to all interested parties
- Implementation dependent – SW/HW bugs, congestion in the control plane.

Protocol dependent failures

- Dependent on the label distribution protocol used.
- Correct protocol behavior yielding a failure scenario.

LDP failure scenarios – loss of synchronization between LDP and the IGP

- Misconfiguration – forget to turn on LDP on a new interface.
- The IGP is using the new interface, but LDP is not.
- The behavior depends on the distribution mode of LDP (independent vs. ordered).

LDP independent control

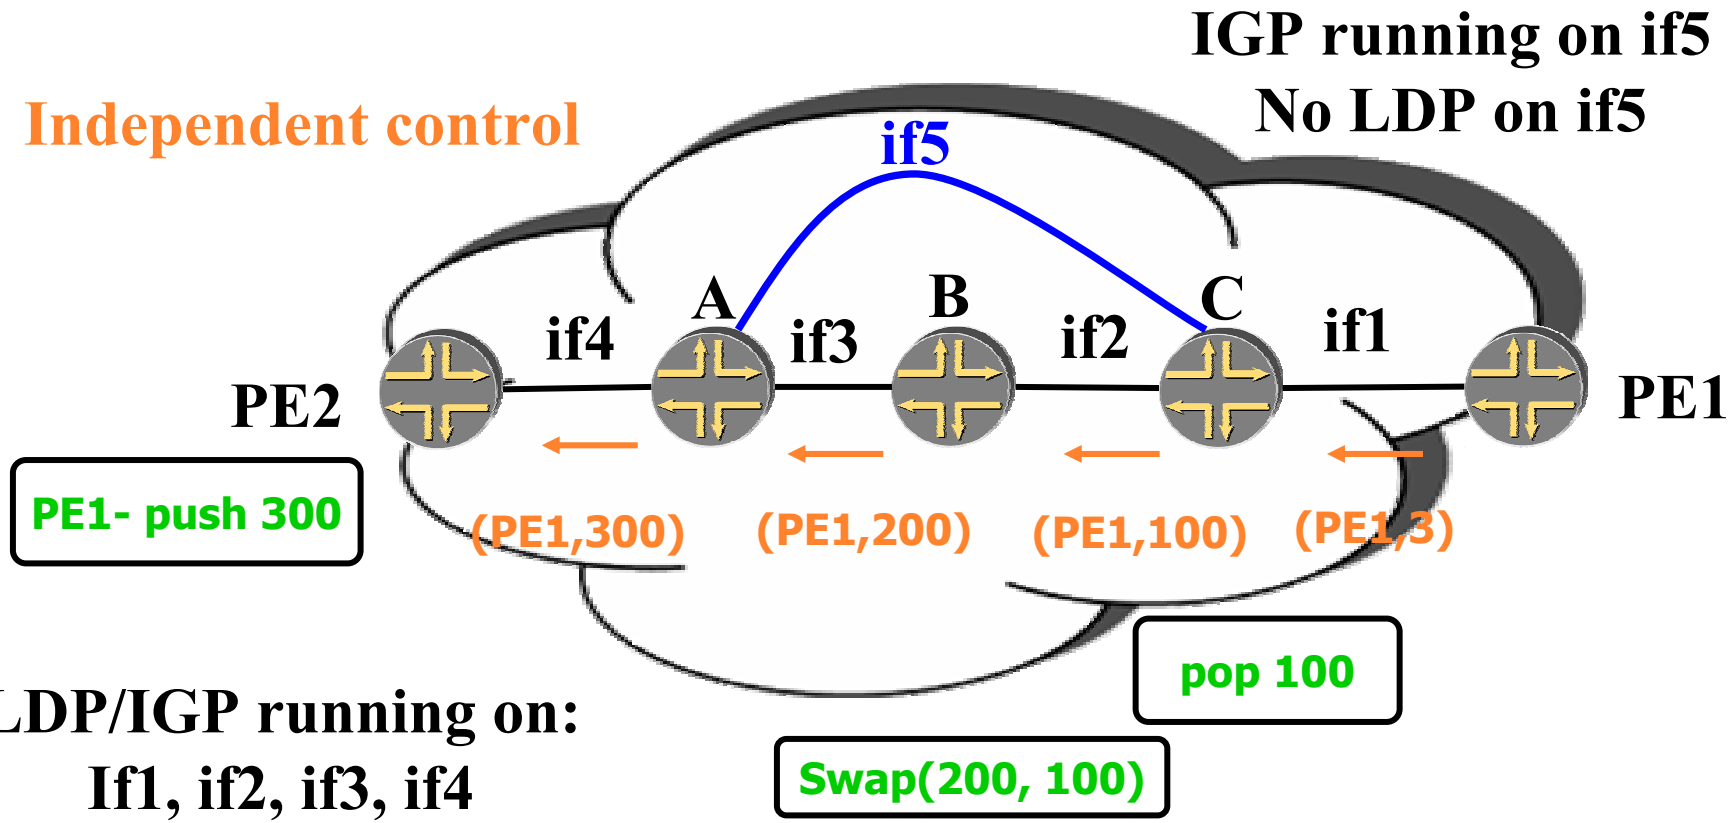

LDP ordered control

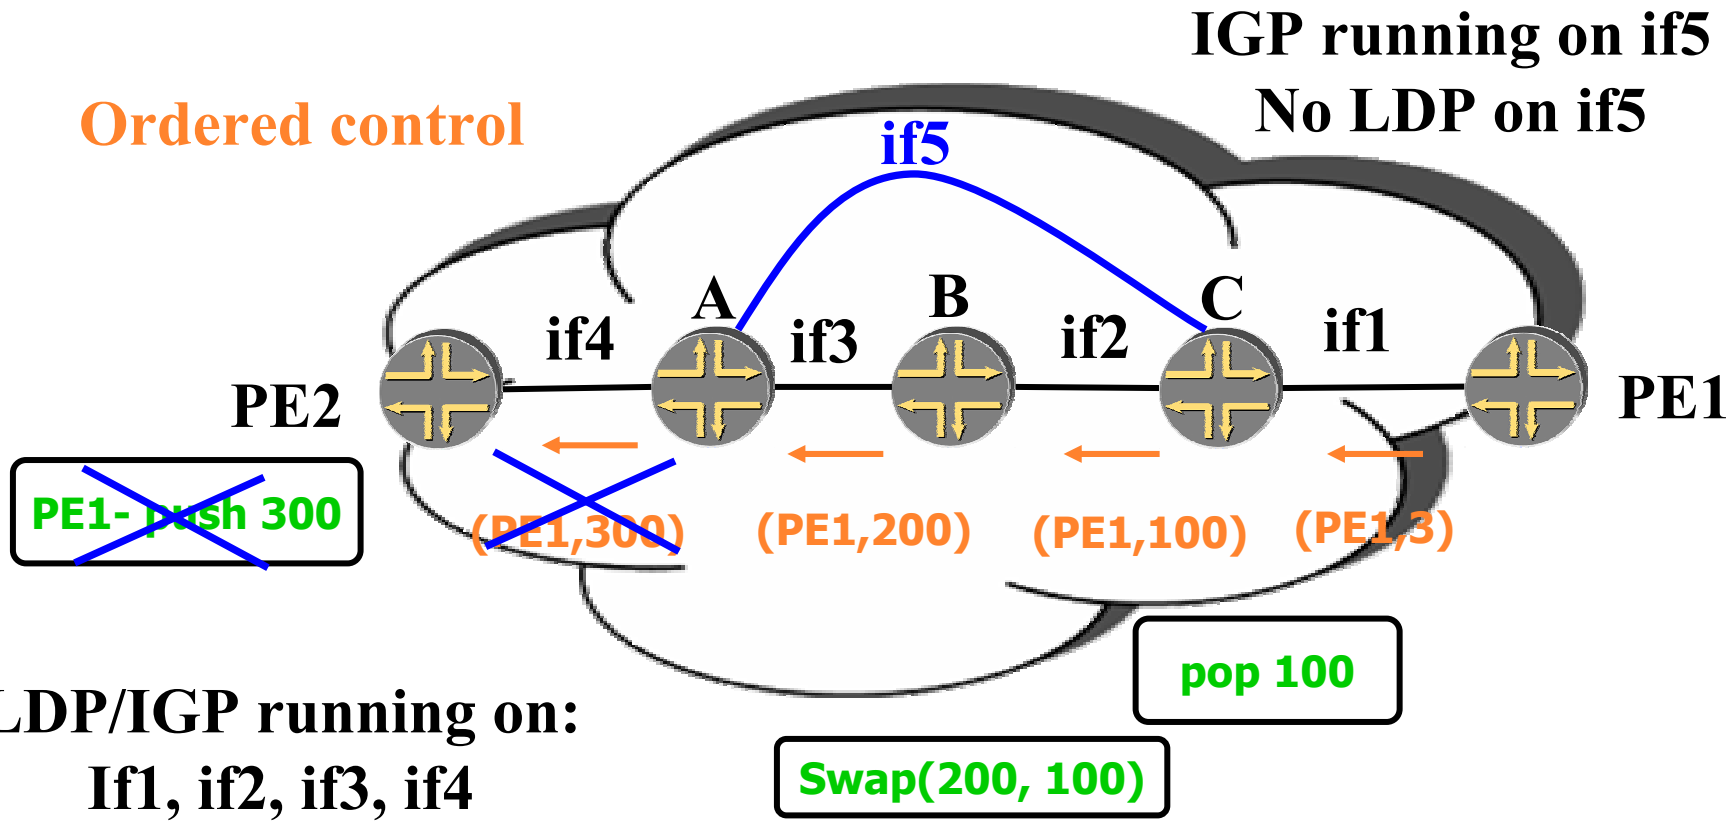

LDP failure scenario
Failure characterization

LDP independent control

- Silent failure, the LSP stays up although there is no forwarding state in the middle of the network
- Unbound in time

LDP ordered control

- Non-silent failure, the LSP will go down, routes resolving over it will become unresolved
- Duration irrelevant, since the problem is reported

LDP failure scenarios - configuration error What can be done?

LDP independent control

- What can vendors do?
  Provide mechanisms to prevent the error from happening
  - Configuration checkers
  - Ways to automate configurations
- What can operators do?
  Poll that the LSP is up

LDP ordered control

- What can vendors do?
  Report the failure via syslog, trap, etc…
- What can operators do?
  Watch for such errors, and as soon as they are detected, fix the problem.

Application dependent failures

- Applications such as L3VPN, pseudo-wires, rely on the existence of an LSP.
- Types of failure:
  - PE-PE tunnel failure not detected at the PE
  - PE-PE tunnel failure not propagated to the application: e.g. no knowledge of the LSP going down at the application level.

L3VPN failure scenario – PE-PE failure not detected at the PE

- BGP VPN routes are required to resolve over LSPs. If the LSP fails, the VPN routes become unresolved.
- If there is routing protocol between the PE and the CE, only the routes that are resolved are exported to the CE.
- What happens when there is a failure of the LSP between the PEs? Depends of the nature of the failure…

L3VPN failure scenario – PE-PE tunnel failure not detected at the PE

Silent failure of the LSP

- The VPN routes continue to stay resolved
- The CE continues to send traffic, the PE forwards the traffic, the traffic is dropped in the middle of the network
- What can operators do? Periodic polling of the VPN routes.

LSP failure known at the PE

- The VPN routes become unresolved, the PE withdraws the routes from the CE
- The CE is aware of the failure, will not attempt to send traffic
- What can operators do? Monitor alarms on the PE and the CE

Pseudo-wire failure propagation

- Assume that the LSP failure is detected at the PE.
- Based on the emulated service, the PEs may send native indications over the related attachment circuits to notify the end points of the fault condition.

  draft-ietf-pwe3-oam-msg-map-00.txt

Pseudo-wire failure propagation (cont)

- What to do if the emulated service doesn't have well defined OAM procedures ?
- What can vendors do? May bring the entire interface down if all the attachment circuits related to the interface are affected.
- What can operators do? Don't rely on the failure propagation, instead use BFD as a failure detection mechanism at the service level.

Implementation dependent failures

1) Software/hardware bugs causing forwarding table corruption
   - Incorrect update of the forwarding table
   - Memory corruption of a forwarding table
     - Causes a silent failure, with traffic loss or traffic misrouting, for an unbound amount of time.
2) Congested control plane not updating the forwarding plane fast enough
     - Causes a silent failure, with traffic loss, for a bound (but variable) amount of time

SW/HW bugs – What can vendors do?

1) Avoid the problem!
2) Implement a consistency check between the forwarding plane and the control plane to detect and report mismatches.
   - Issues: computational resources, false negatives, etc

Congestion in the control plane – What can vendors do? (cont)

3) Avoid the problem – make sure the router has adequate control plane resources for its role
4) Improve processing under load
   - Assign priorities to different tasks, prioritize route resolution
4) Graceful degradation
   - Avoid session flaps caused by a busy control plane - offload hello processing to lower layers, make sure keepalives are sent, etc.
5) Detect and report a busy control plane

So… are my LSPs up?

So… are my LSPs up?

- LSP failures are not always caused by bugs
- Failures are difficult to detect – silent failures, short-lived failures, etc.
- The best medicine is prevention
  - Vendors and operators share in the effort of improving the resilience of MPLS networks.
  - New work in the standards bodies to address some of the problems.

Thank you!