

BGP Graceful Shutdown

draft-ietf-grow-bgp-gshut-02

Pierre Francois

Bruno Decraene

Cristel Pelsser

Clarence Filsfils

Keyur Patel

UCL-BE

France Telecom - Orange

IIJ

Cisco Systems

Cisco Systems

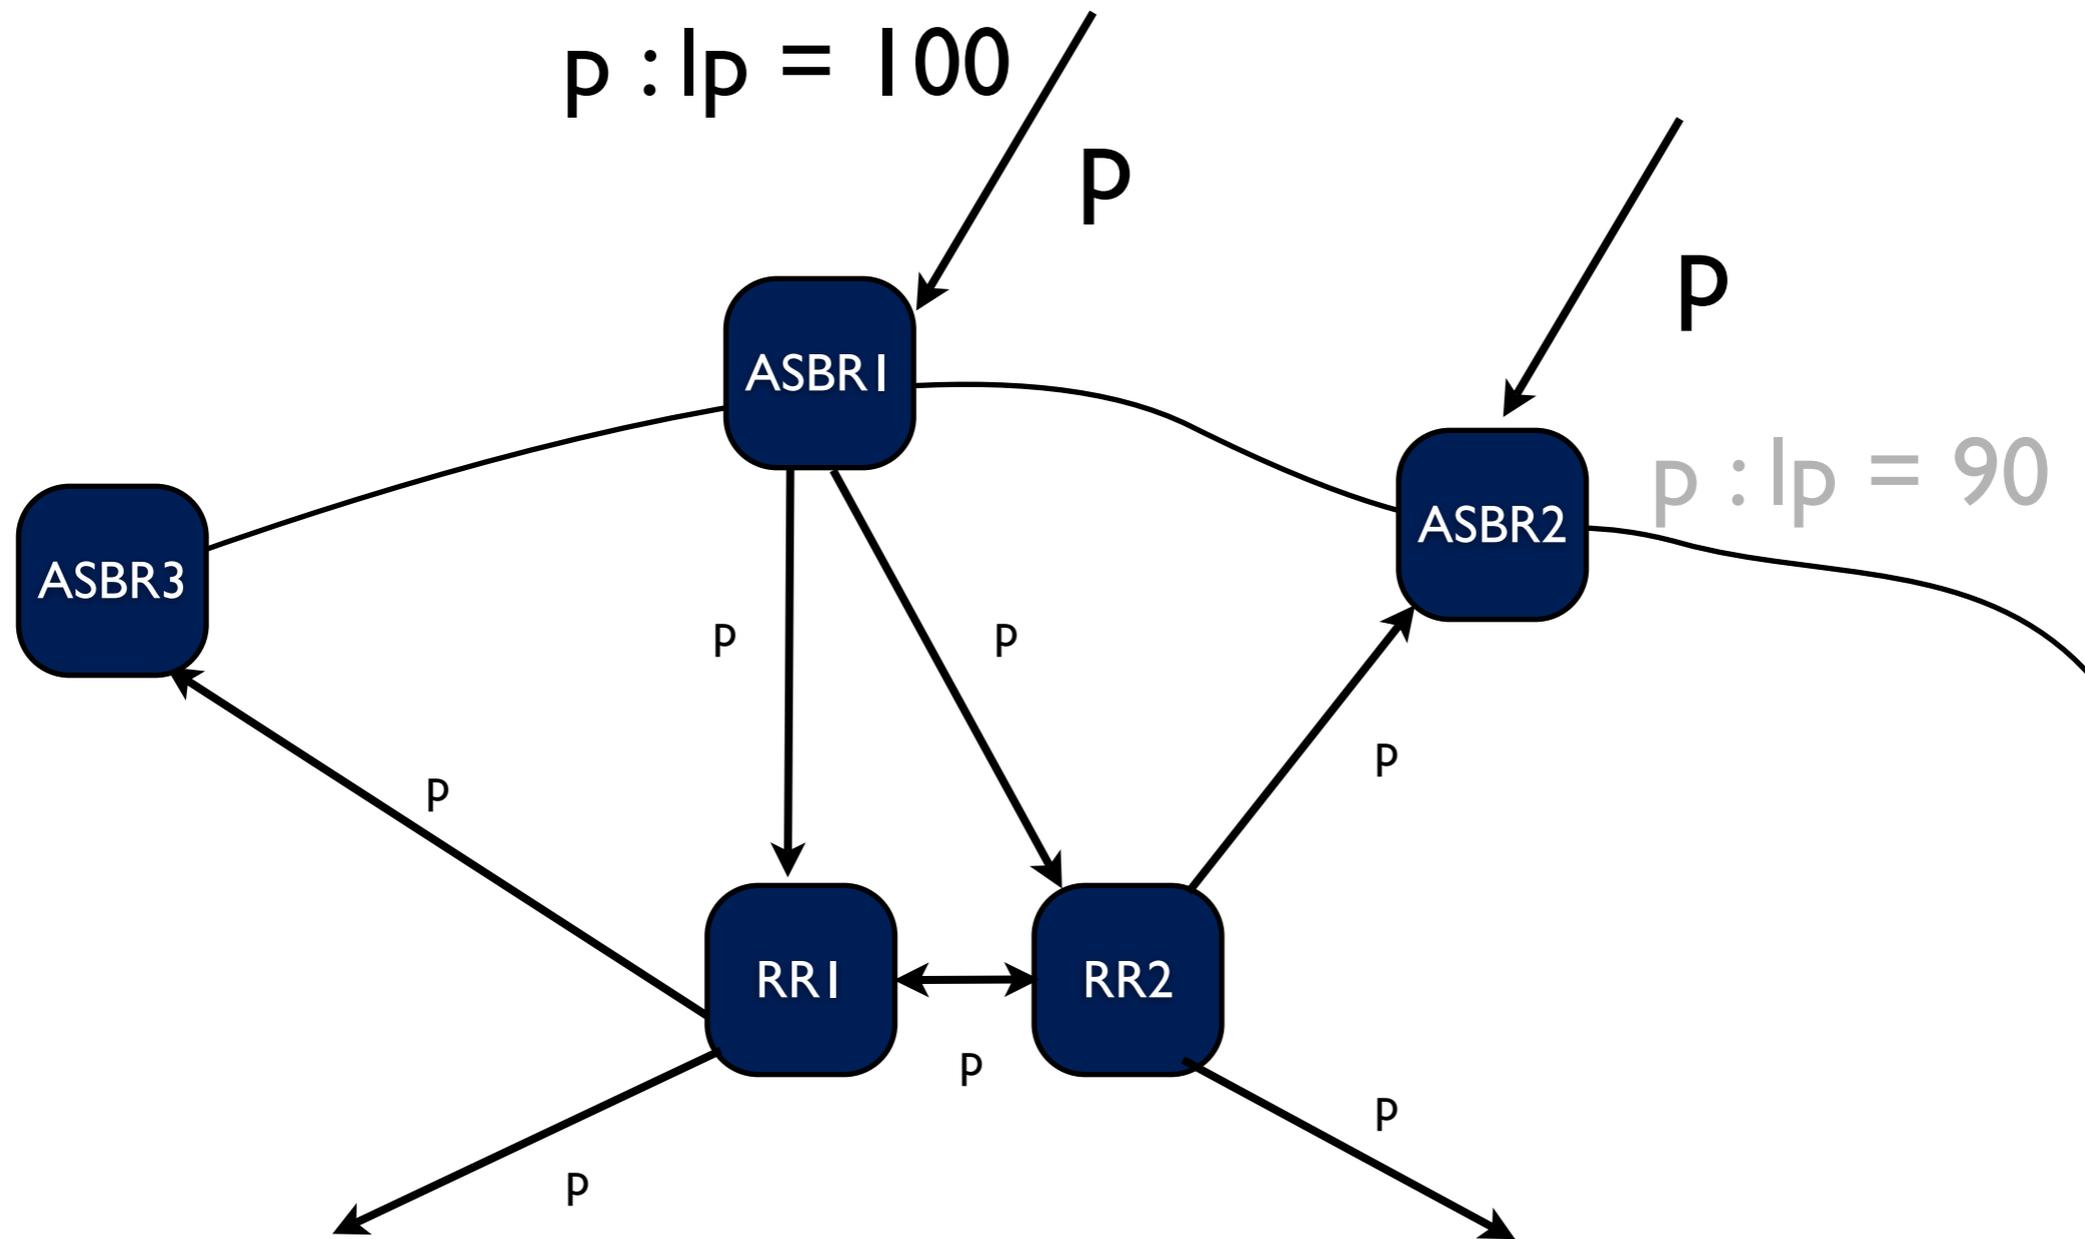
What is it about ?

- Avoid traffic loss upon planned BGP session shutdown
- Operational procedures
 - Pre-configuration set
 - At maintenance time
- *New:* Open option for a “simplified” procedure
- *New:* Vendor support considerations
 - Less configuration burden

LoC

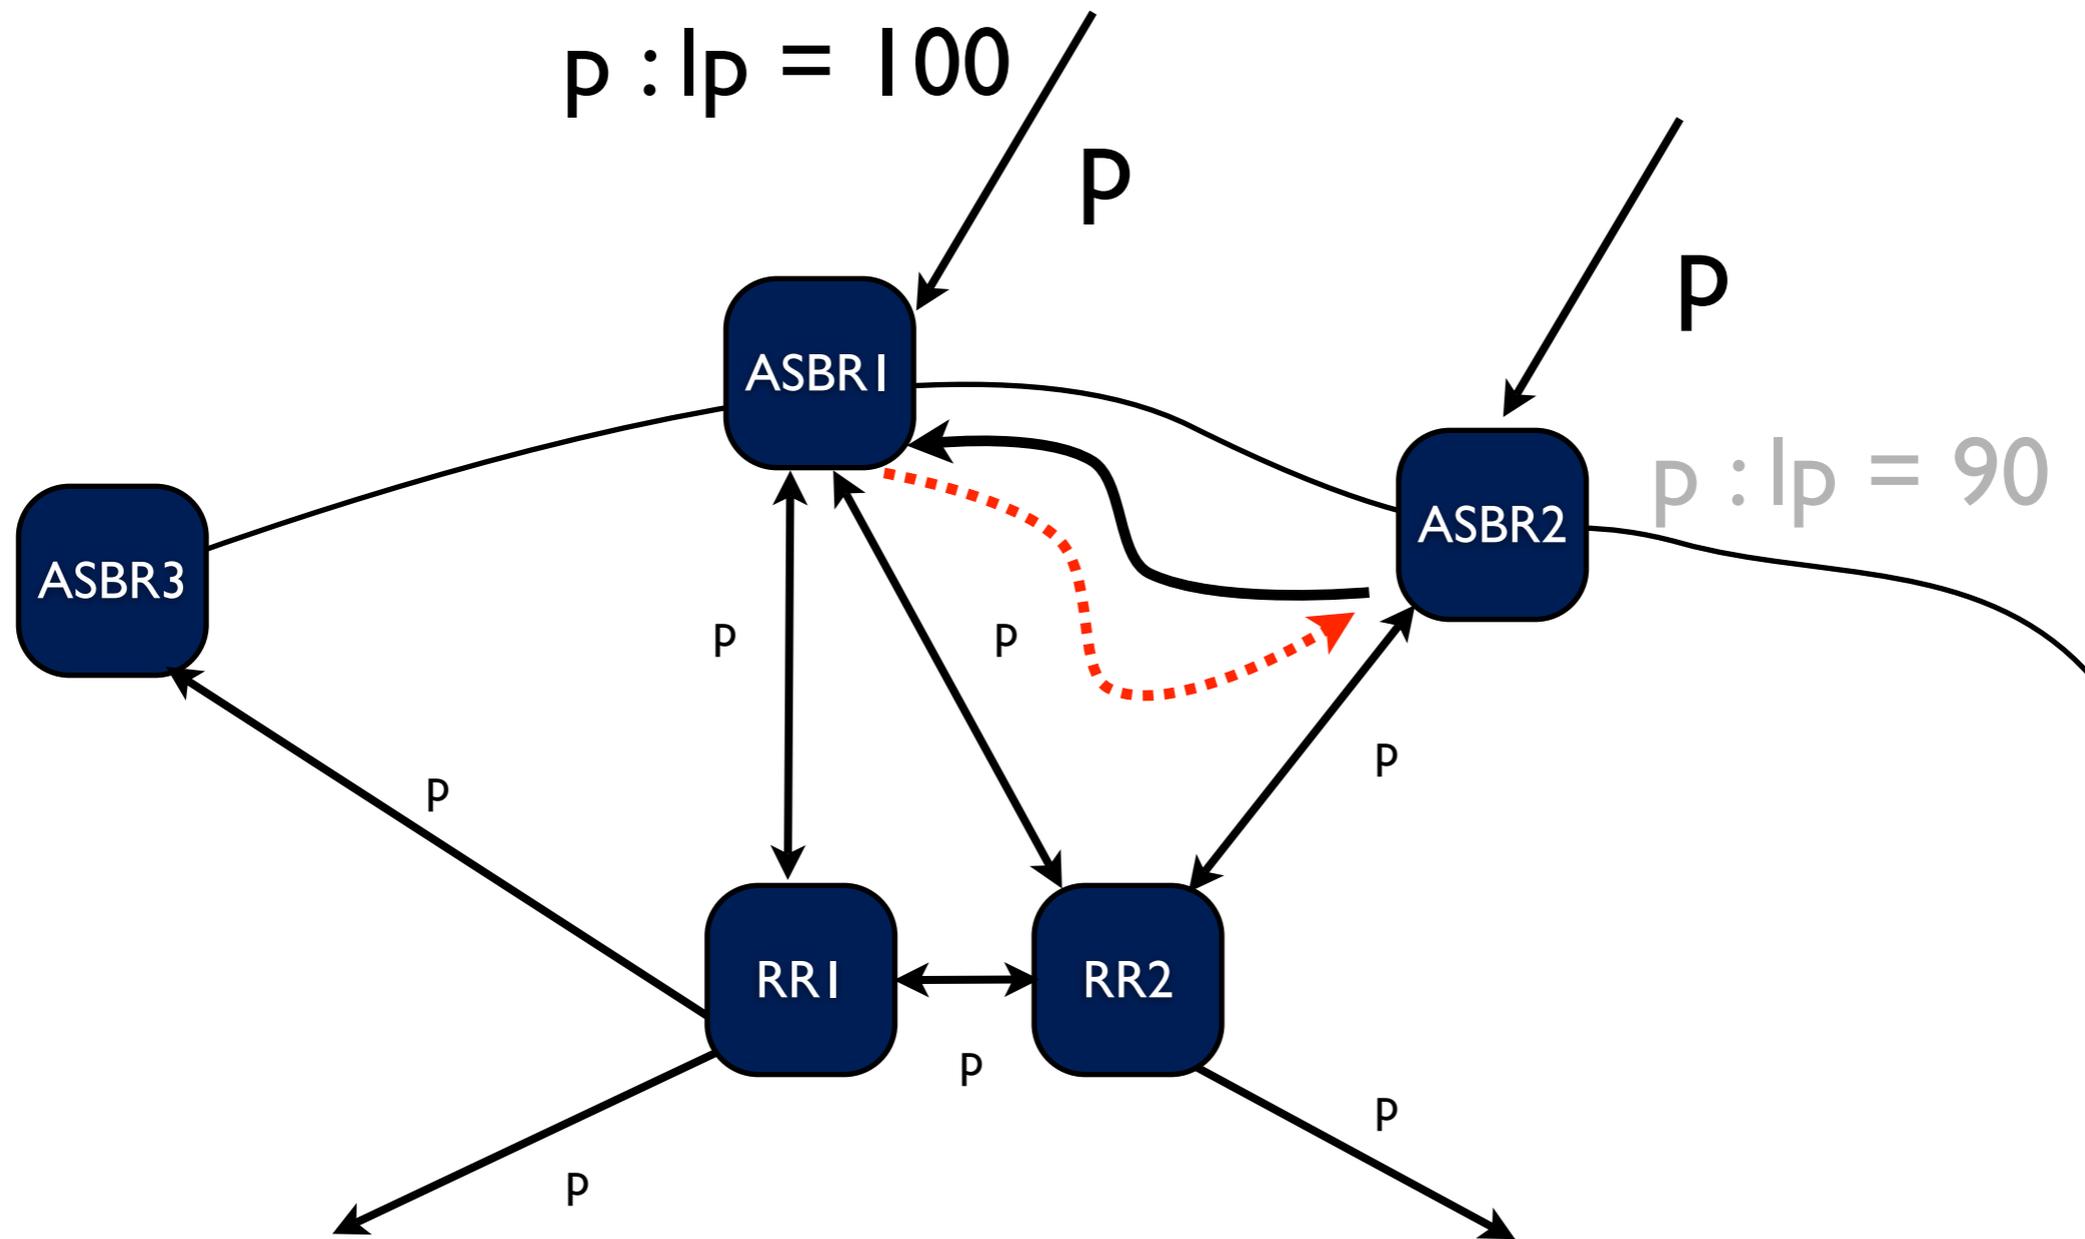
- Outbound :
Path hiding in iBGP leads to LoC
Transient inconsistencies lead to LoC
- Inbound :
Same problem
Need to trigger G-Shut at neighboring AS

LoC upon maintenance Path hiding



LoC upon maintenance

Transient inconsistencies

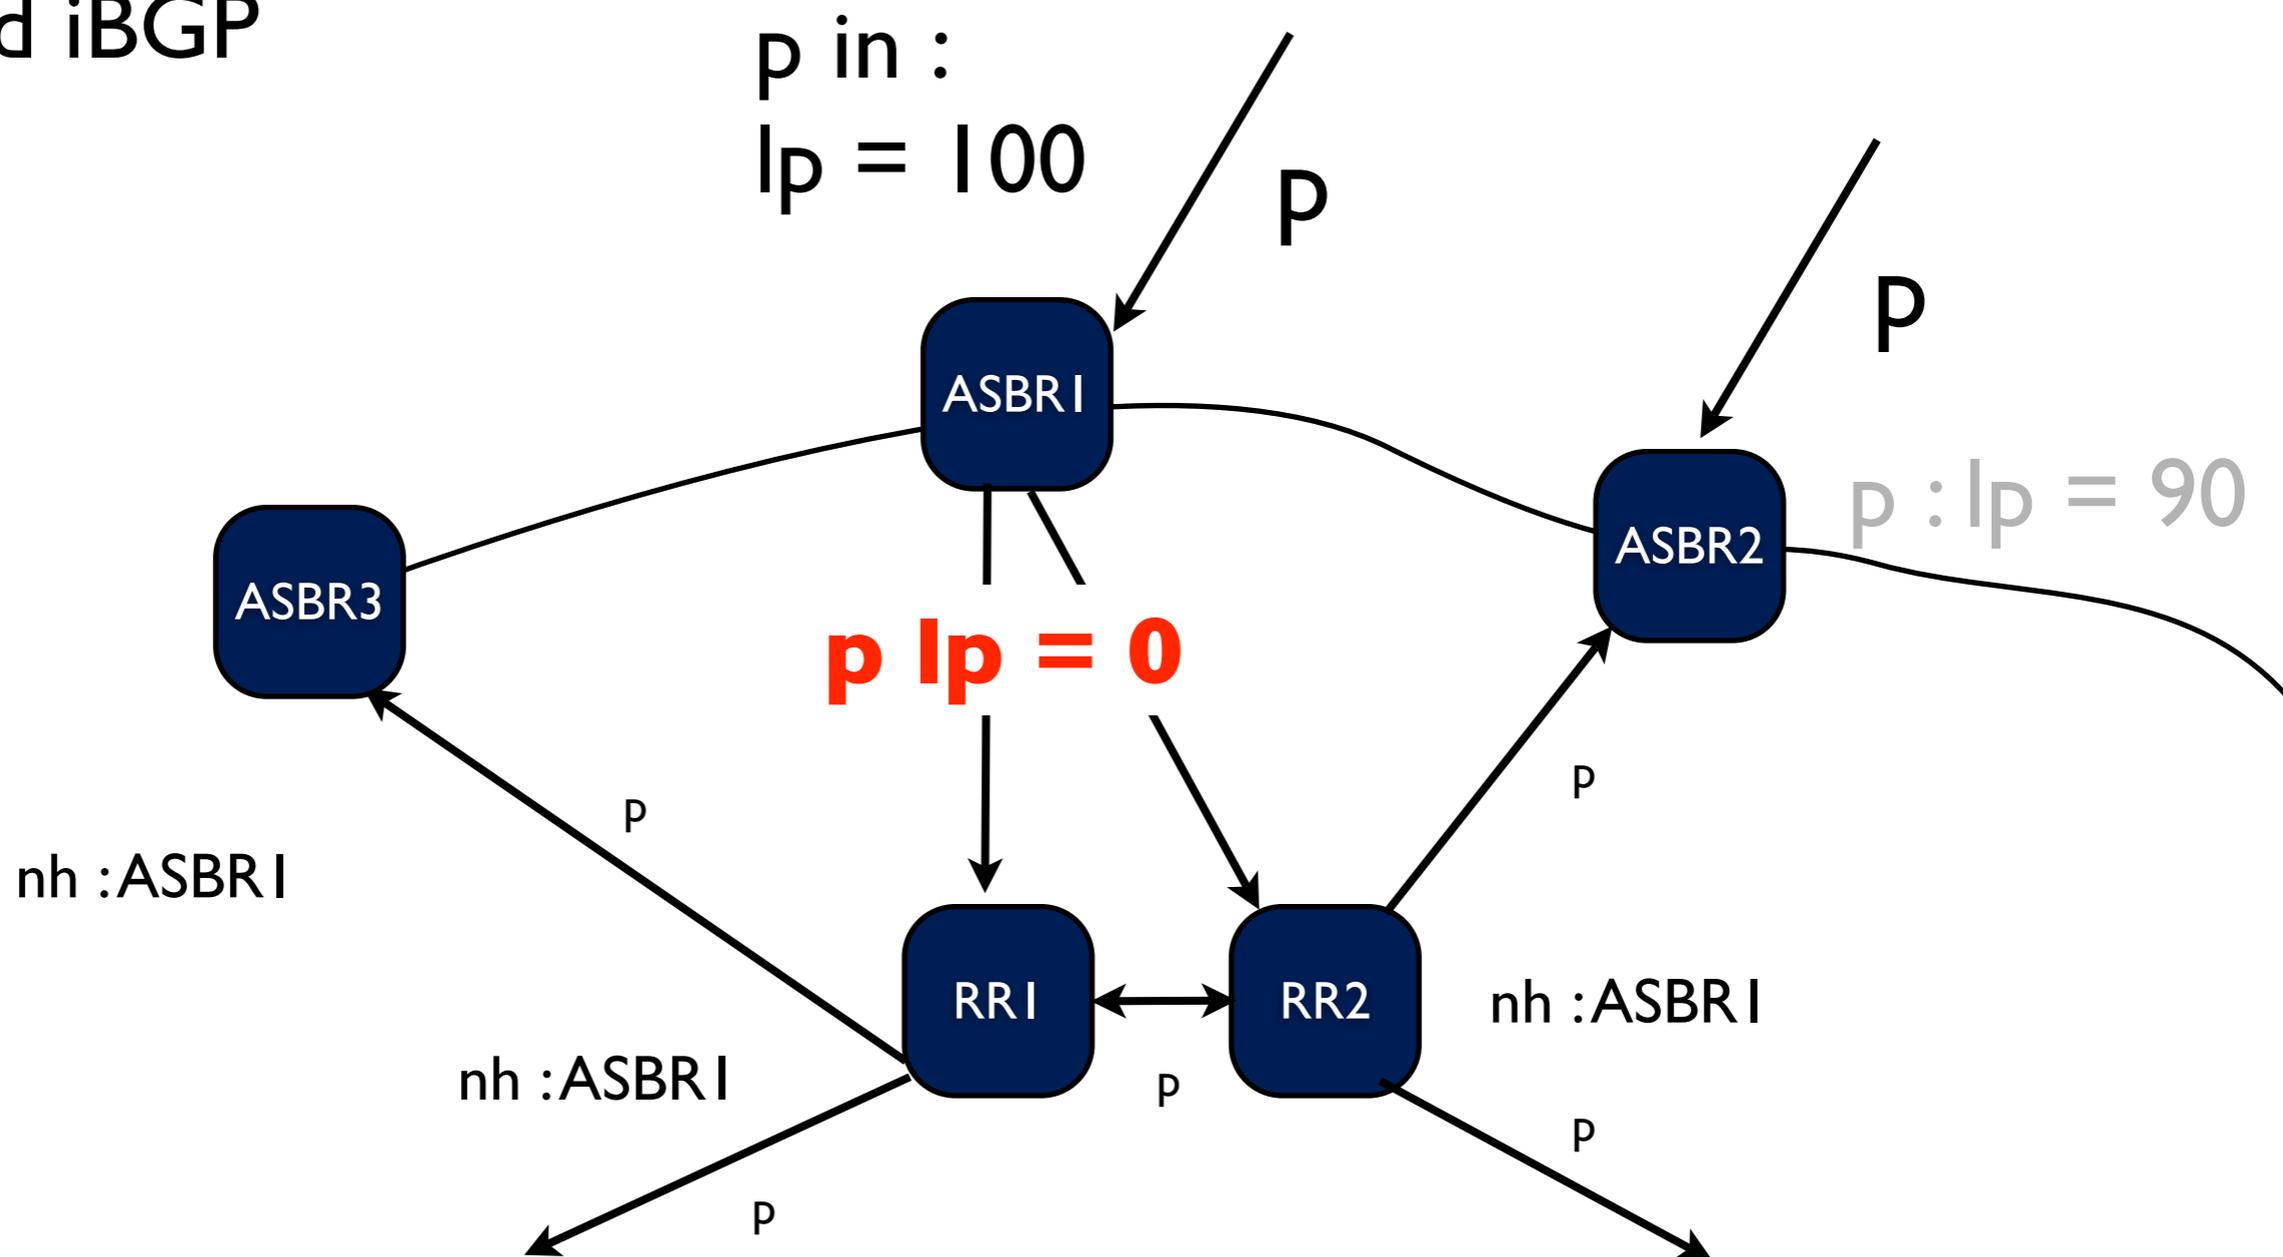


Outbound traffic Outbound LP

Decrease LP on
outbound iBGP
session

Wait

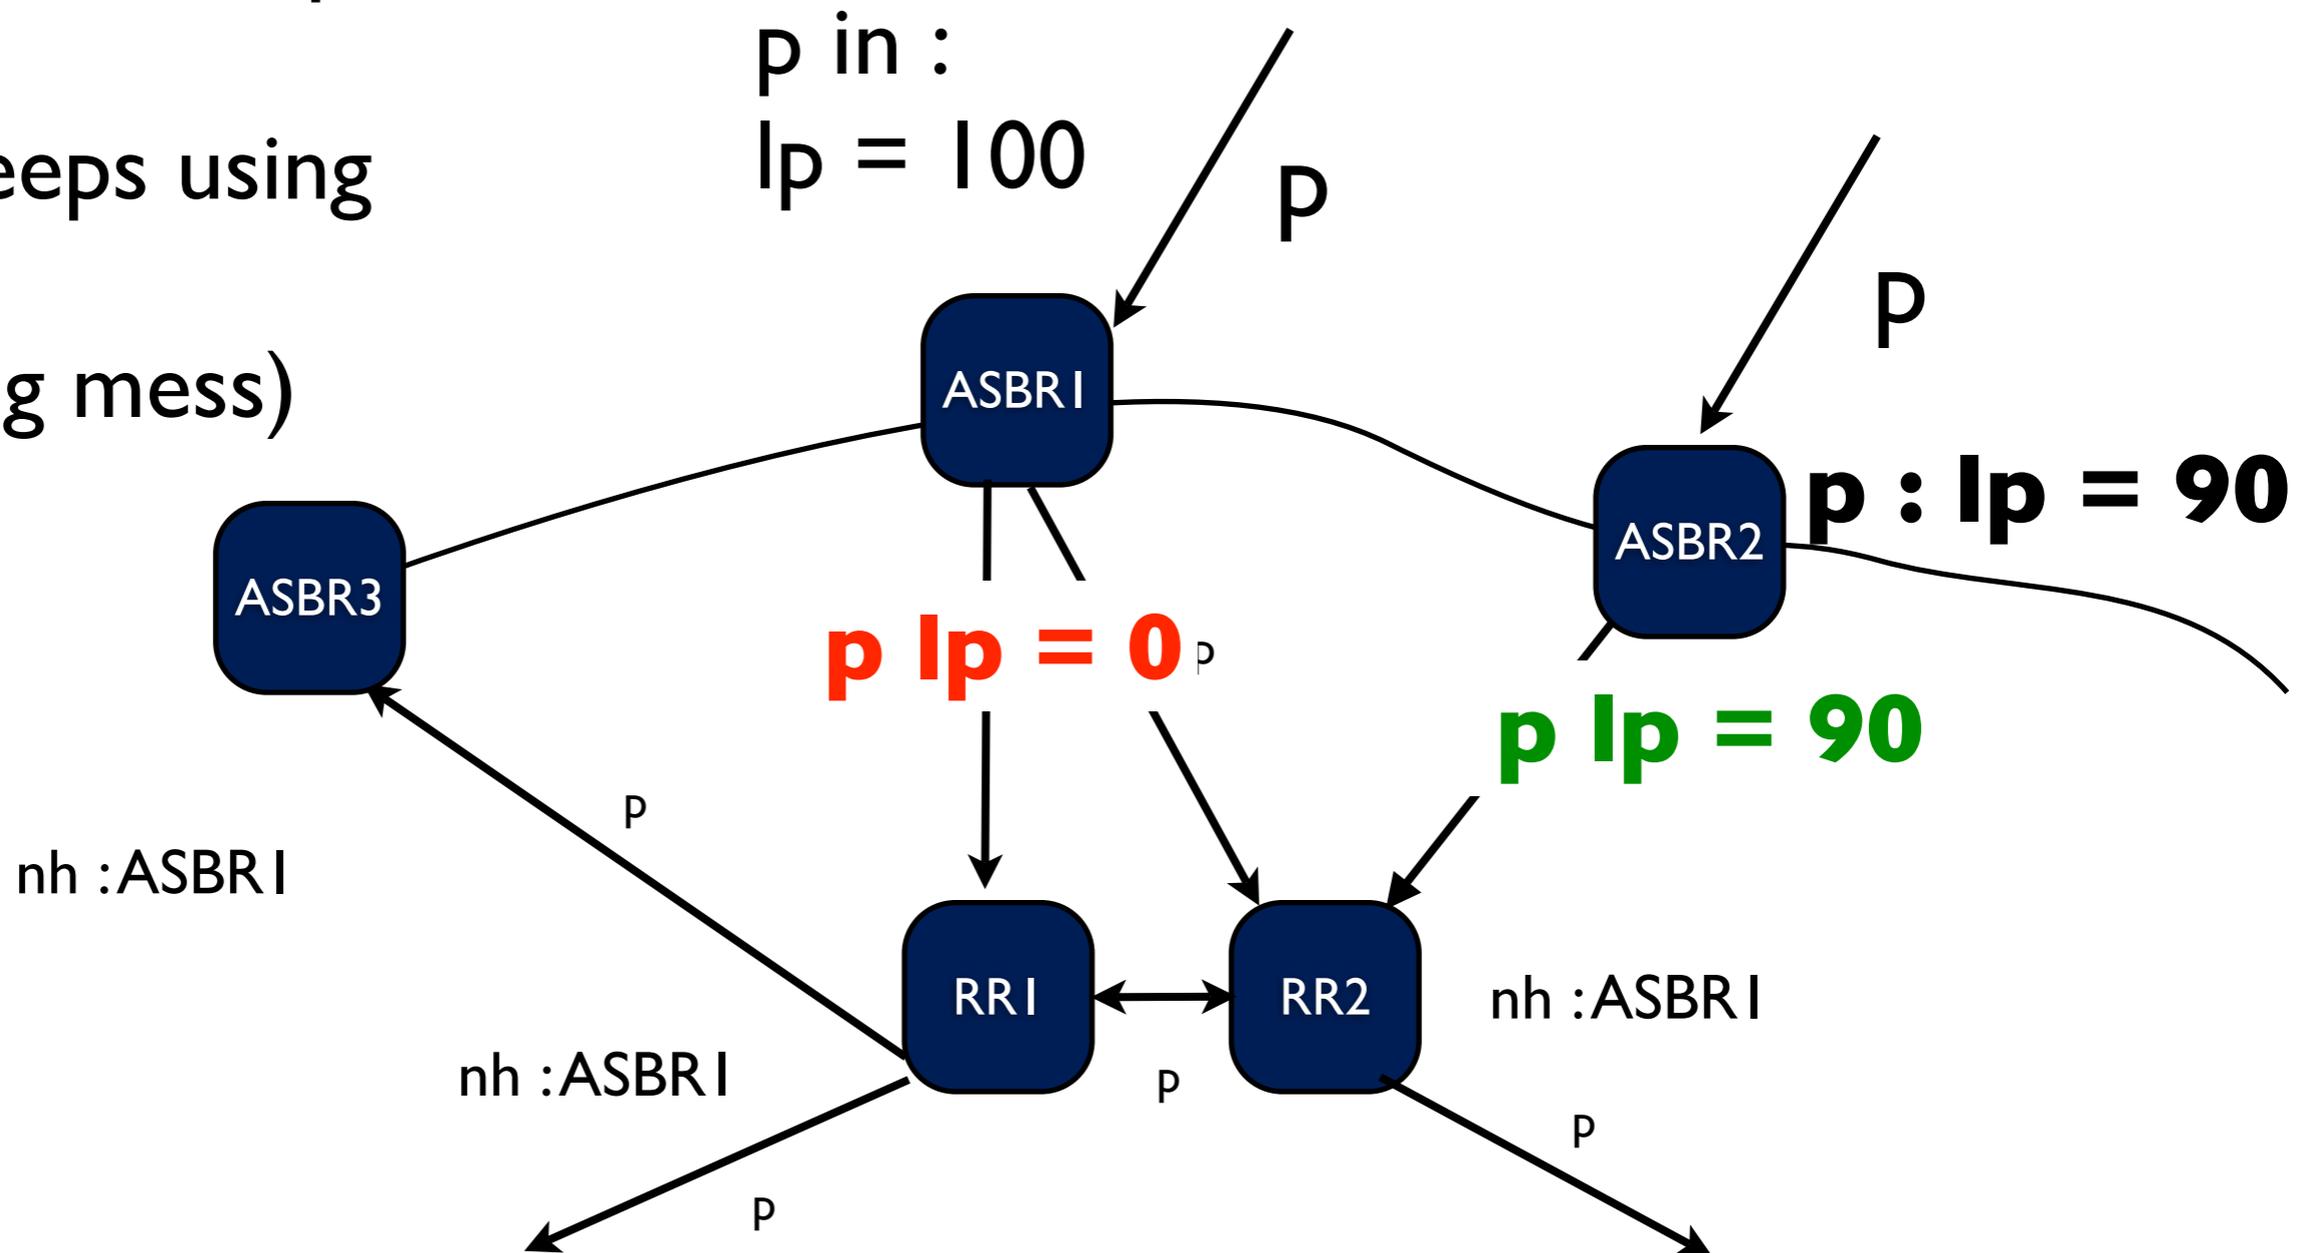
Shut



Outbound traffic Outbound LP

No more lack of path

ASBR1 keeps using
primary
(no fwding mess)



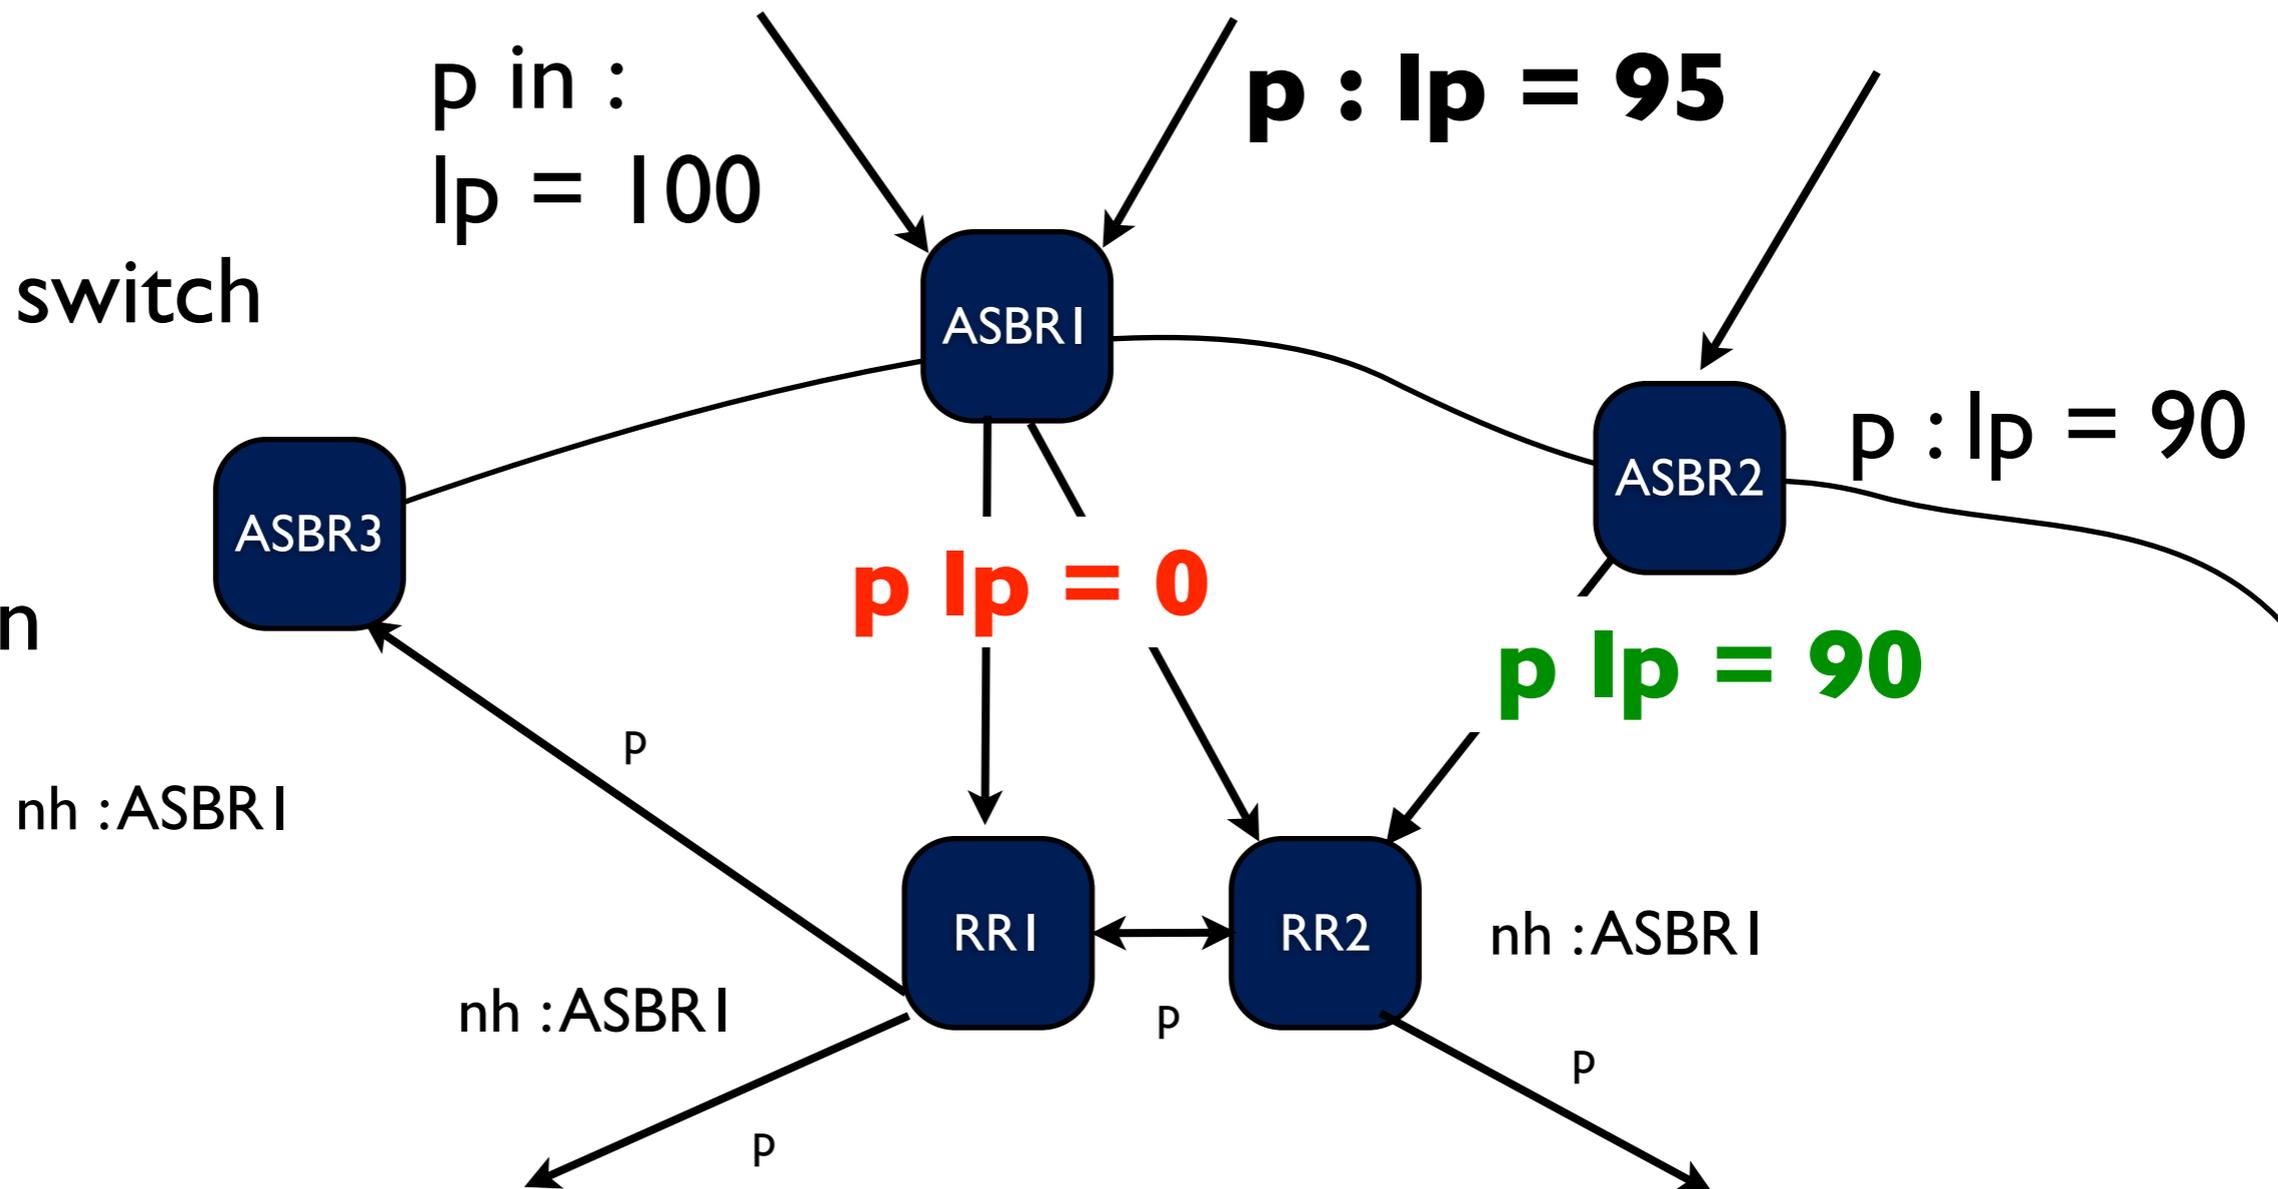
Outbound traffic

Outbound LP : glitch

Decrease LP on
outbound iBGP
session

Transient switch
to ASBR2

Update on
Shut



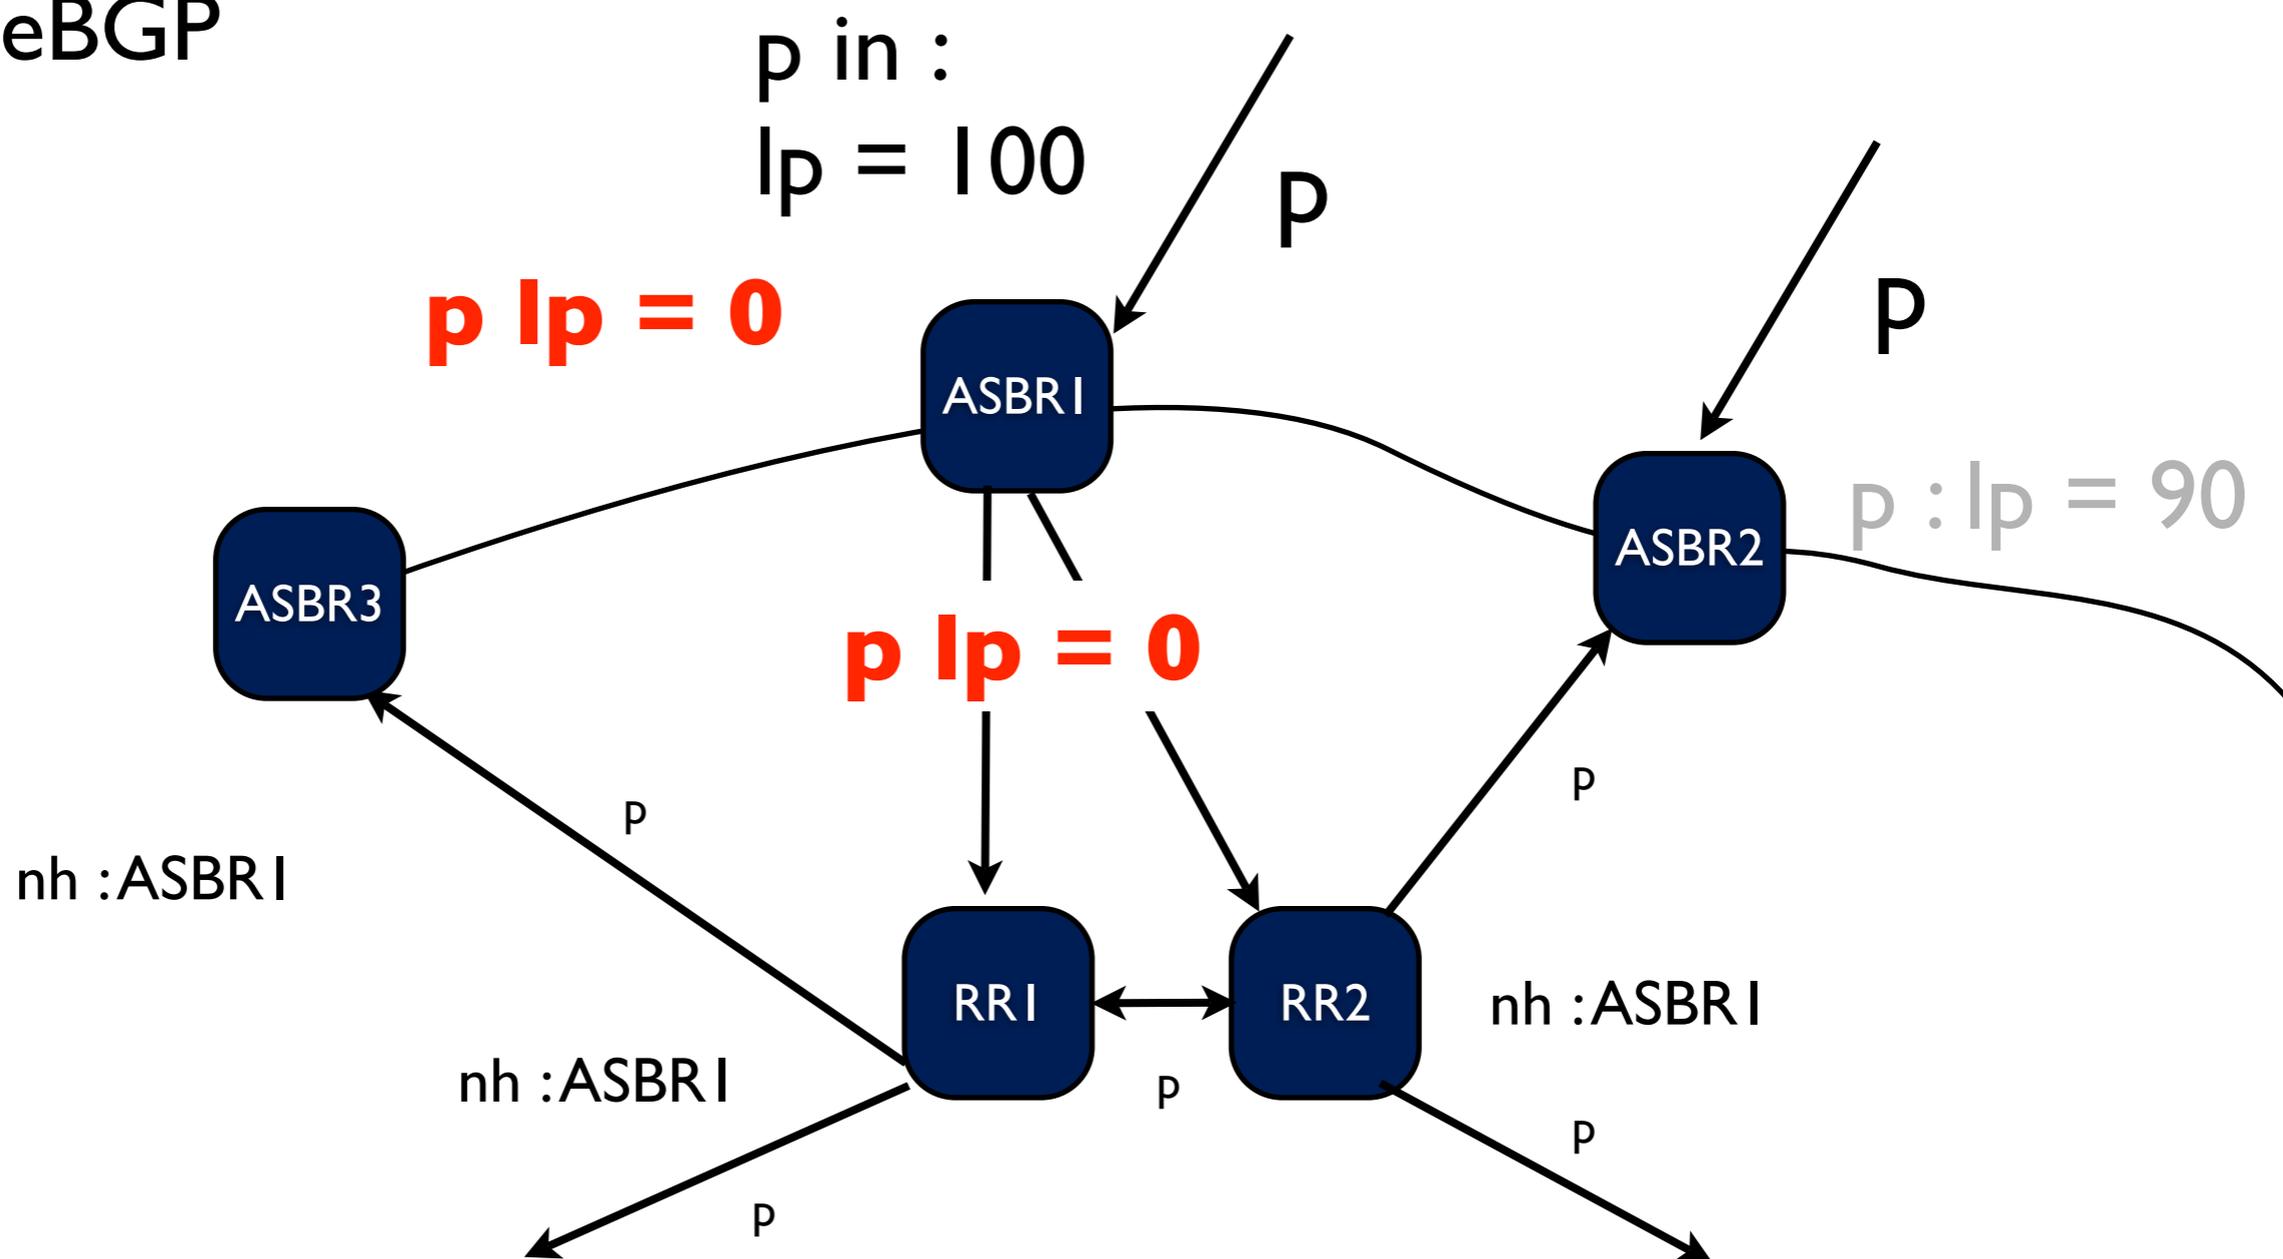
Outbound traffic

Simple : Inbound LP

Decrease LP on
inbound eBGP
session

Wait

Shut

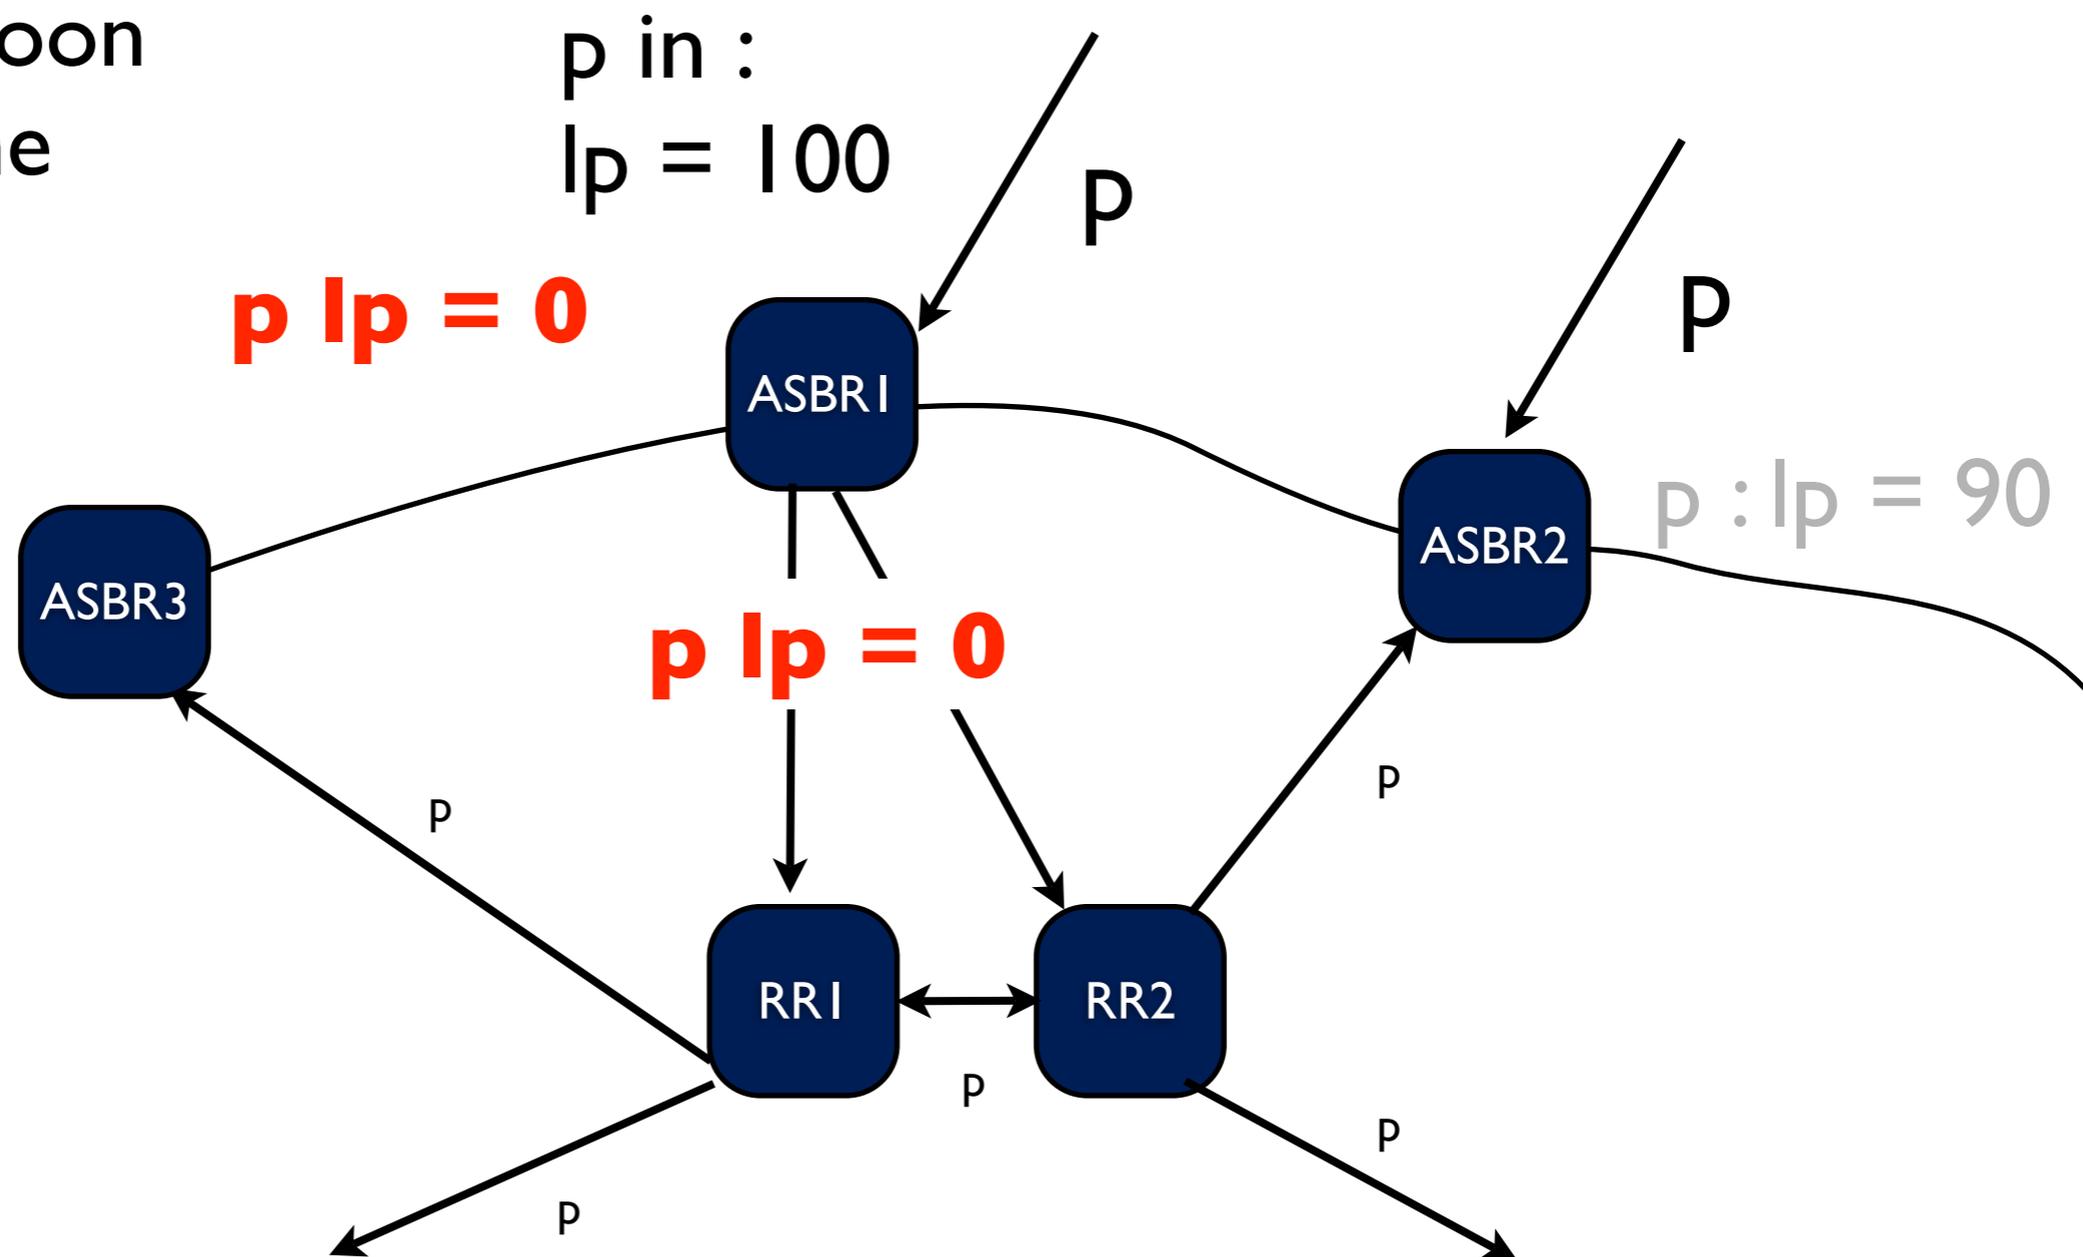


Outbound traffic Inbound LP : glitches

ASBR1 switches to
alternates as soon
as it knows one

Early primary
withdraw

fwd mess if
alternate was
not installed
(Ext. best)



Outbound traffic

Inbound LP glitches: dilemma ?

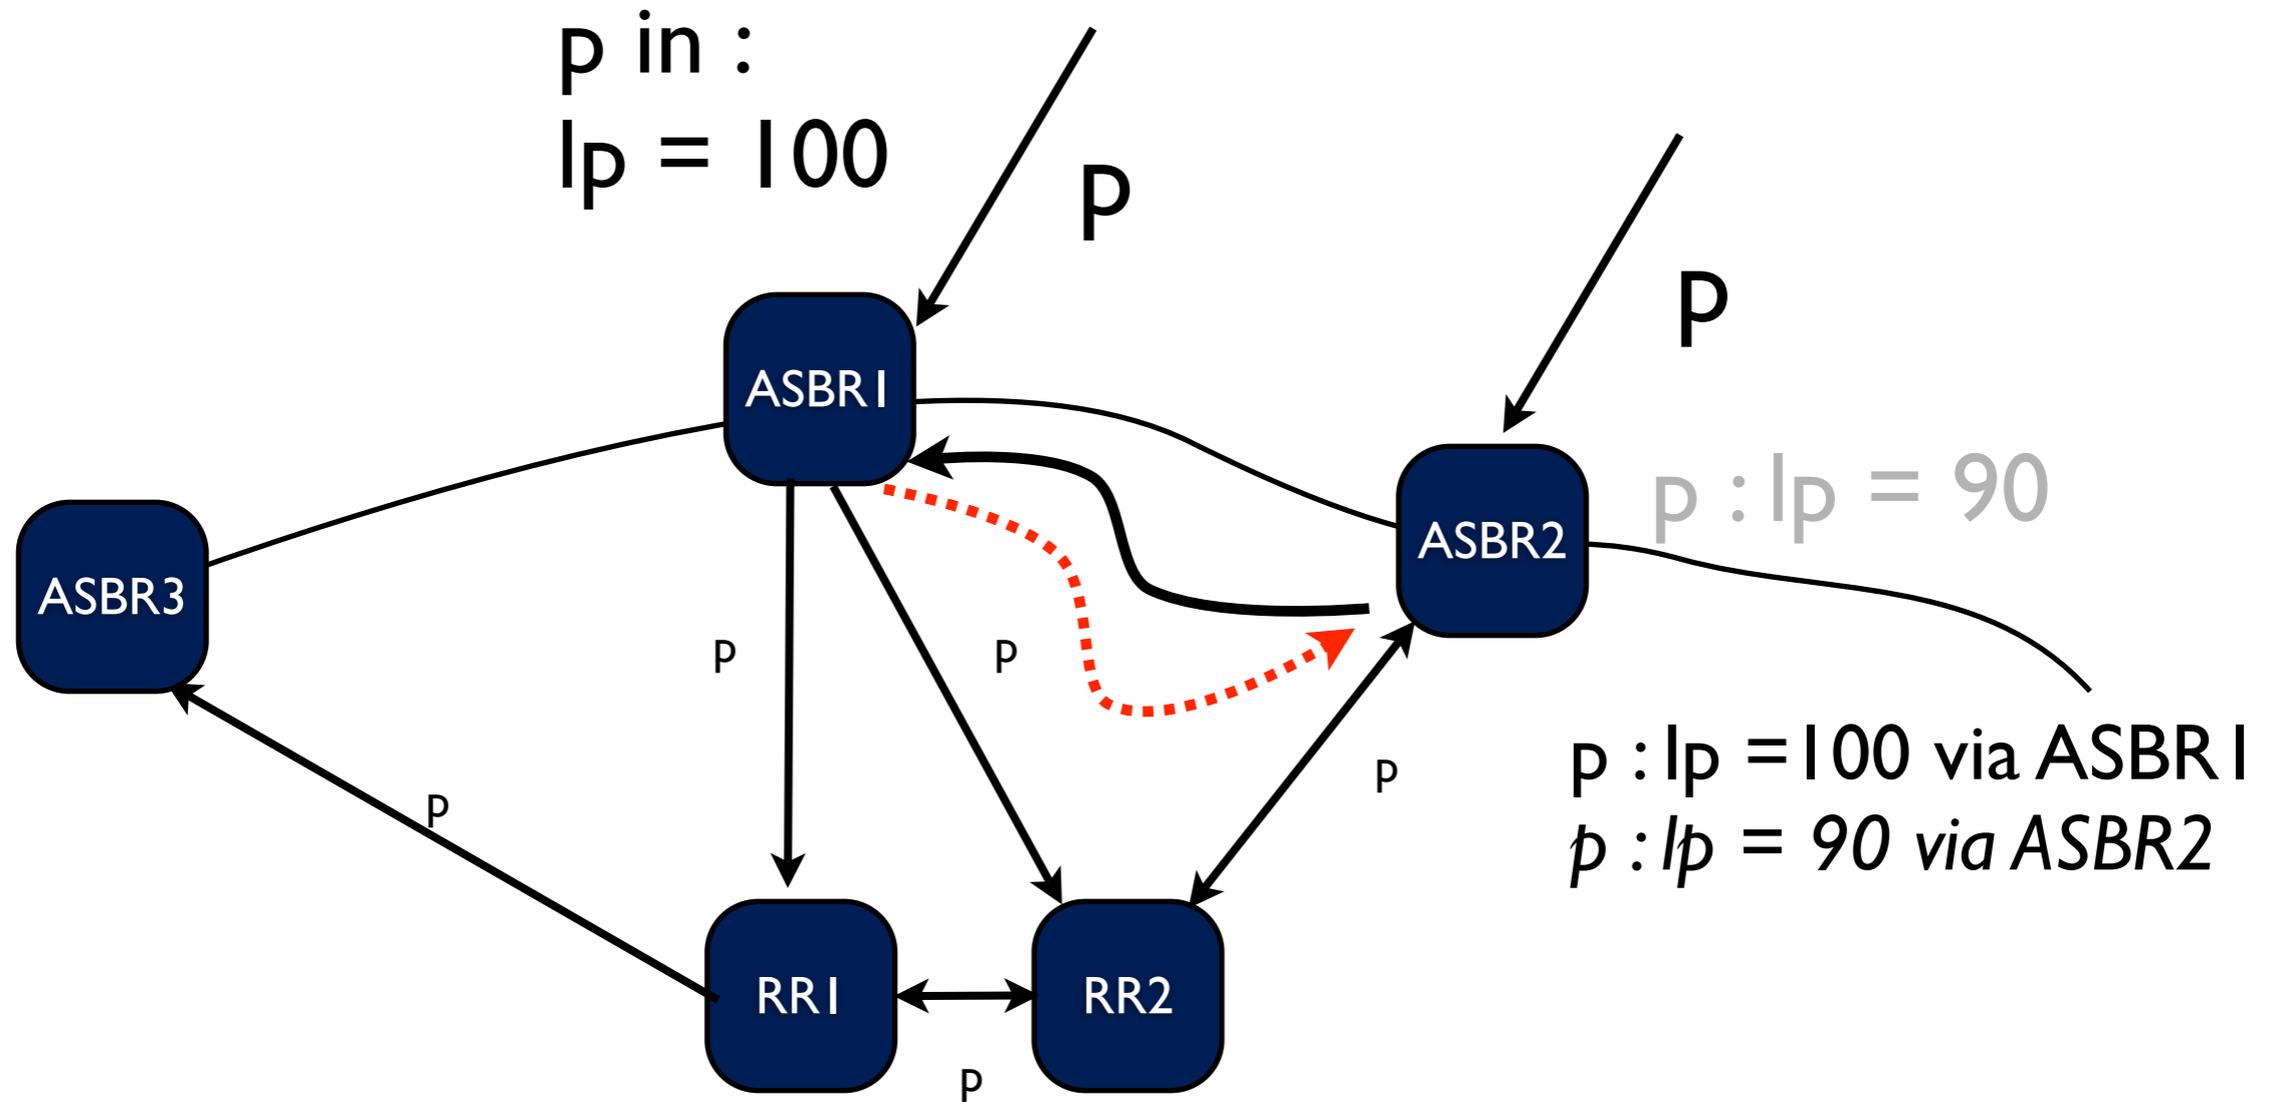
- Solving path hiding problem with Ext Best leads to transient forwarding loops with inbound LP

Outbound traffic

Inbound LP glitches: dilemma ?

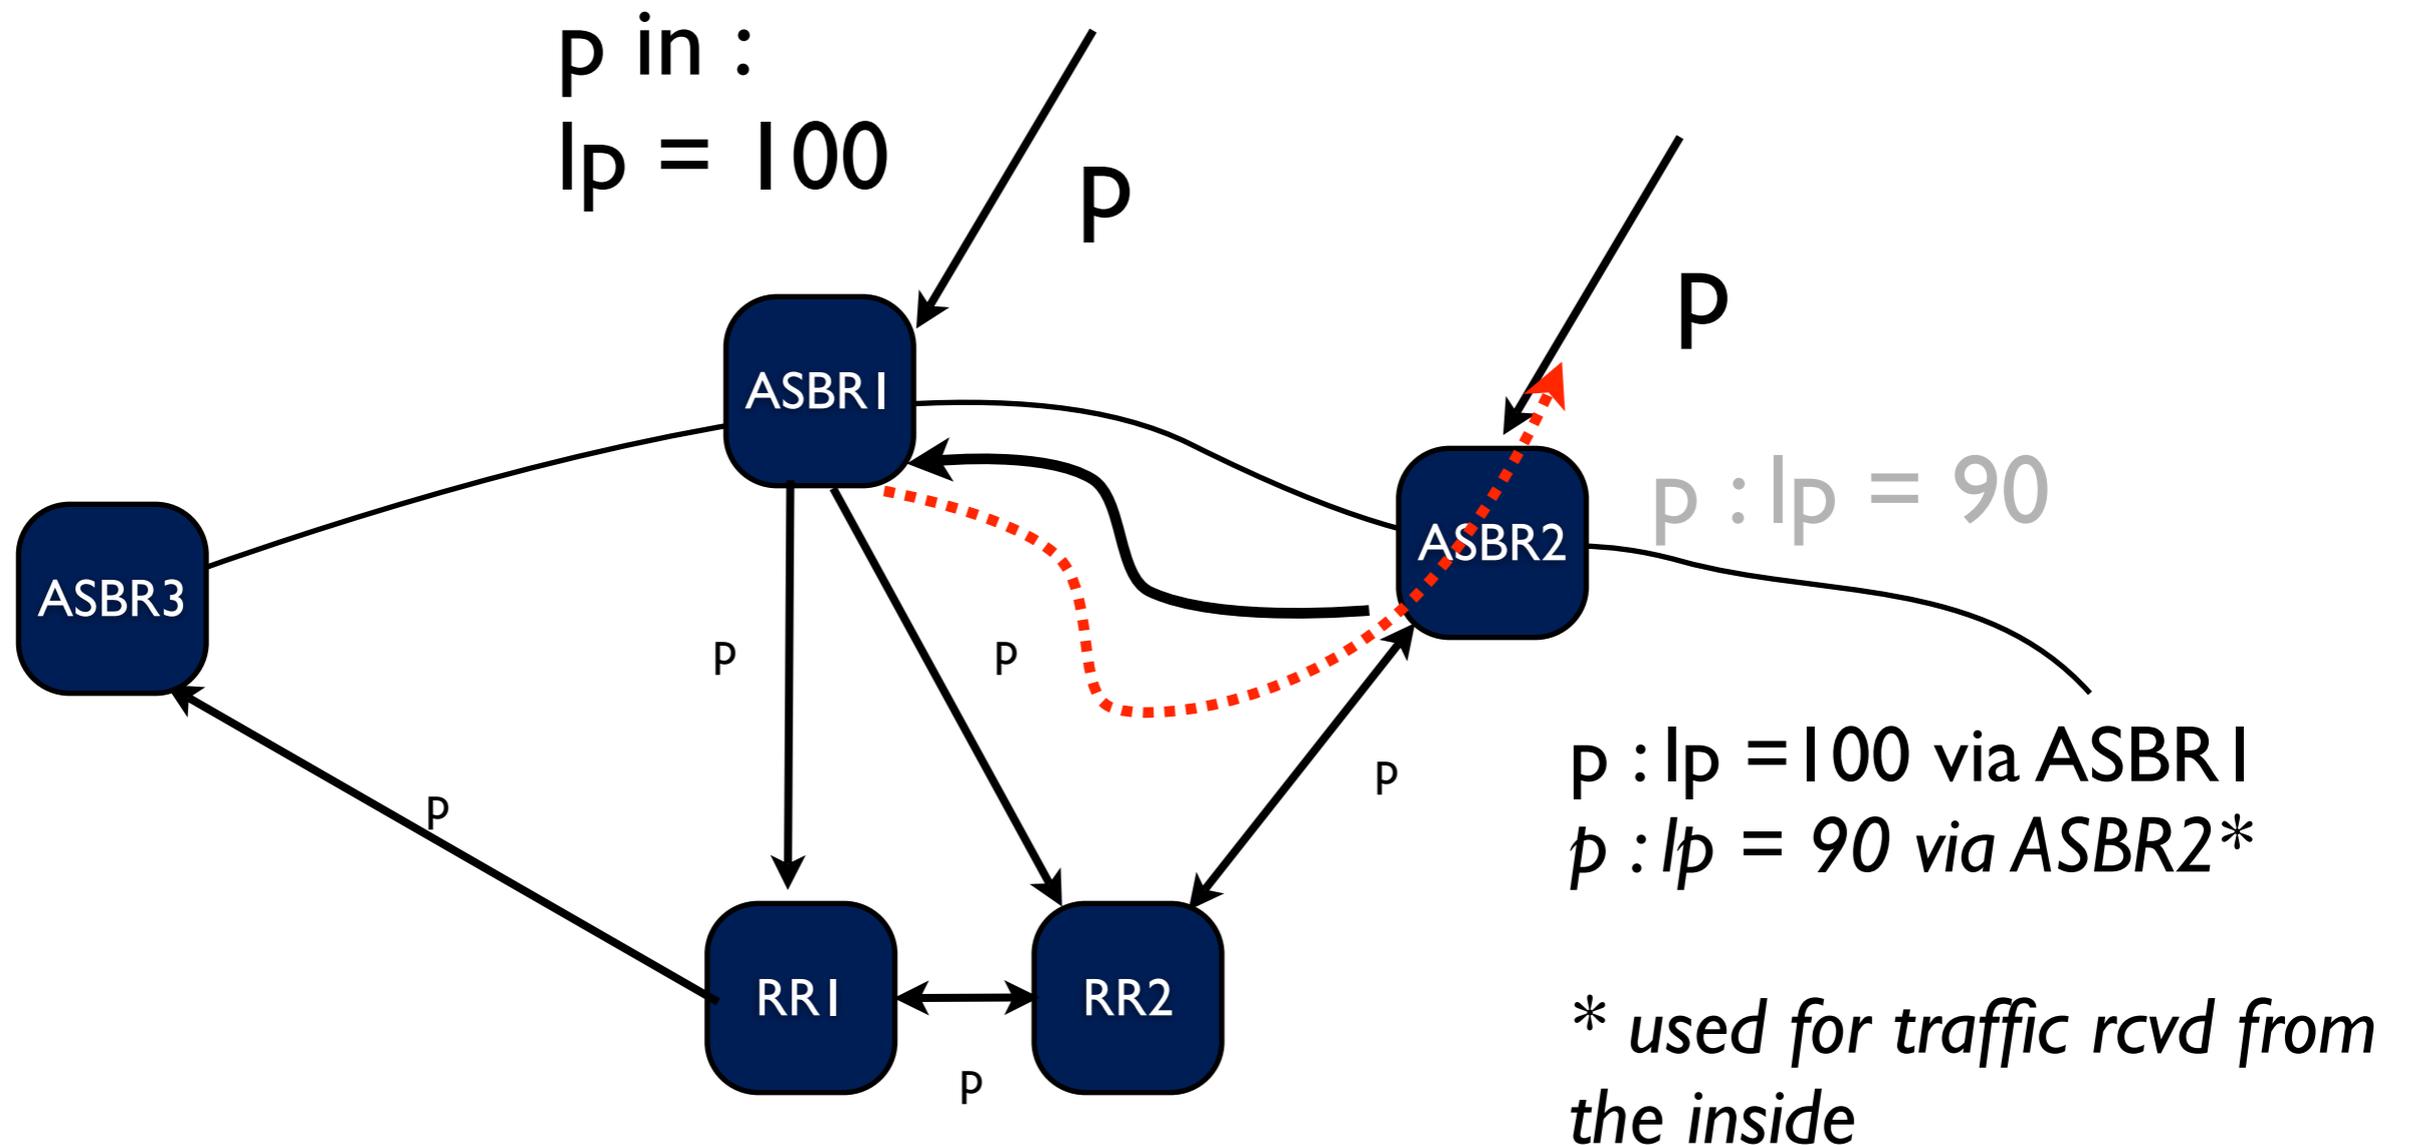
- Solving path hiding problem with Ext Best leads to transient forwarding loops with inbound LP Unless you have a proper PIC implementation !

Best External behavior



$p : l_p = 100$ via ASBR1
 $p : l_p = 90$ via ASBR2

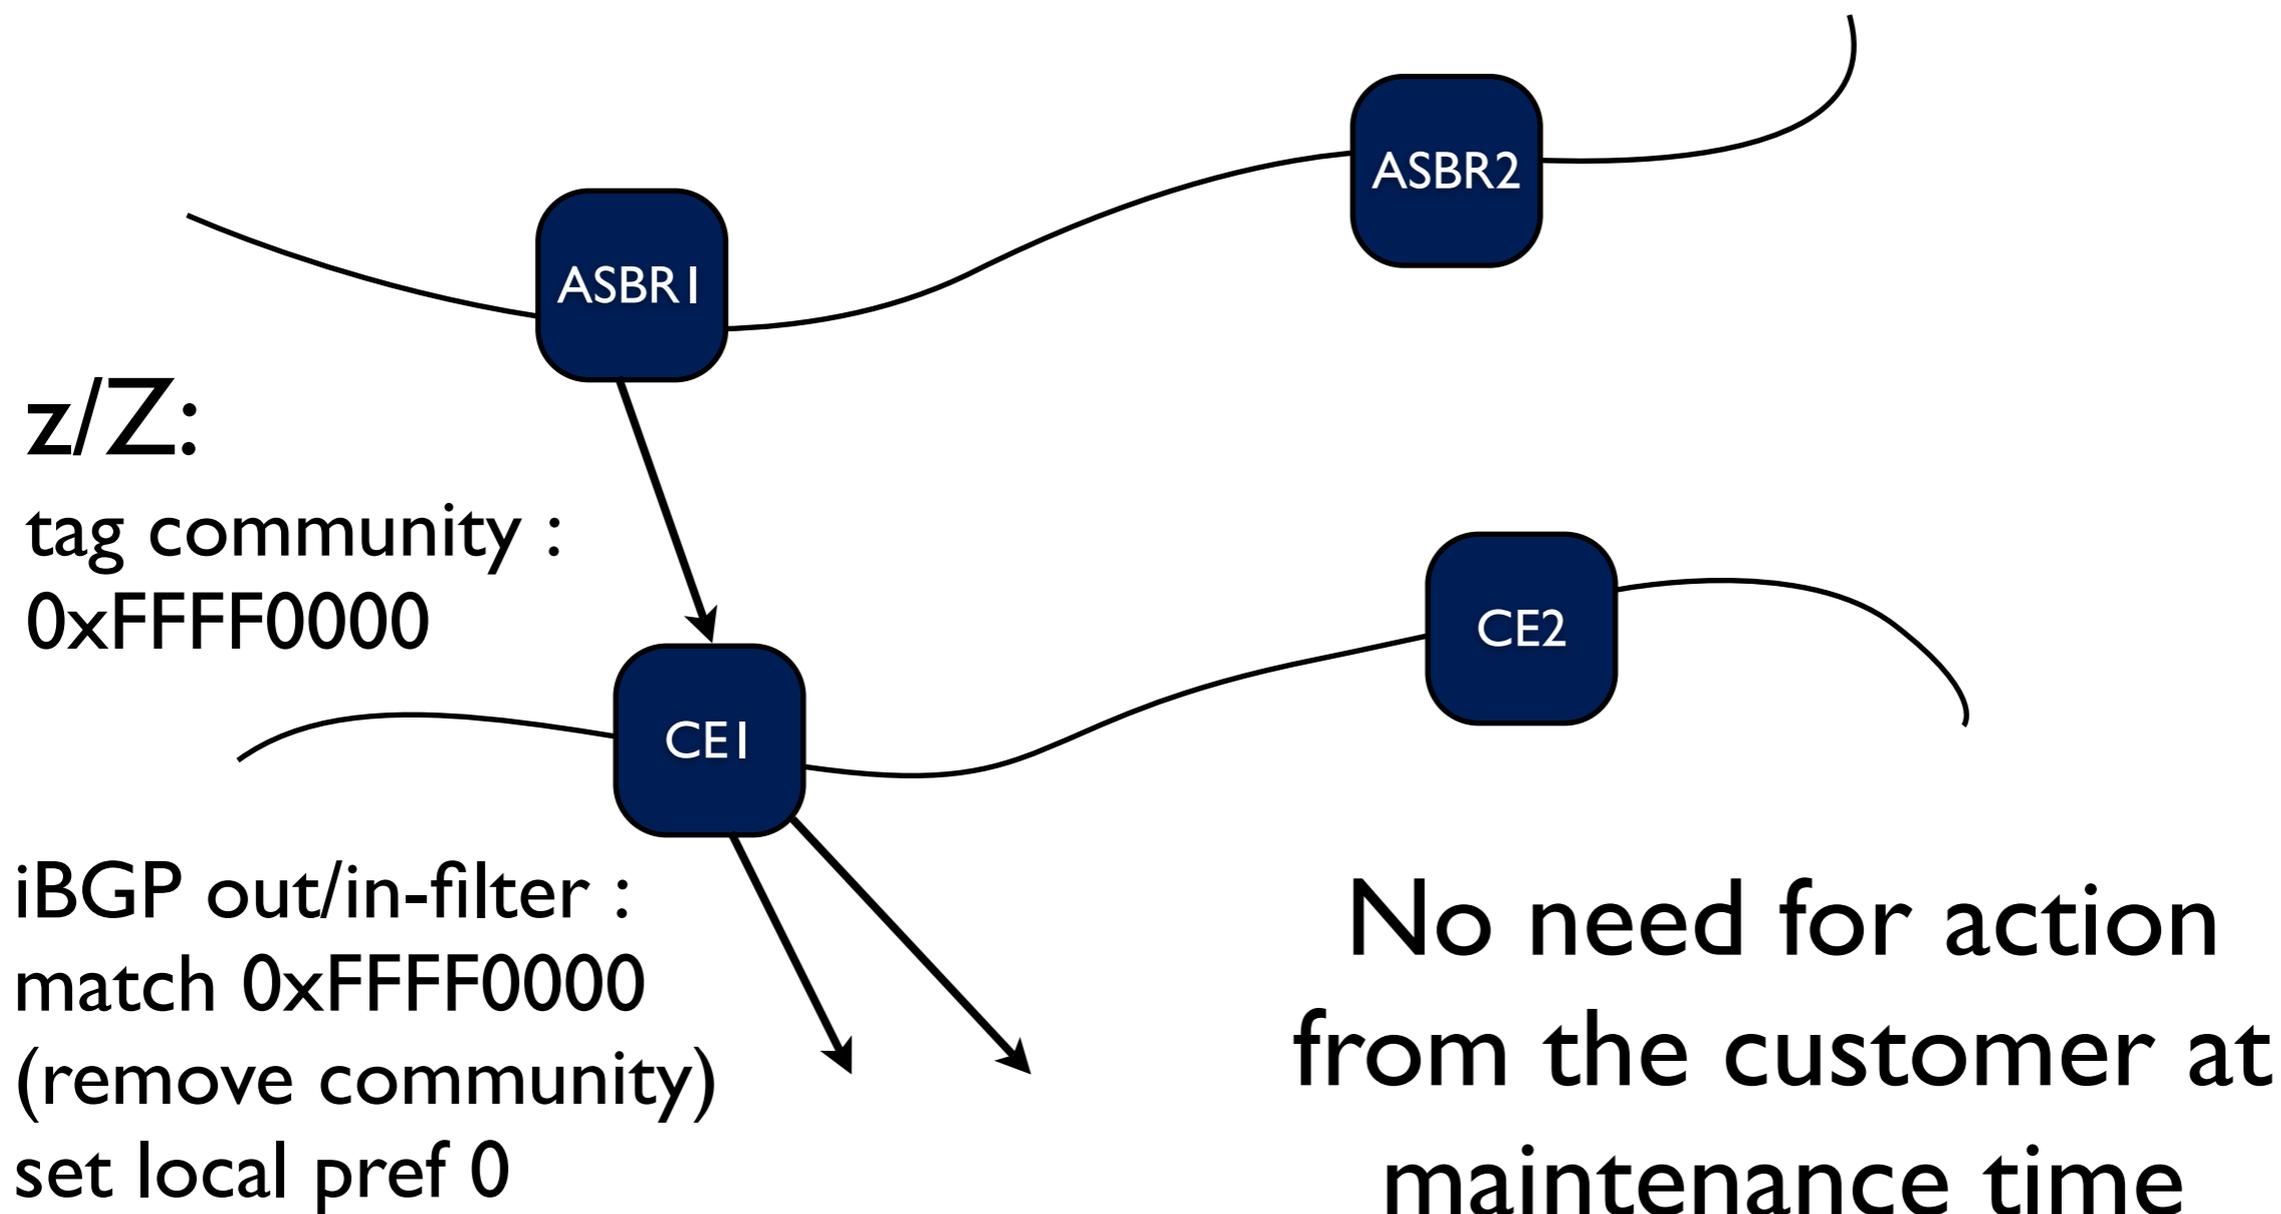
Clever Best External behavior



p : lp = 100 via ASBR1
p : lp = 90 via ASBR2

Inbound traffic

Trigger outbound g-shut at the other side of the peering link



Inbound traffic

Boring aspect of GSHUT community

- Need to handle its non-transitive essence
- Option to use a *non transitive* extended community
`draft-decraene-idr-reserved-extended-communities-00`

Next steps

- Radical draft verbosity decrease
- Simplified procedure when obtained FCFS non transitive extended community
- Needs IDR agreement