

BGP Graceful Shutdown

draft-francois-bgp-gshut-01.txt

Pierre Francois

UCL-BE

Bruno Decraene

France Telecom - Orange

Cristel Pelsser

NTT

Clarence Filsfils

Cisco Systems

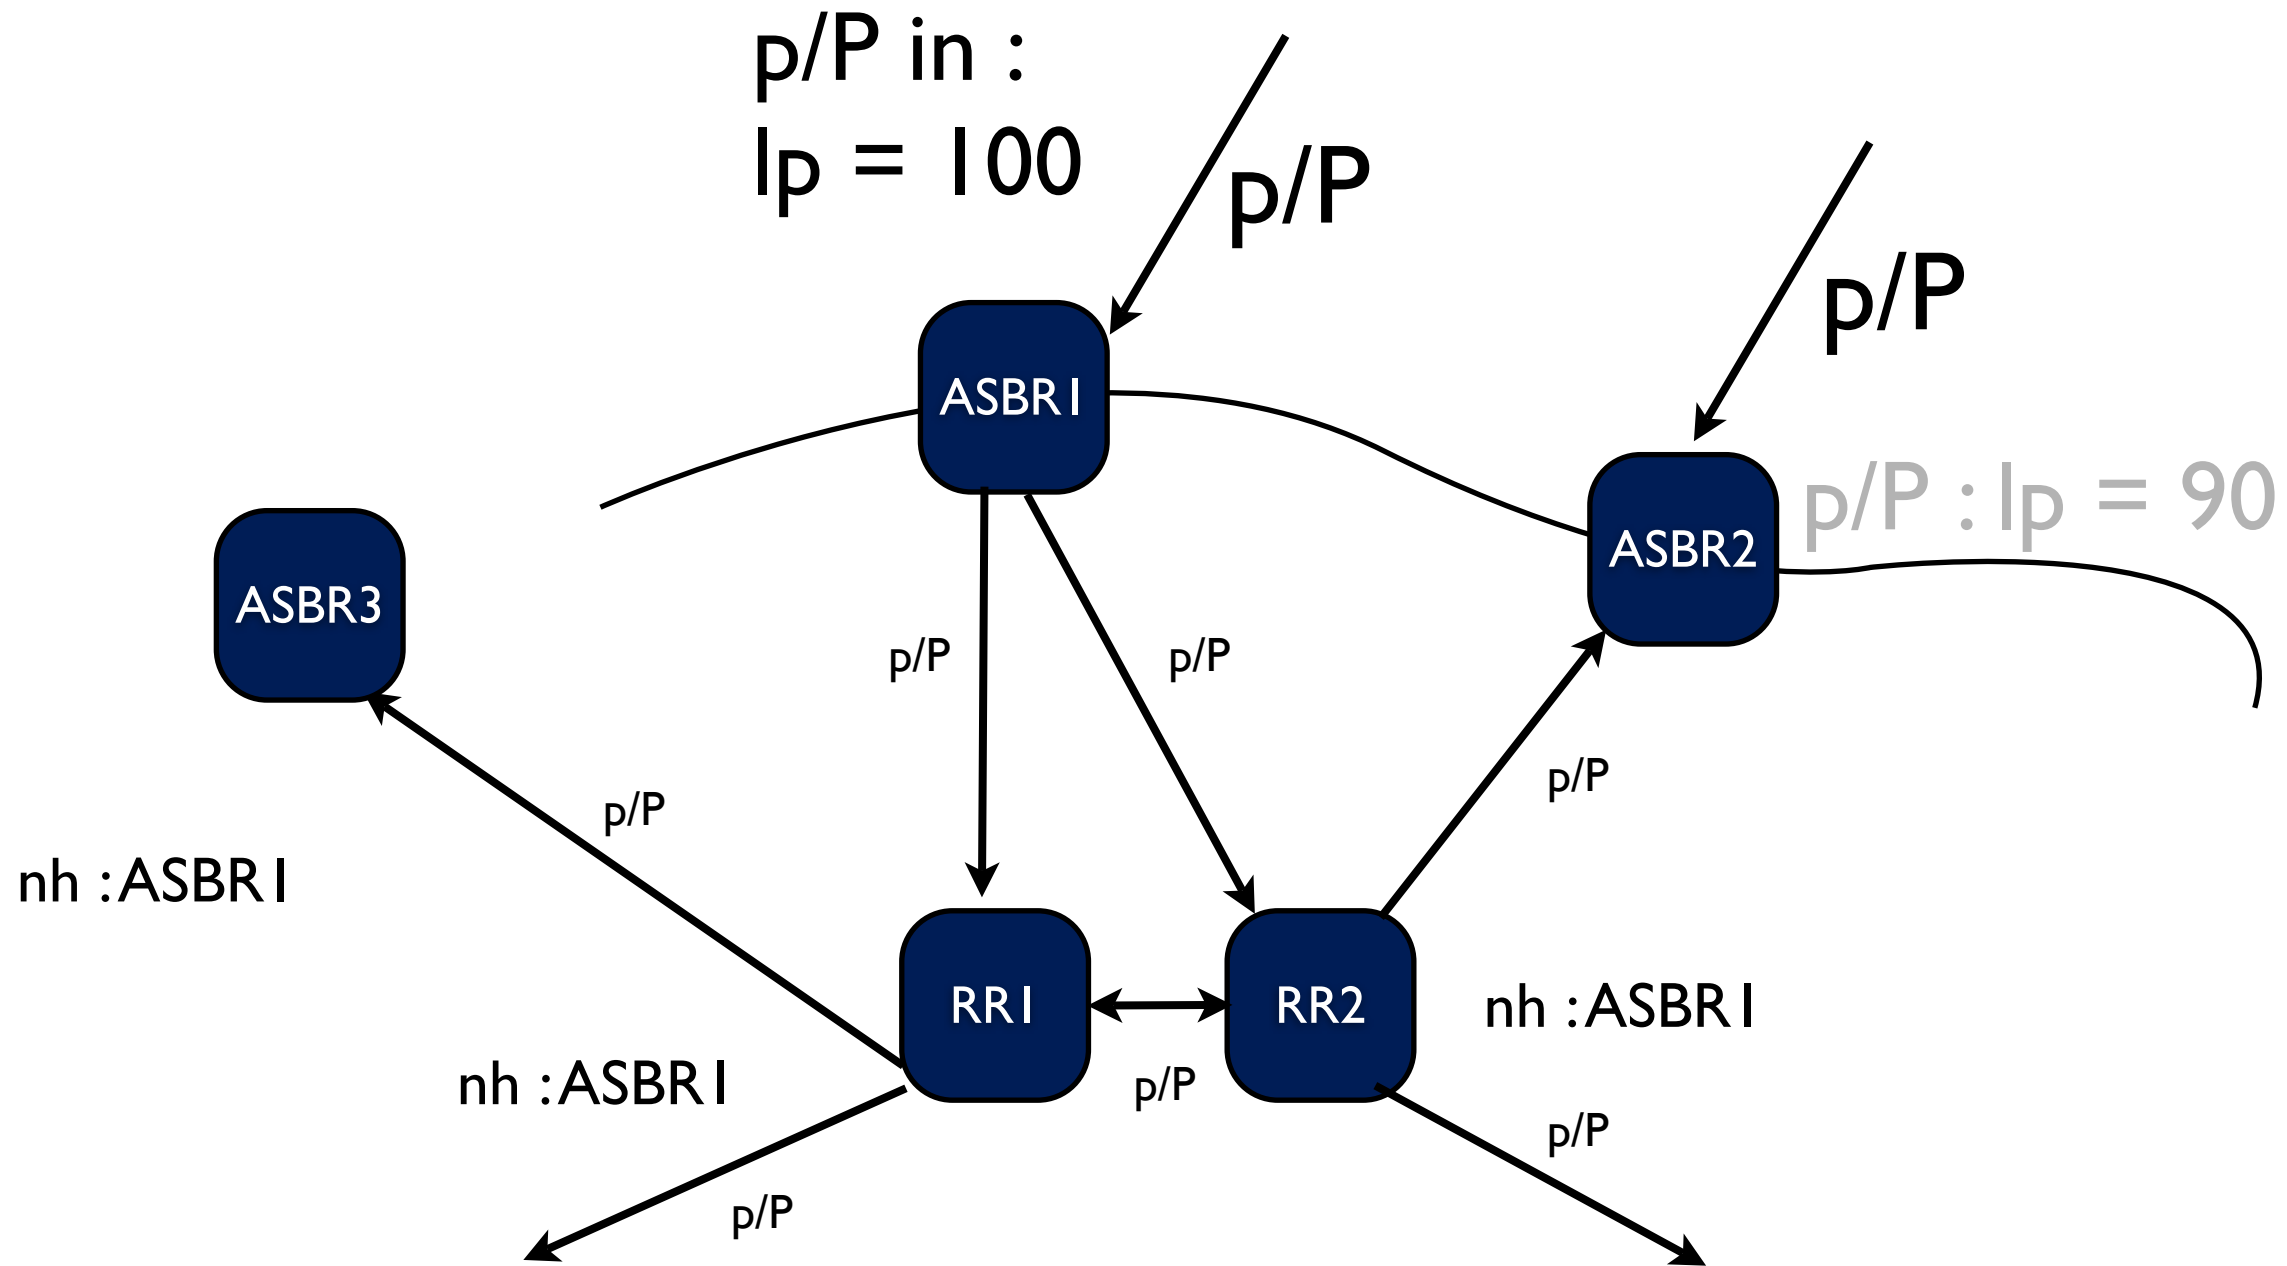
Agenda

- Brief reminder
- Changes since -00

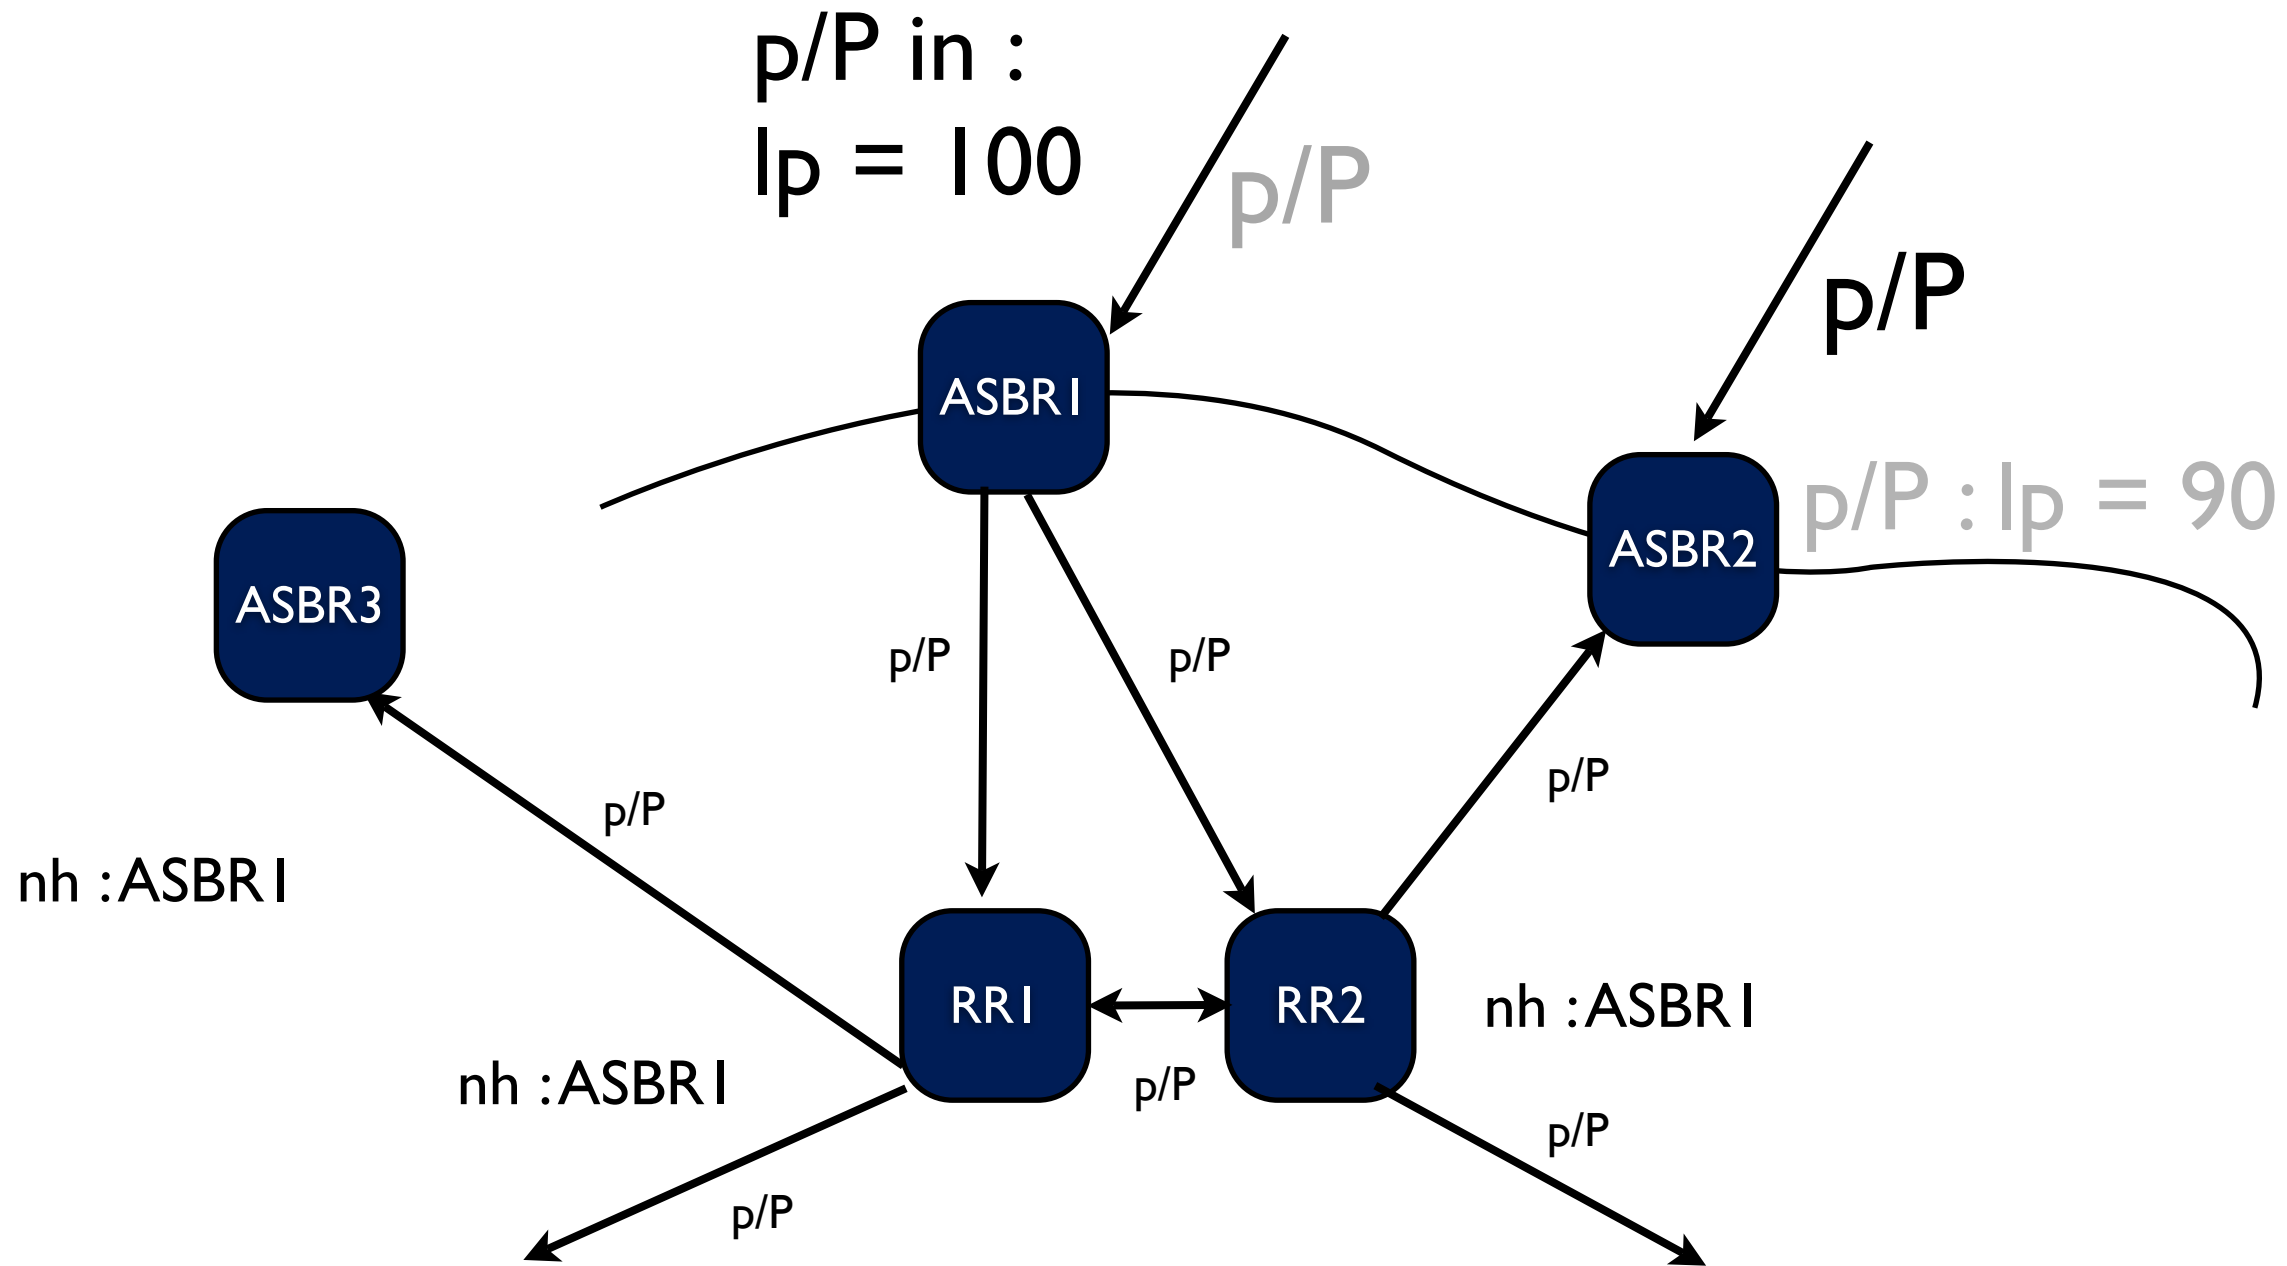
Goal

- Avoid traffic loss upon manual eBGP session shutdown
- Due to transient lack of path during convergence within the AS
- Operational procedures

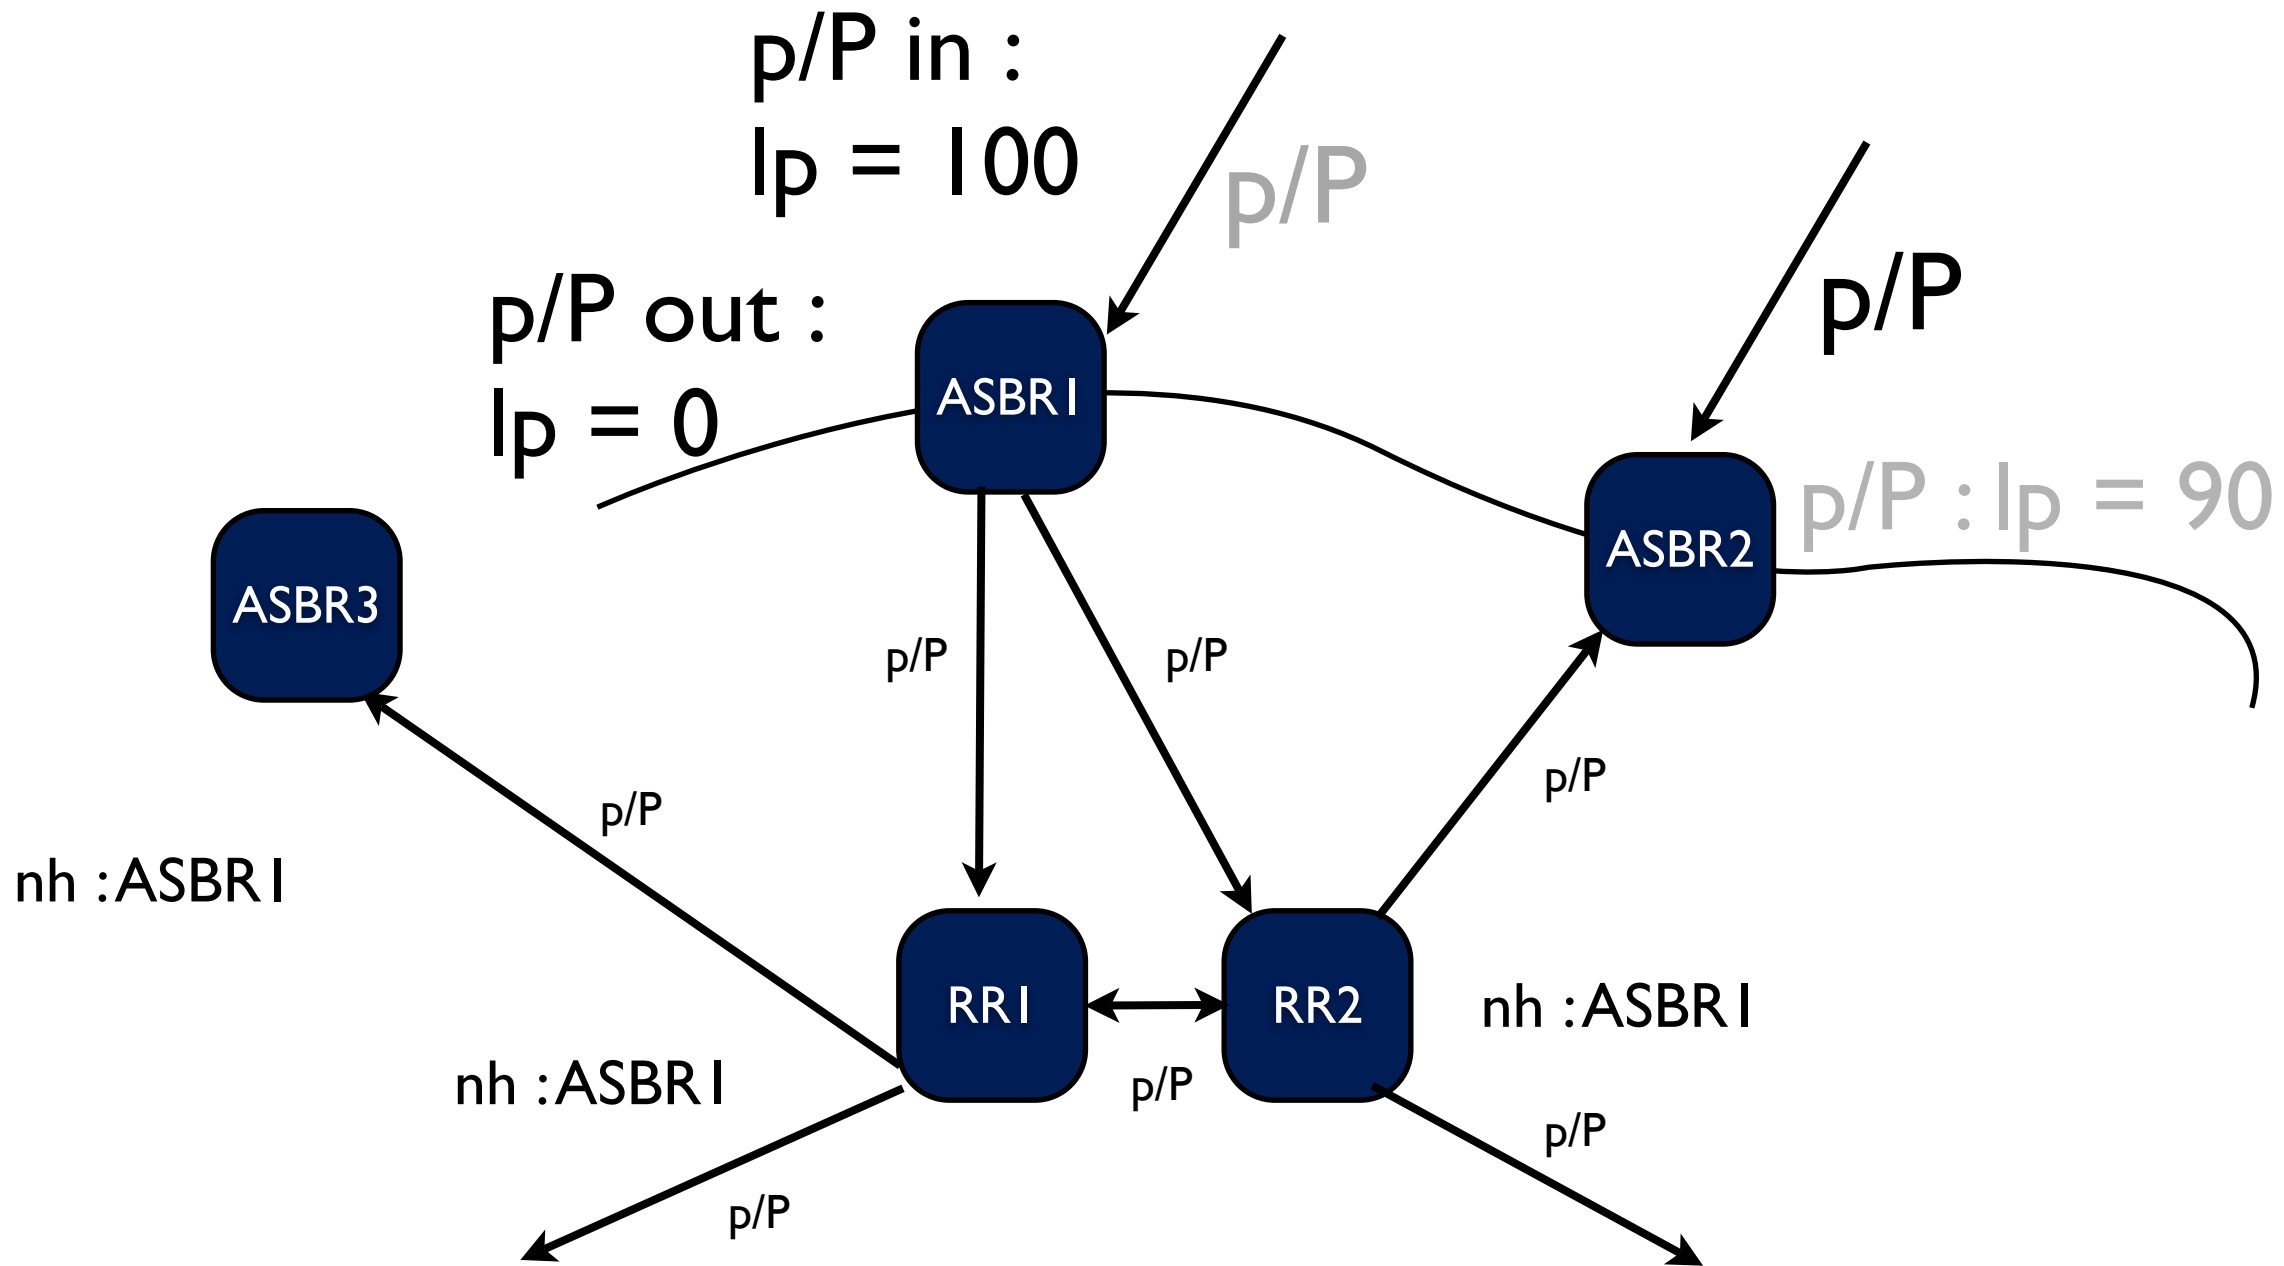
G-Shut : Outbound traffic



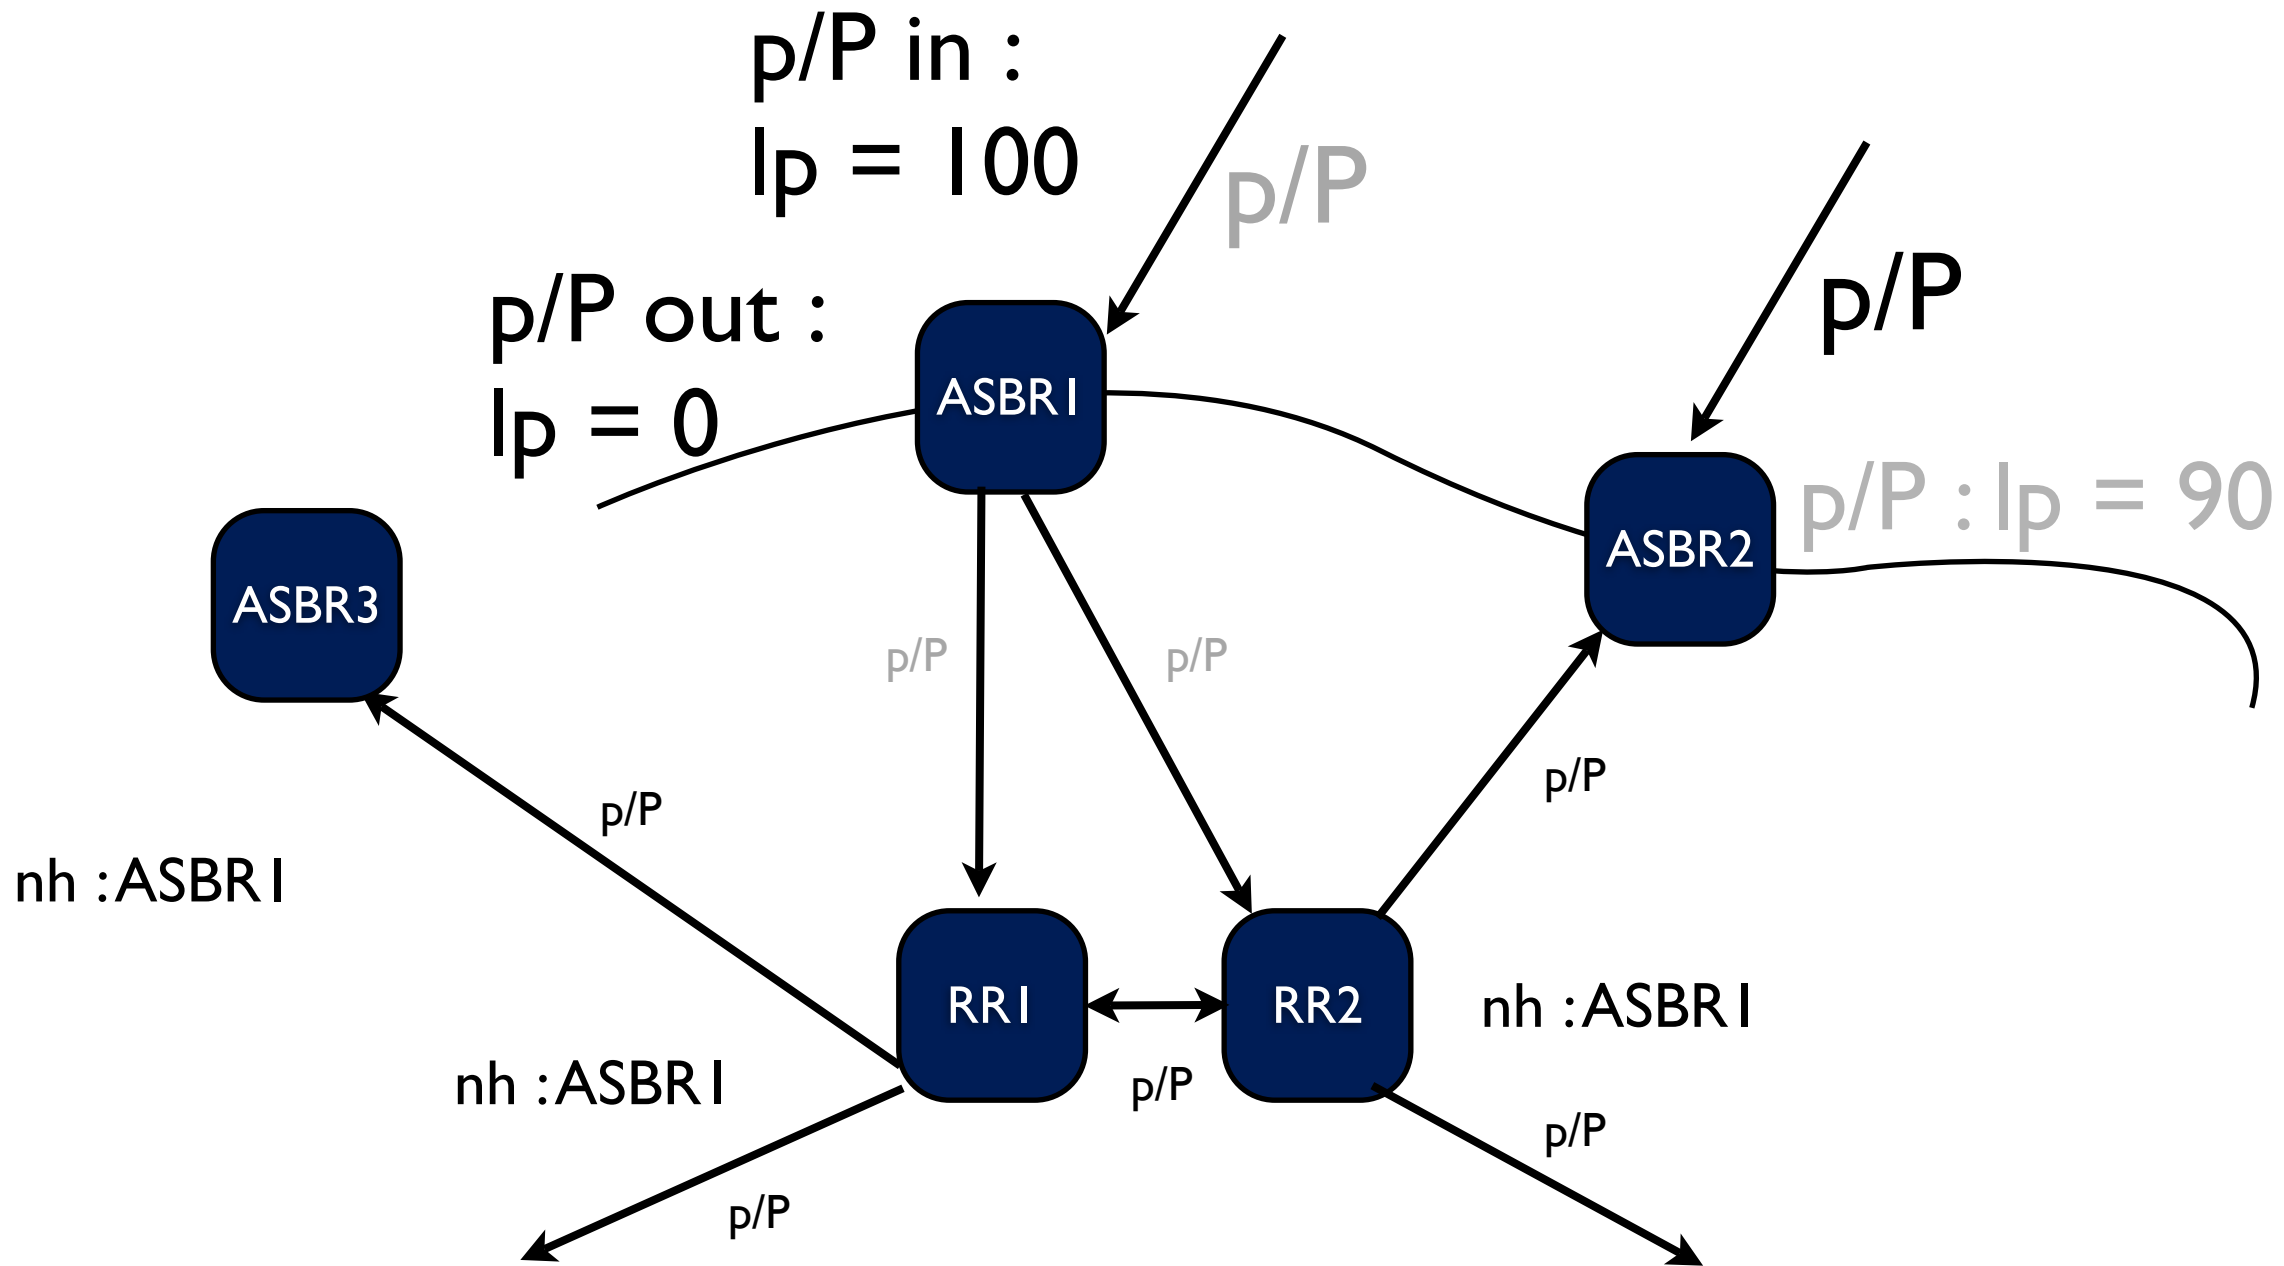
G-Shut : Outbound traffic



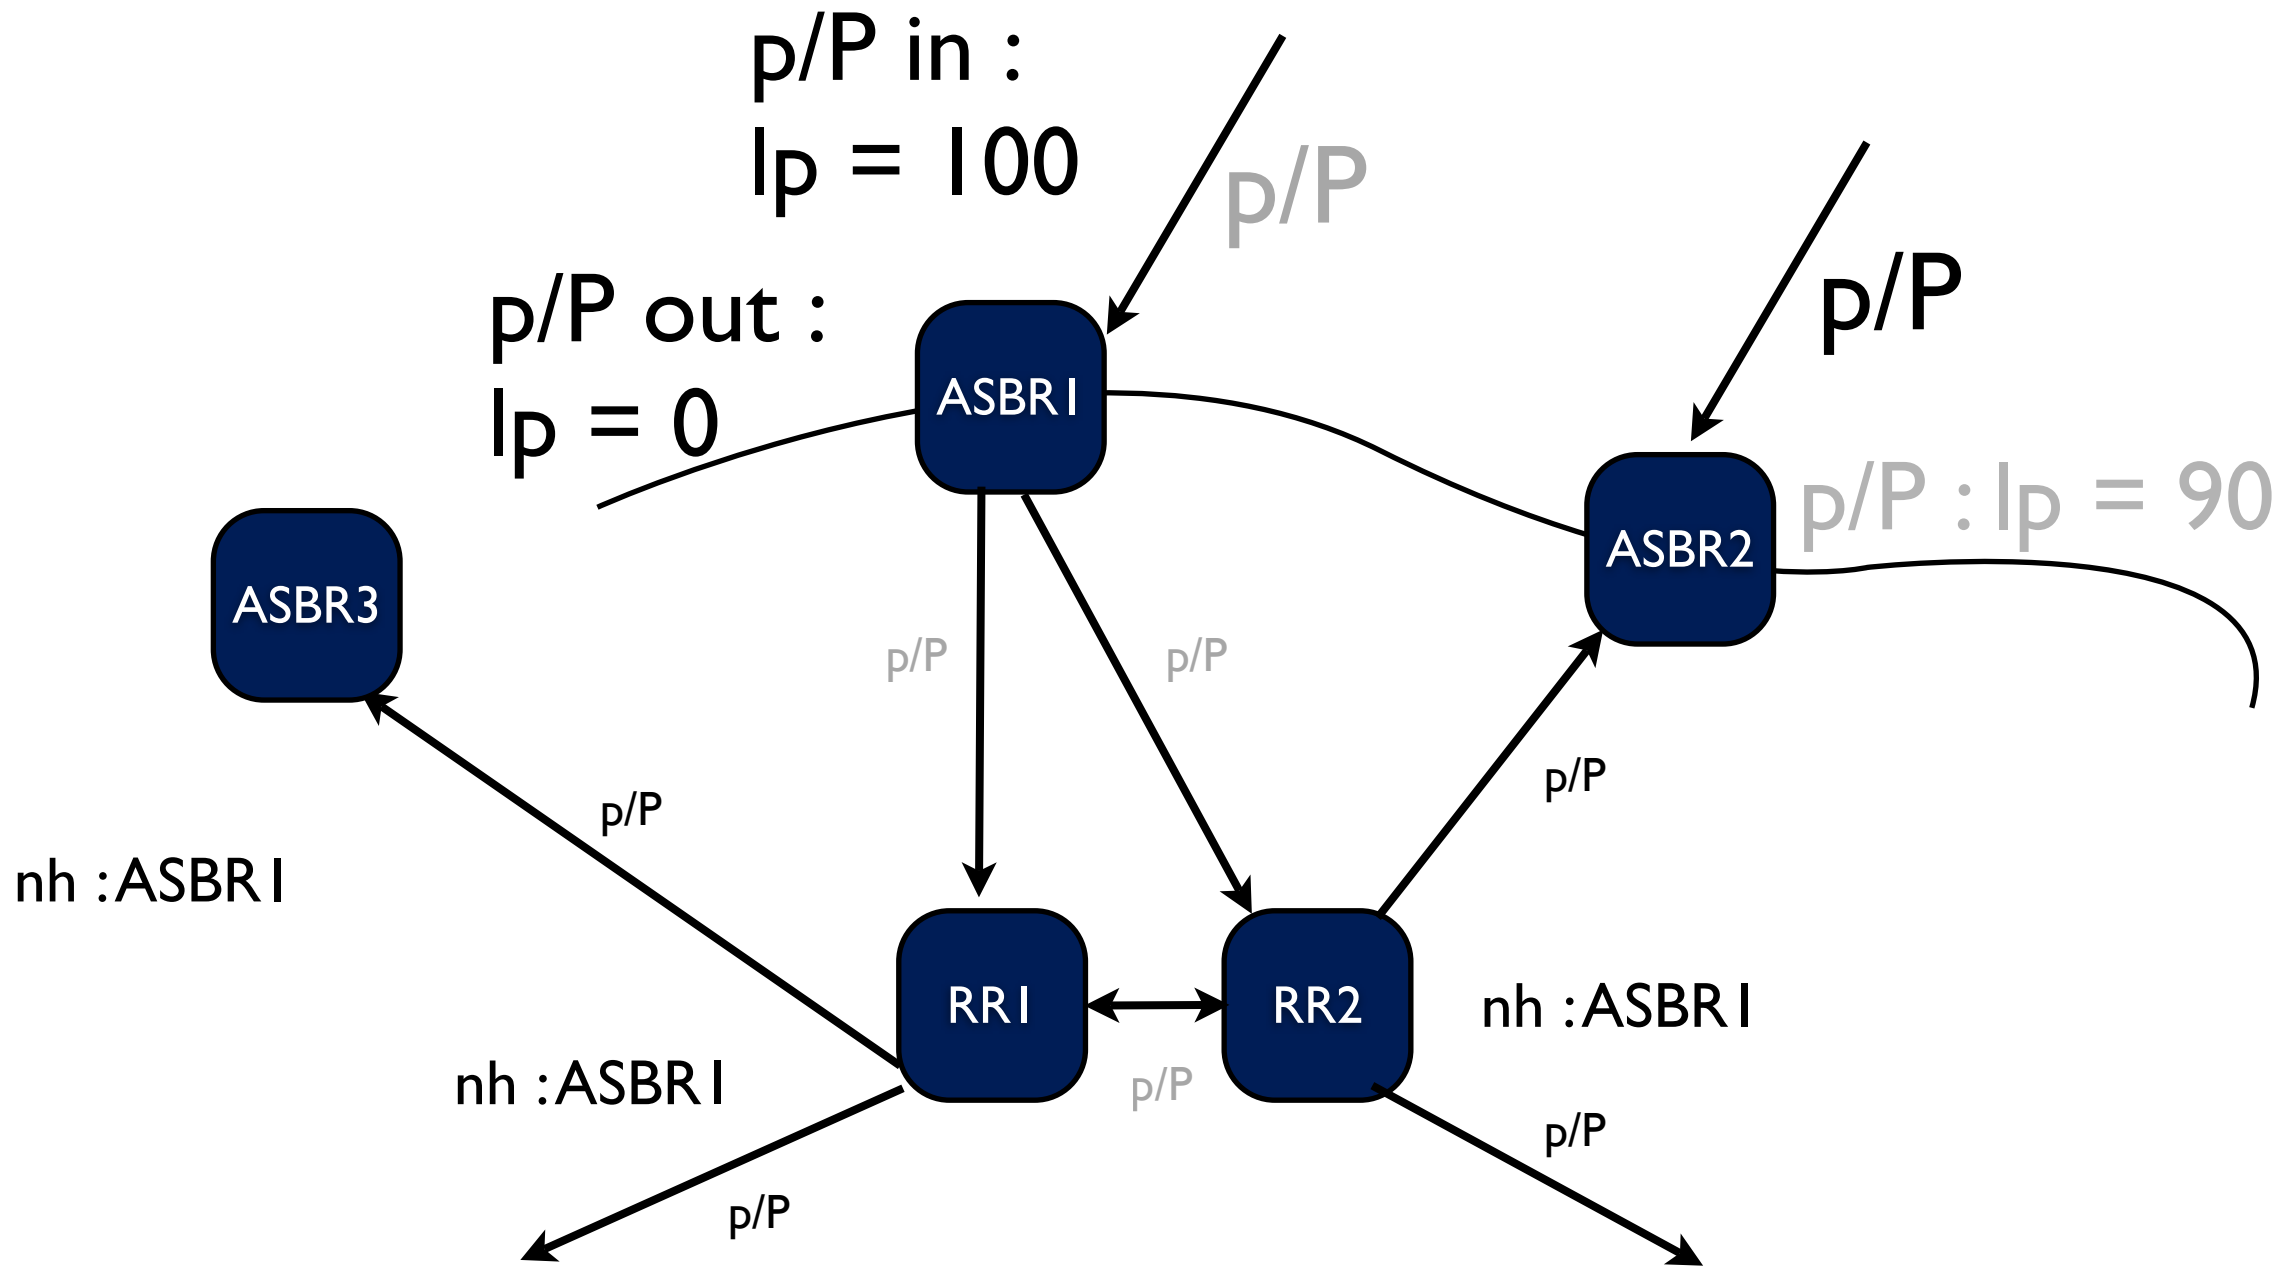
G-Shut : Outbound traffic



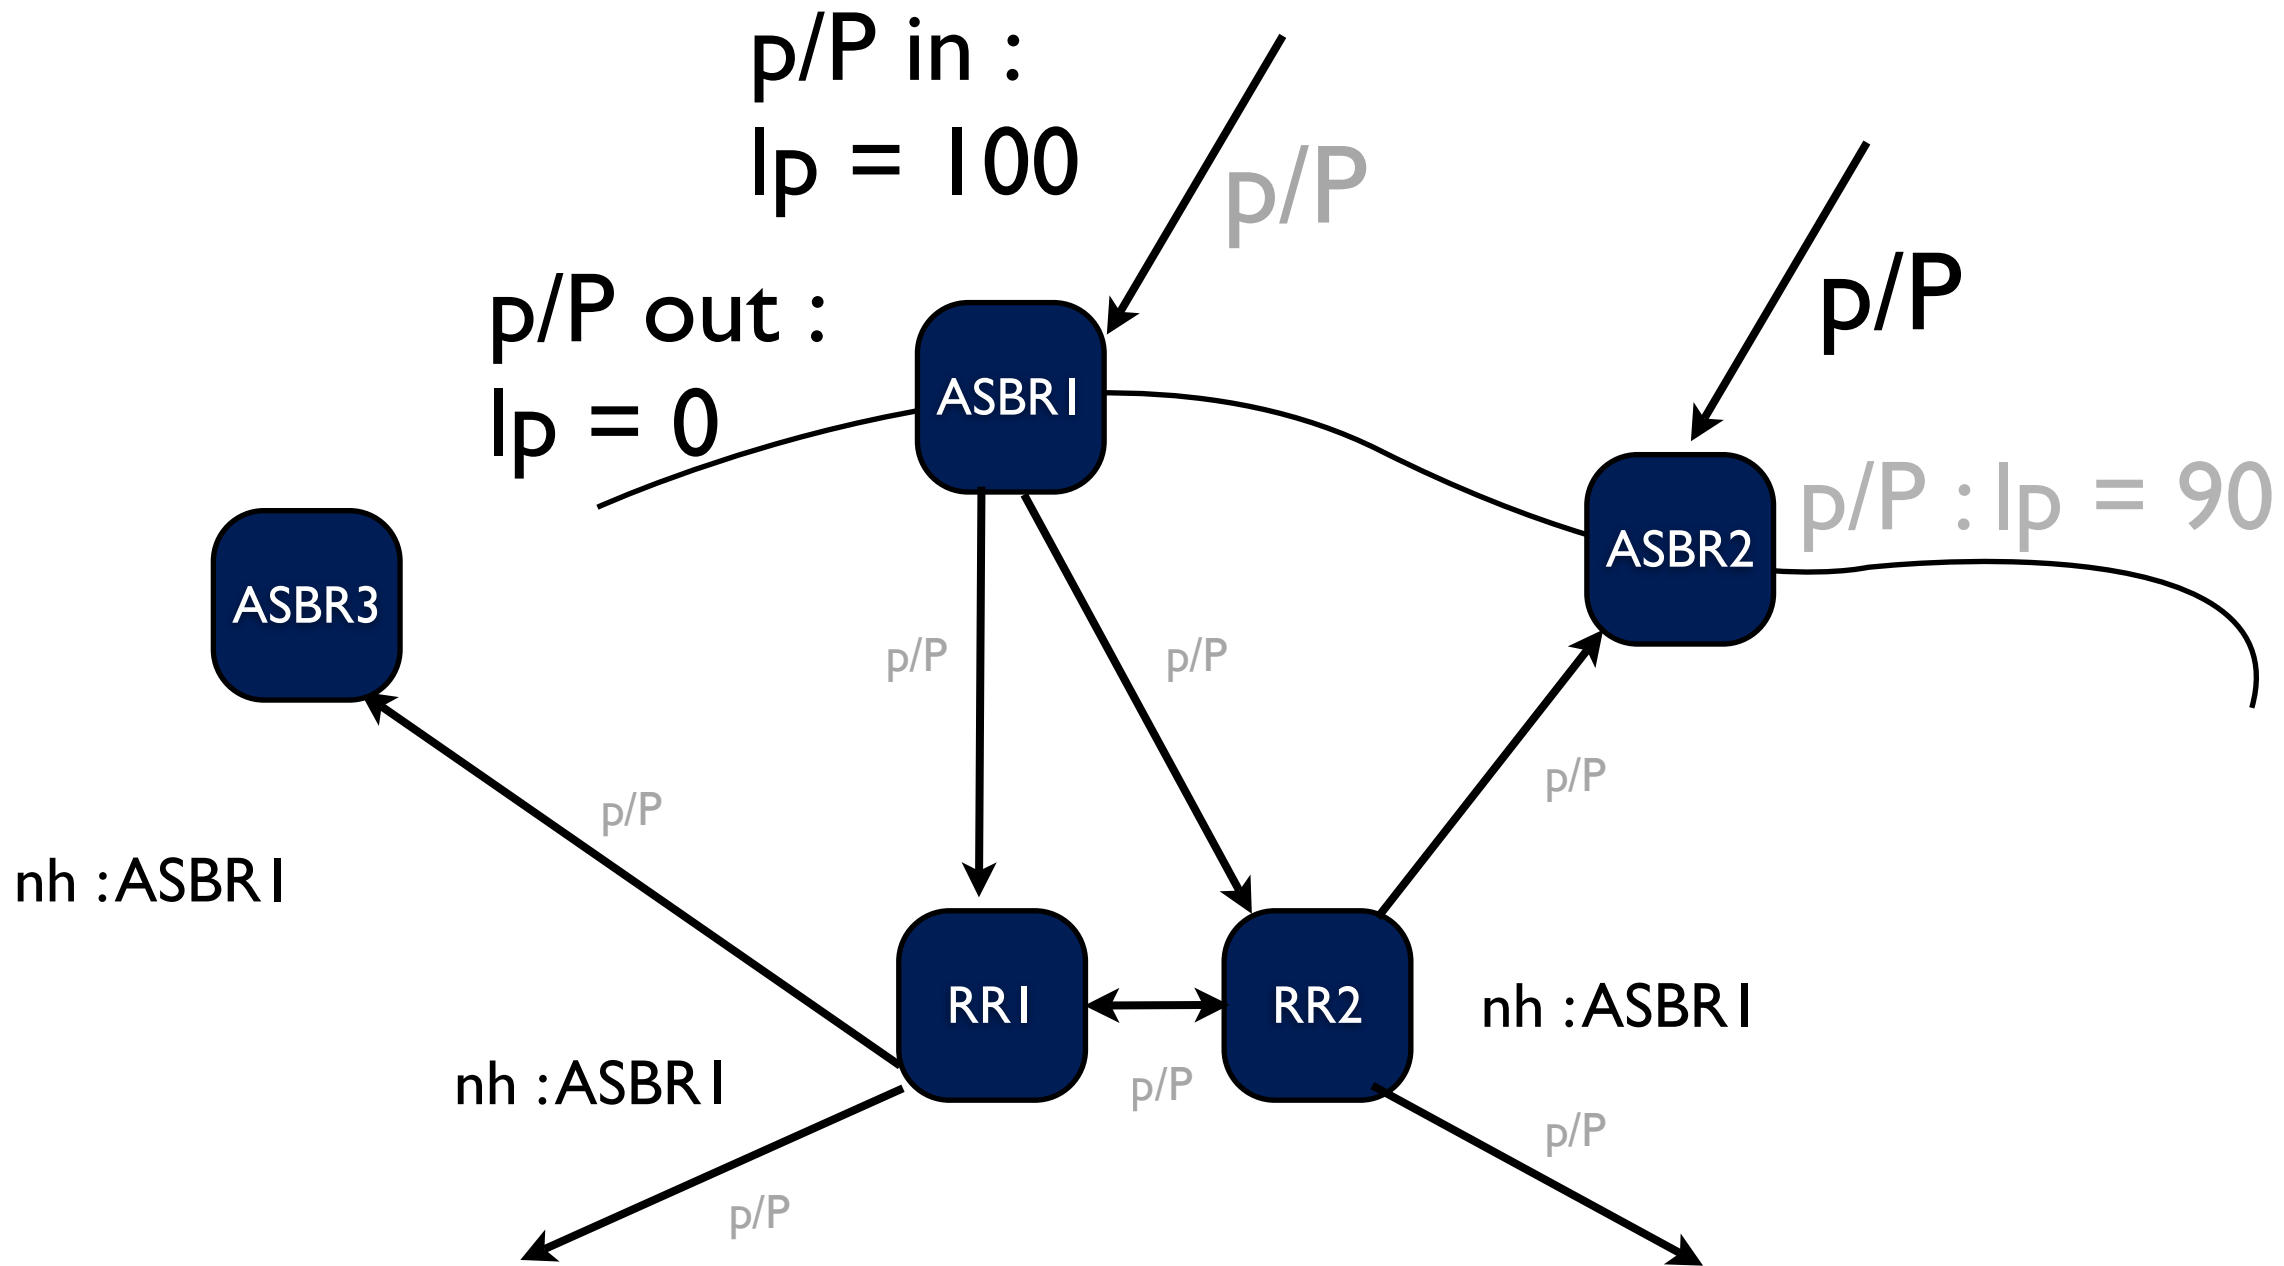
G-Shut : Outbound traffic



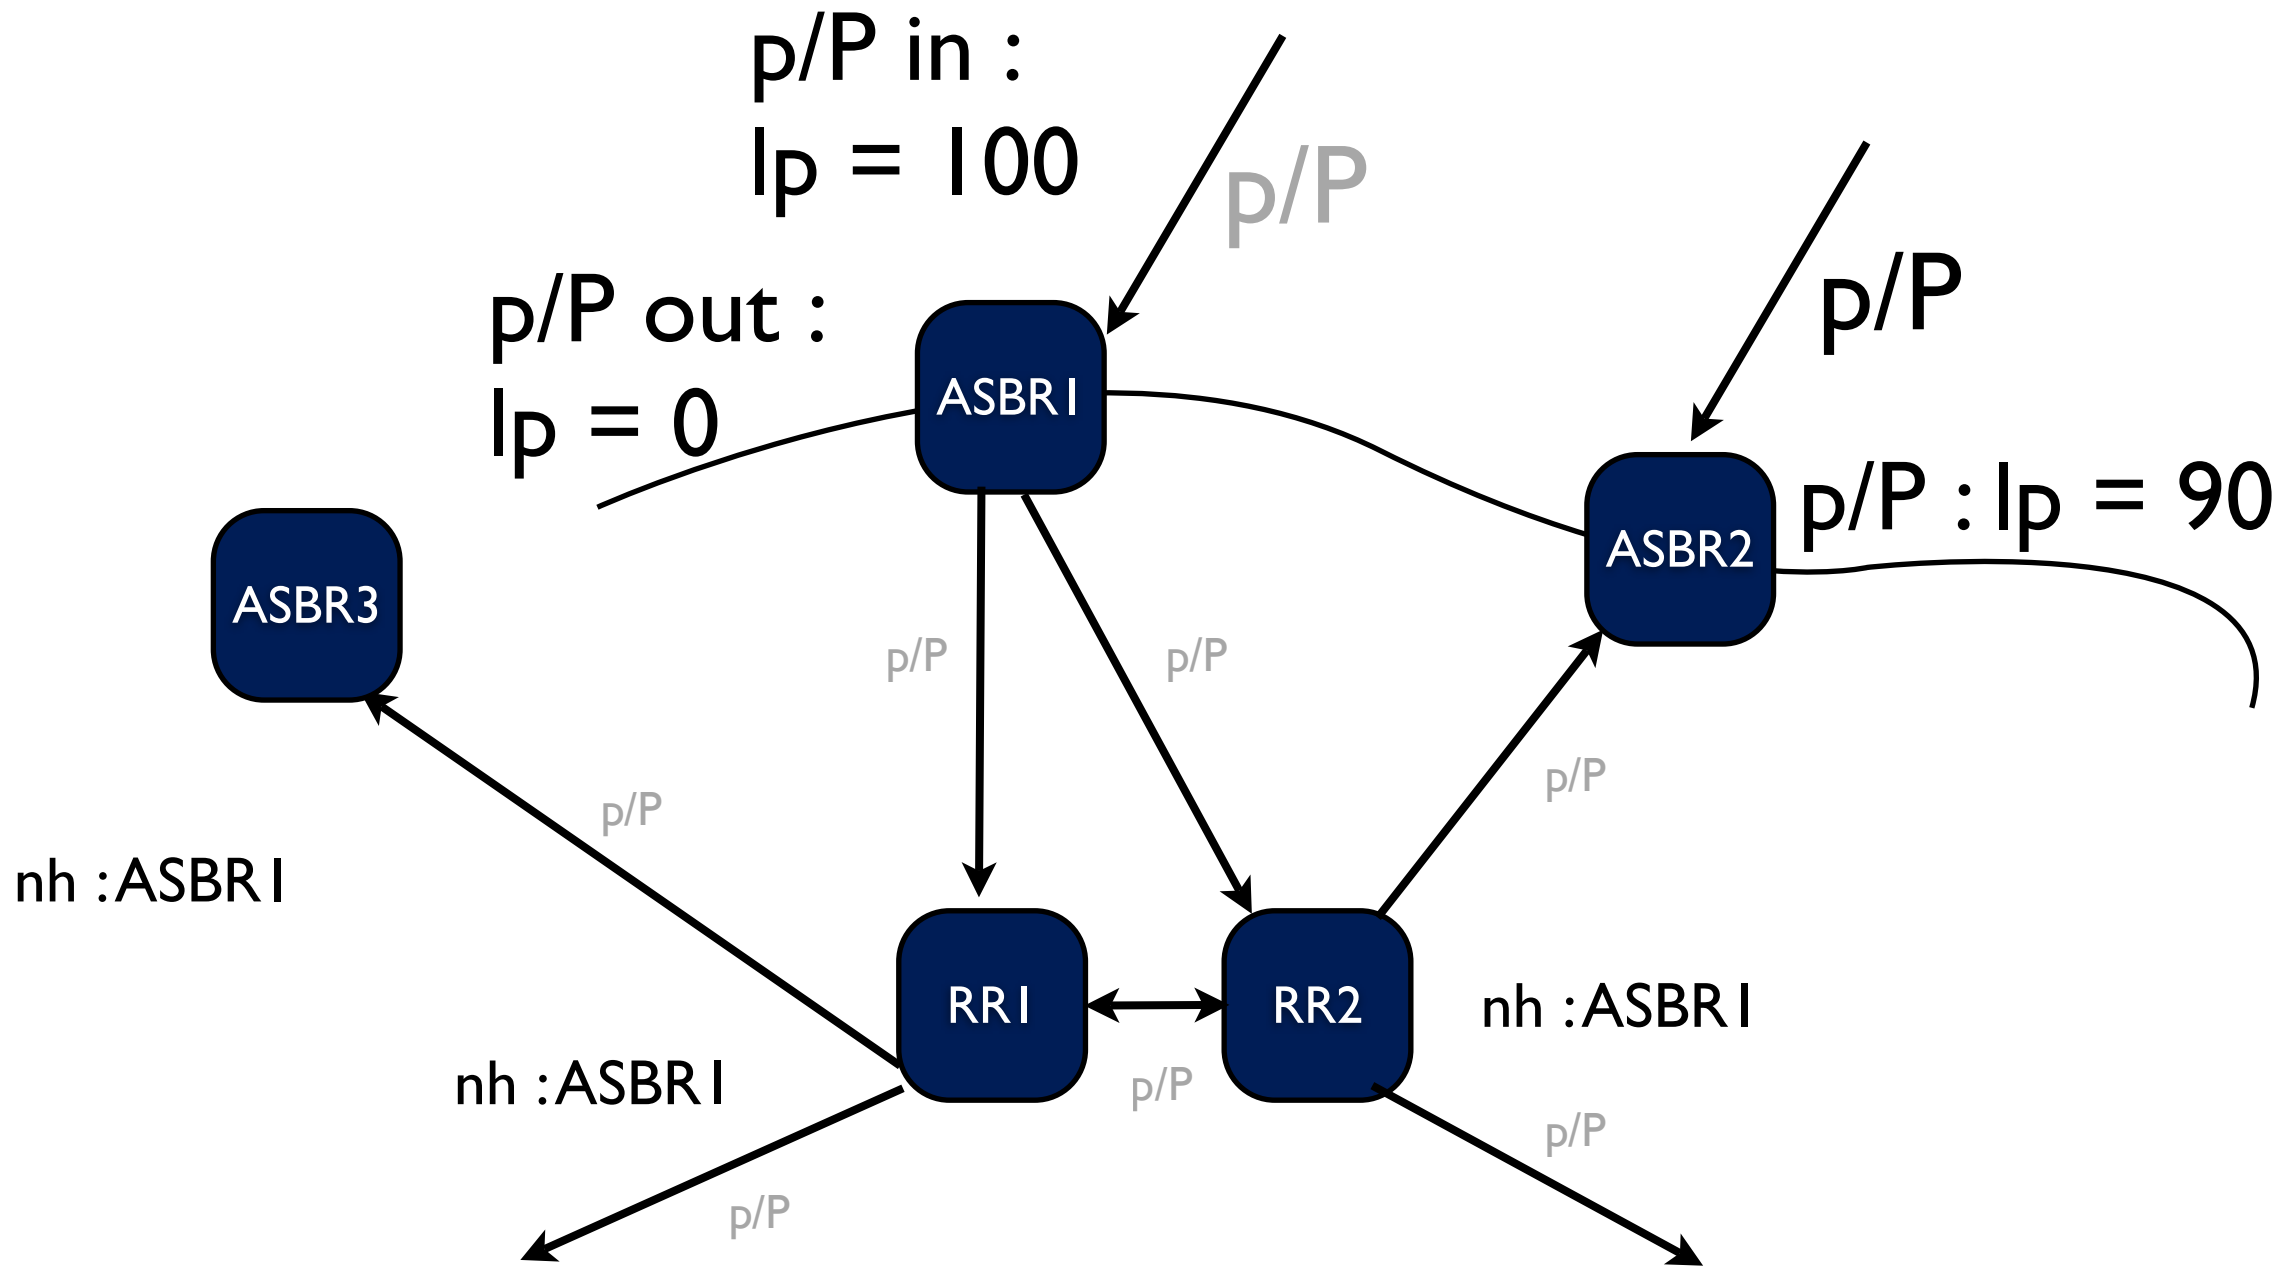
G-Shut : Outbound traffic



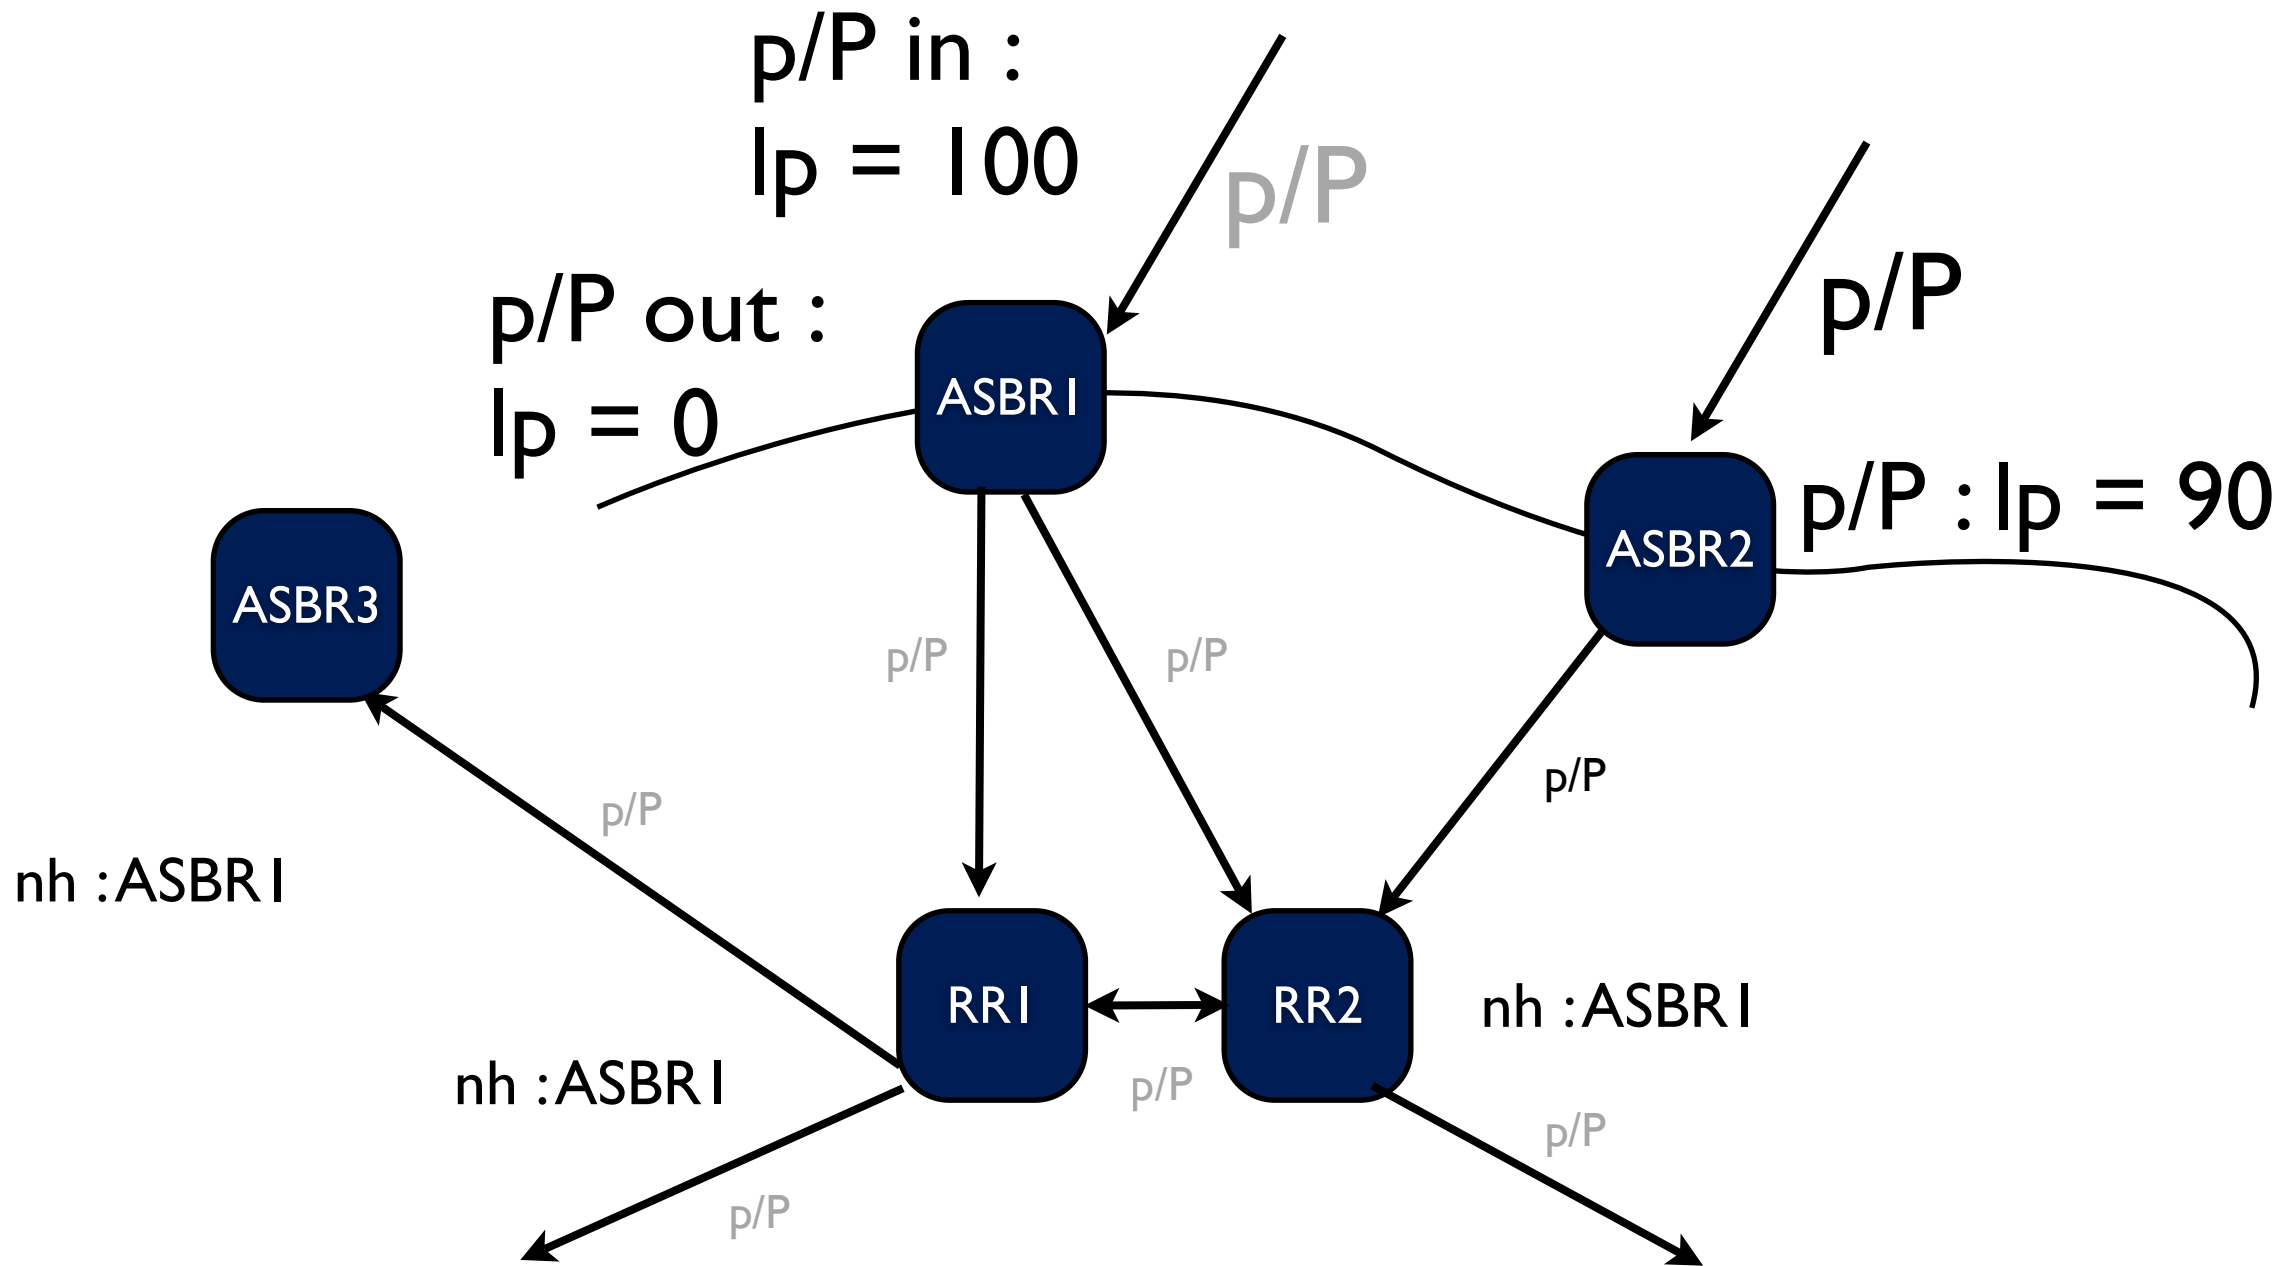
G-Shut : Outbound traffic



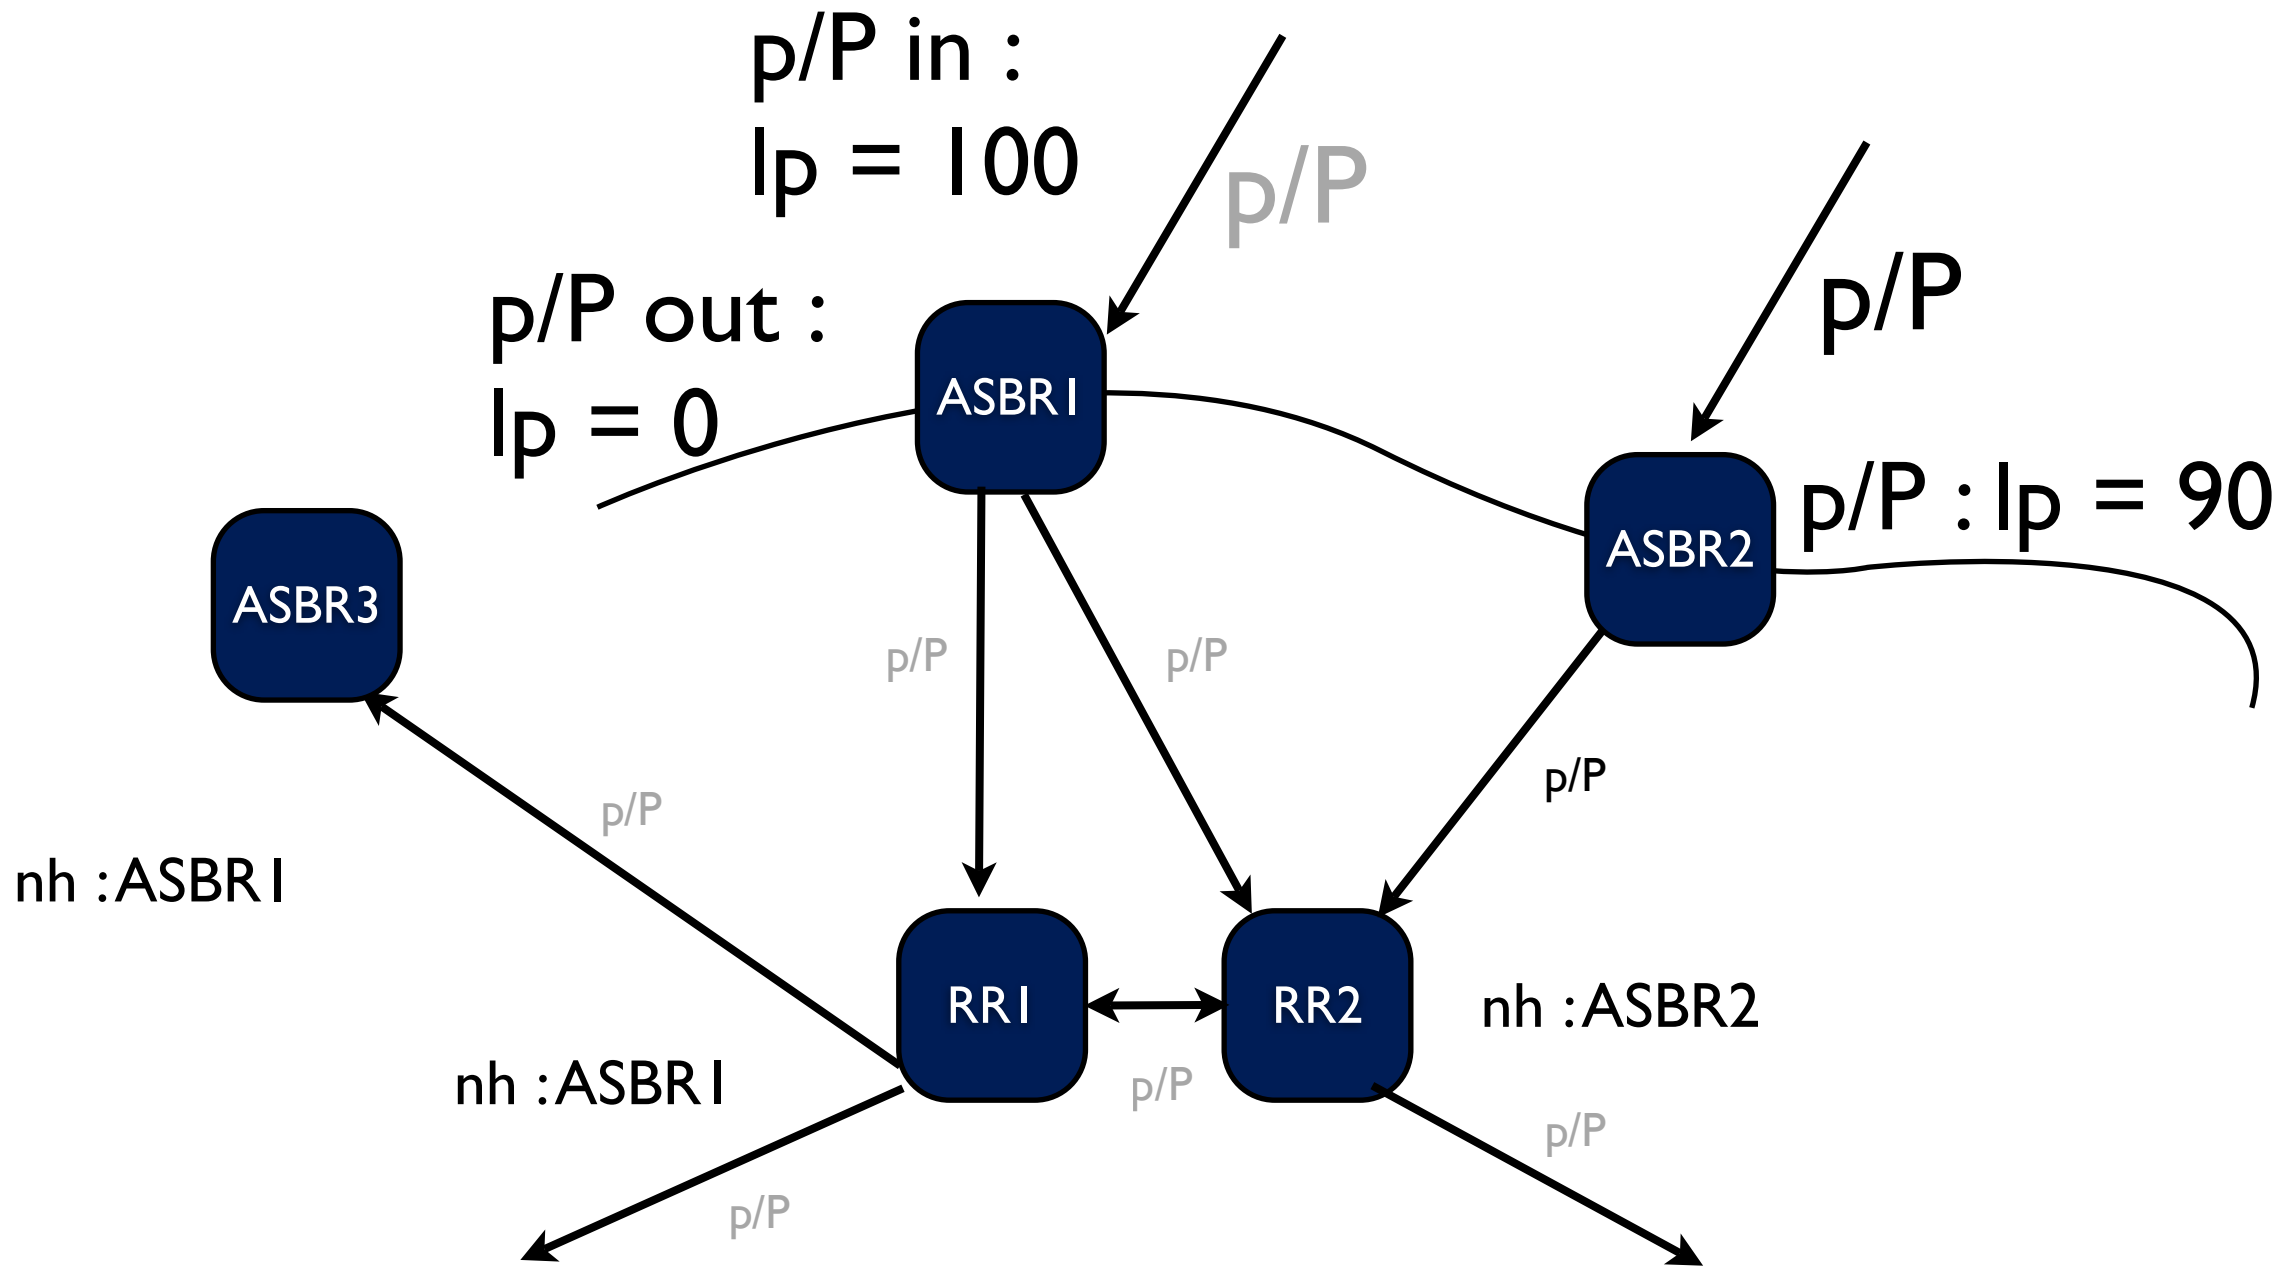
G-Shut : Outbound traffic



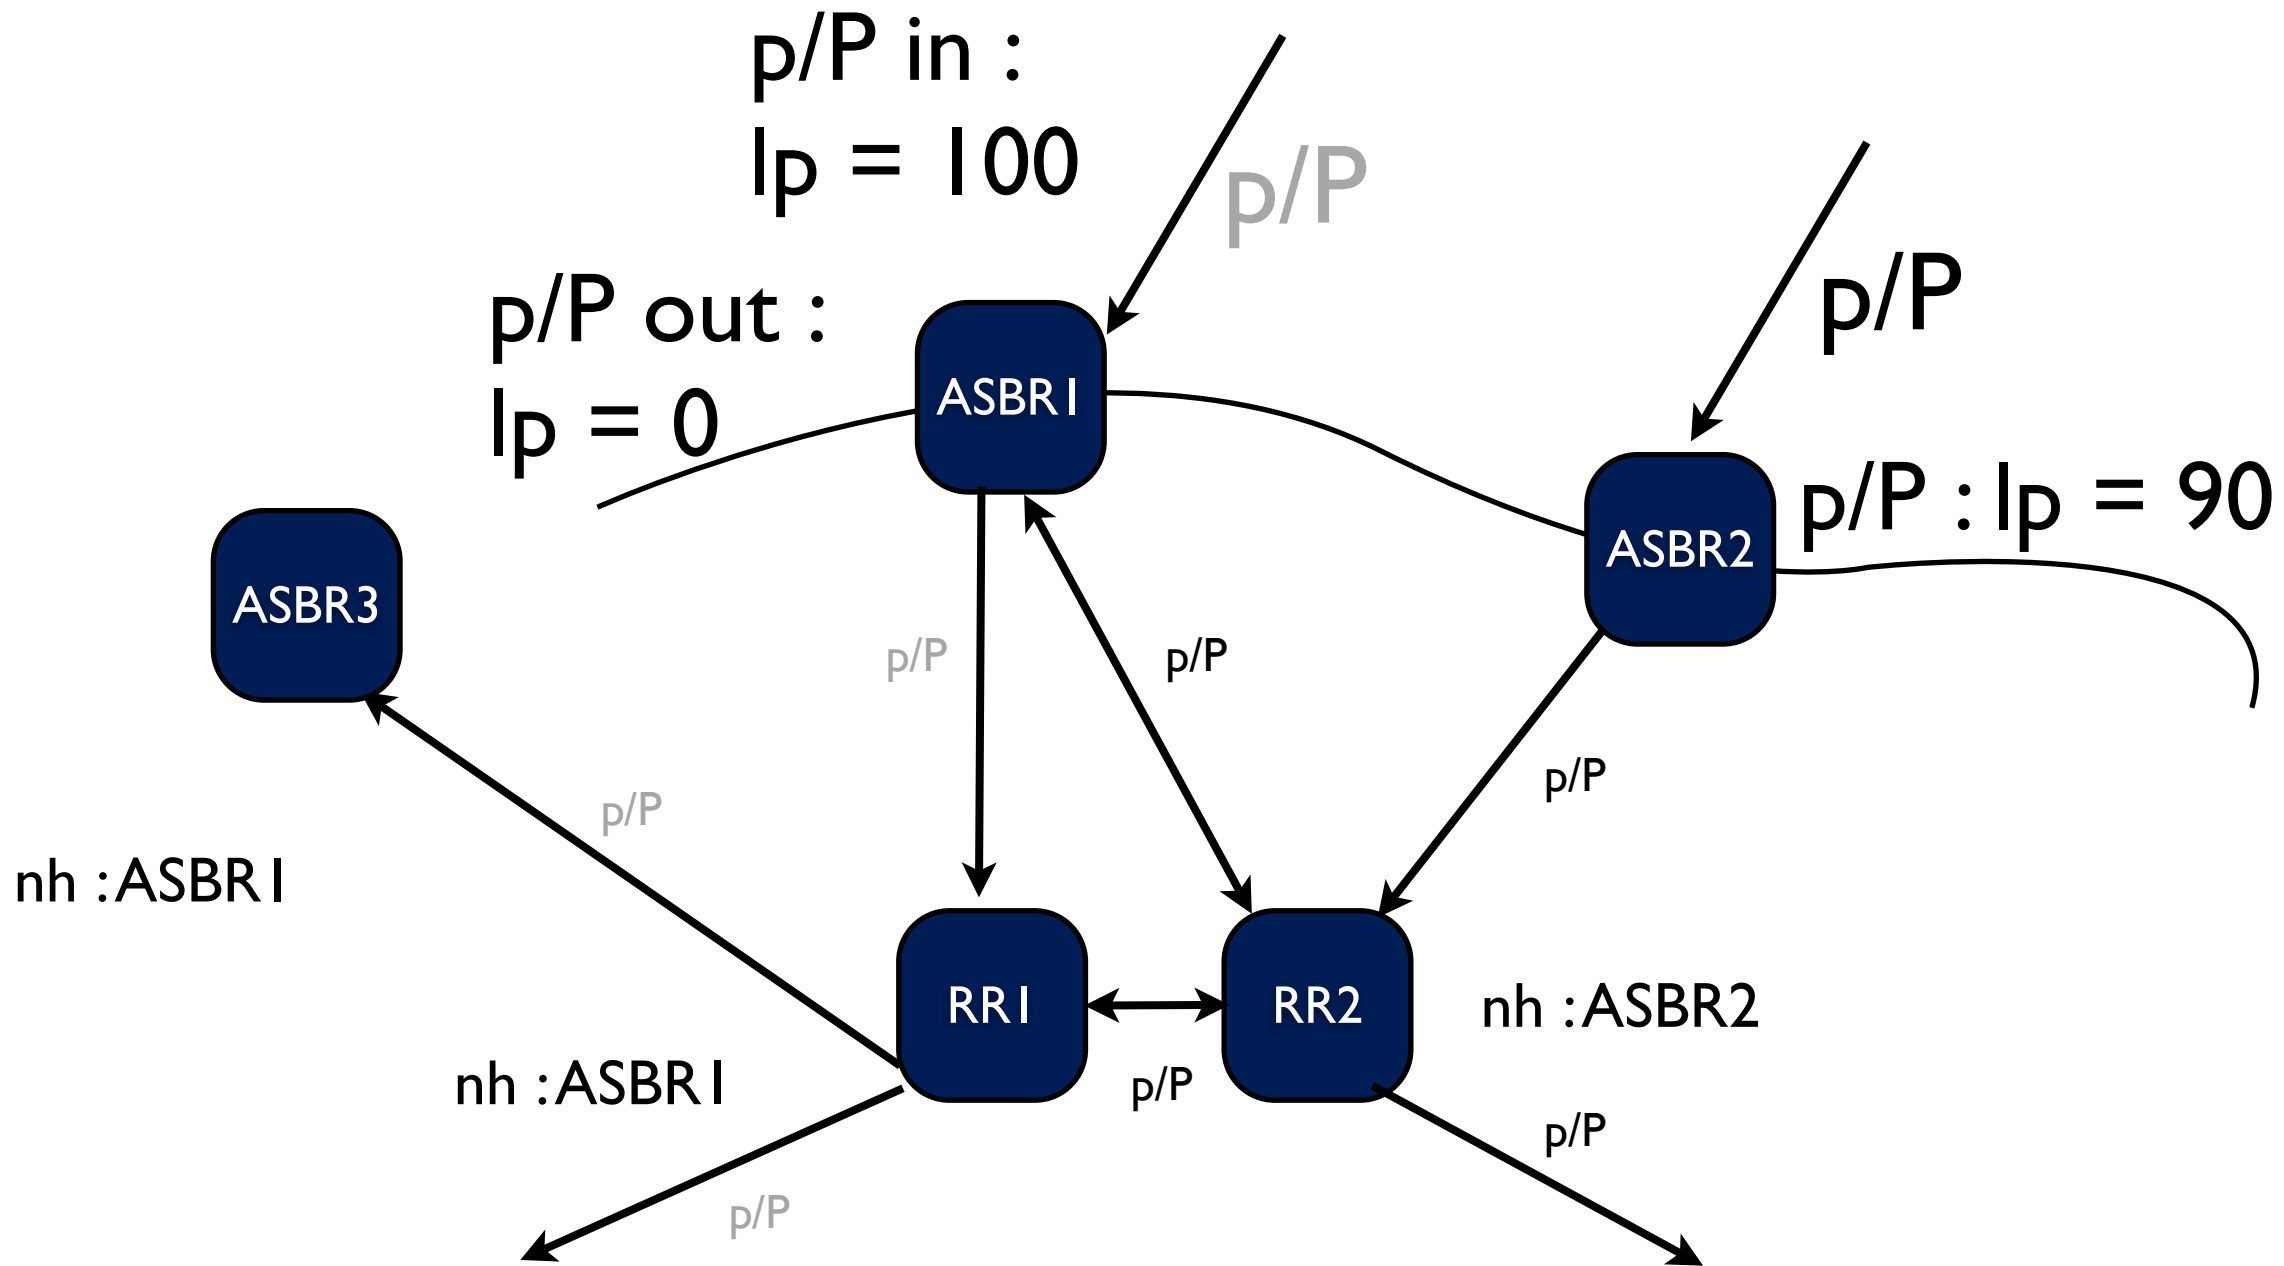
G-Shut : Outbound traffic



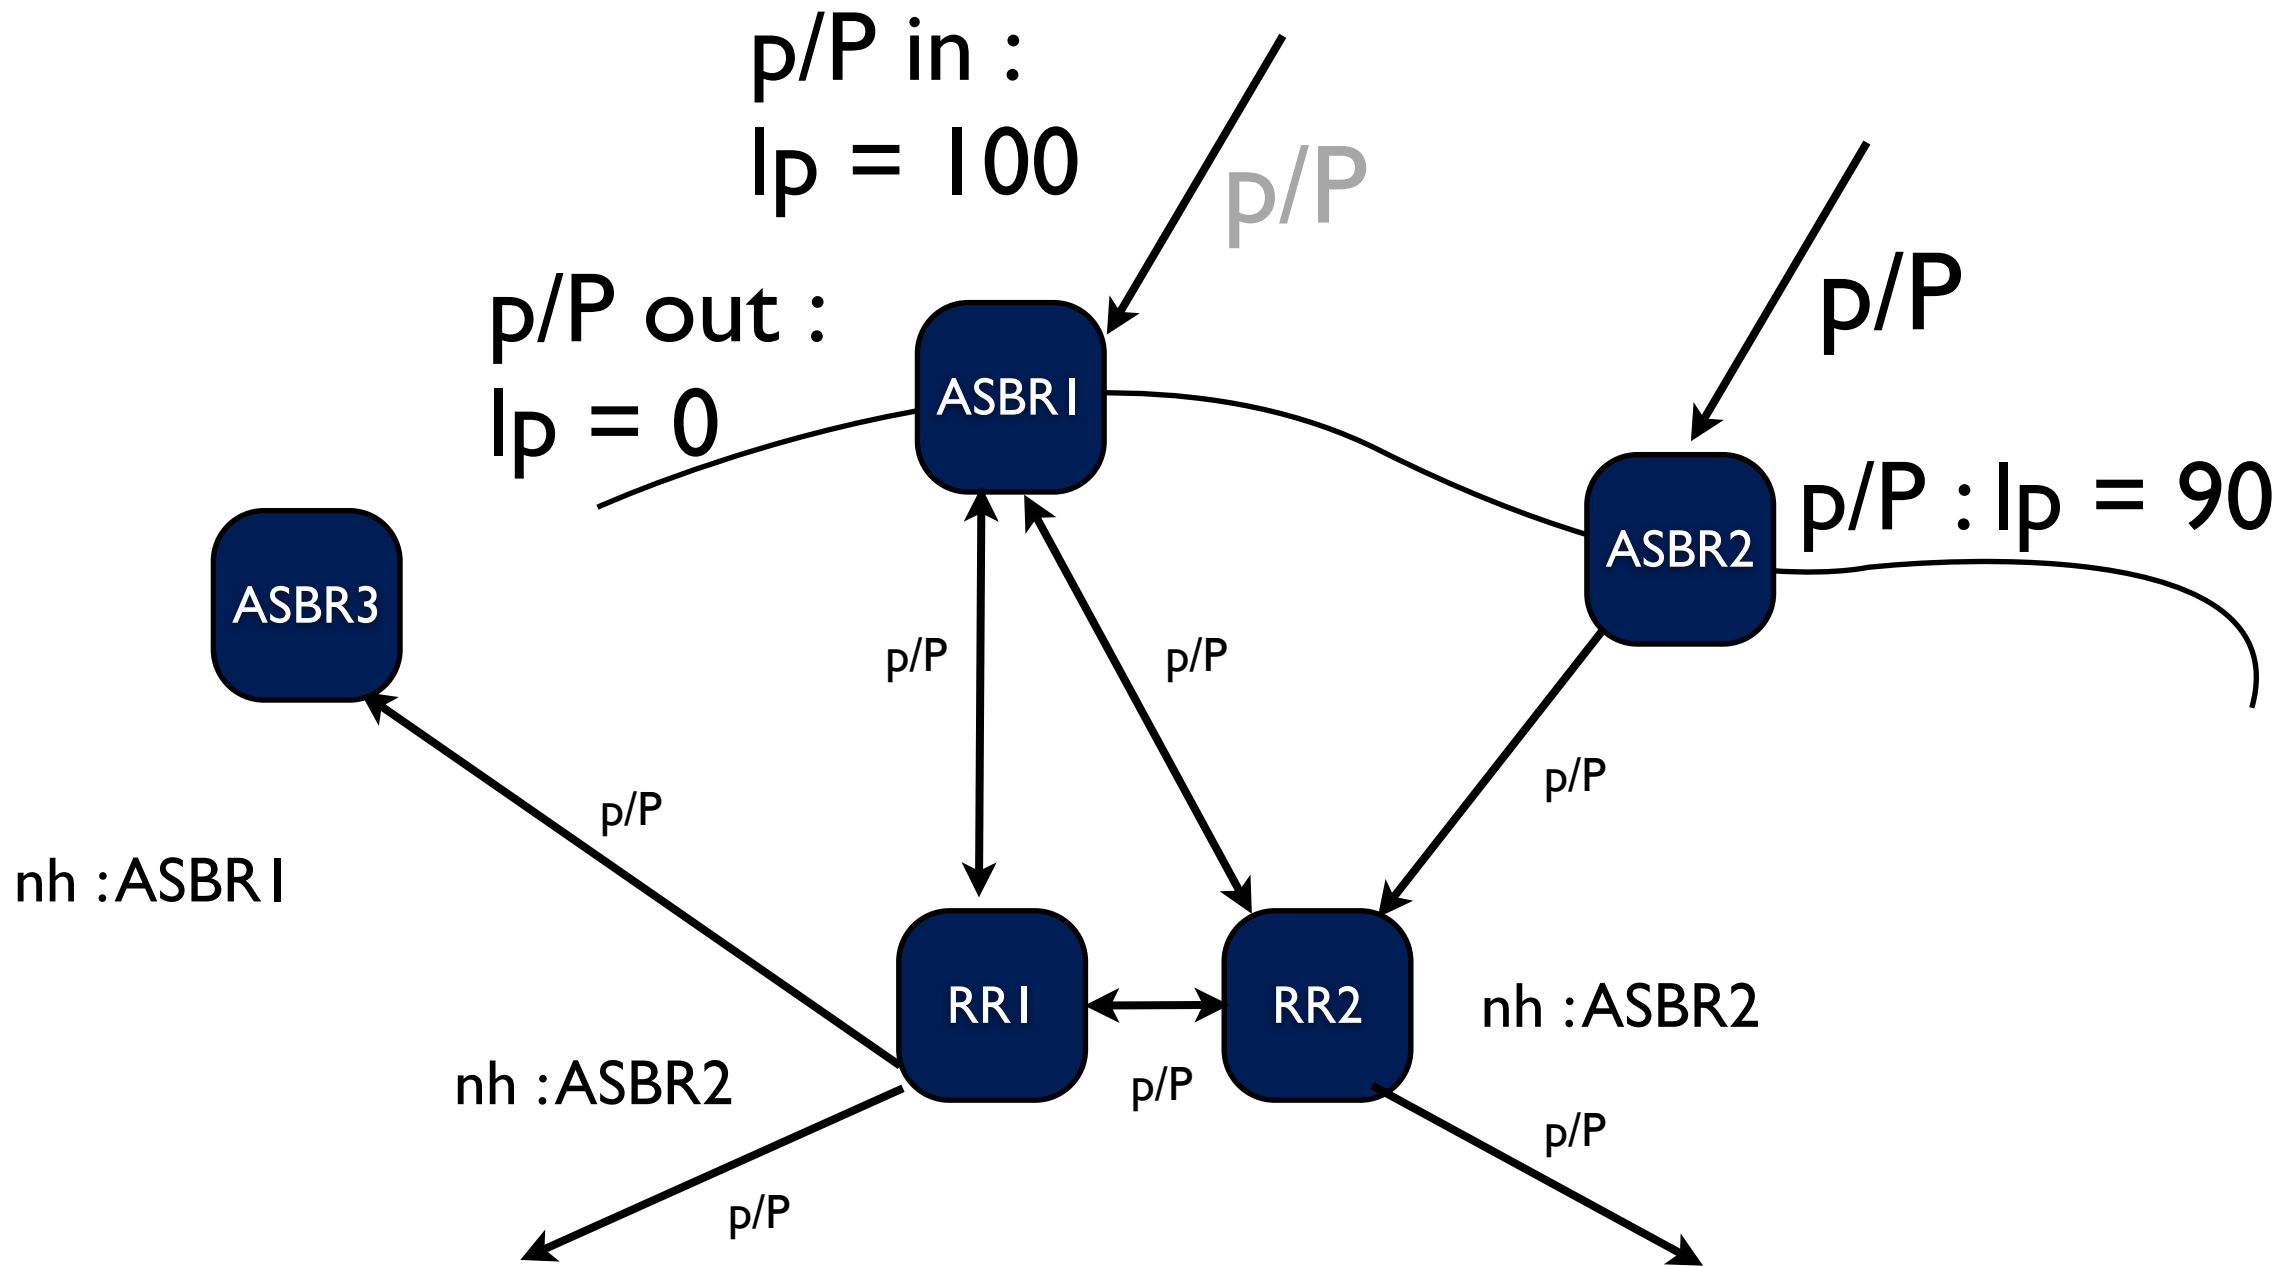
G-Shut : Outbound traffic



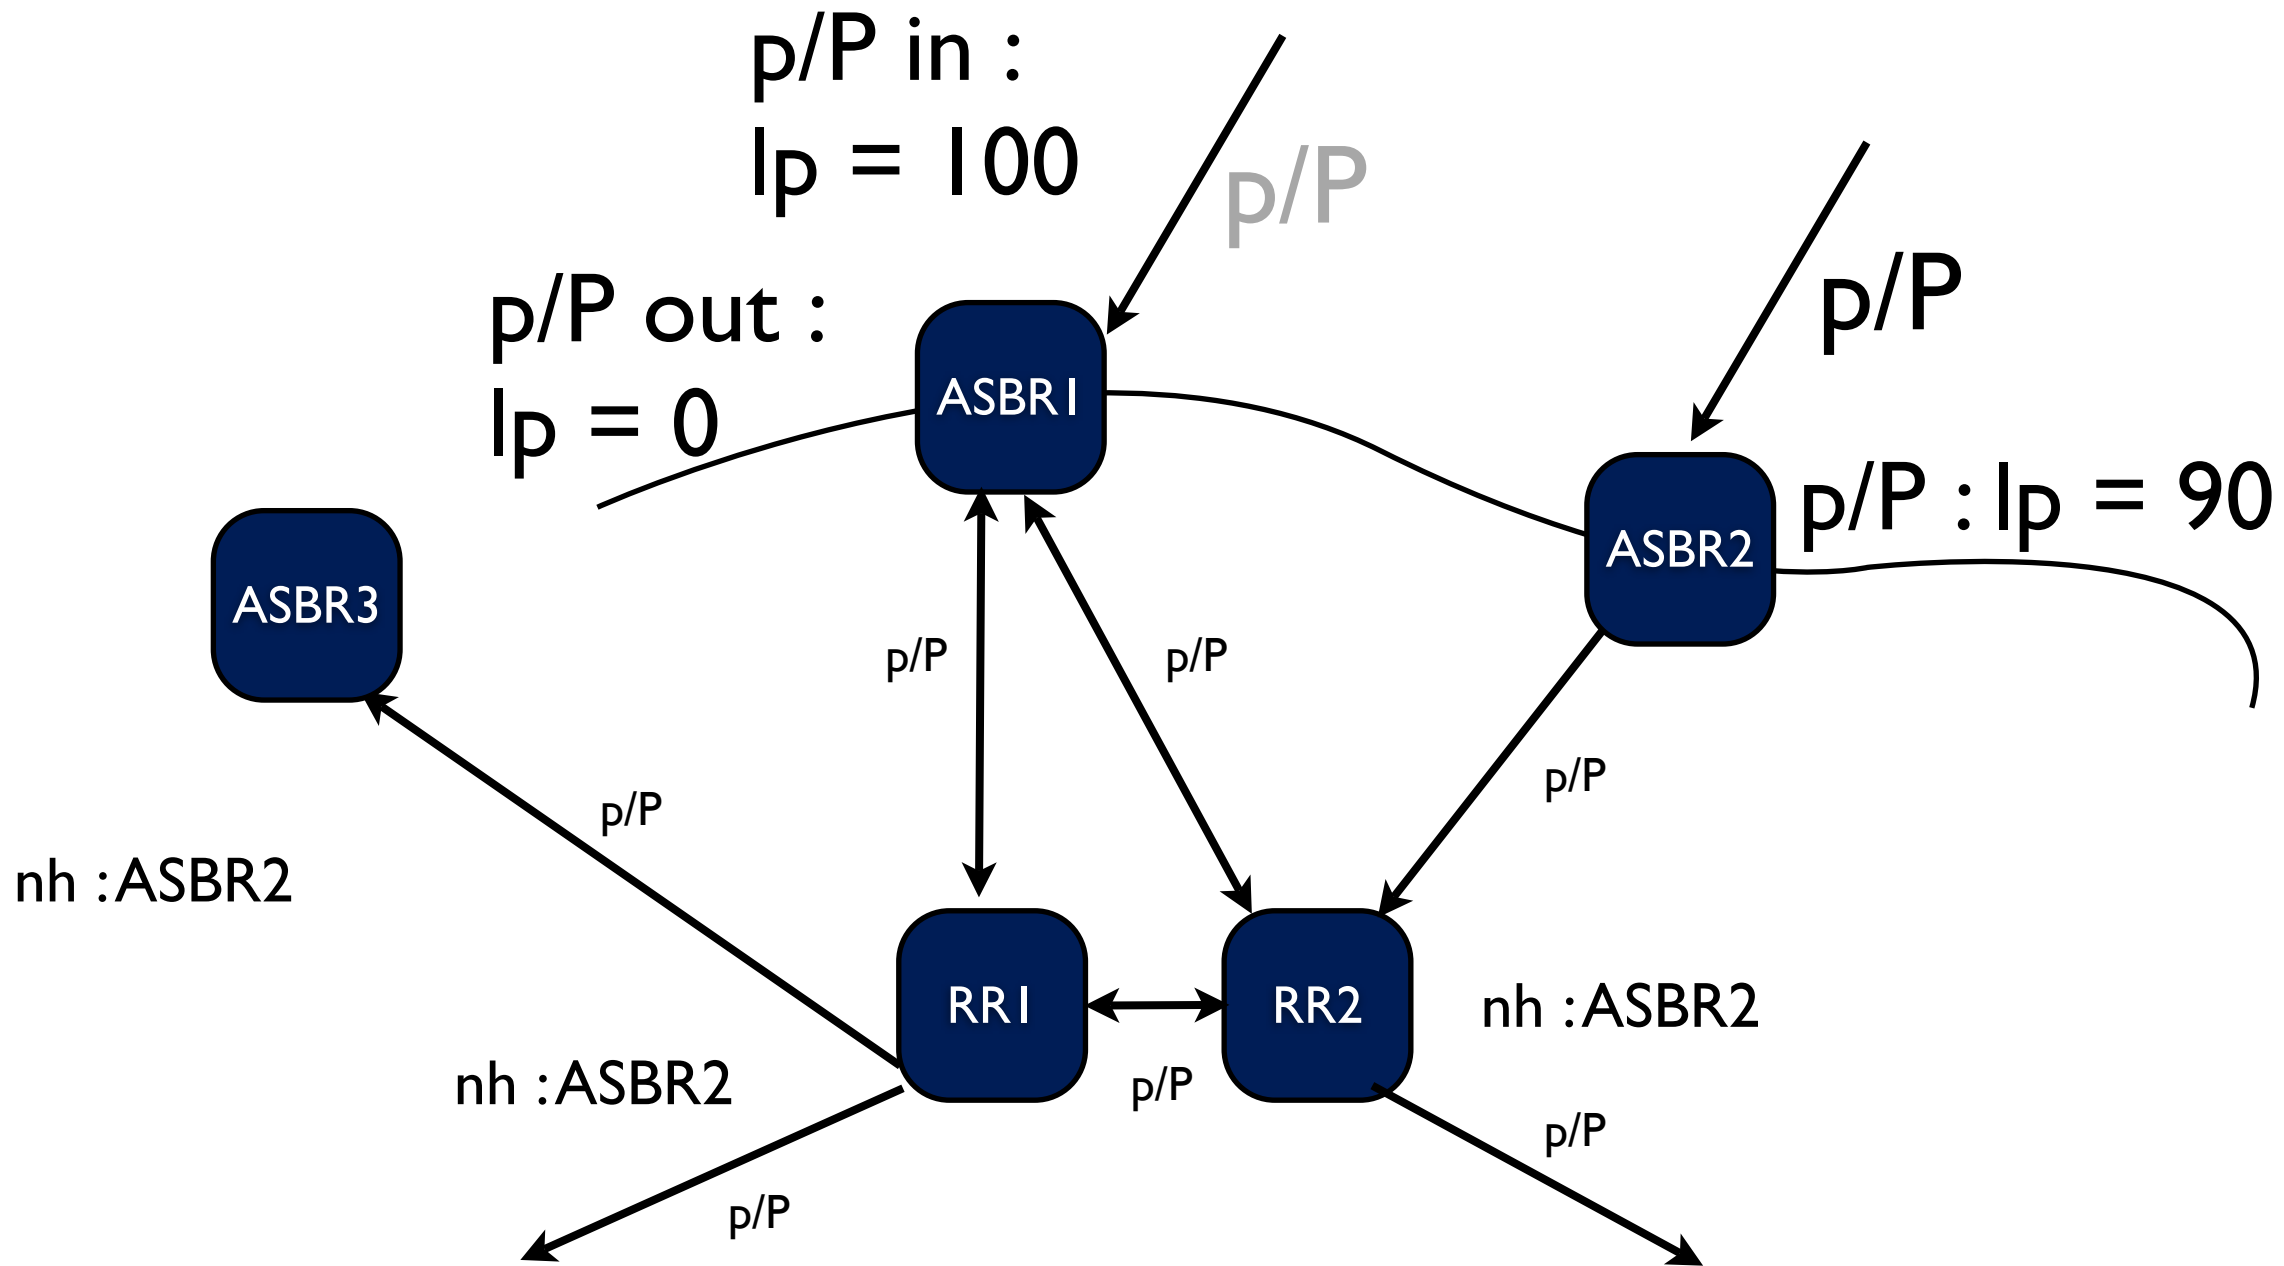
G-Shut : Outbound traffic



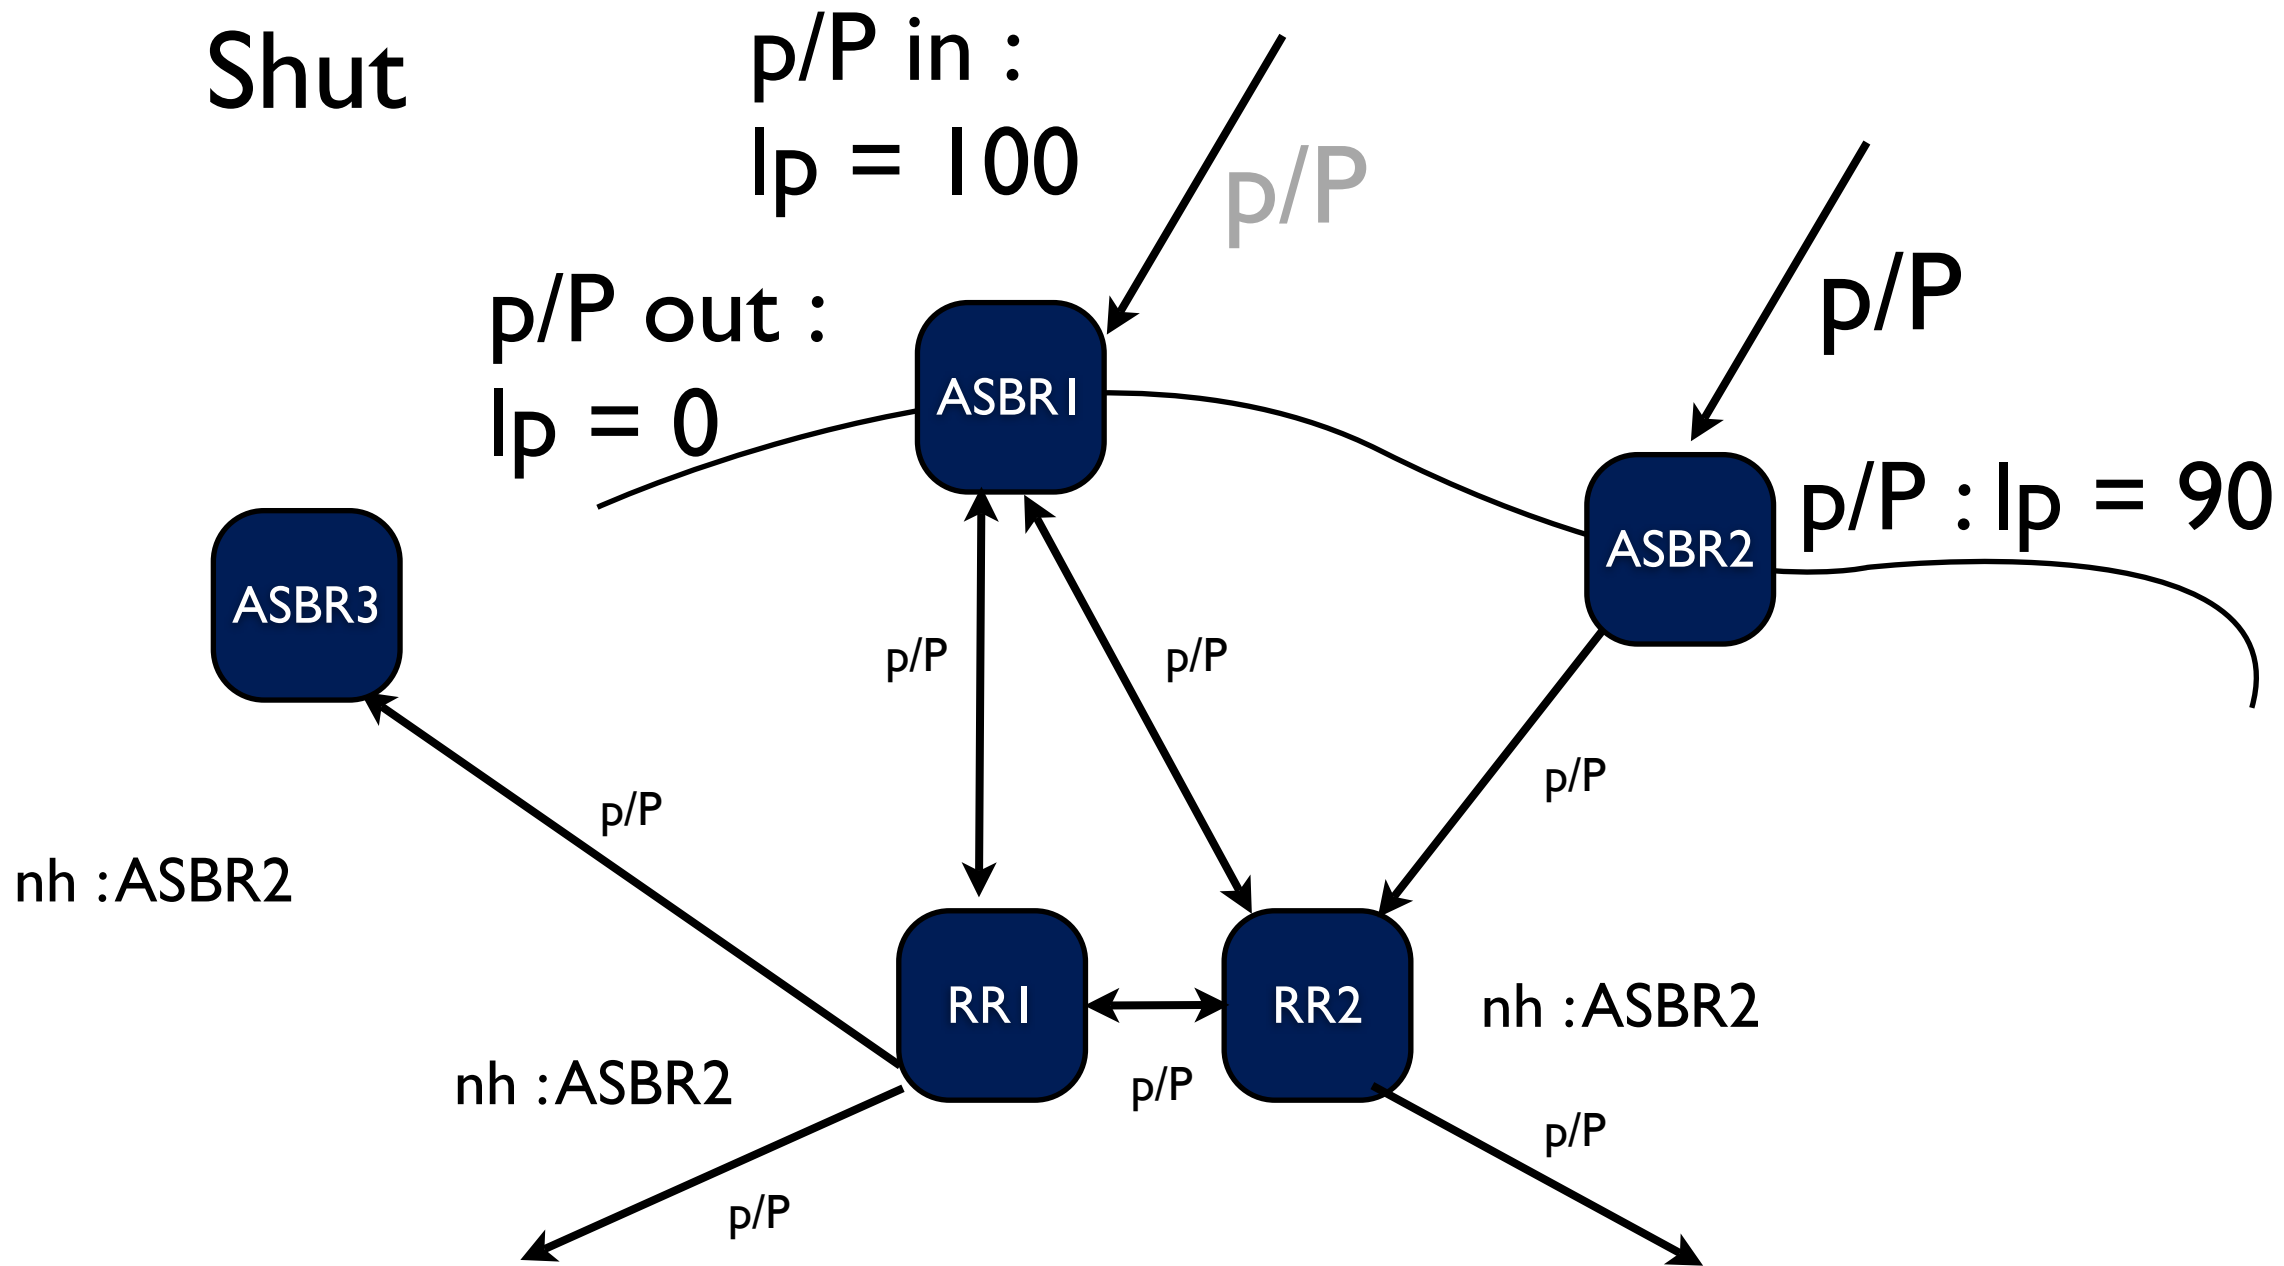
G-Shut : Outbound traffic



G-Shut : Outbound traffic

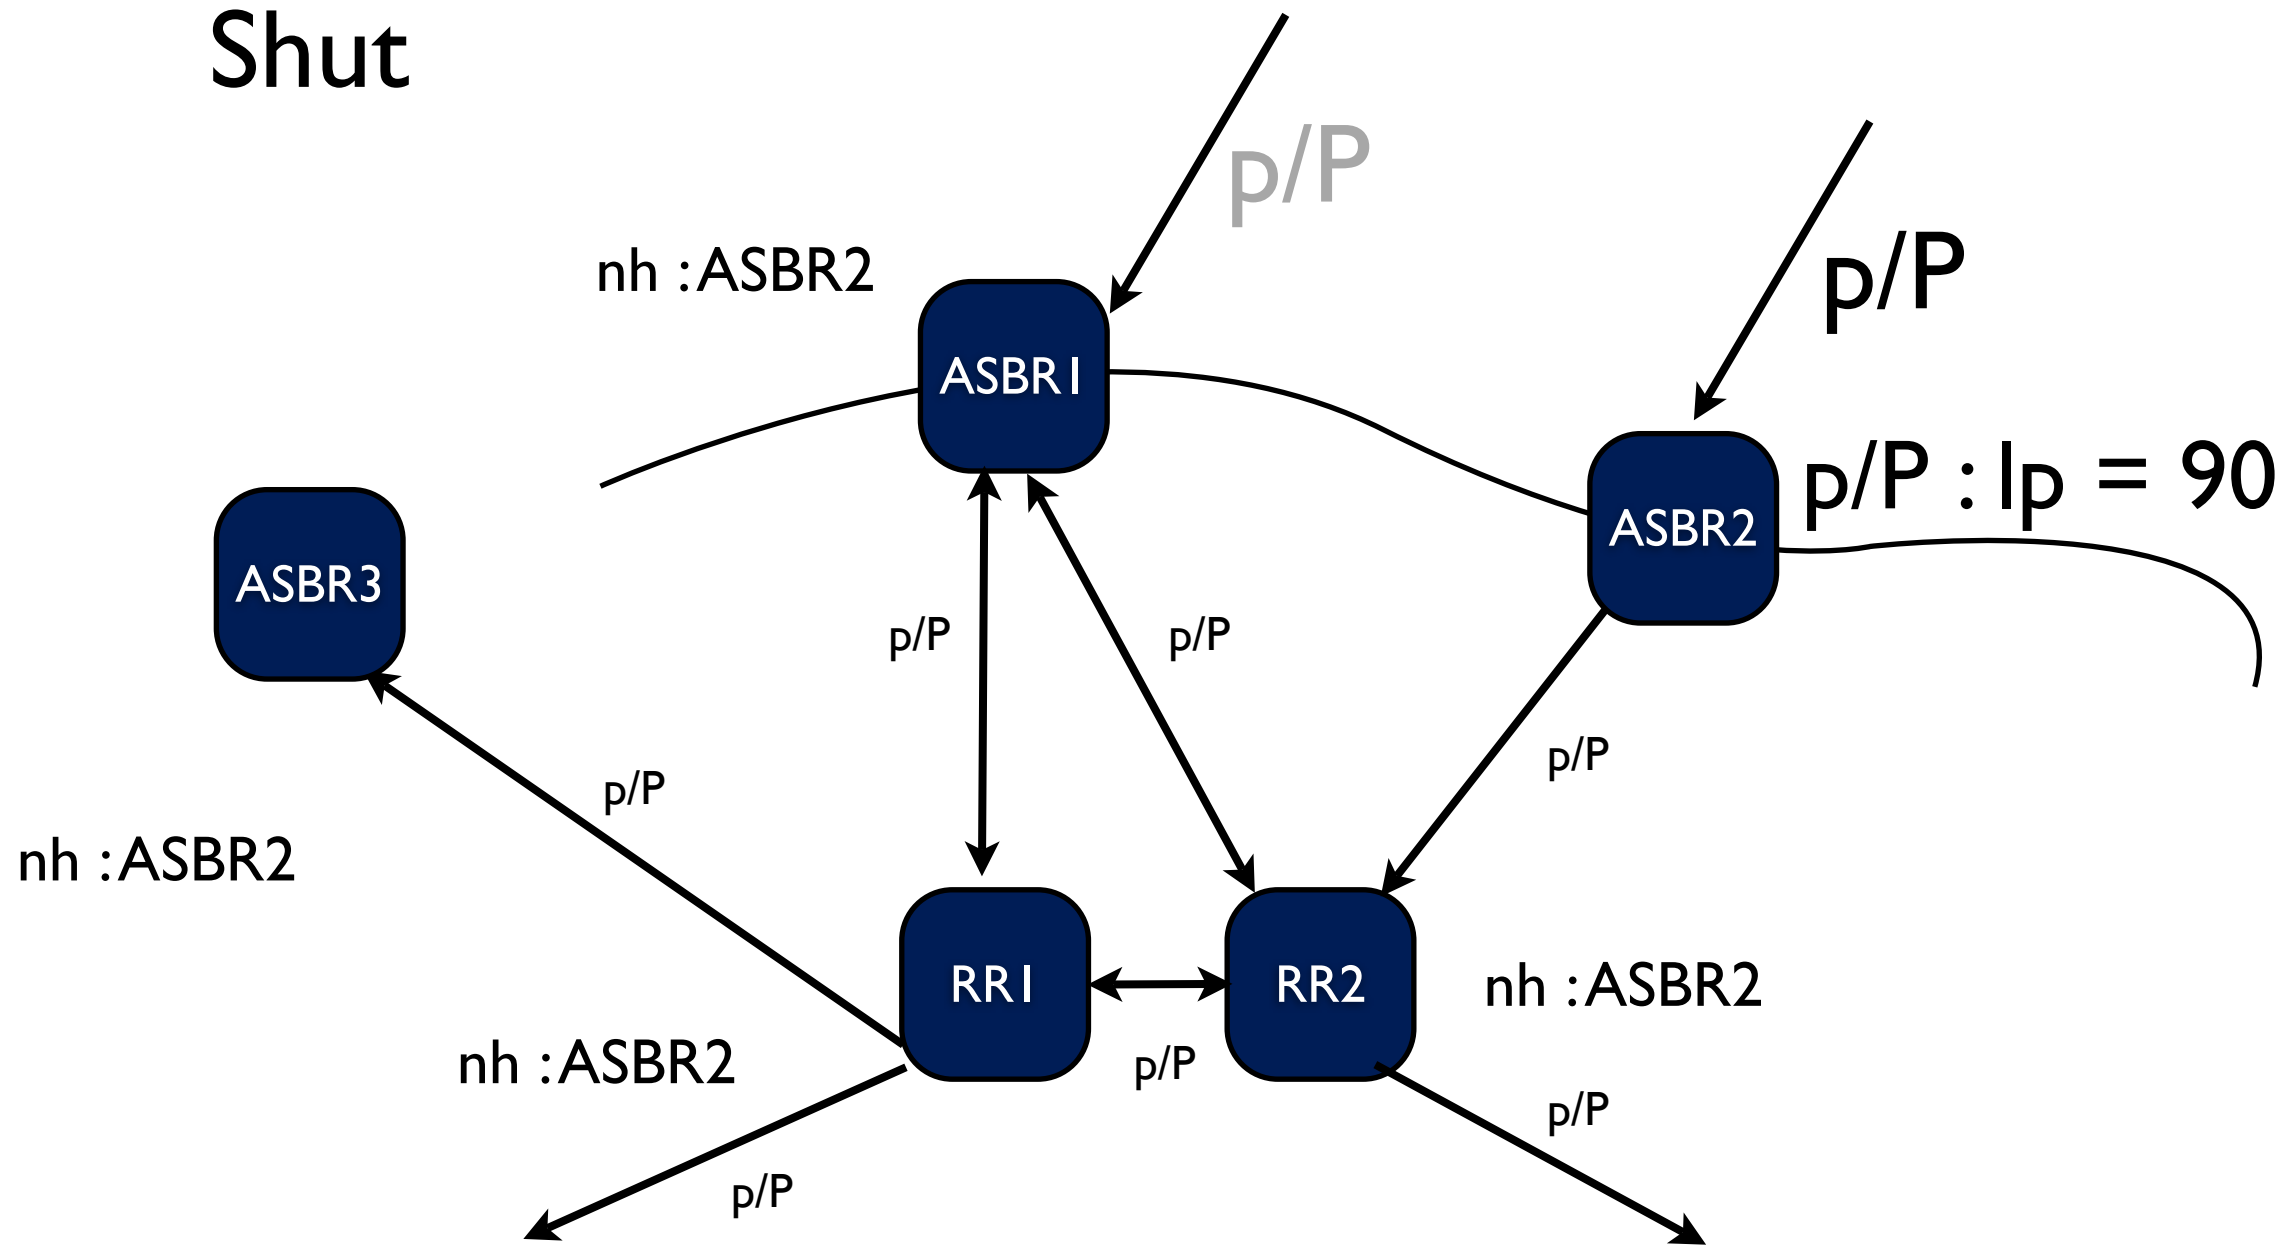


G-Shut : Outbound traffic



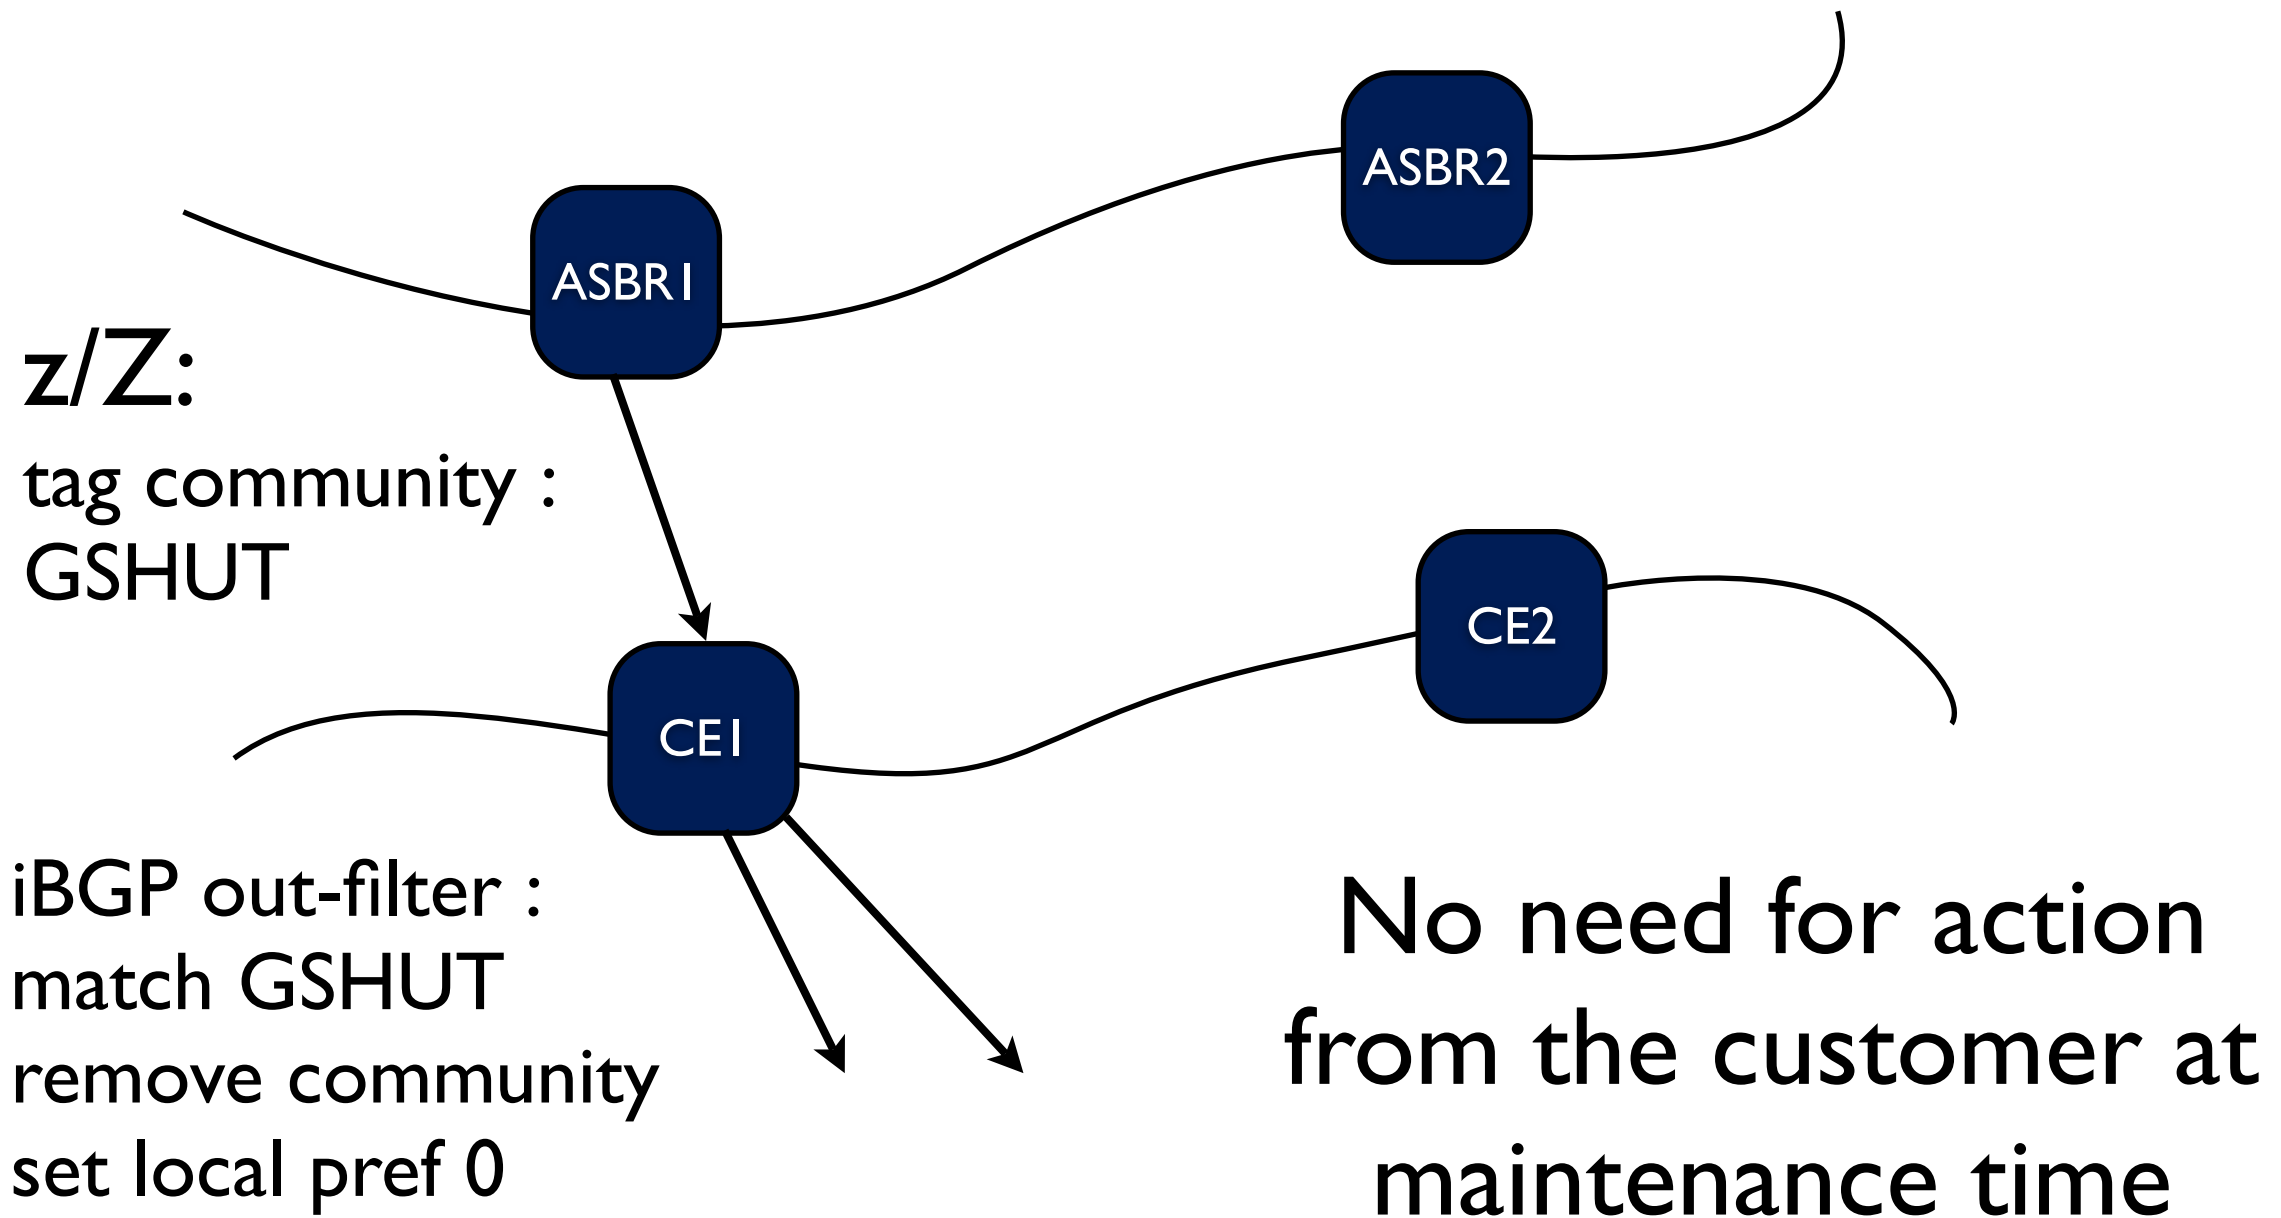
G-Shut : Outbound traffic

Shut



Inbound traffic

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



Deployment simplicity

- Incremental deployment
 - per eBGP peering session
- Incremental gain
- Can be implemented now by ISPs through configuration
 - Vendors could help make it simpler

Why documenting this ?

- Gather input and improvements from the community
- Inform about the procedure
 - If you want to implement g-shut
 - If your peers want to implement g-shut
- Let vendors think about how to ease this behavior
- Discuss the value of that community (next slide)

The g-shut community

- One g-shut community value per customer/peer/provider
 - difficult to manage
 - error prone
- We want to **reserve** one for this purpose

Changes

- Status : Informational
- Can an Informational Document describe what is the goal of a reserved FCFS community (0xFFFF0000) ?

Changes(2)

- Take comments from -00 into account
 - Link up
 - LOC
 - What you can do about it
 - Cases where LP of iBGP paths are overridden by ingress ASBR
 - Open the option to use the g-shut community over iBGP

What do we get ?

- 0 packet loss for link down and up cases
- Some corner case timings require best-external (intra-cluster-best) support
- still provides very significant loss reduction without support of best-external

Next steps

- Actually reserve the community
- WG Doc ?

?