

BGP Graceful Shutdown

draft-ietf-grow-bgp-gshut-01

Pierre Francois

UCL-BE

Bruno Decraene

France Telecom - Orange

Cristel Pelsser

IJ

Clarence Filsfils

Cisco Systems

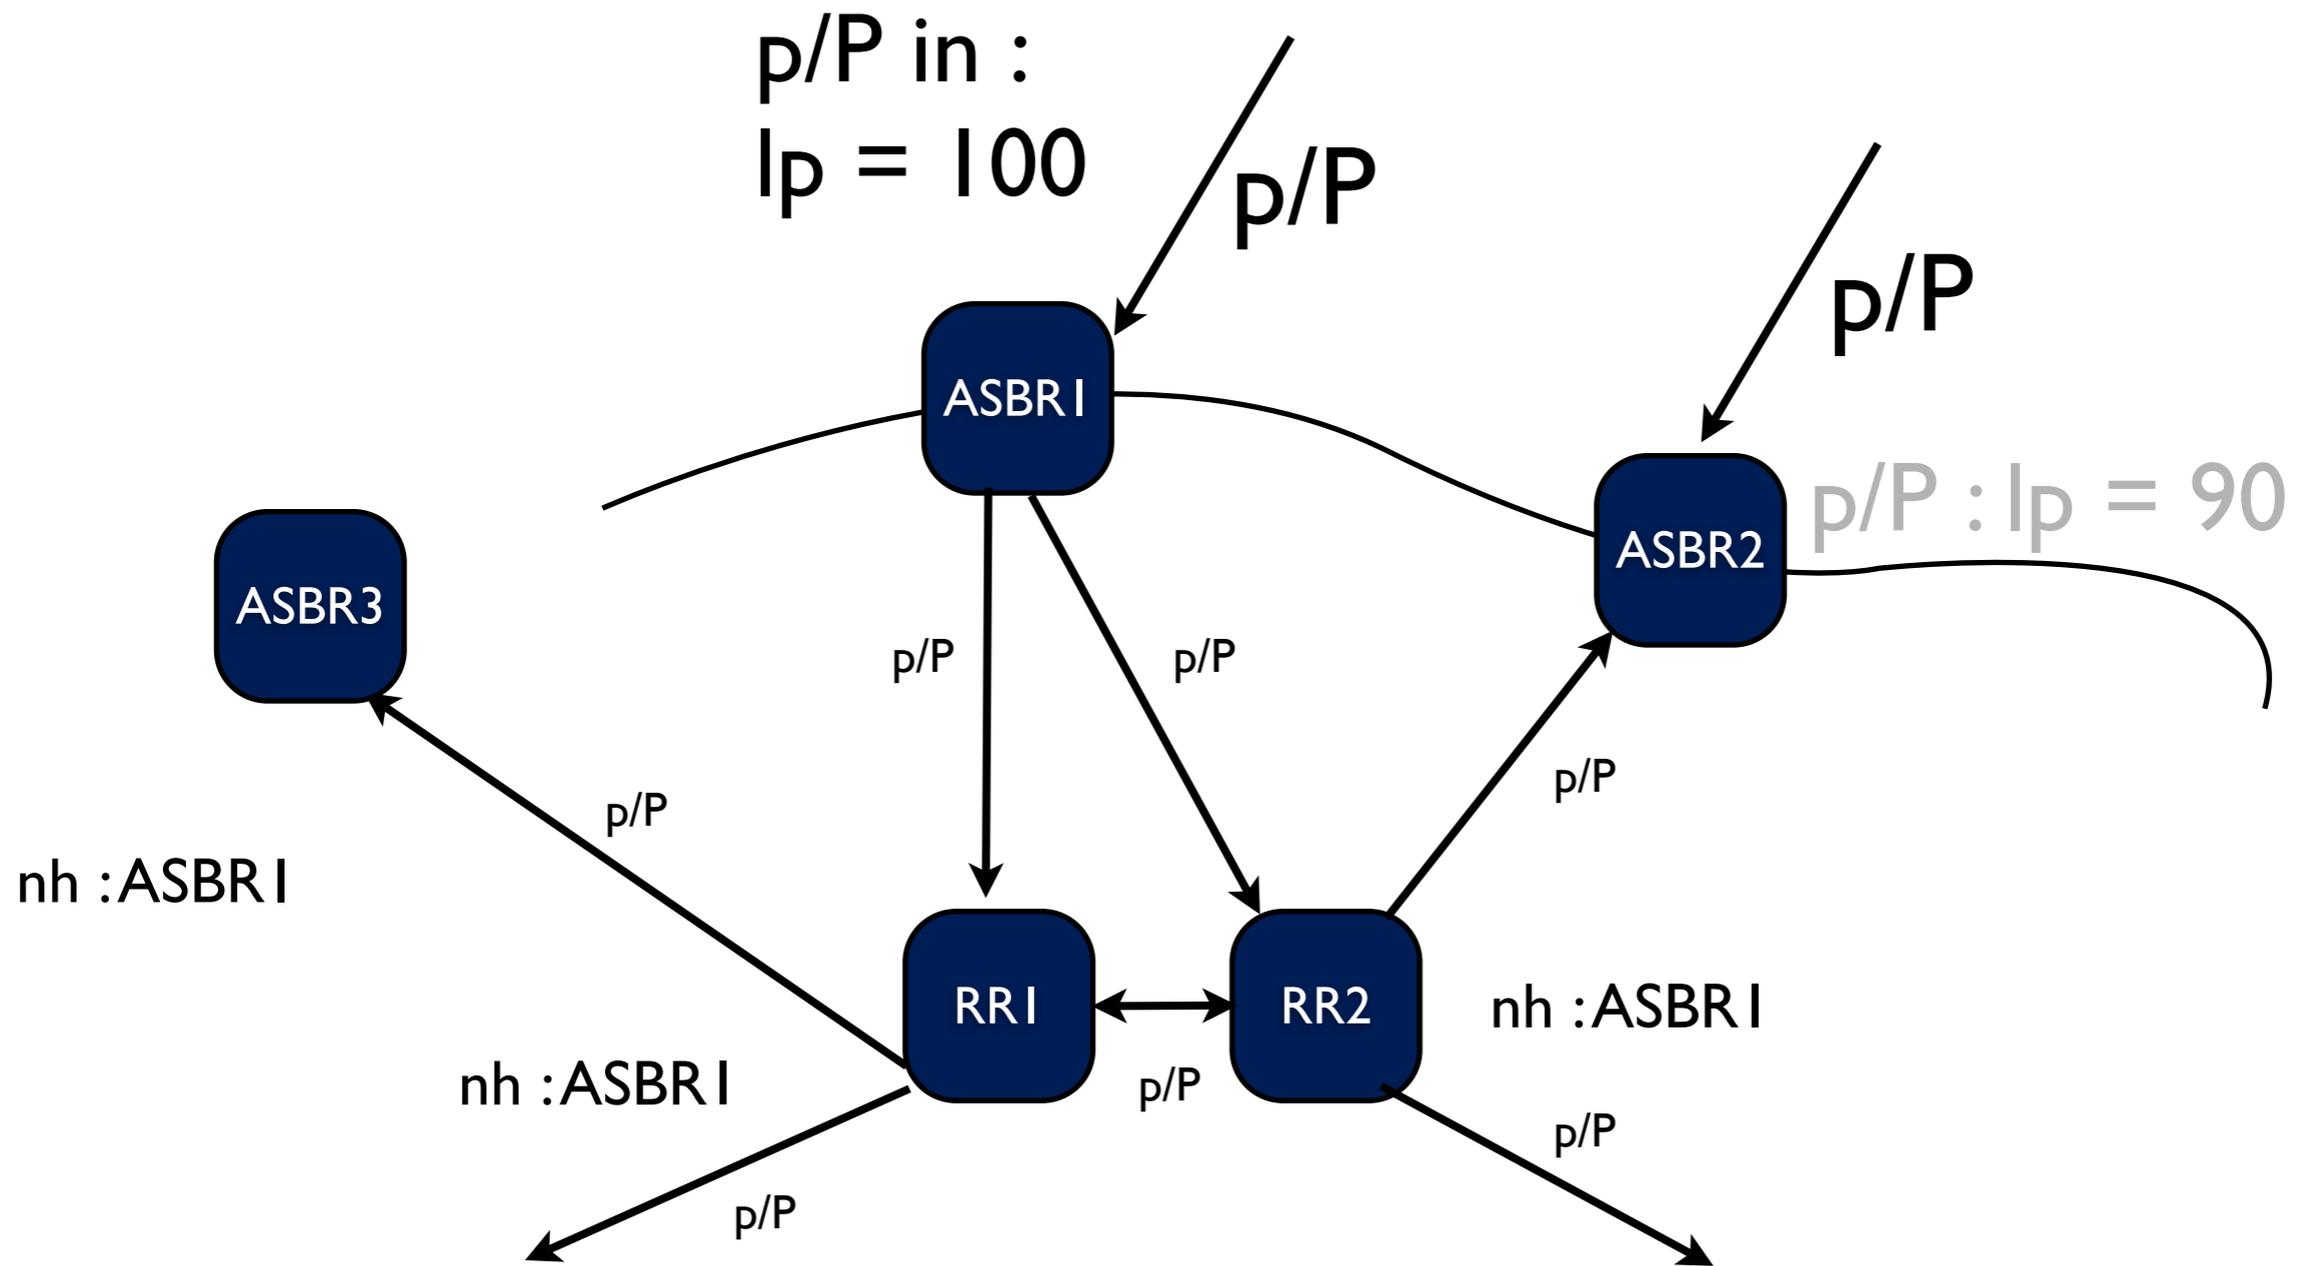
Agenda

- Quick reminder
- The G-shut community value
- G-shut behavior concealment
- Local-pref overriding

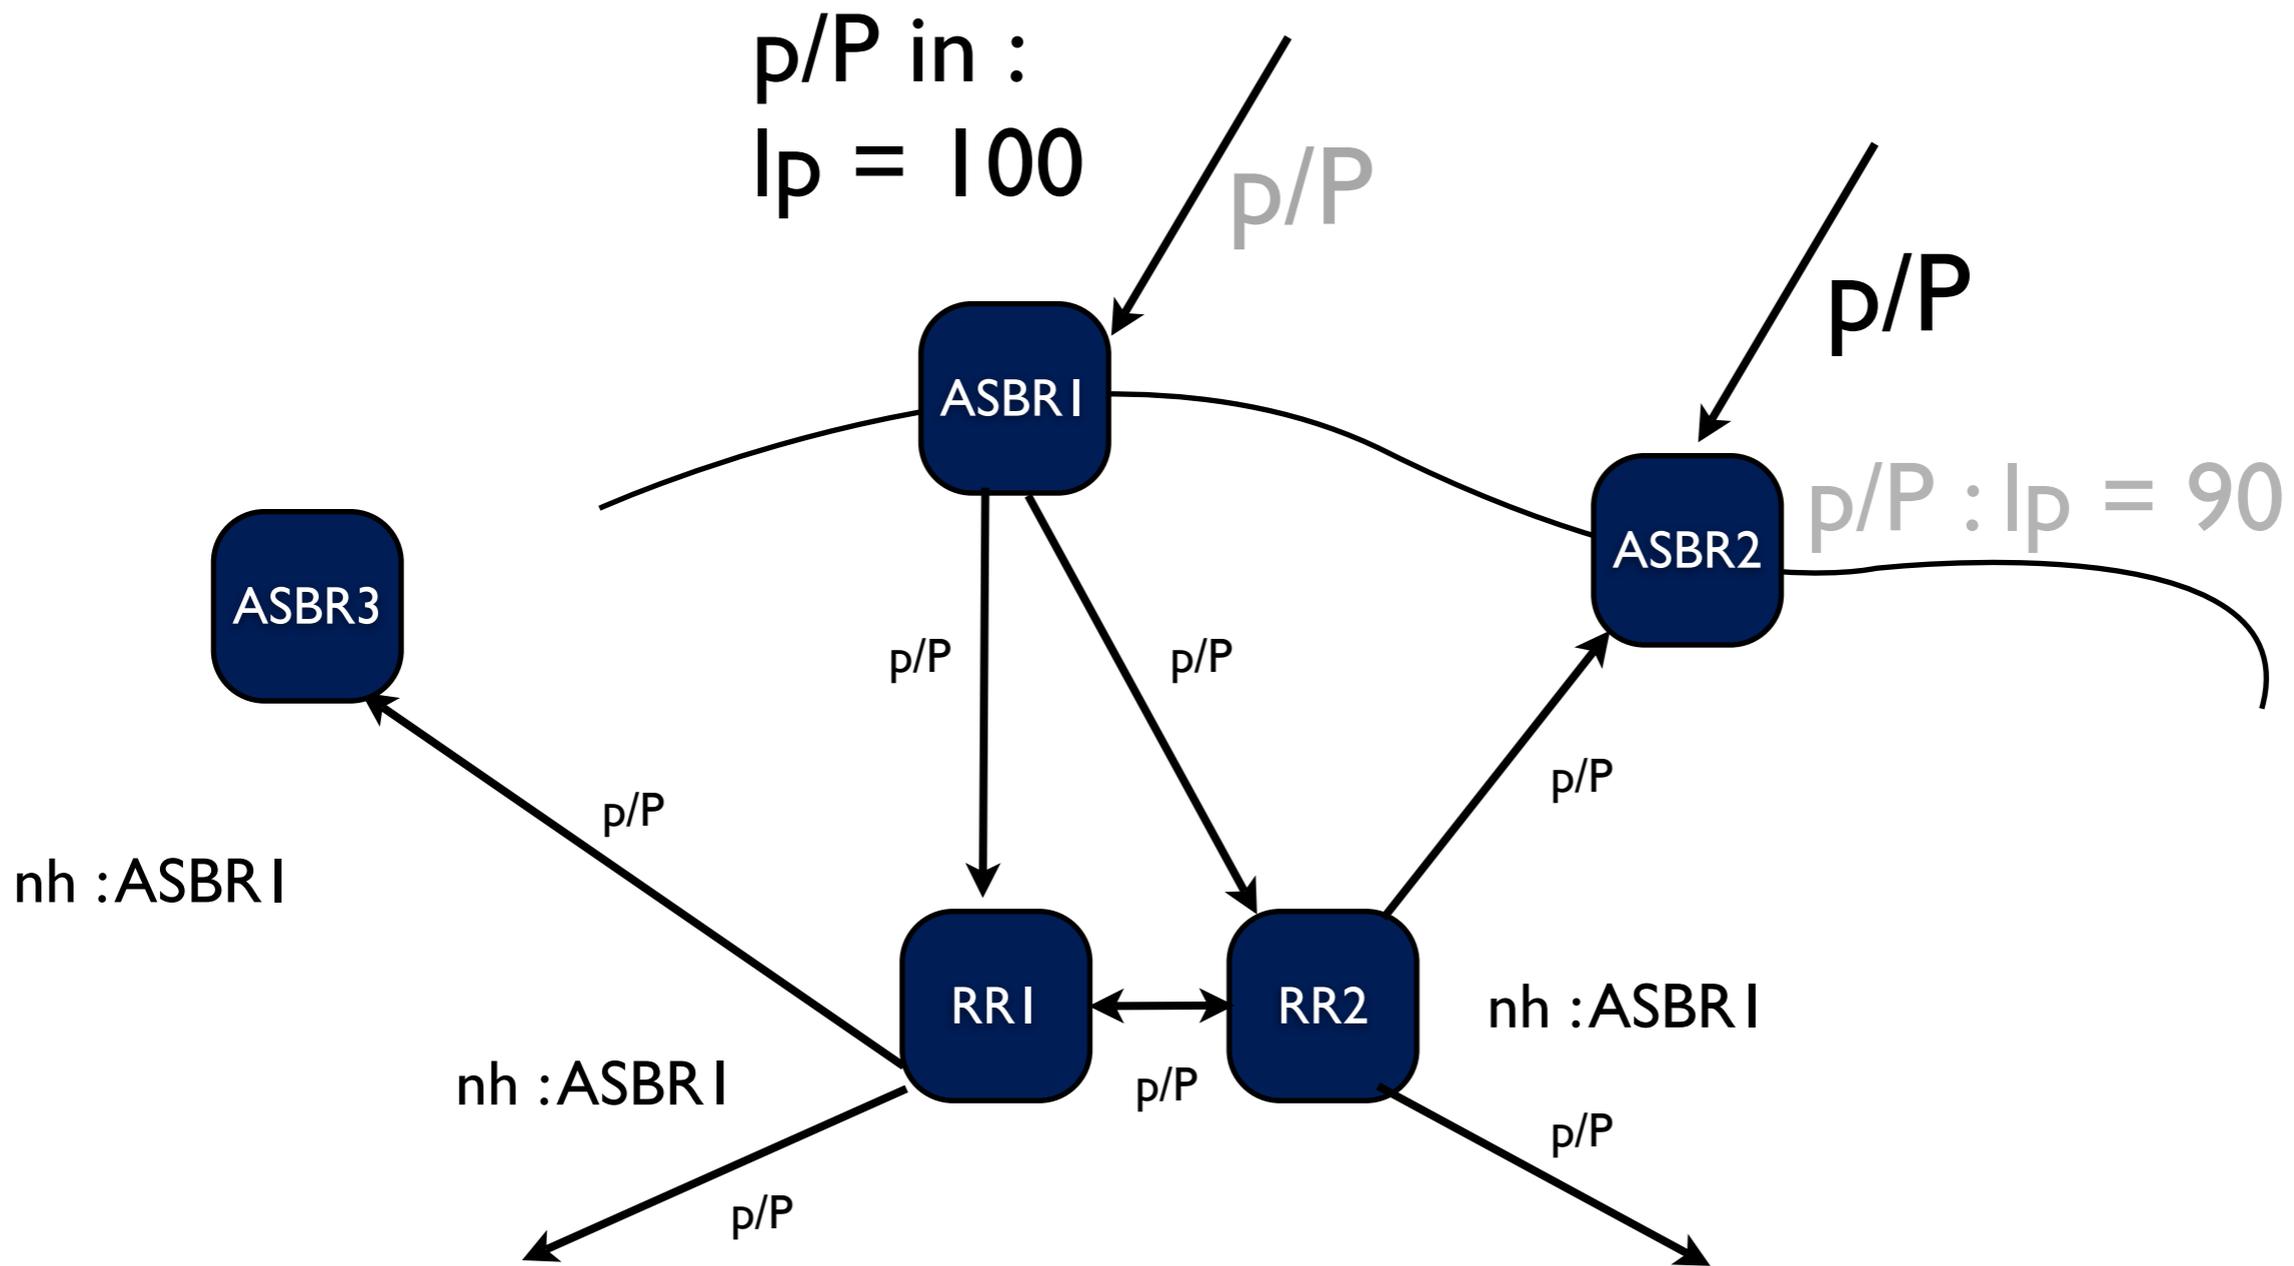
BGP G-Shut

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

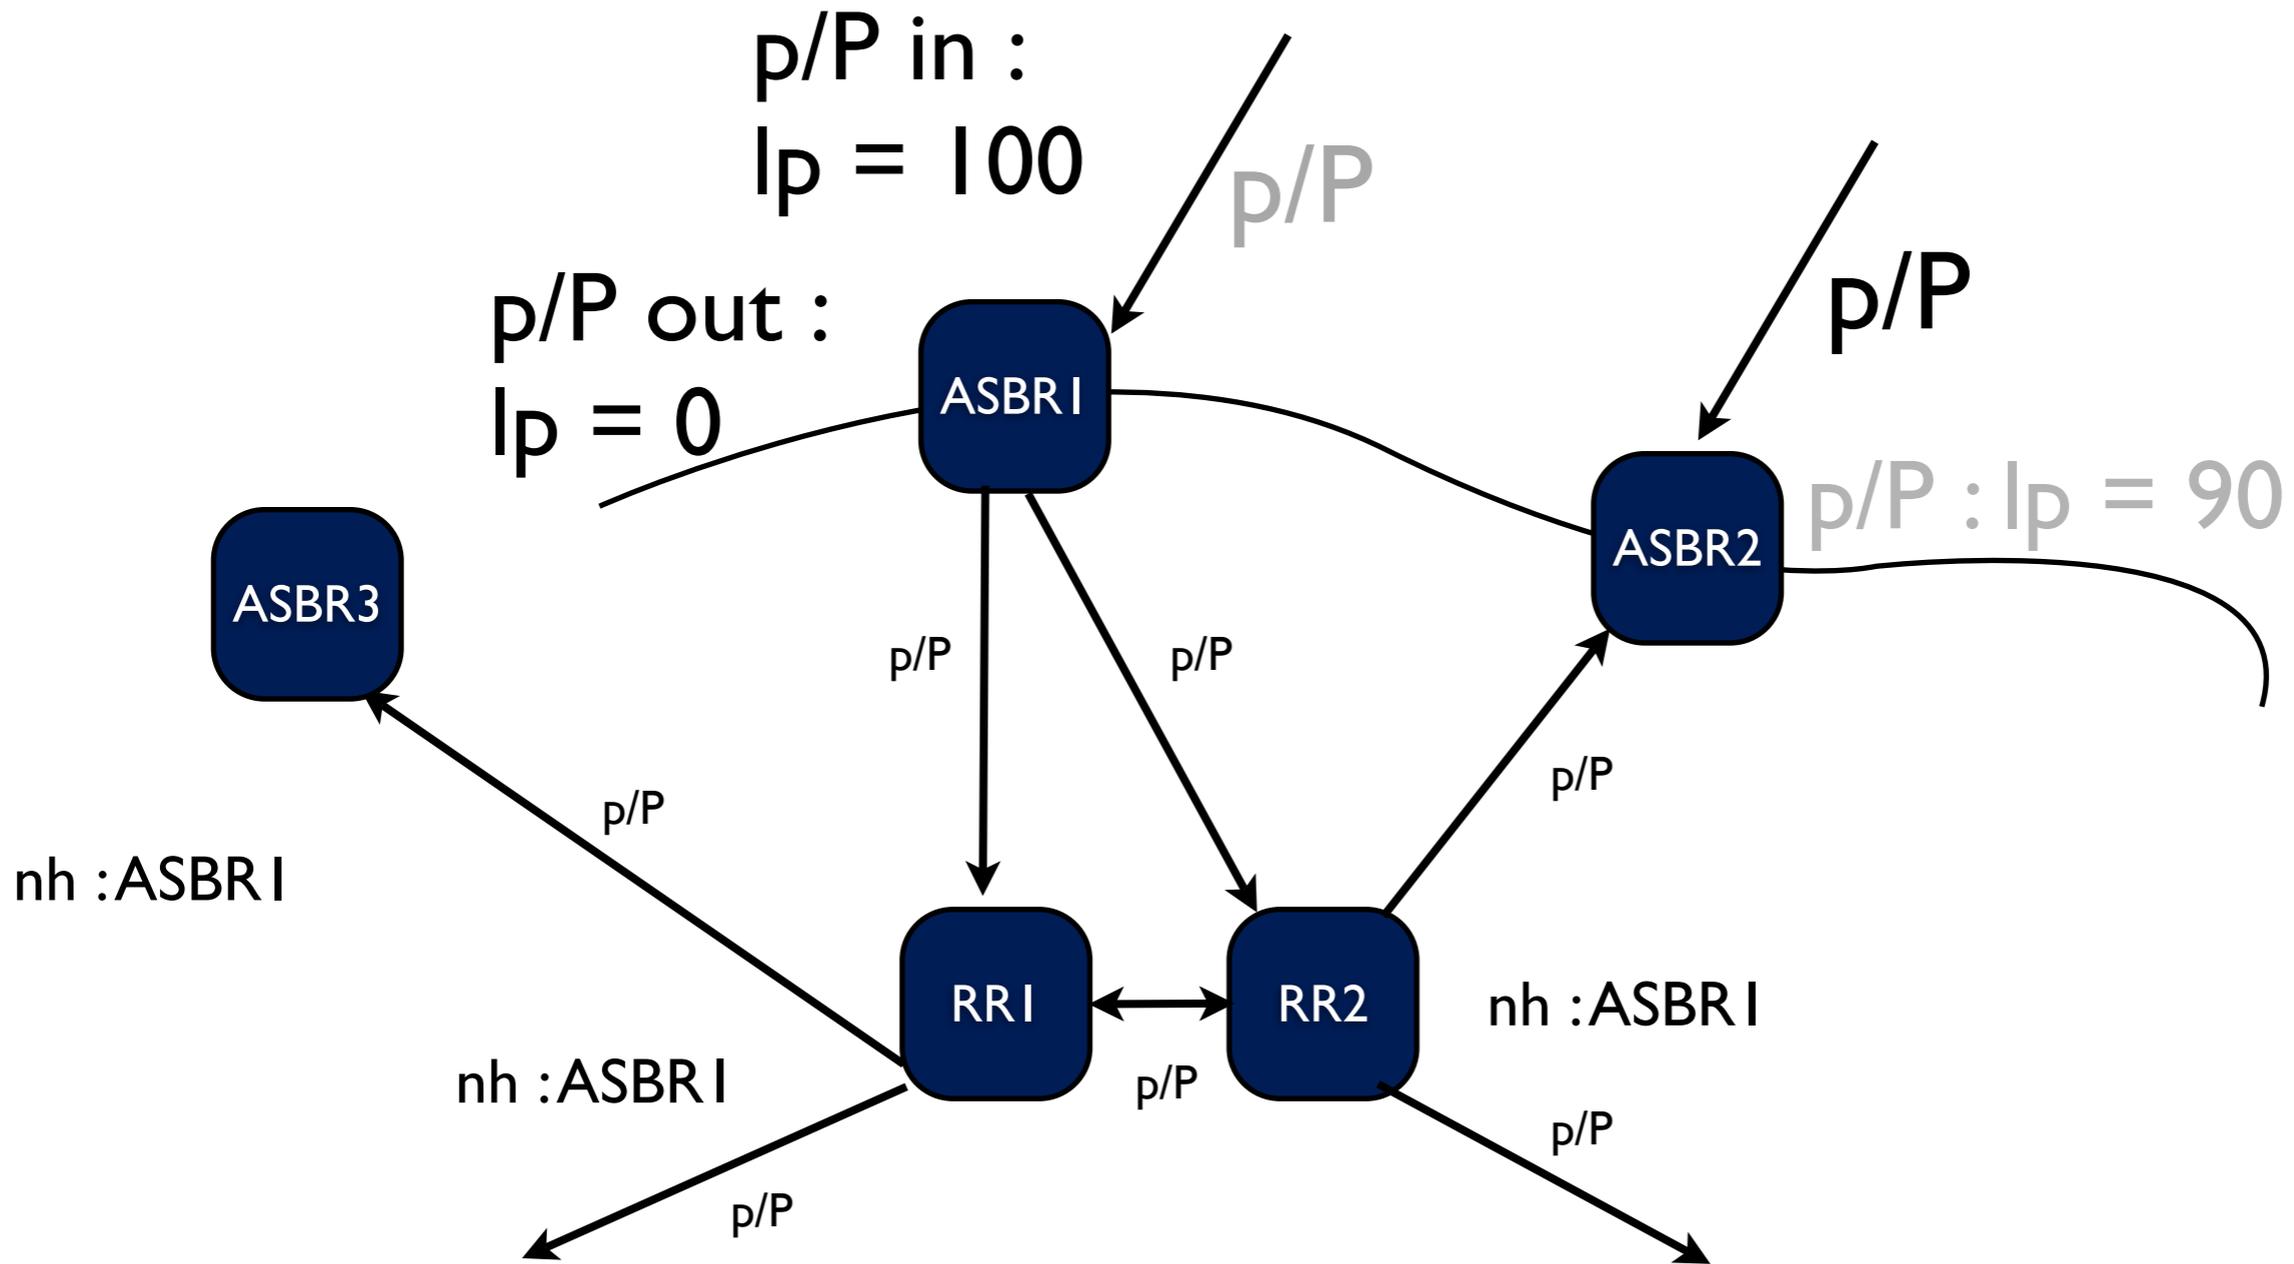
Outbound traffic



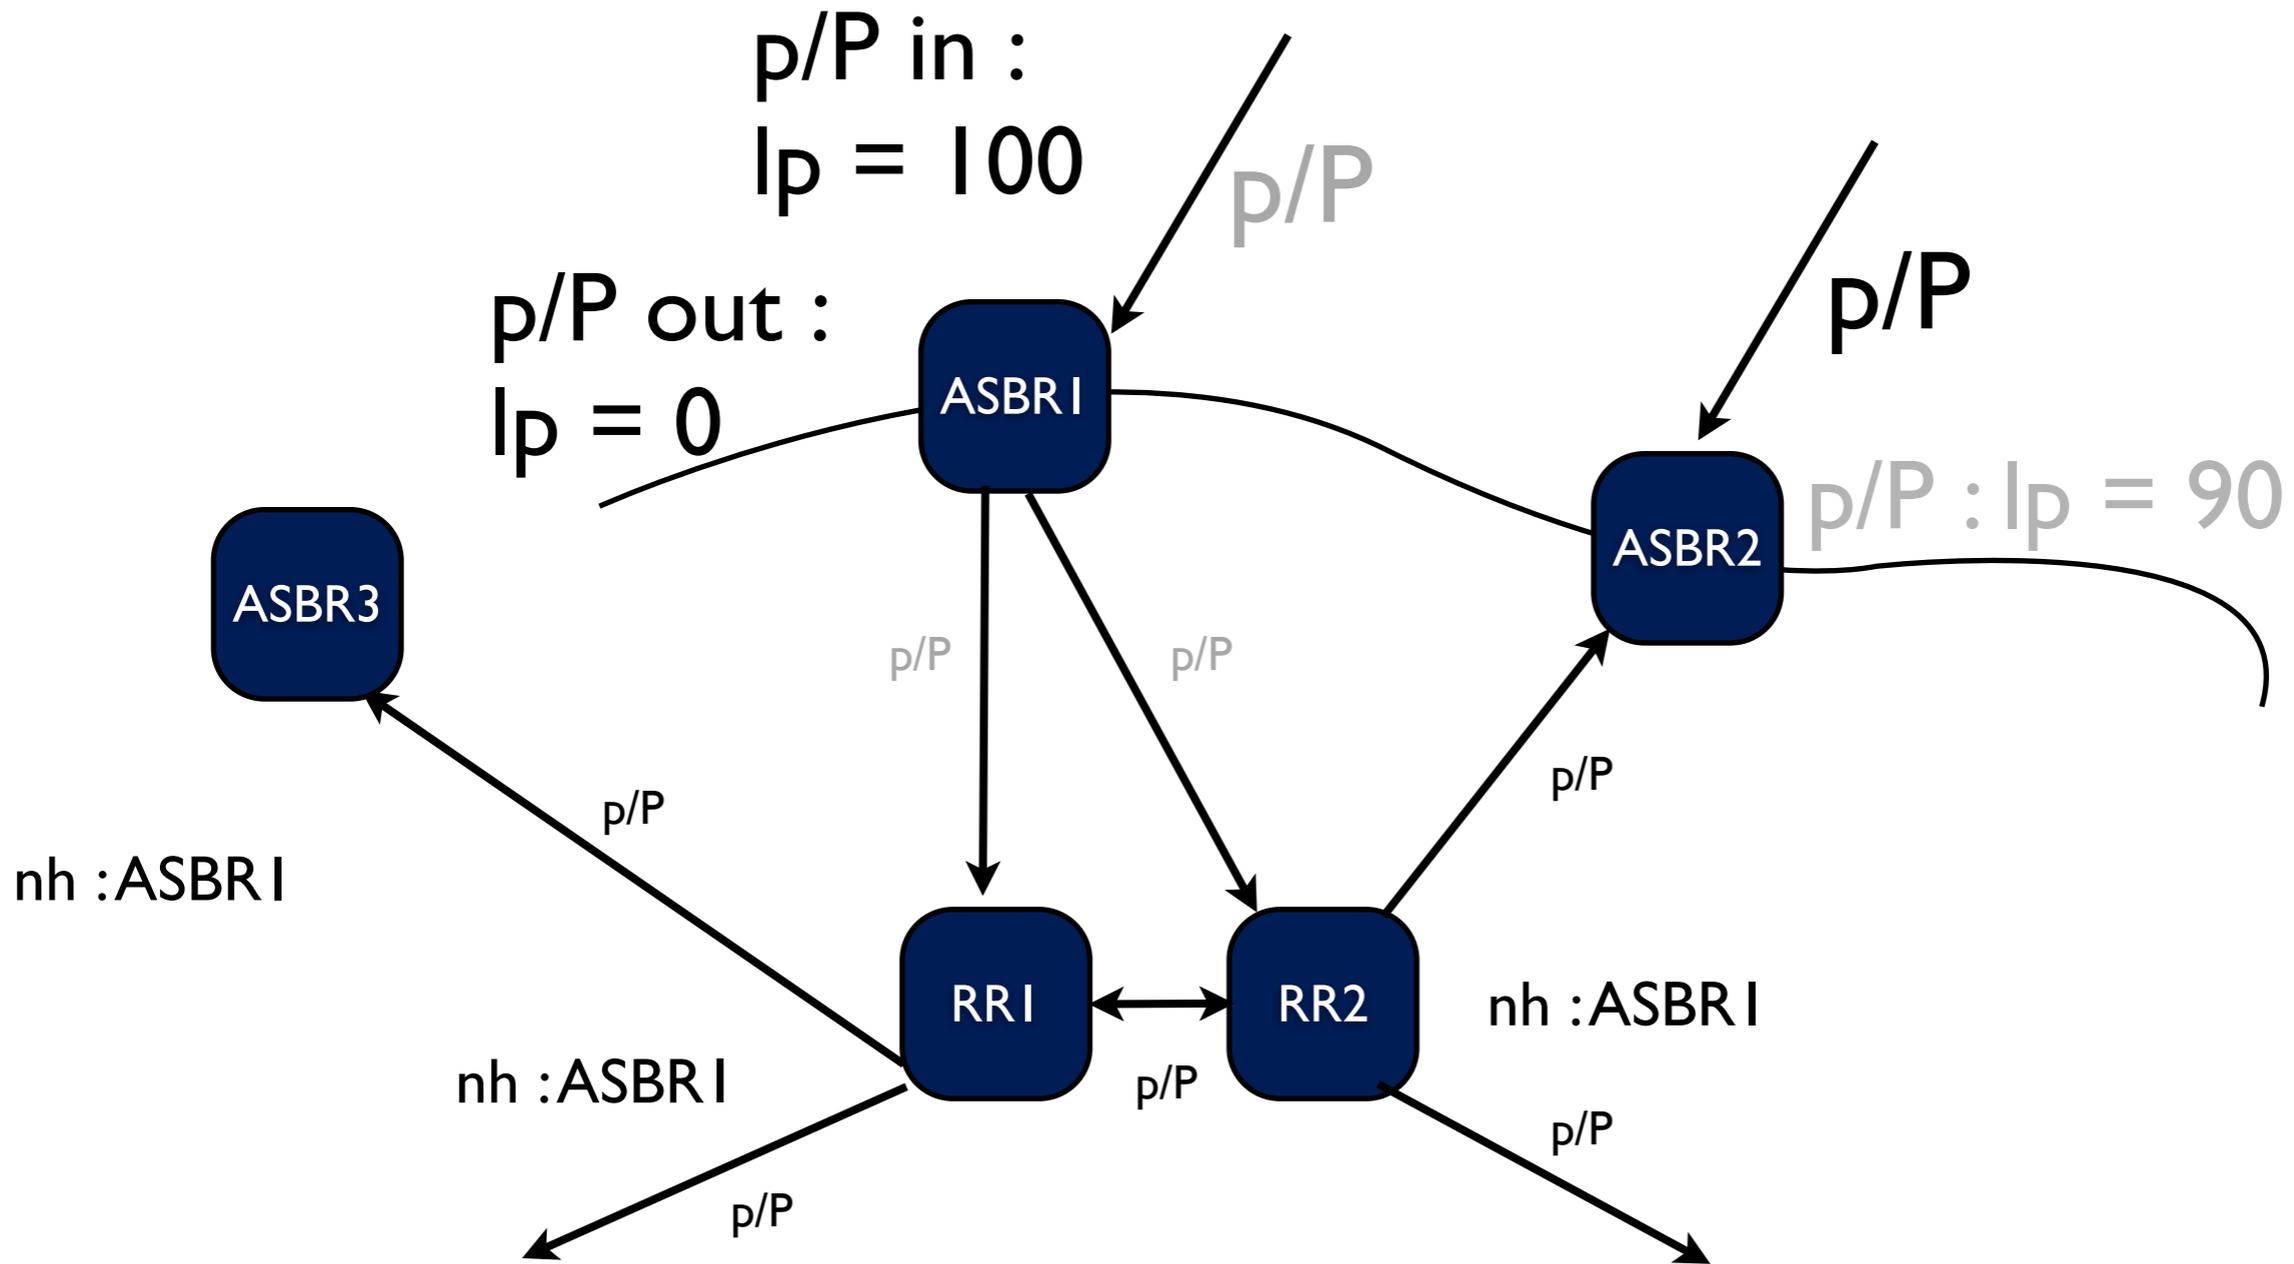
Outbound traffic



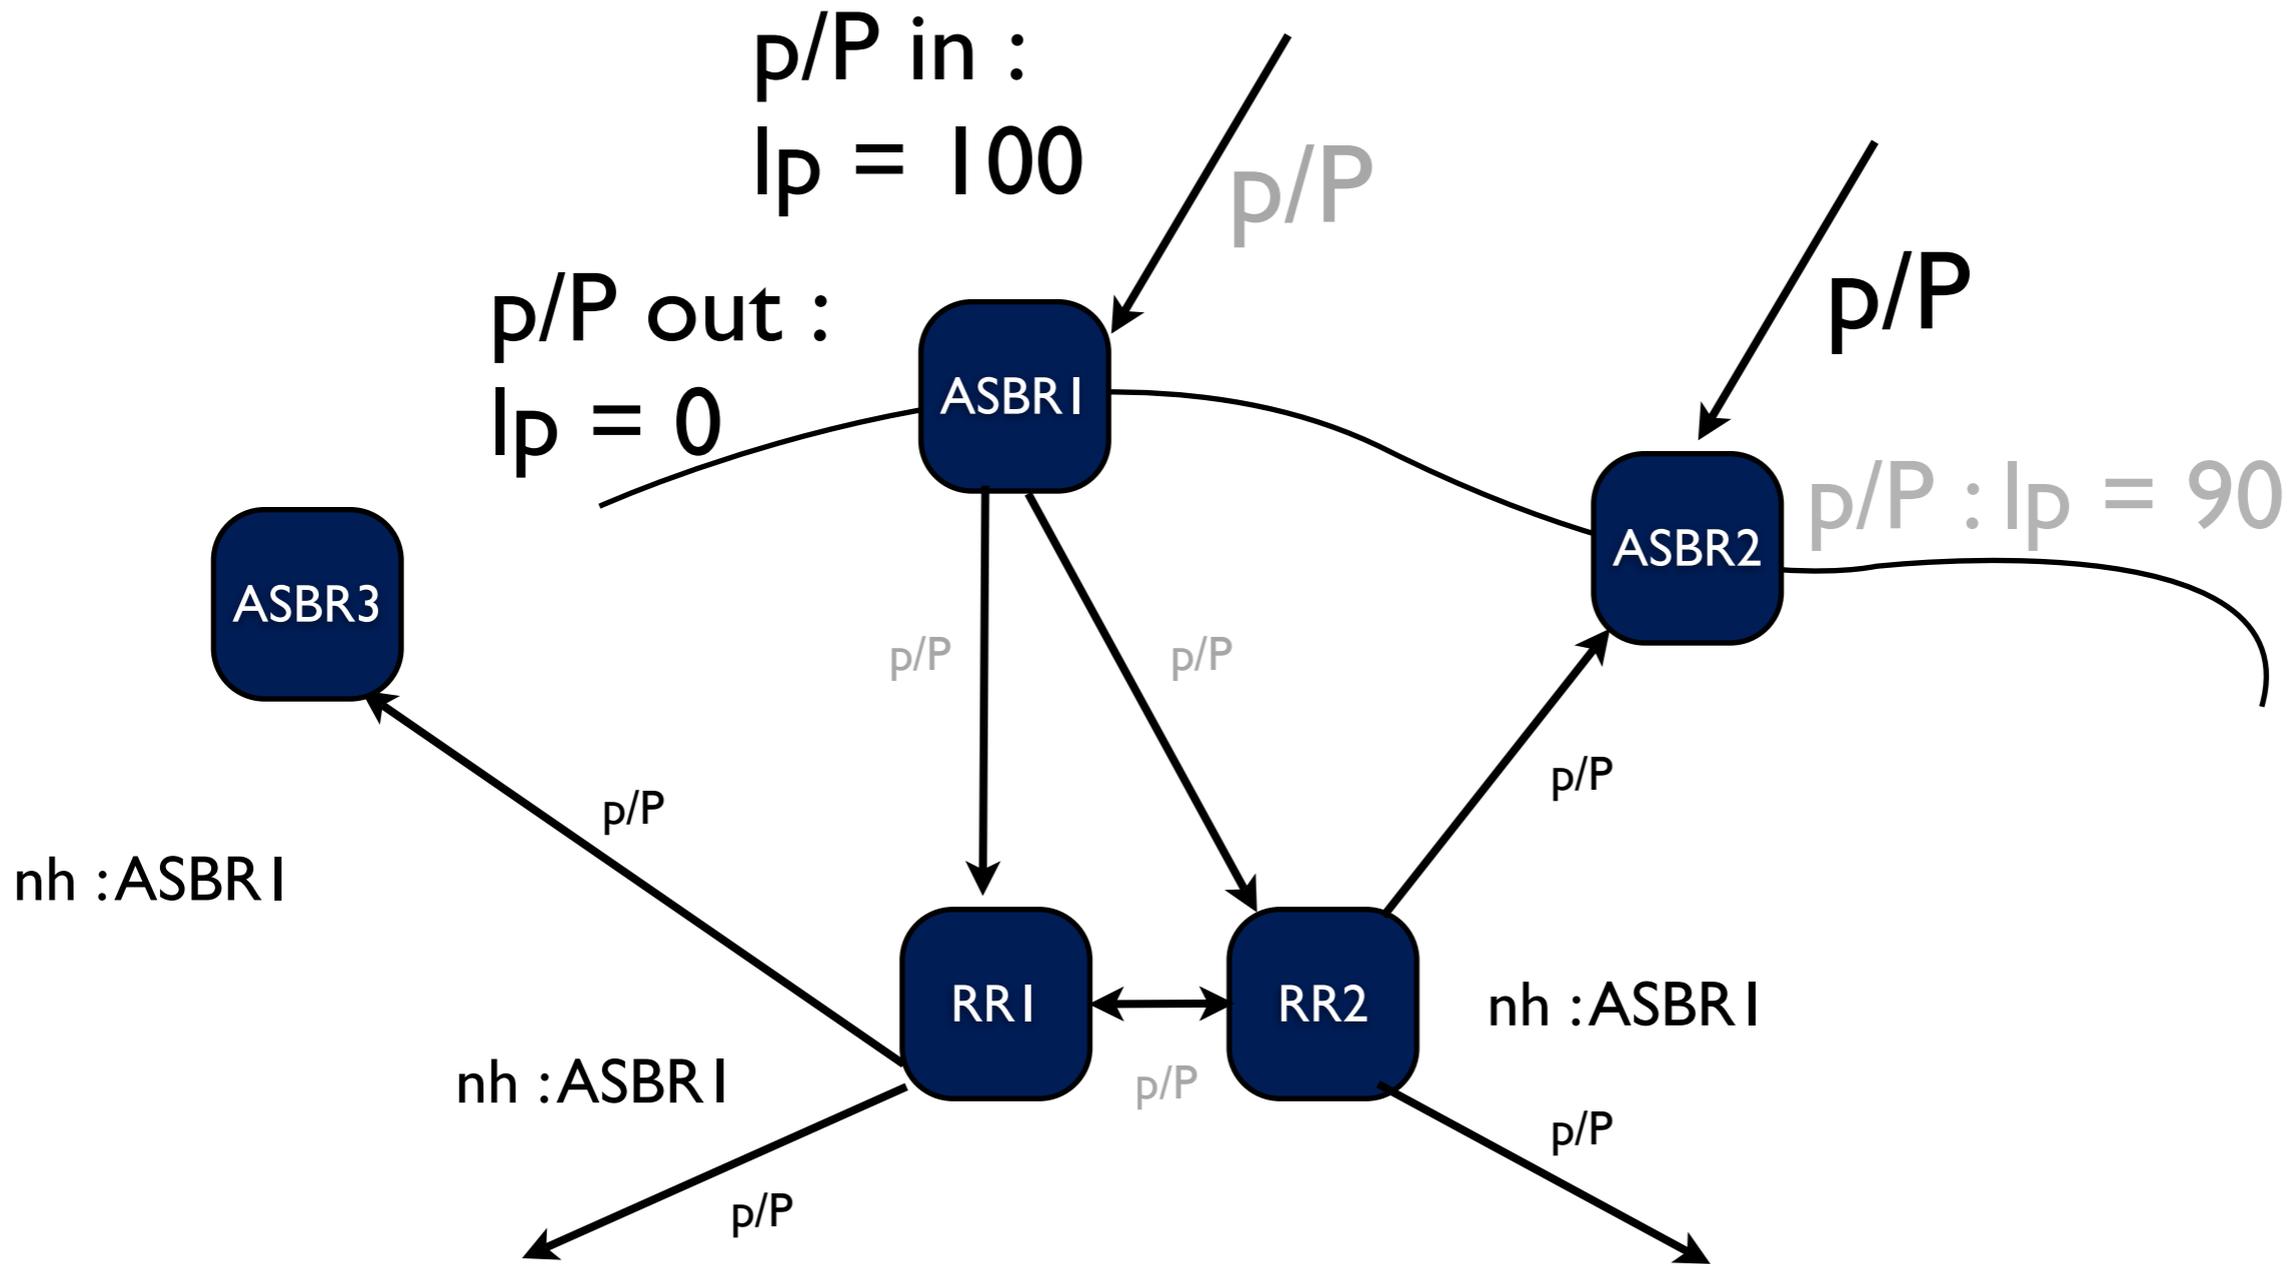
Outbound traffic



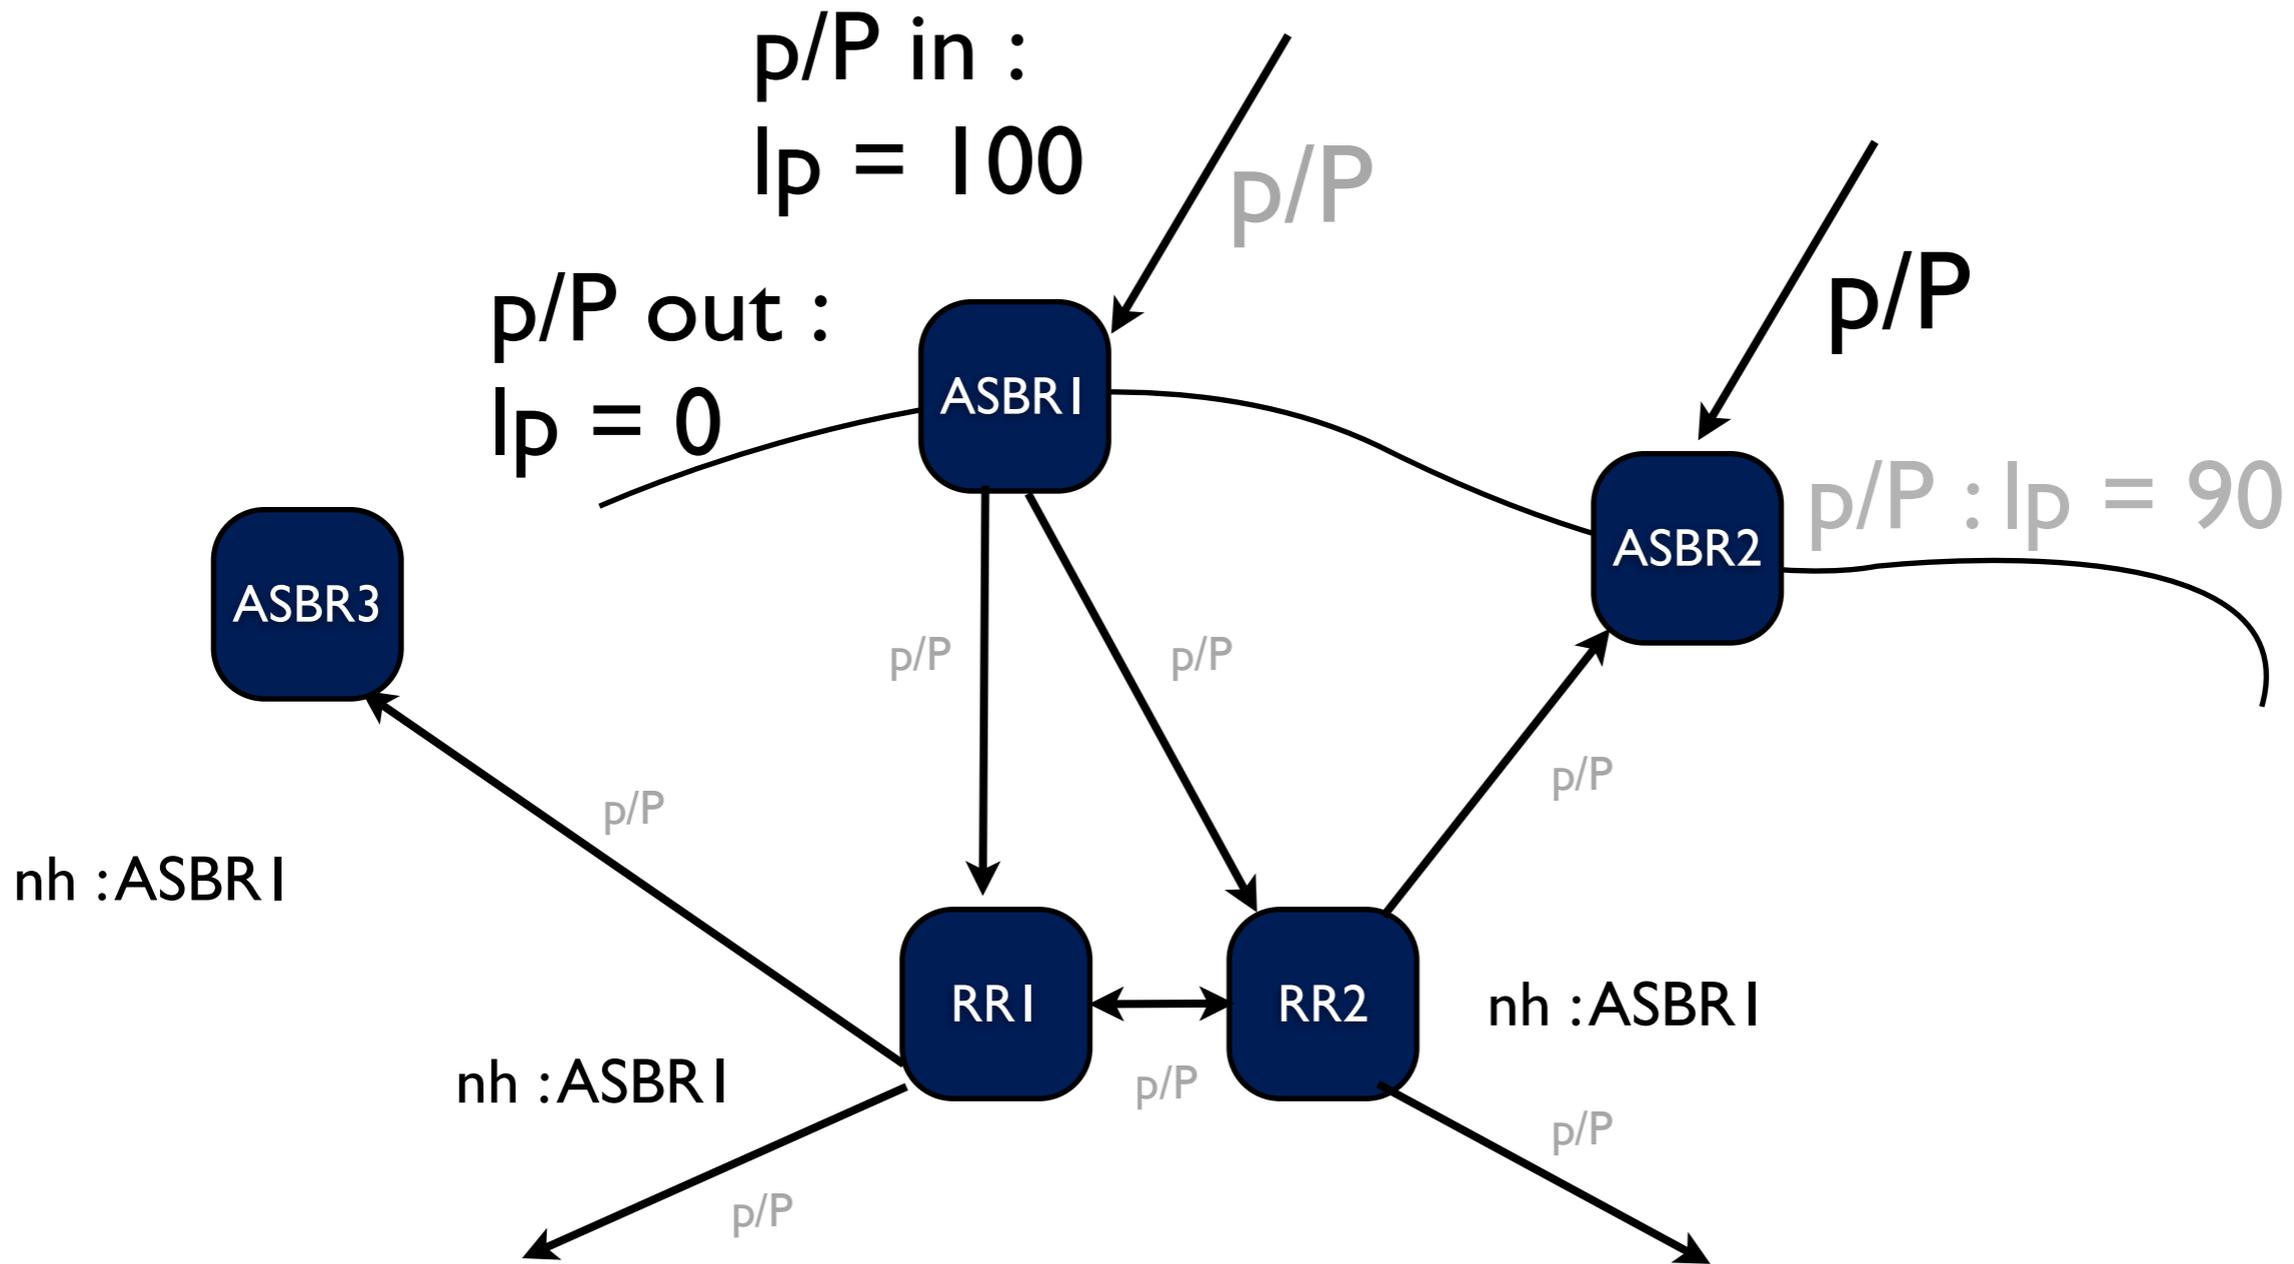
Outbound traffic



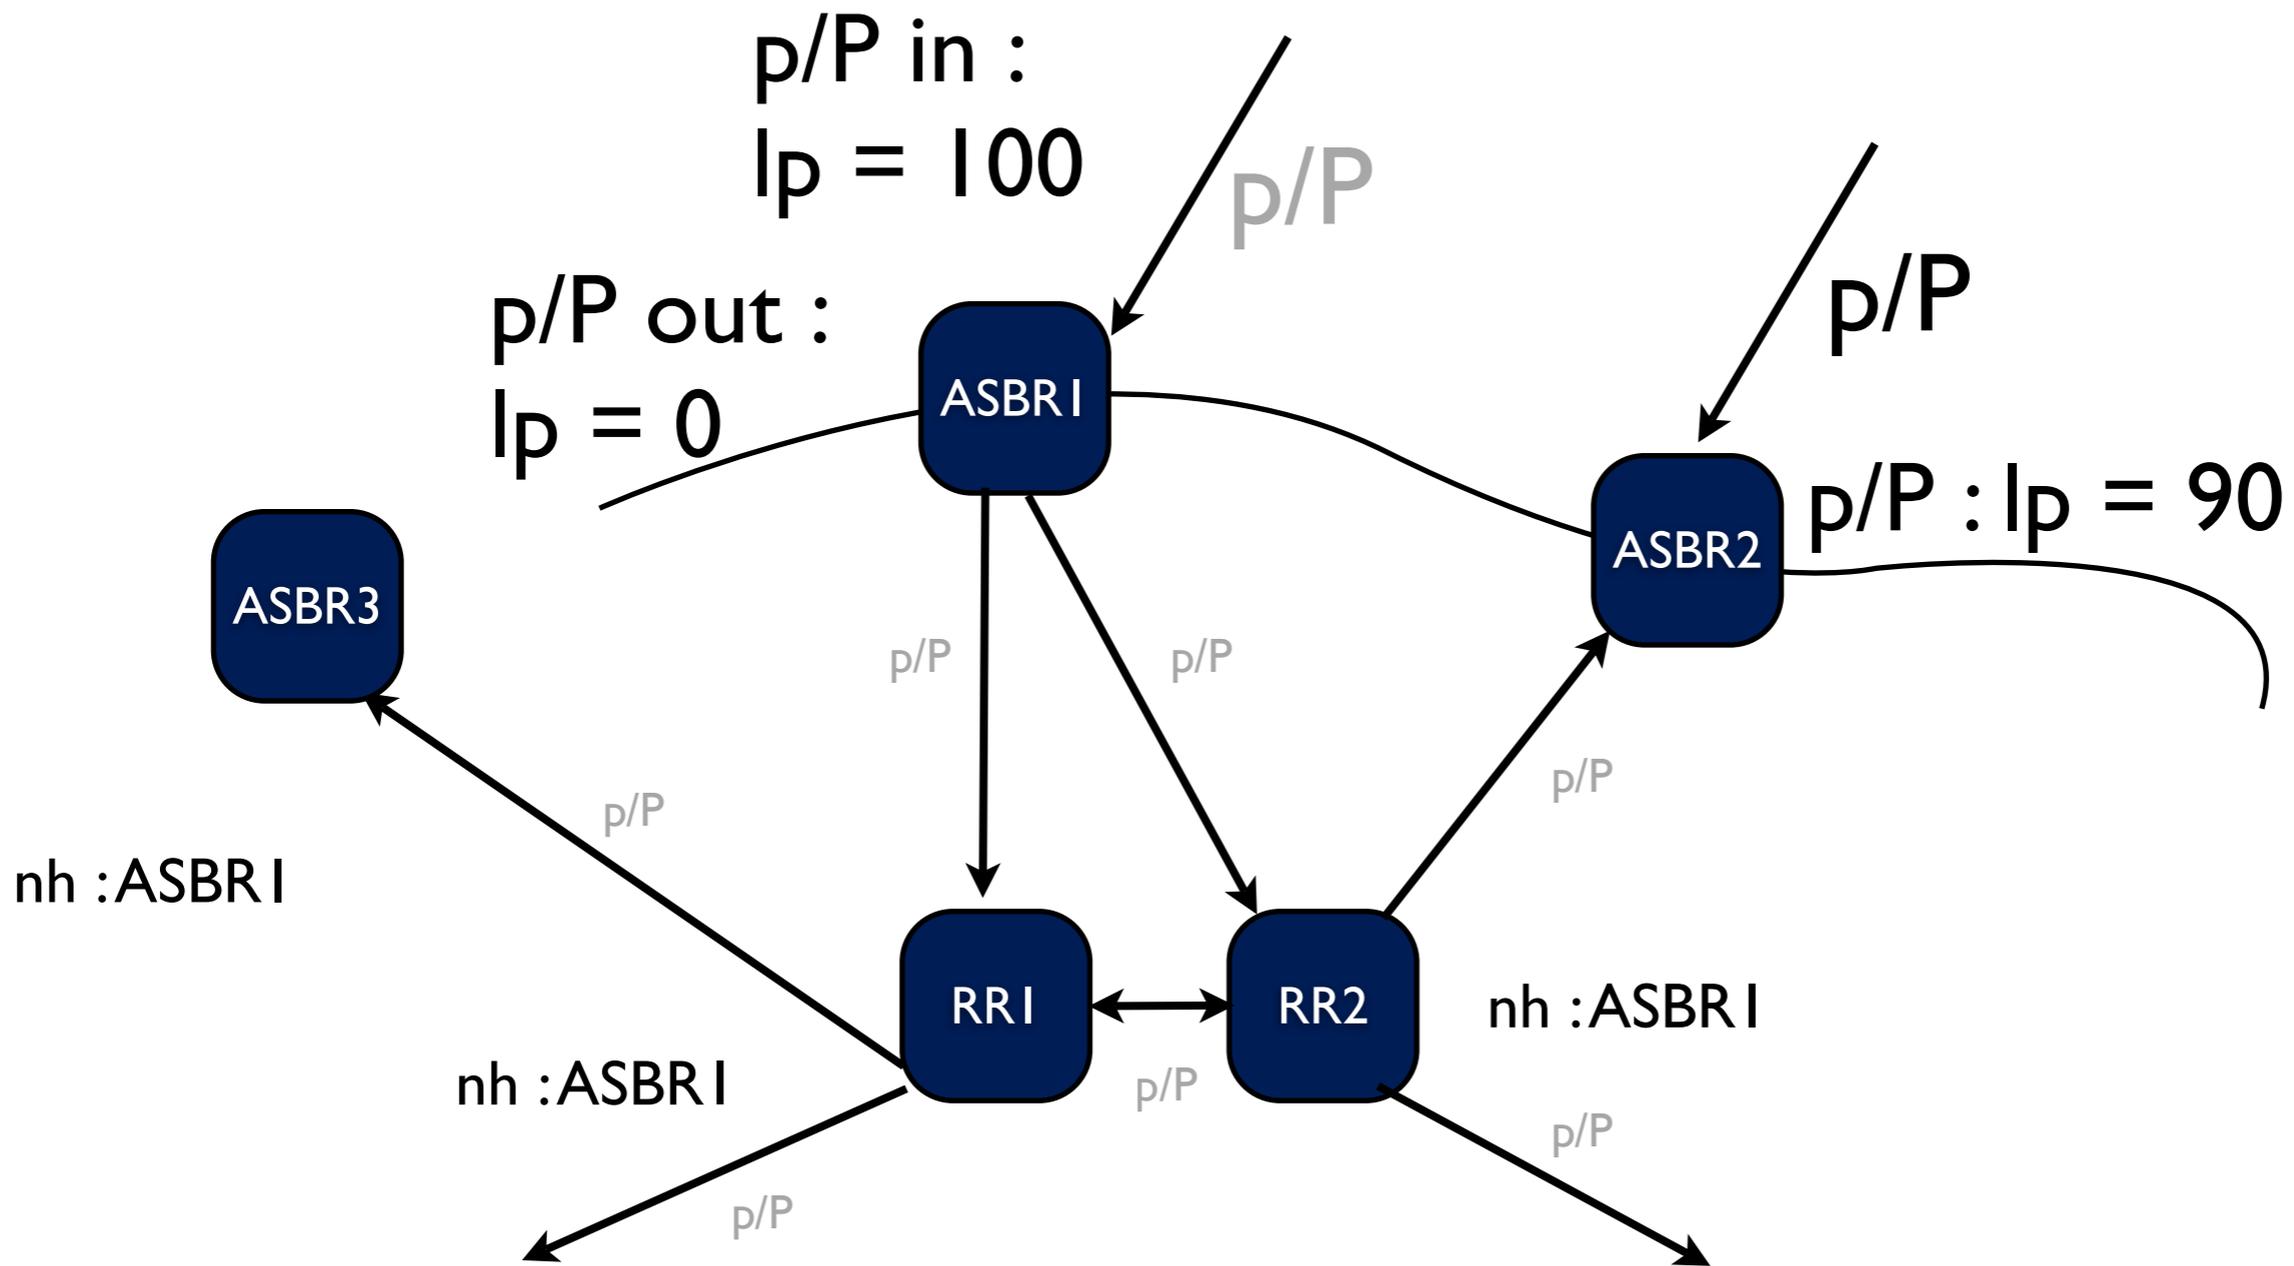
Outbound traffic



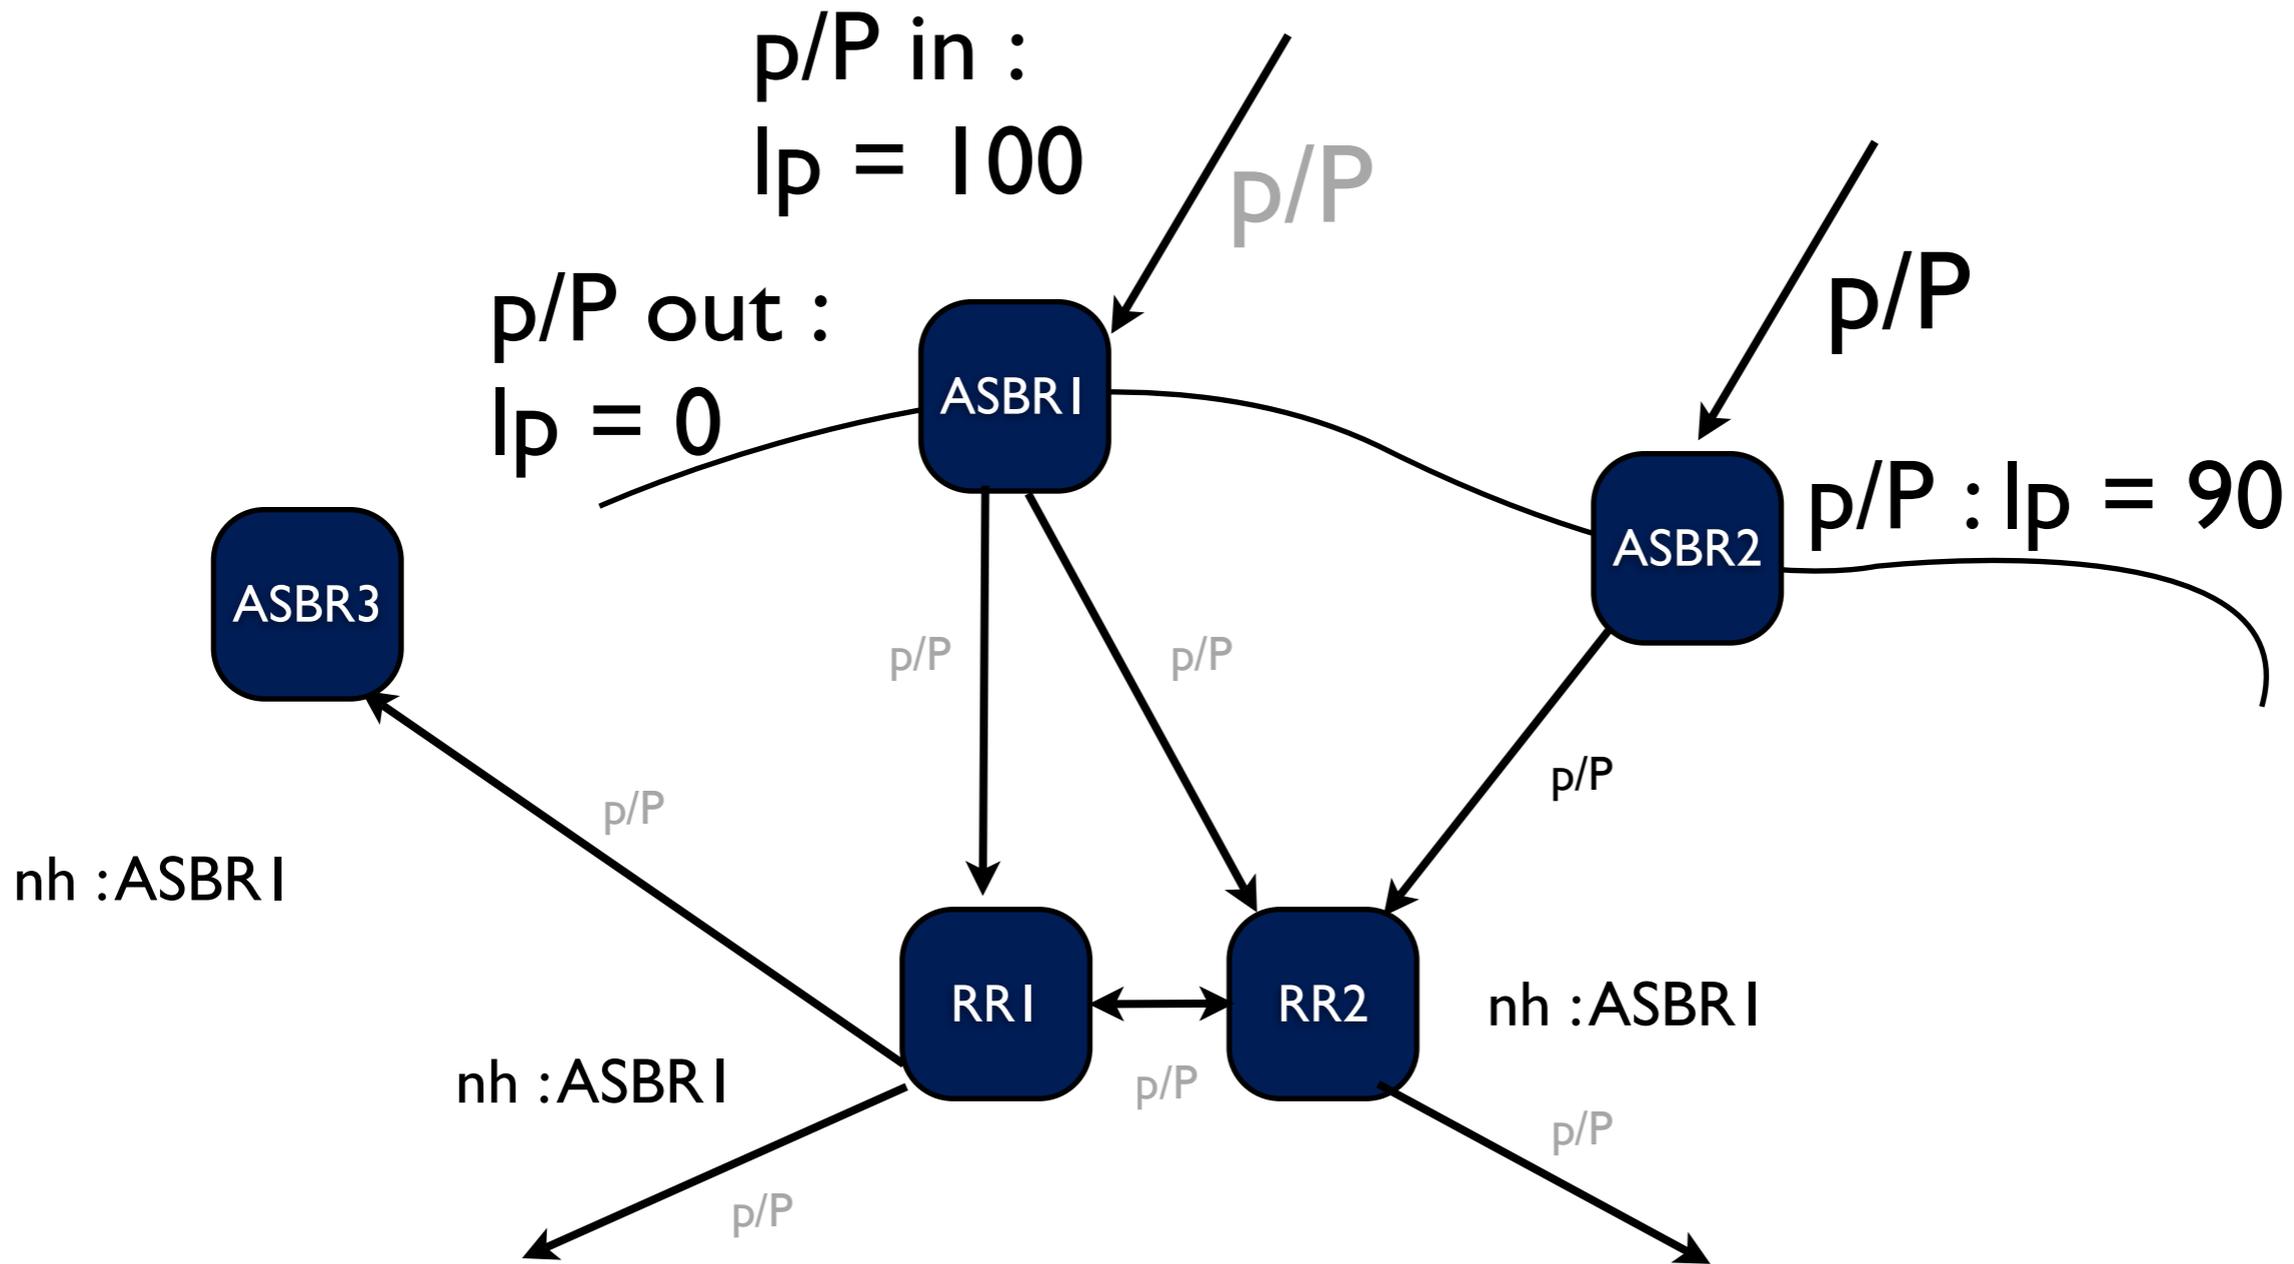
Outbound traffic



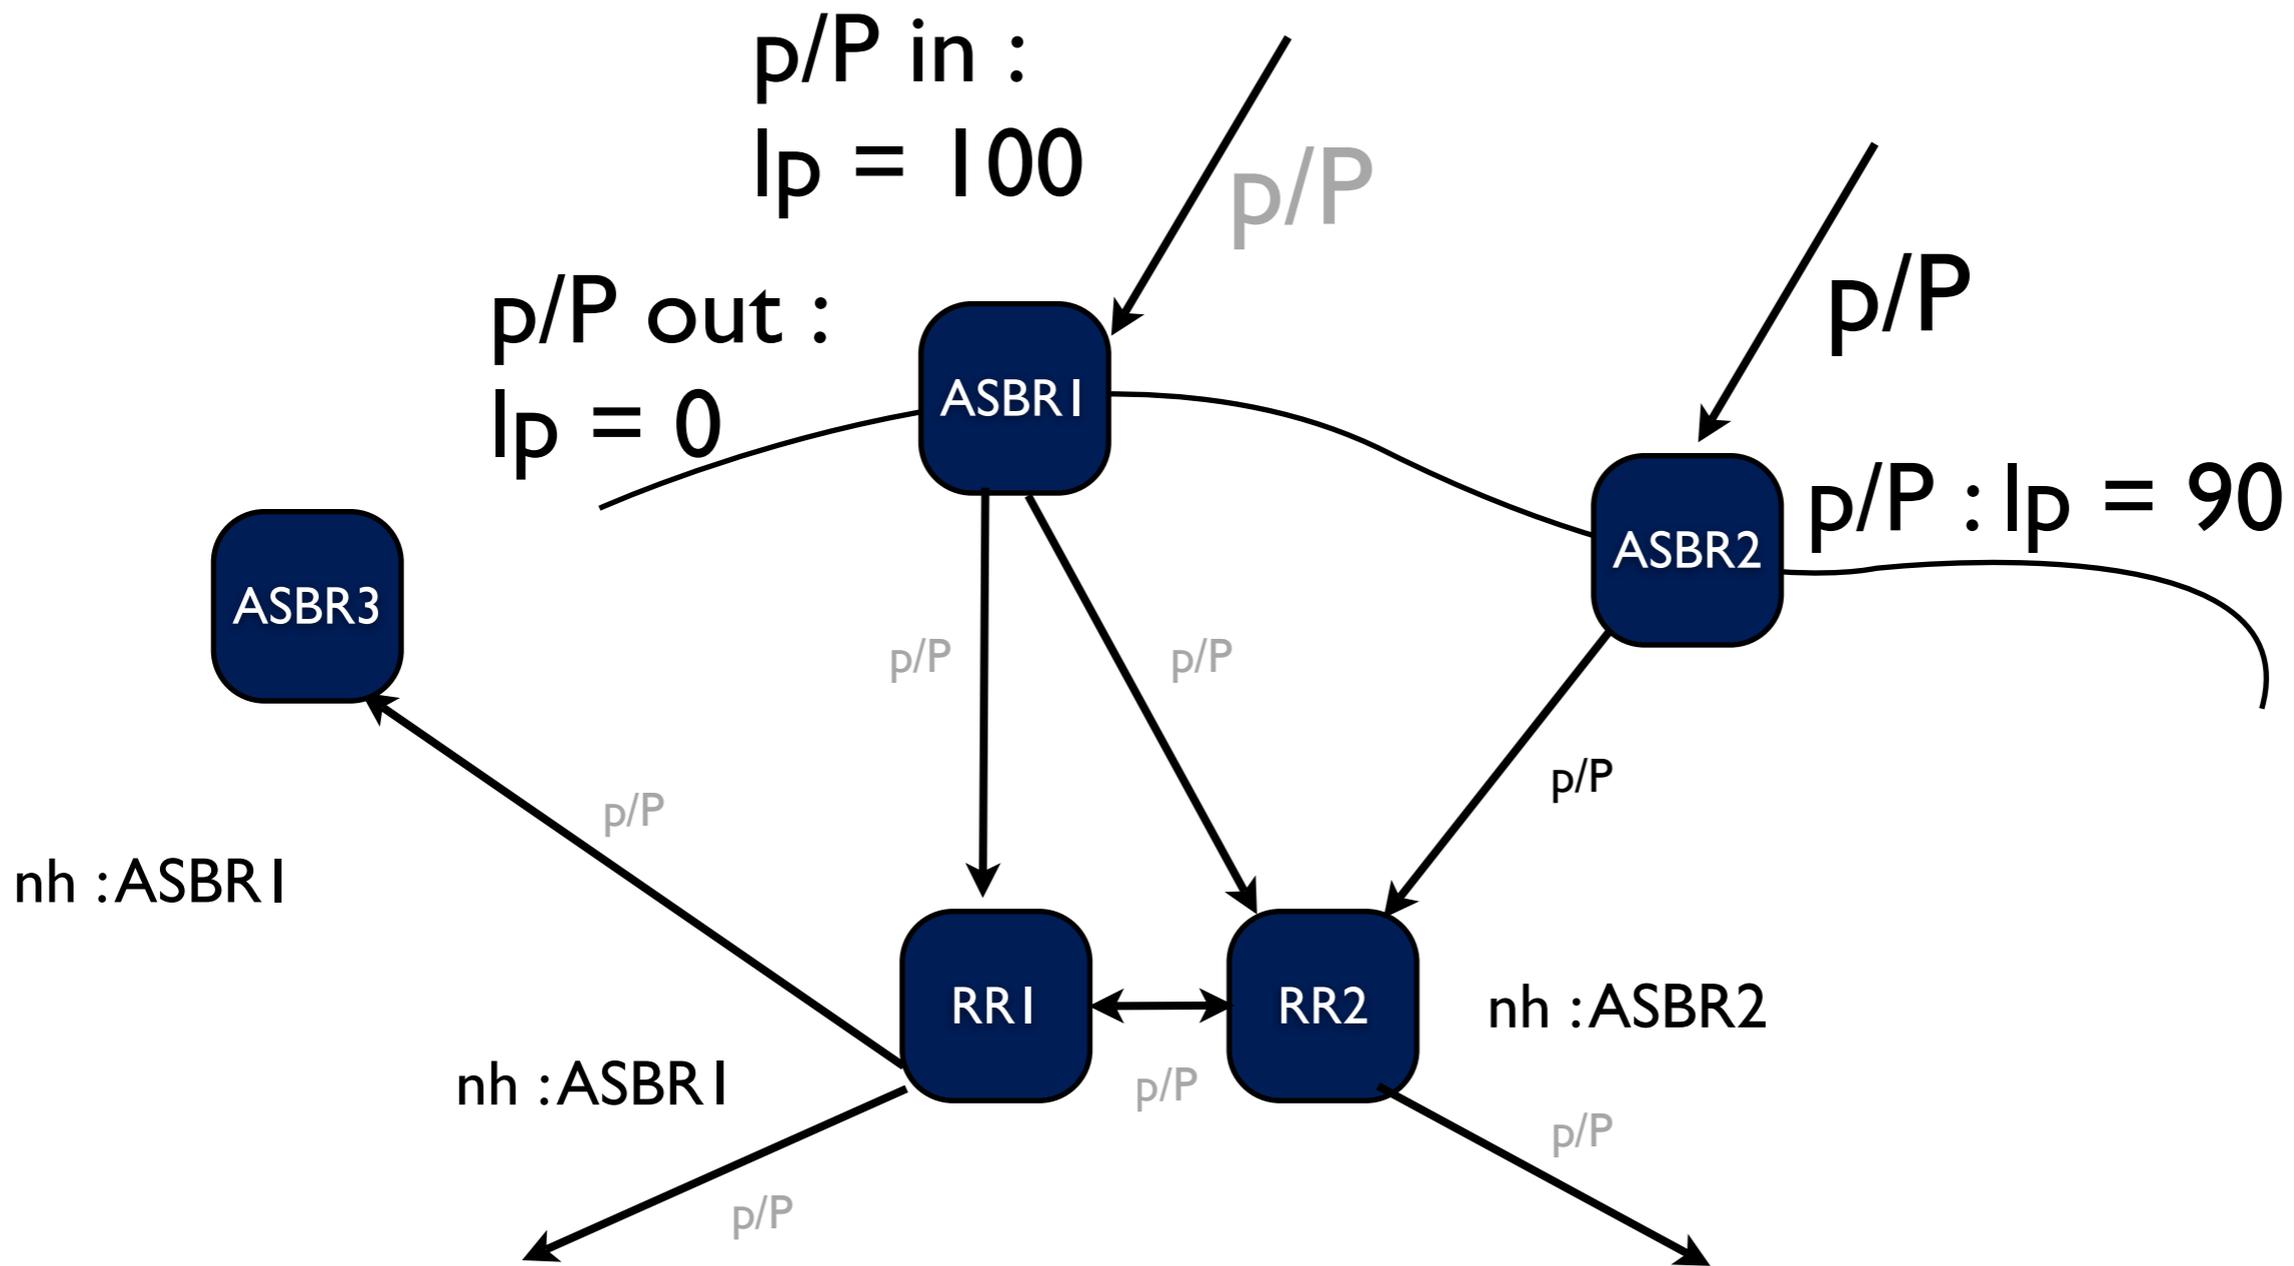
Outbound traffic



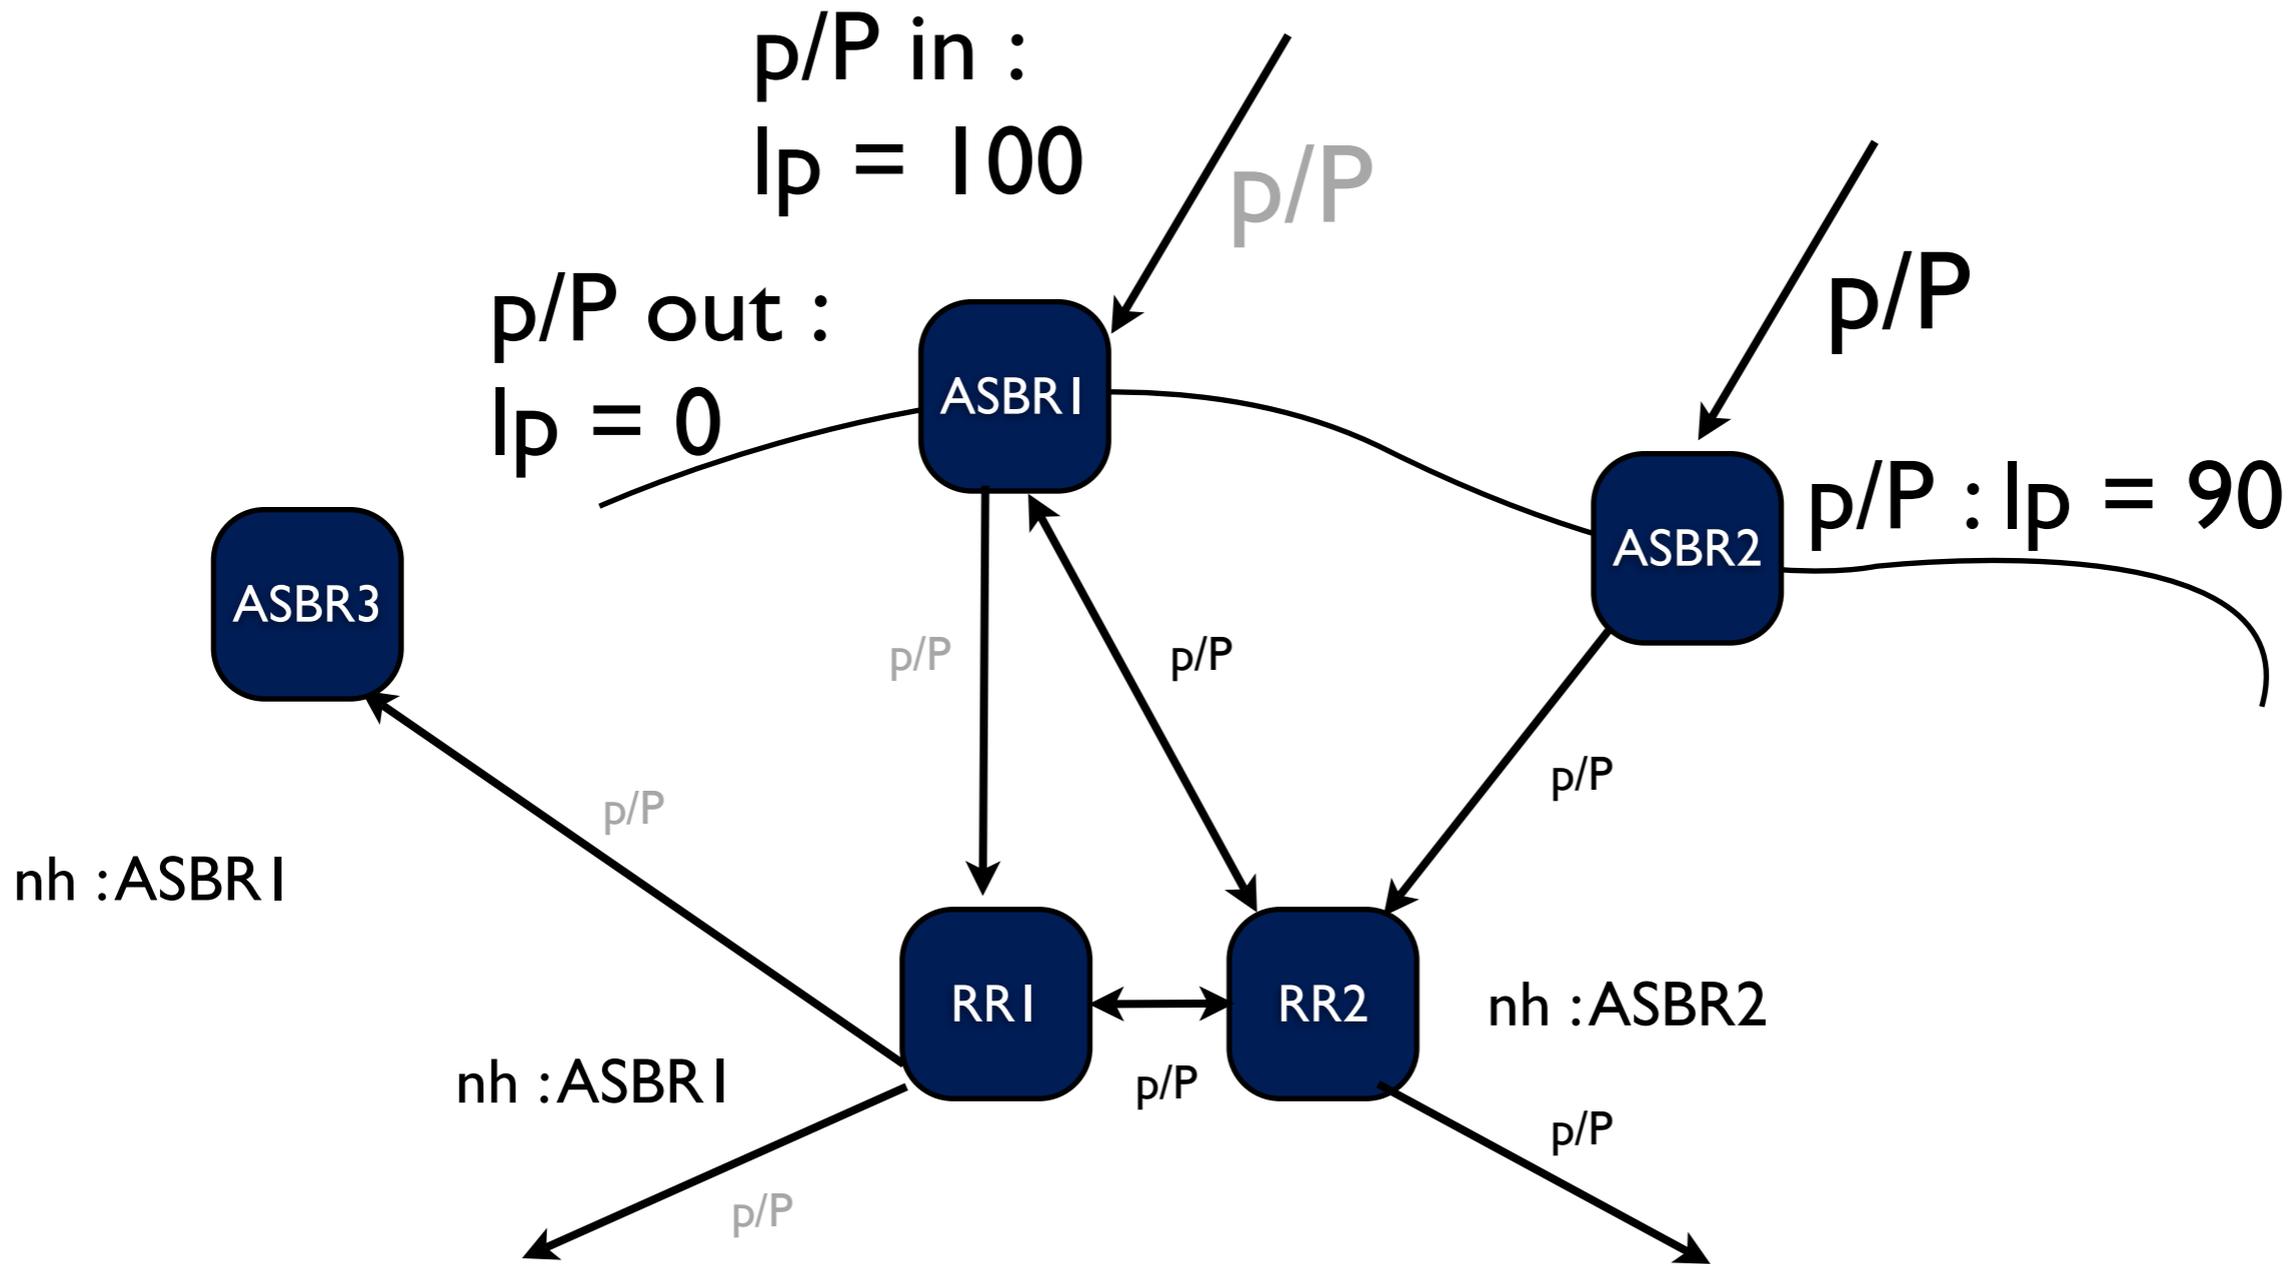
Outbound traffic



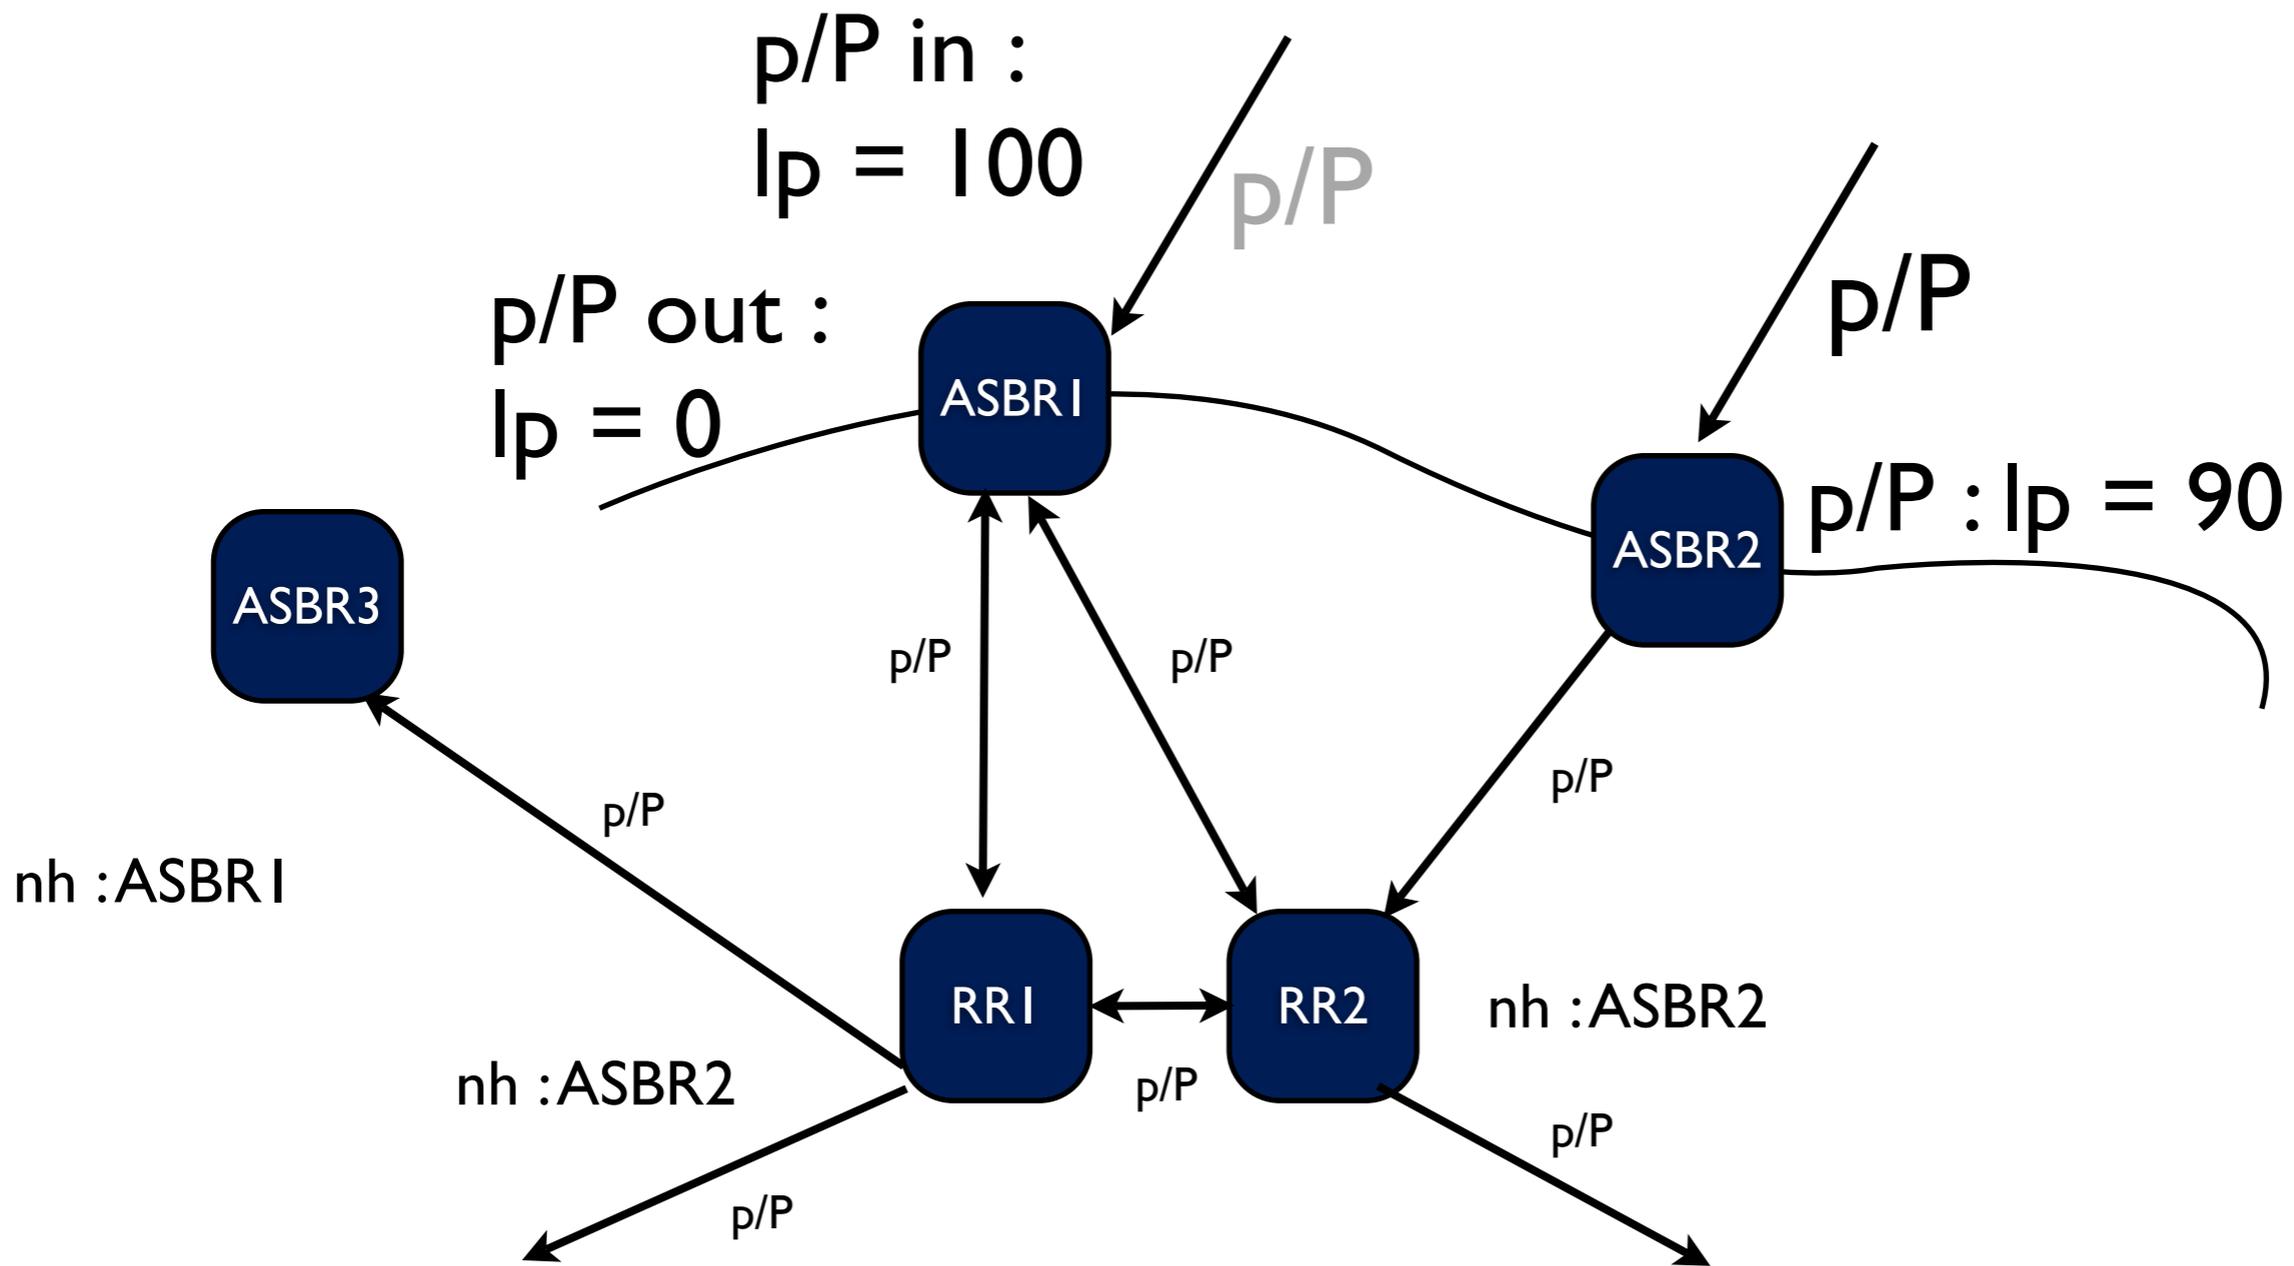
Outbound traffic



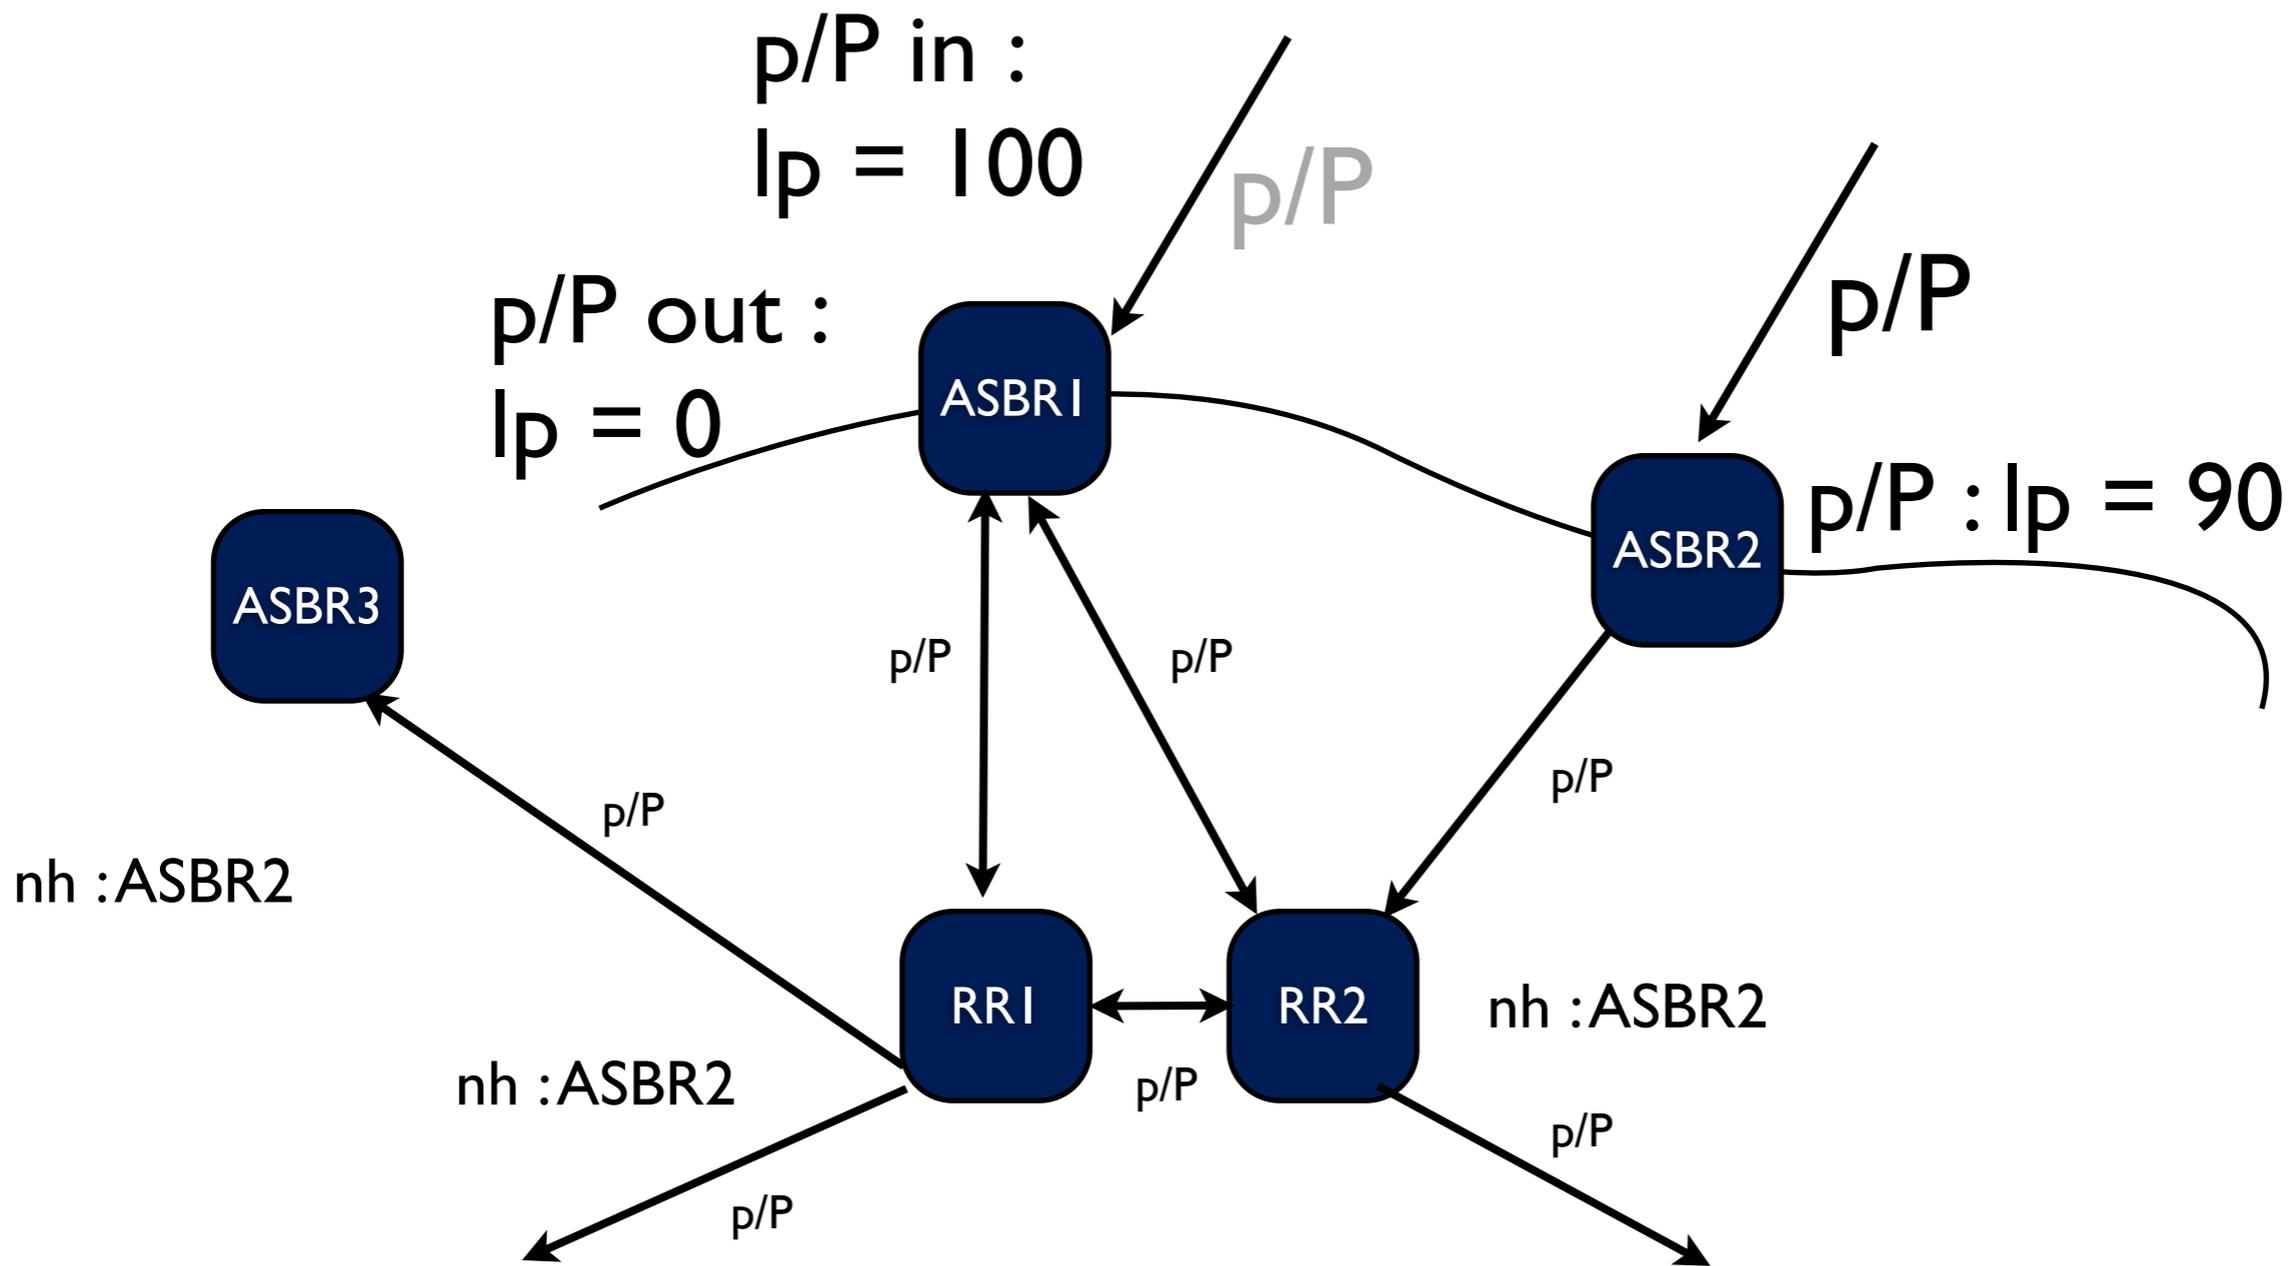
Outbound traffic



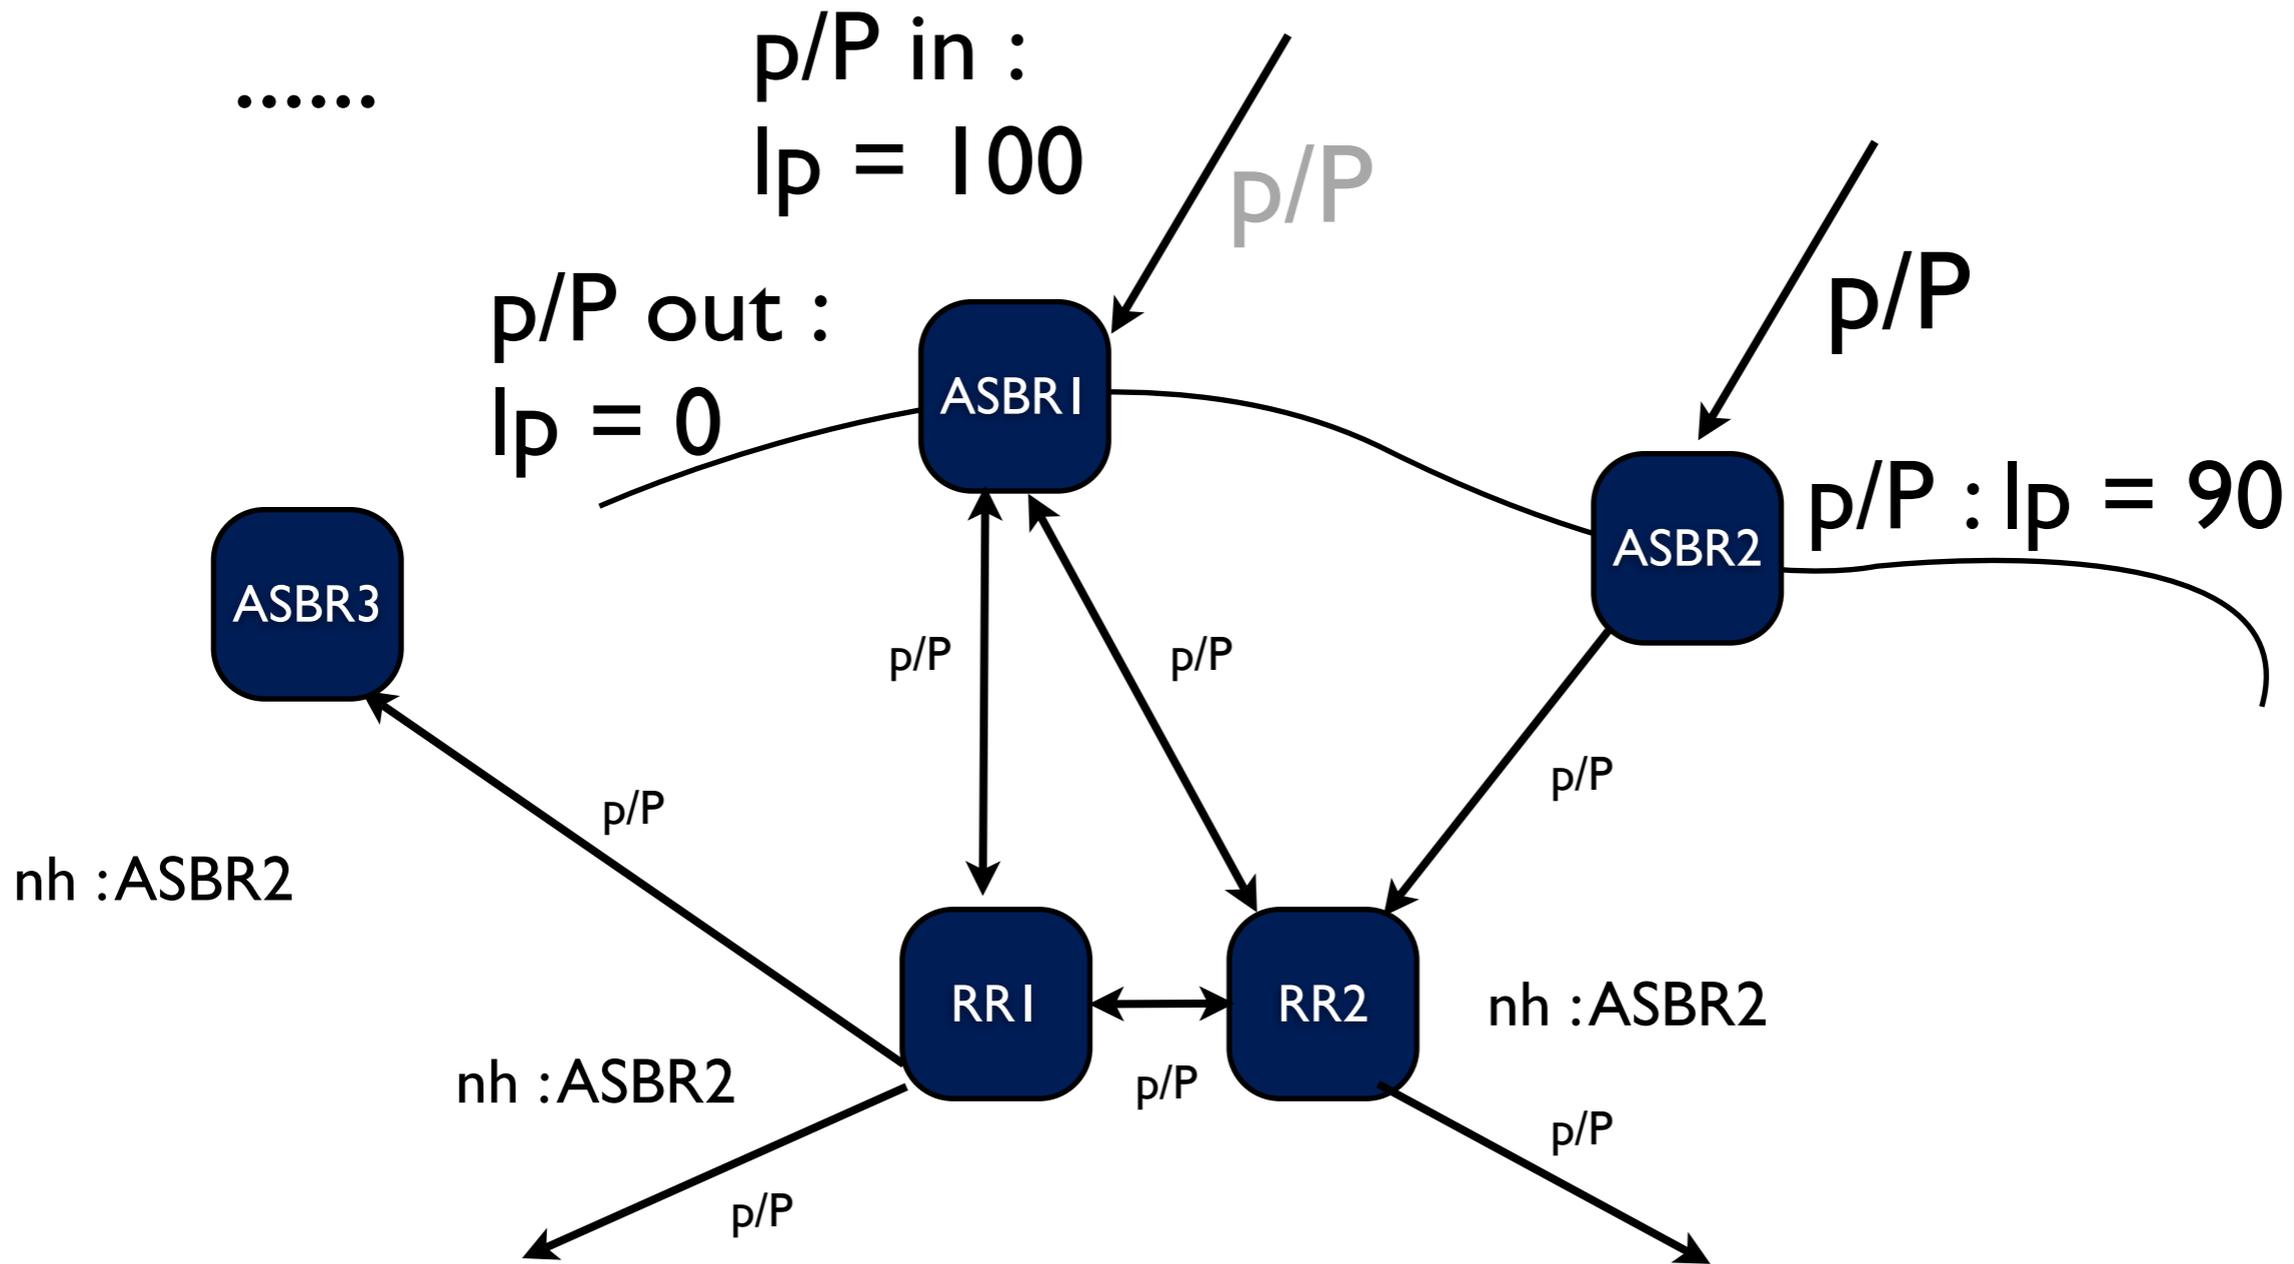
Outbound traffic



Outbound traffic

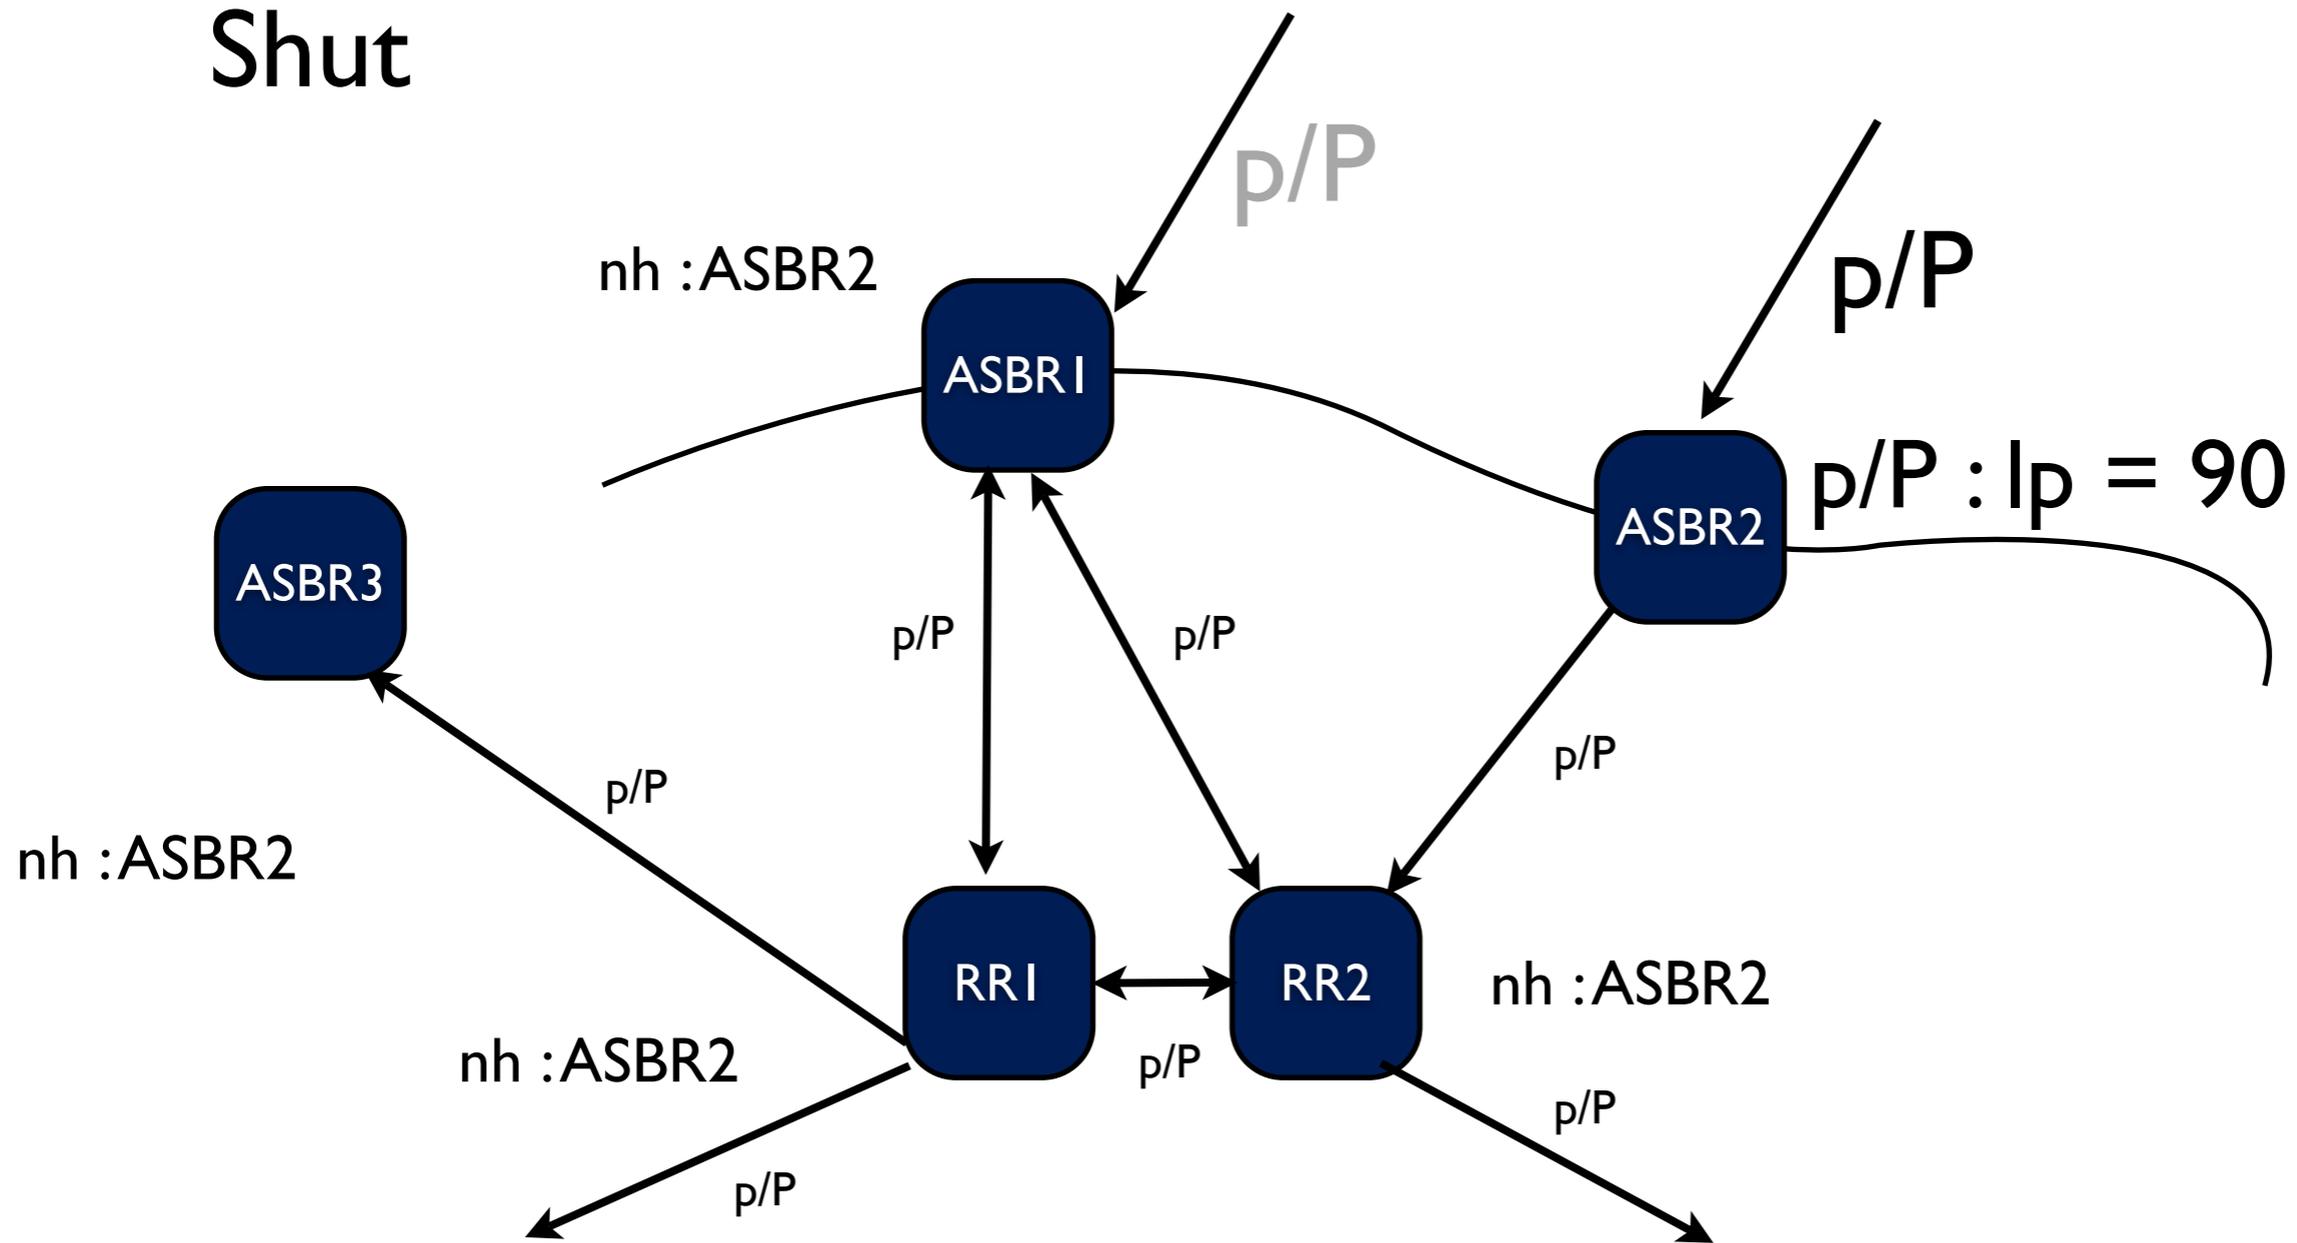


Outbound traffic



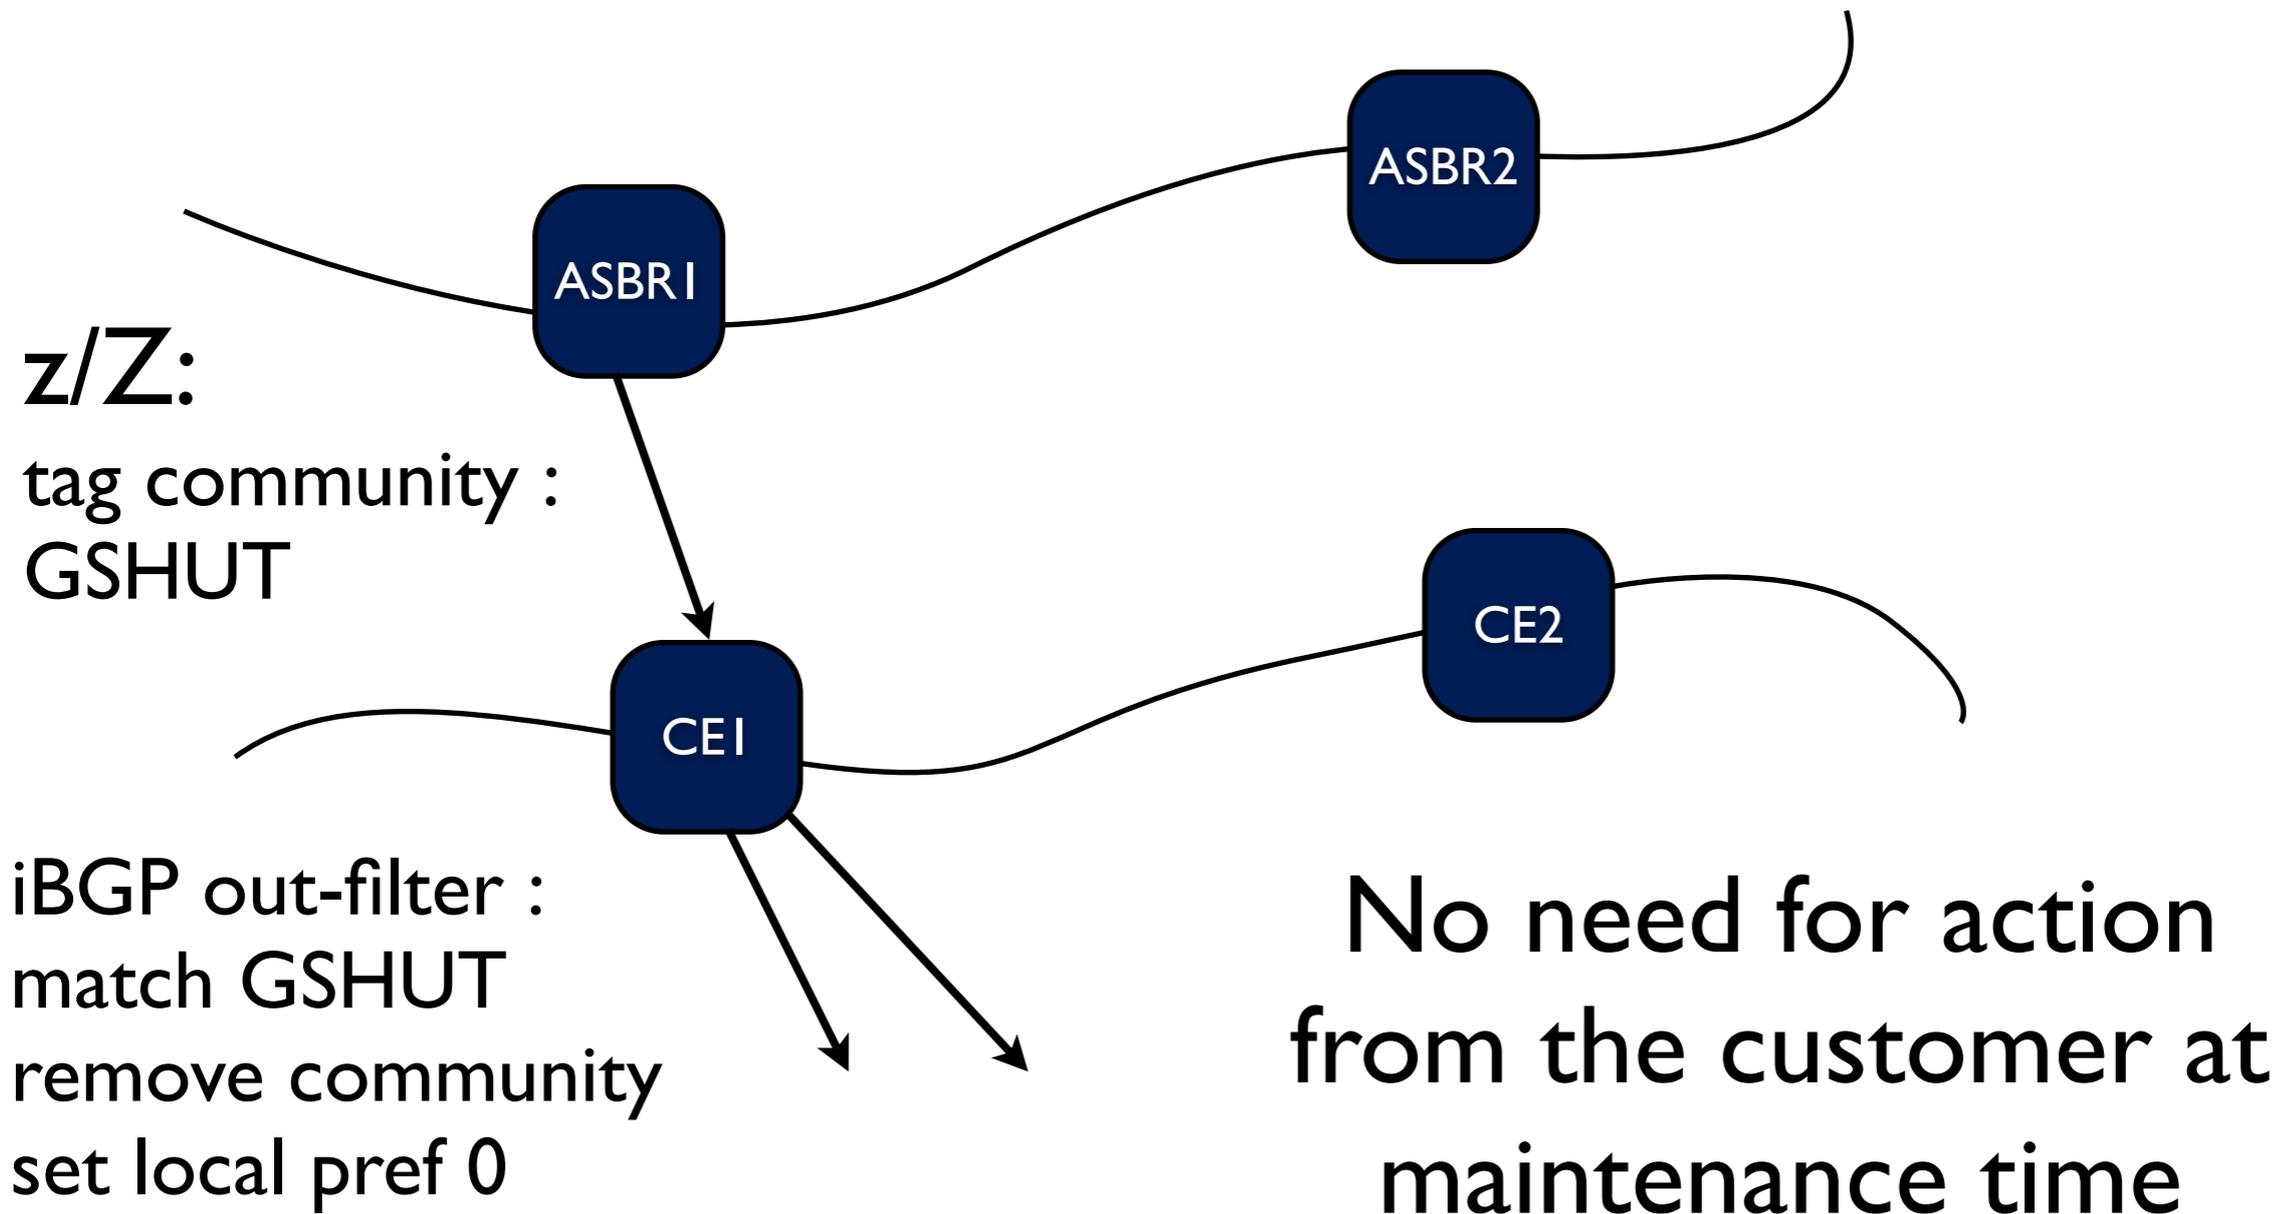
Outbound traffic

Shut



Inbound traffic

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



0xFFFF0000

- One g-shut community per neighbor AS is a pain
 - neighbor-specific config
 - help from vendors less practical
- 0xFFFF0000 value has been reserved for this
 - see <http://www.iana.org/assignments/bgp-well-known-communities/>

G-shut concealment

- Avoid the g-shut community to be propagated through the whole Internet upon maintenance of a link
- Convergence is kept local to the AS
- Prevents misuses of the community

G-shut concealment

- match FFFF0000
remove community
set local pref 0
on reception from eBGP neighbors that do not support gshut or to whom the service is not provided
- match FFFF0000
remove community
set local pref 0
on iBGP out-filters

Extended community ?

- “If the GSHUT community is an extended community, it SHOULD be chosen non-transitive. In that case, some clarification is required about the handling of such communities”
- draft-decraene-idr-rfc4360-clarification-00

What would we get ?

- No need to re-implement community non-transitivity of by configuration on ASBRs
- No default behavior of 4360 can lead to misuses of the solution

Local pref overriding

- If some speakers override LP on iBGP sessions
- Use of an AS specific community over the iBGP sessions