Internet dove gli altri non arrivano.

eolo

The BLU Project

Past, Present, and Future

The “BLU people”

F. Alberti

P. Biasoli

S. Ceccato

M. Citterio

A. Milani

Outline

Internet dove gli altri non arrivano.

1. Challenges

2. Tools

3. Solutions

About EOLO

Internet dove gli altri non arrivano.

- At the beginning (2006):
  - Broadband for rural areas (in northern Italy)
- Today:
  - Coverage: >90% northern Italy, > 50% central Italy
  - 230K customers (4% market share)
  - From 30Mbps to 1Gbps and white-label wholesale
- Each month:
  - 13K new customers

FWA – Anatomy of a Base Transceiver Station

Internet dove gli altri non arrivano.

To the top

- Backhaul (PTP links)
- Client access (PMP links)

To the ground

- Routers, switches
- Energy stations

About EOLO

- Fixed Wireless Access (FWA) ISP
- Core in Milan
- Backhaul connections by means of **wireless** links established between *Base Transceiver Station*

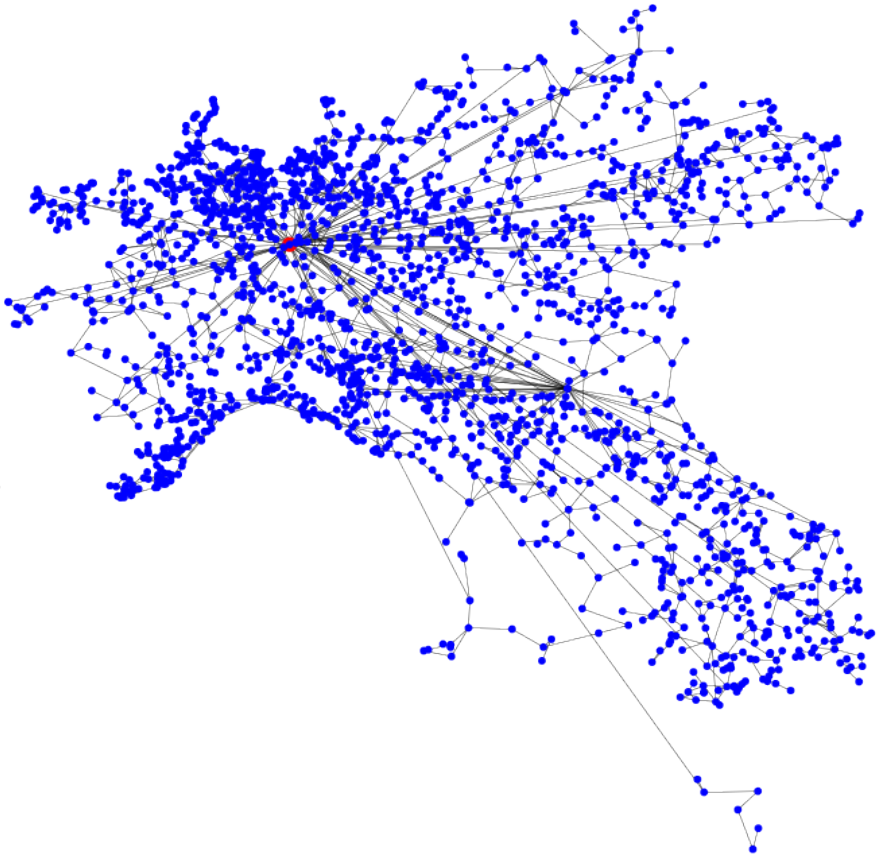

(A very) Distinguishing feature

Opportunistic approach

- New link established whenever two BTSs “see eachothers”

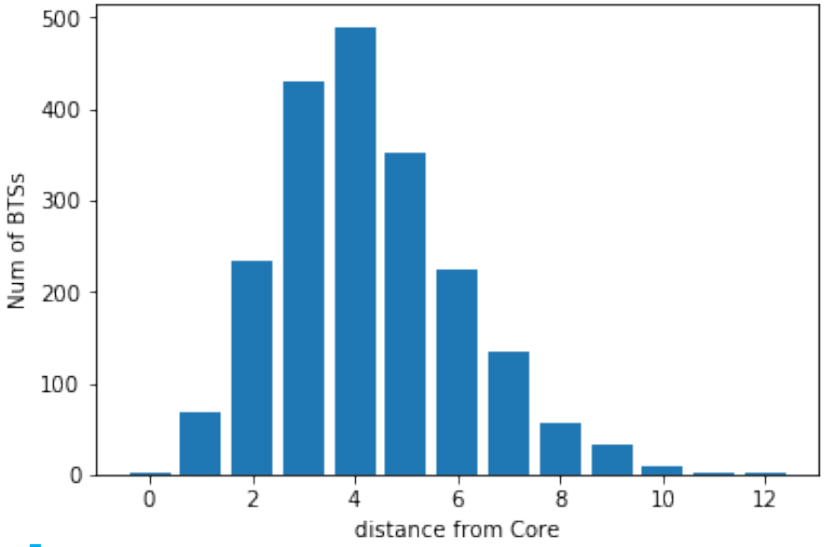

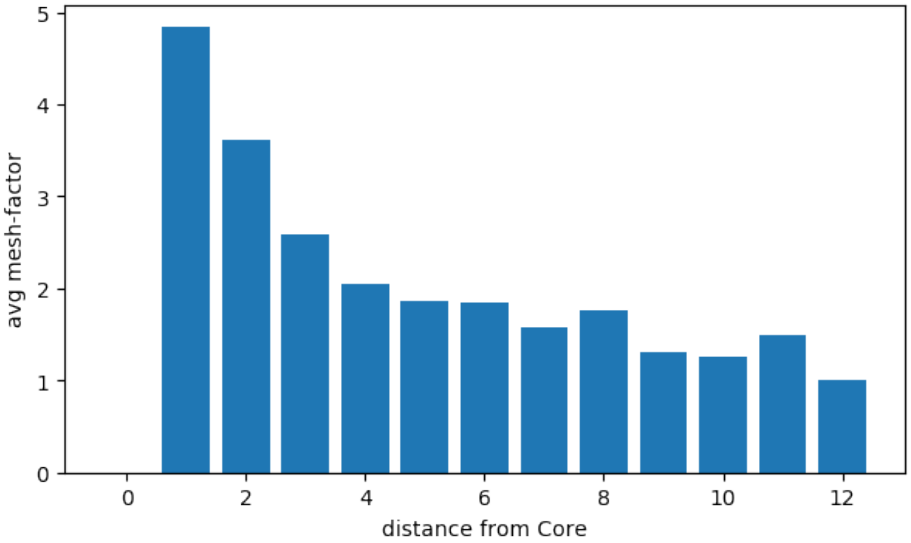

Challenges
Diameter of the network

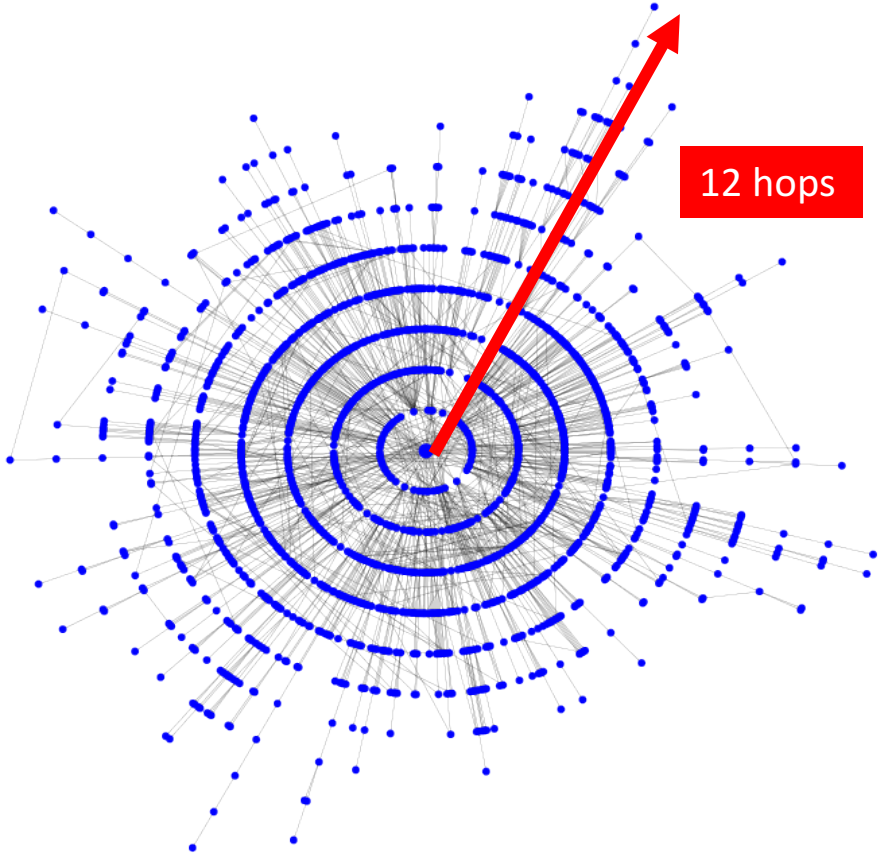

Challenges
Topology changes

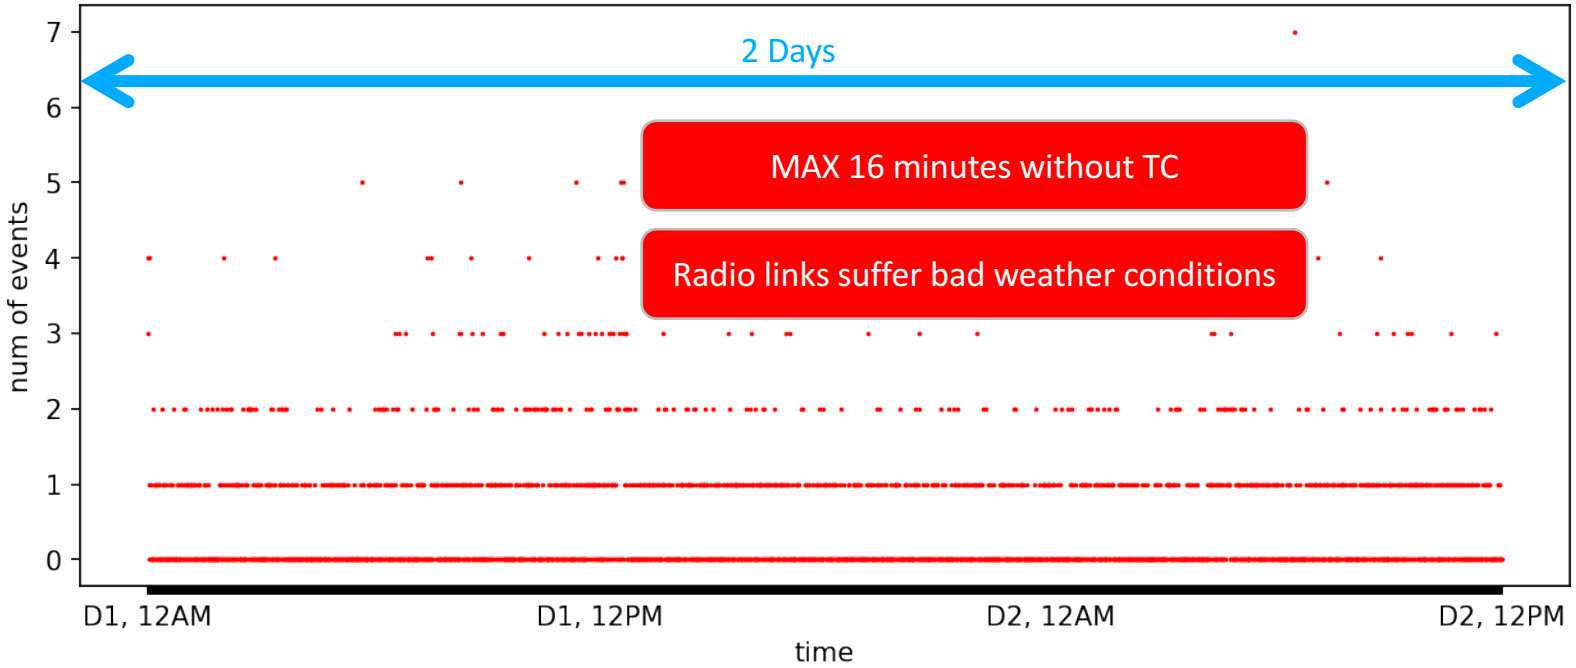

Our solution: the BLU project

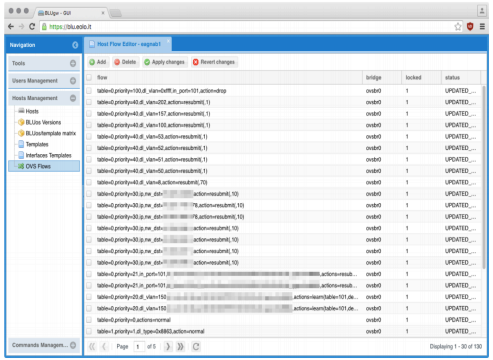

BLU routers

BLUos

BLU-GW

The BLU project
Routers

- TileGX architecture
- 72 core CPU
- 16 GB DDR3 RAM
- 128GB SSD
- Low power: 90W avg, 150W max
- 24x1G ports, 2x10G ports

- IEEE 1613 compliant
  - Electromagnetic compatibility
  - Error-free operation in -15C/+75C environment
  - EMI fields of up to 180V/m

The BLU project
BLUos

Based on 6windgate by 6WIND S.A.

- Linux 3.10.61
- OpenVSwitch
- 6WIND userspace network stack
  - Based on DPDK
  - Associate cores to ports (different configuration based on the routers tasks/position)
  - Userspace dataplane – HIGHLY EFFICIENT
- PPP daemon
- DHCP relay
- Quagga (BGP)

EOLO customizations

- OpenVSwitch mac-learning improved
- RFC 3107 inside Quagga BGP
- Automatic set-up/rescue procedure
  - Self-management strategies
  - Corrupted configurations
  - Host temporarily disconnected
- Monitoring daemons
  - LLDP neighboring details
  - Bandwidth (radio links suffer bad weather, interferences, etc.)
  - Tx/Rx traffic on each interface (both BH and Access)

The BLU project
BLU-GW

- Avoid (or, at least, limit) human error
- Full management of our BLU routers
  - Parallel execution of commands
  - Scheduler for night updater
  - Success/failure/rollback
- Collection and analysis of data from the BLU routers
  - Topology updates
  - Statistics, logs and error messages
- Statistics visualization

The BLU project
BLU-GW - Management

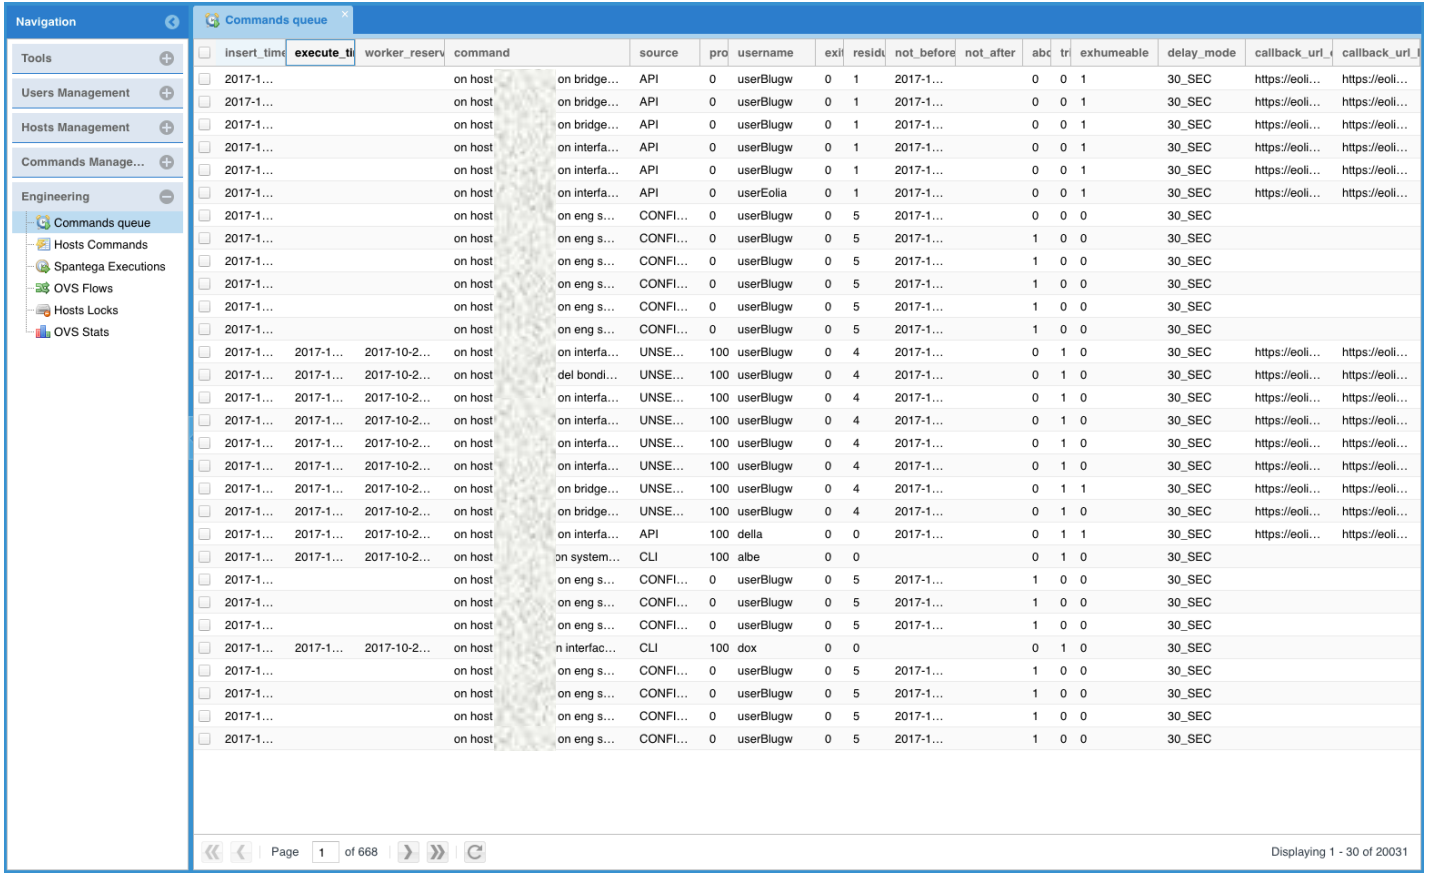

The BLU project
BLU-GW - Management

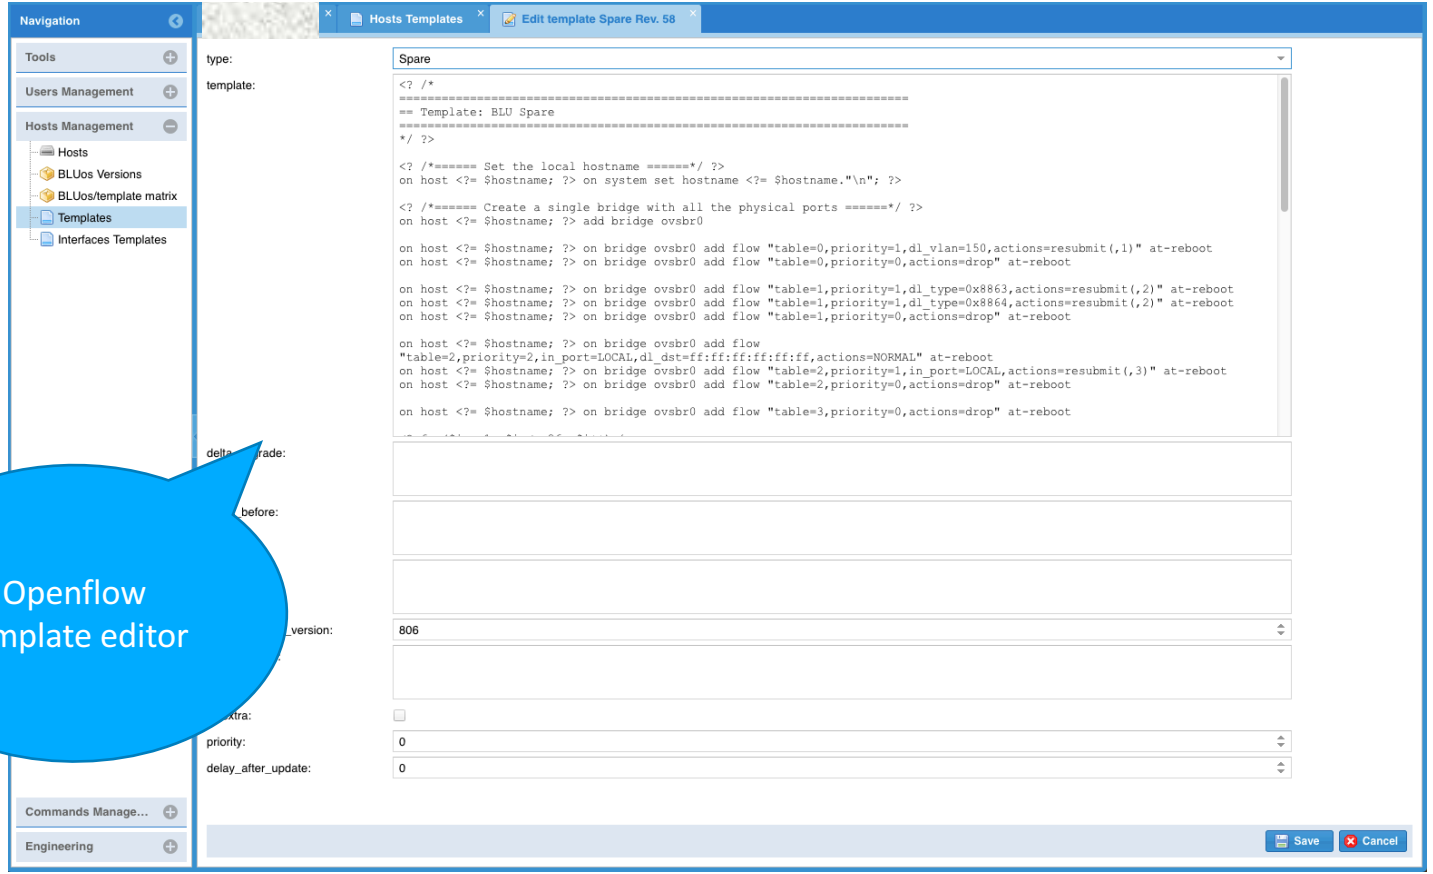

Openflow template editor

The BLU project
BLU-GW - Monitoring

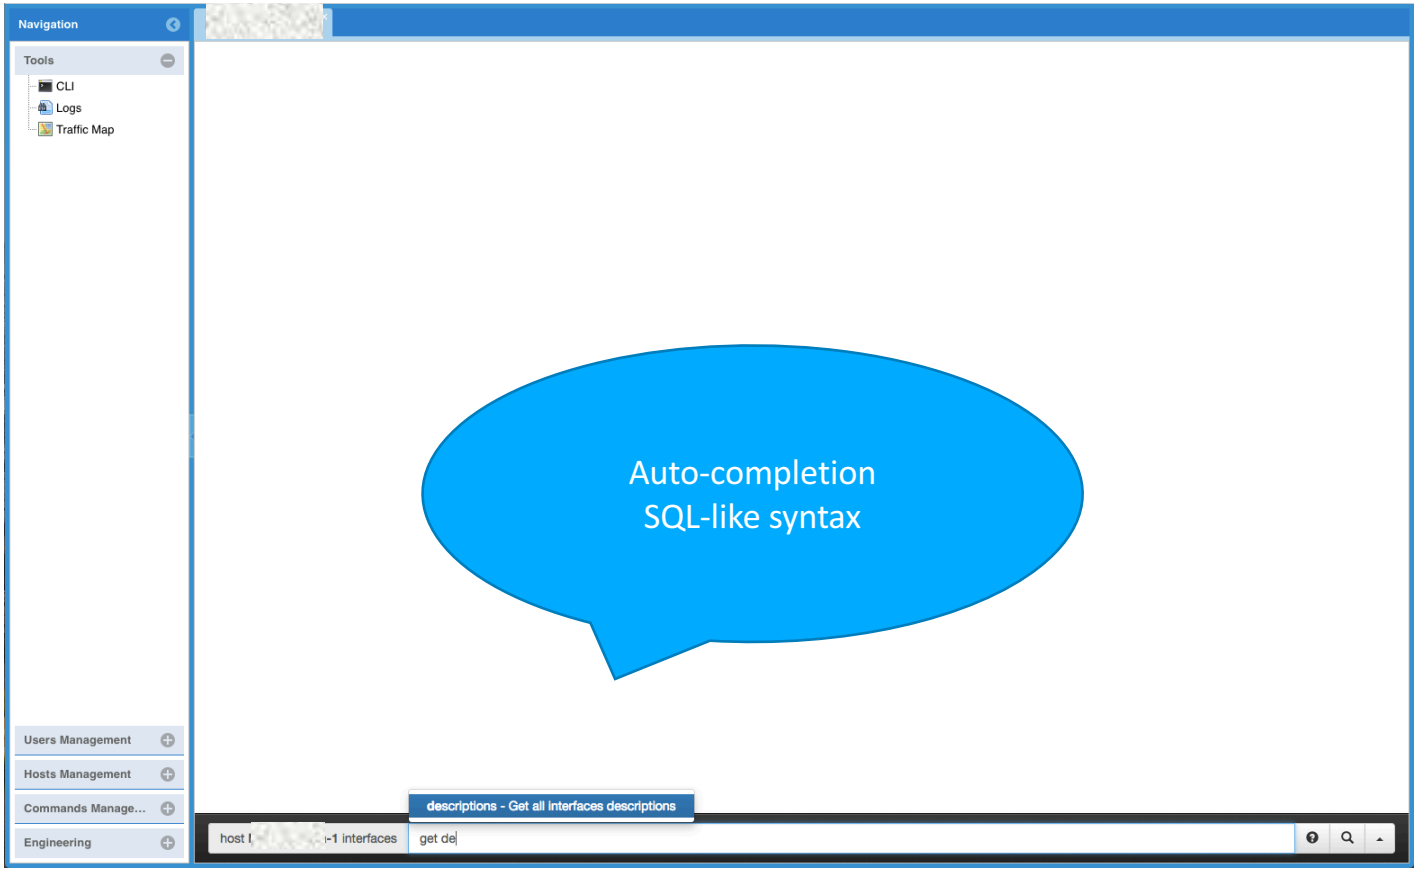

The BLU project
BLU-GW - Monitoring

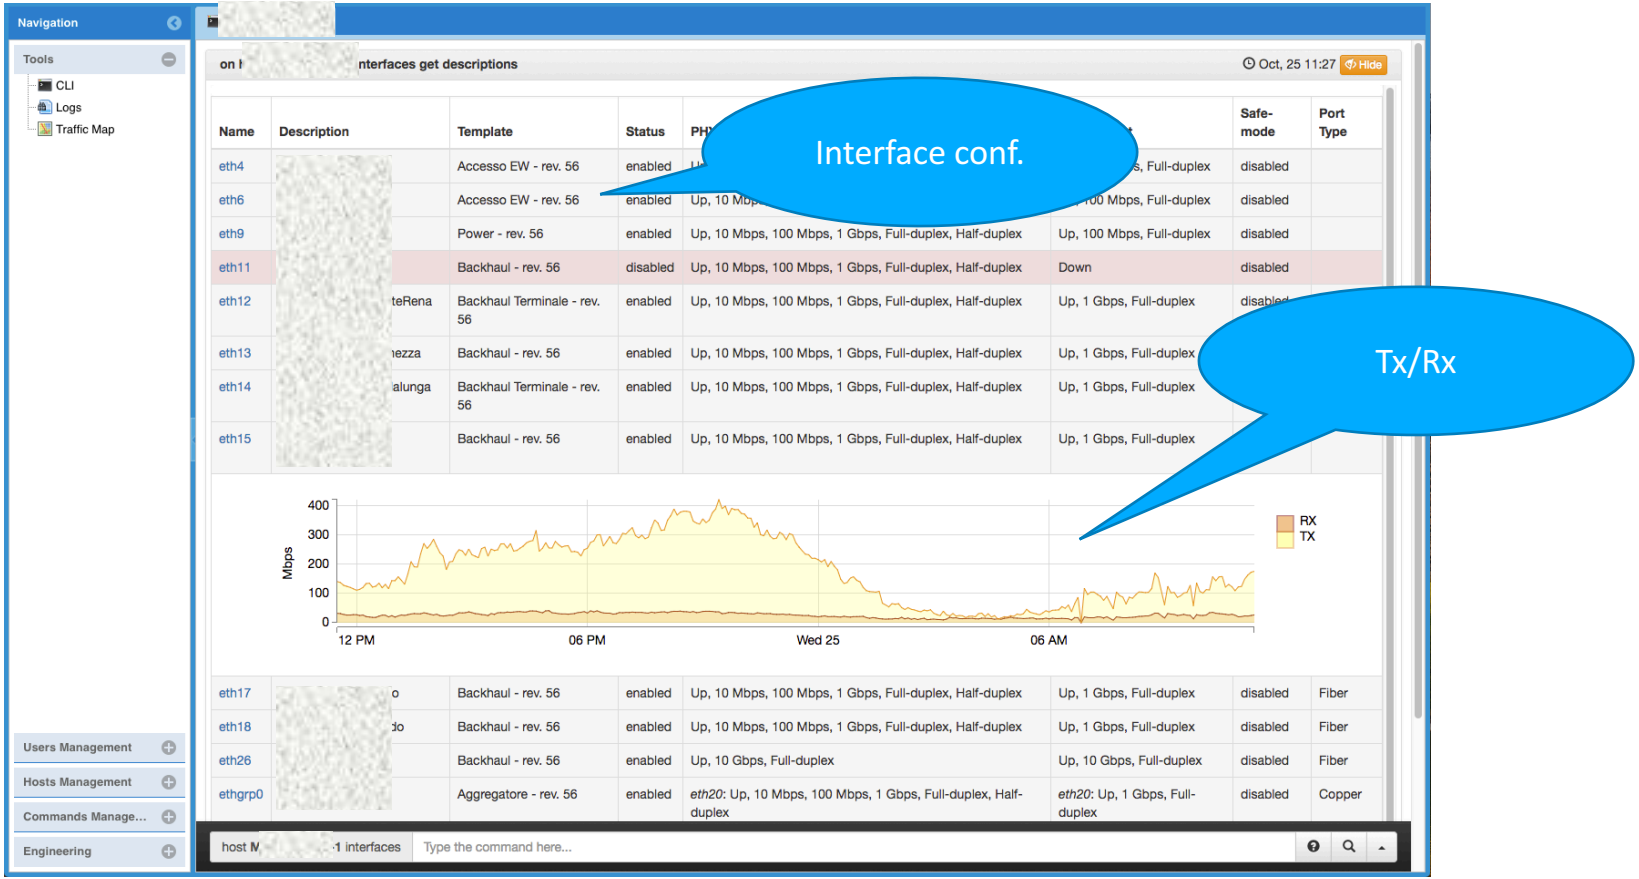

The BLU project
BLU-GW - Monitoring

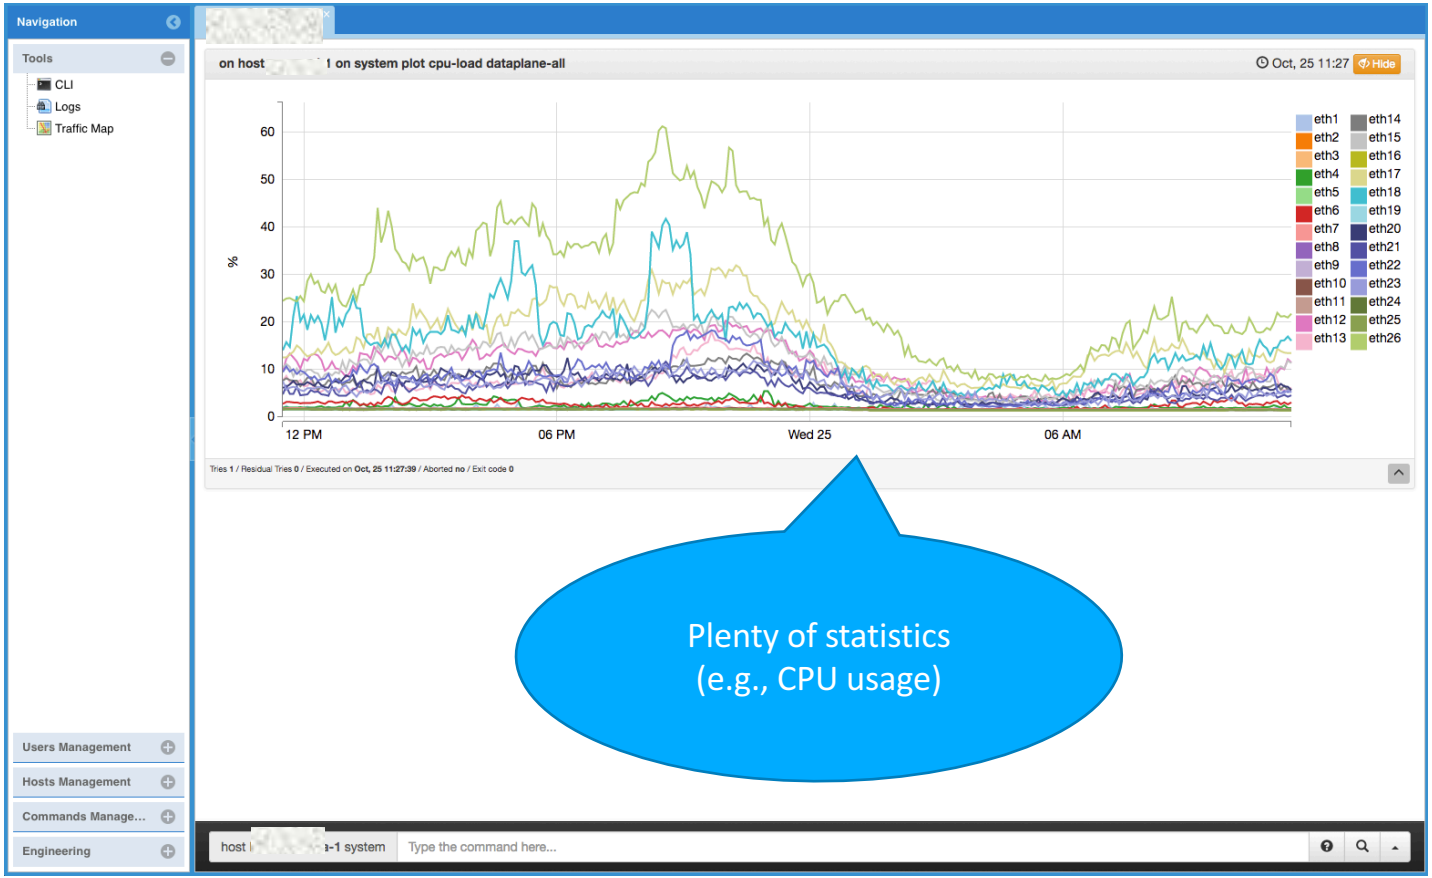

The BLU project
BLU-GW - Monitoring

The BLU project
Phase1 - Challenges

- Backhaul:
  - Wireless links
  - PVST+ L2
- Very high number of mac-address!
- PPP? DHCP?
- Traffic balancing? Overloaded links?

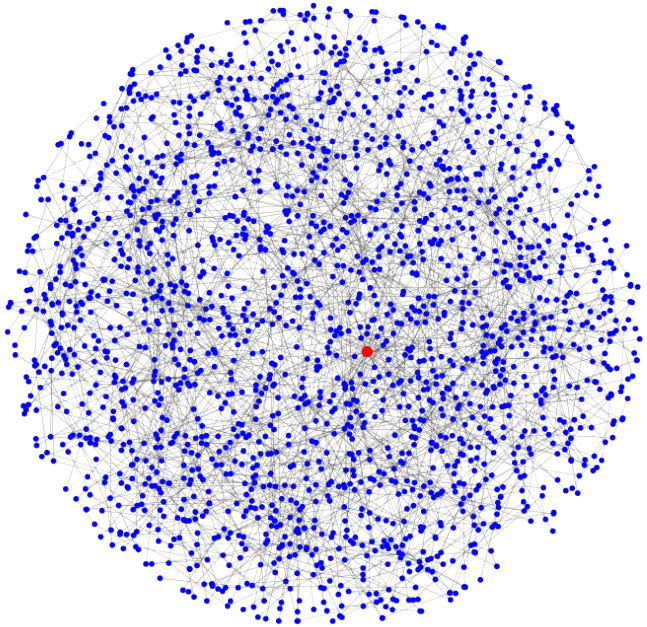

The BLU project
Phase1 - Solution

- Backhaul:
  - Wireless links
  - PVST+ L2
- Very high number of mac-address!
  - PPP server/DHCP relay on the BLU
  - Each BLU advertises to the core its PPP/DHCP IP pools
- Traffic balancing? Overloaded links?
  - Multiple BH vlans (cost changes determine different paths)

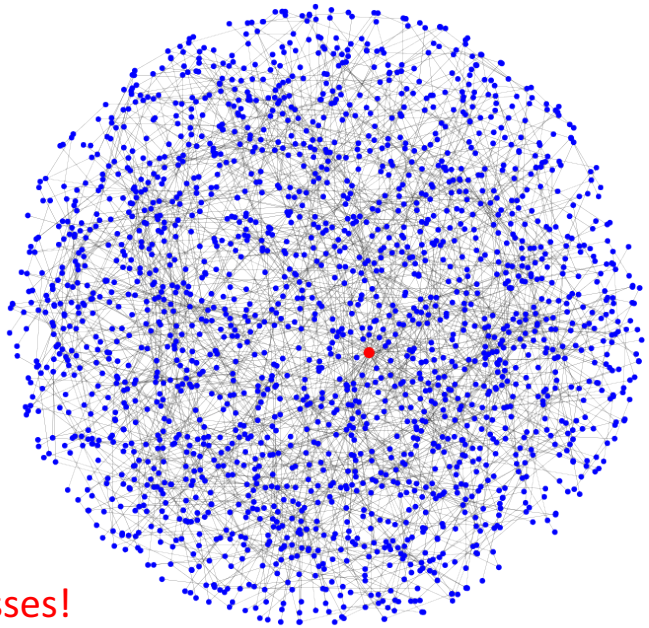

- Cost changes performed manually!
- OVS still suffers the high number (> 1K per BLU) of mac-addresses!
  - Flows-handling performance decreases – traffic gets lost

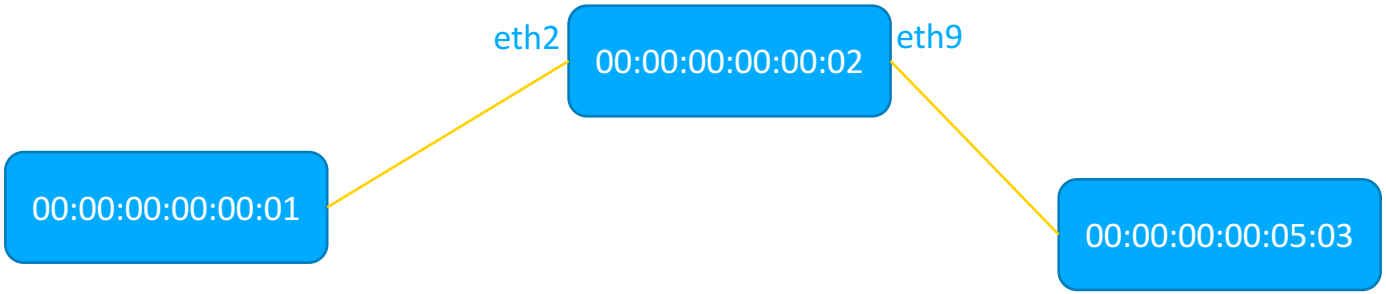


```
in_port(2),eth(src=00:00:00:00:00:01,dst=00:00:00:00:05:03),…,actions:9
in_port(9),eth(src=00:00:00:00:05:03,dst=00:00:00:00:00:01),…,actions:2
```

The BLU project
BLUos – Hacking OVS


```
in_port(2),eth(src=00:00:00:00:00:01,dst=00:00:00:00:05:03),…,actions:9
in_port(2),eth(src=00:00:00:00:01:01,dst=00:00:00:00:05:03),…,actions:9
in_port(2),eth(src=00:00:00:00:02:01,dst=00:00:00:00:05:03),…,actions:9
in_port(2),eth(src=00:00:00:00:03:01,dst=00:00:00:00:05:03),…,actions:9

in_port(9),eth(src=00:00:00:00:05:03,dst=00:00:00:00:00:01),…,actions:2
in_port(9),eth(src=00:00:00:00:05:03,dst=00:00:00:00:01:01),…,actions:2
in_port(9),eth(src=00:00:00:00:05:03,dst=00:00:00:00:02:01),…,actions:2

in_port(9),eth(src=00:00:00:00:05:03,dst=ff:ff:ff:ff:ff:ff),…,actions:2,3,4,5
in_port(9),eth(src=00:00:00:00:01:03,dst=ff:ff:ff:ff:ff:ff),…,actions:2,3,4,5
```

The BLU project
BLUos – Hacking OVS

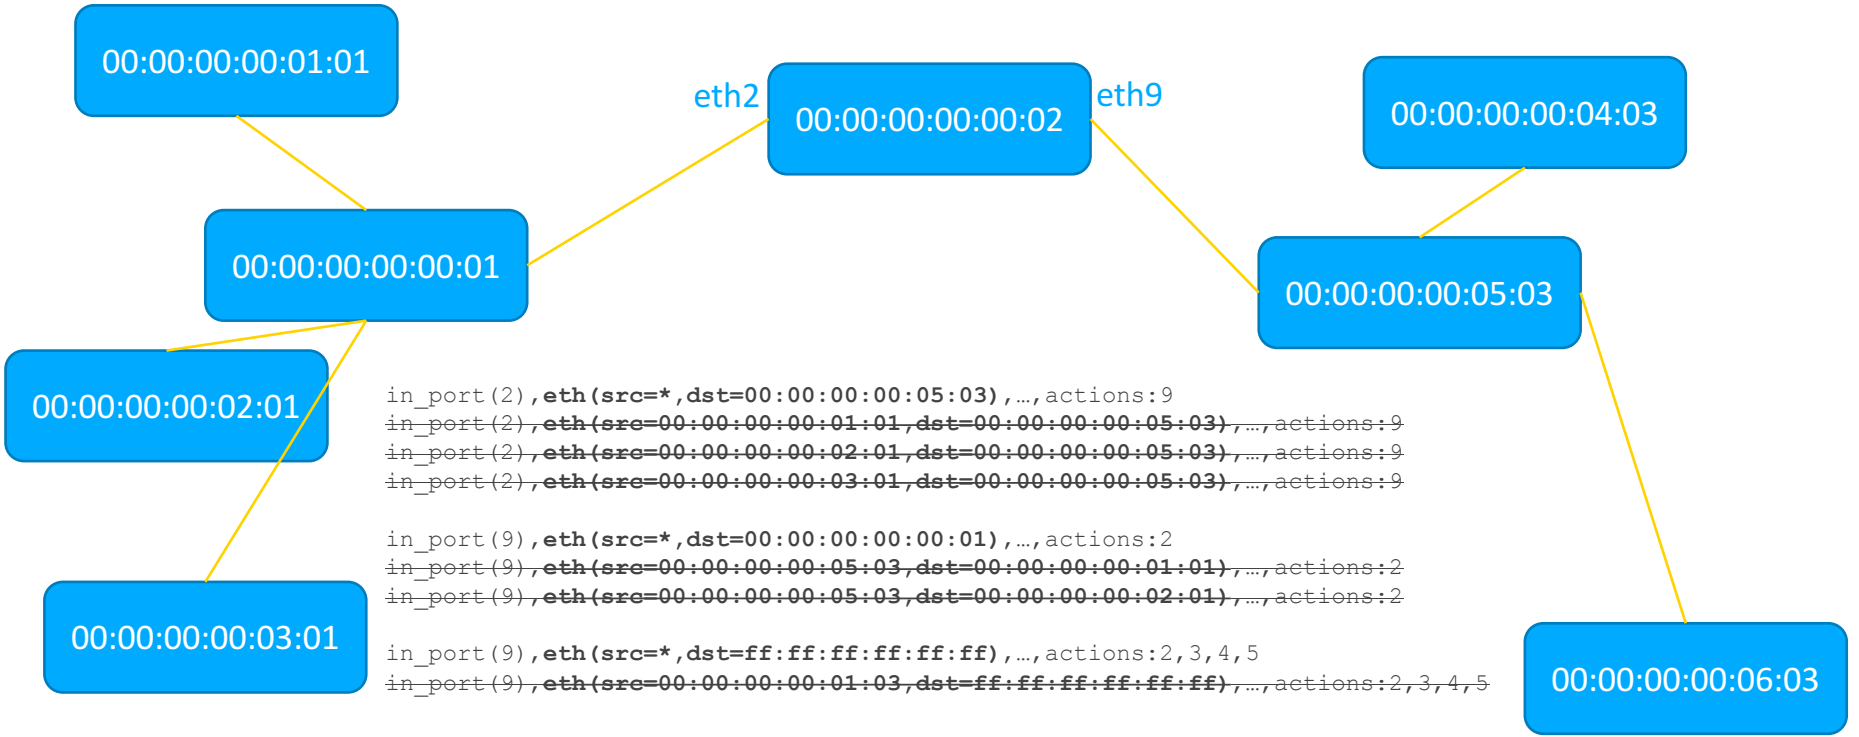


```
in_port(2),eth(src=*,dst=00:00:00:00:05:03),…,actions:9
~~in_port(2),eth(src=00:00:00:00:01:01,dst=00:00:00:00:05:03),…,actions:9~~
~~in_port(2),eth(src=00:00:00:00:02:01,dst=00:00:00:00:05:03),…,actions:9~~
~~in_port(2),eth(src=00:00:00:00:03:01,dst=00:00:00:00:05:03),…,actions:9~~

in_port(9),eth(src=*,dst=00:00:00:00:00:01),…,actions:2
~~in_port(9),eth(src=00:00:00:00:05:03,dst=00:00:00:00:01:01),…,actions:2~~
~~in_port(9),eth(src=00:00:00:00:05:03,dst=00:00:00:00:02:01),…,actions:2~~

in_port(9),eth(src=*,dst=ff:ff:ff:ff:ff:ff),…,actions:2,3,4,5
~~in_port(9),eth(src=00:00:00:00:01:03,dst=ff:ff:ff:ff:ff:ff),…,actions:2,3,4,5~~
```

The BLU project
BLUos – Hacking OVS

Internet dove gli altri non arrivano.

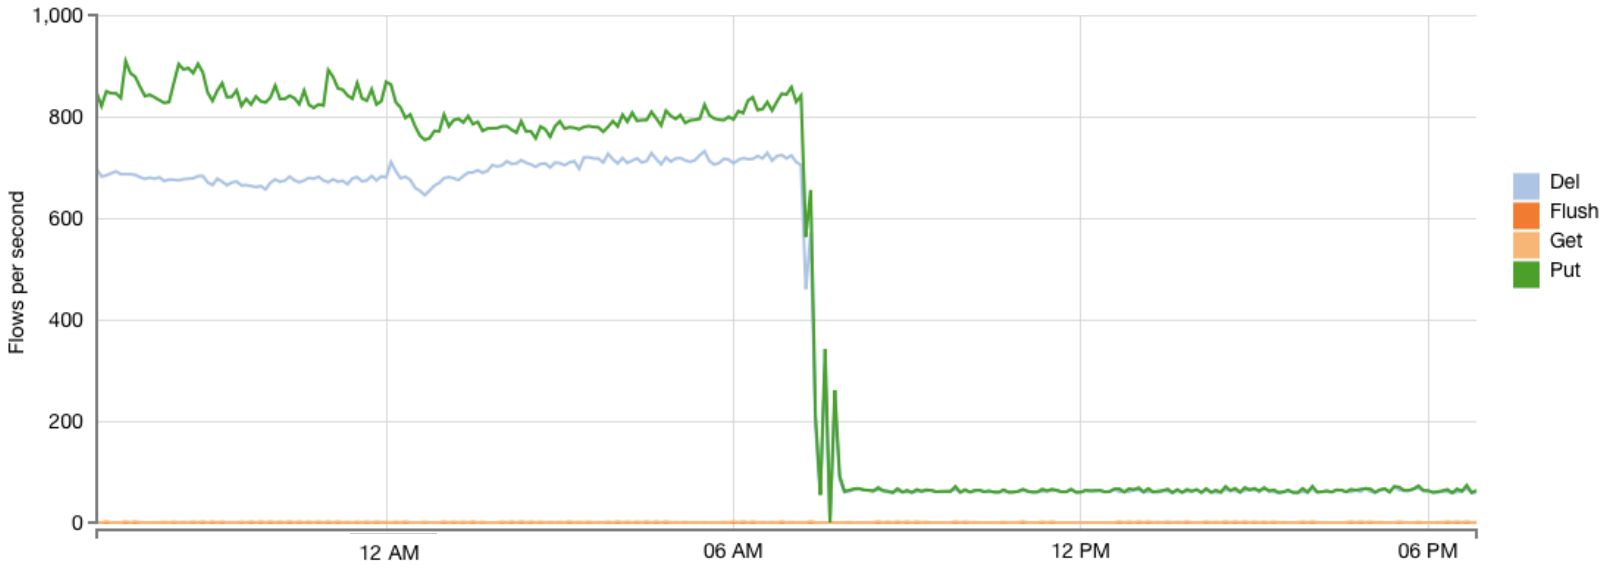

The BLU project
Phase2

- BLUGW knows
  - Link capacity
  - Residual BH link capacity
  - BLUs' "desired bandwidth" (access ifaces bandwidth)
- A centralized SDN controller decides the routing
  - Residual capacity gets optimized
  - Informs the BLU on how to handle traffic
- Openflow rules for MPLS label switching
- RFC3107 for (downlink) traffic labeling

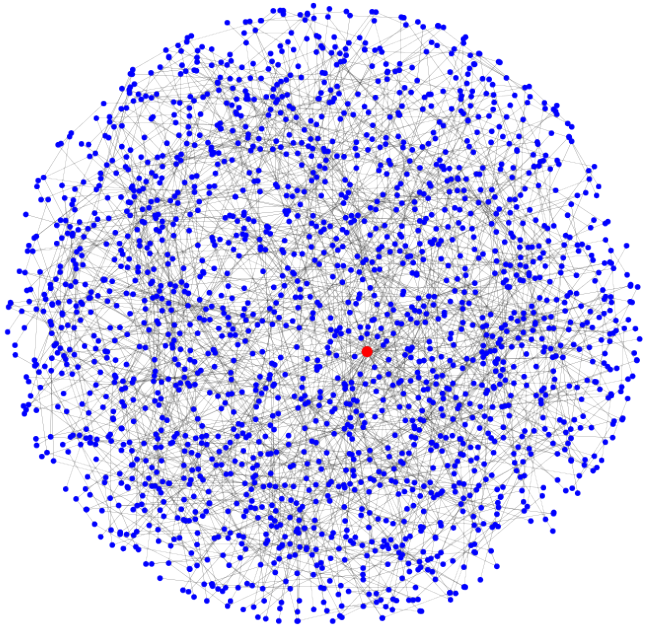

The BLU project
Phase2

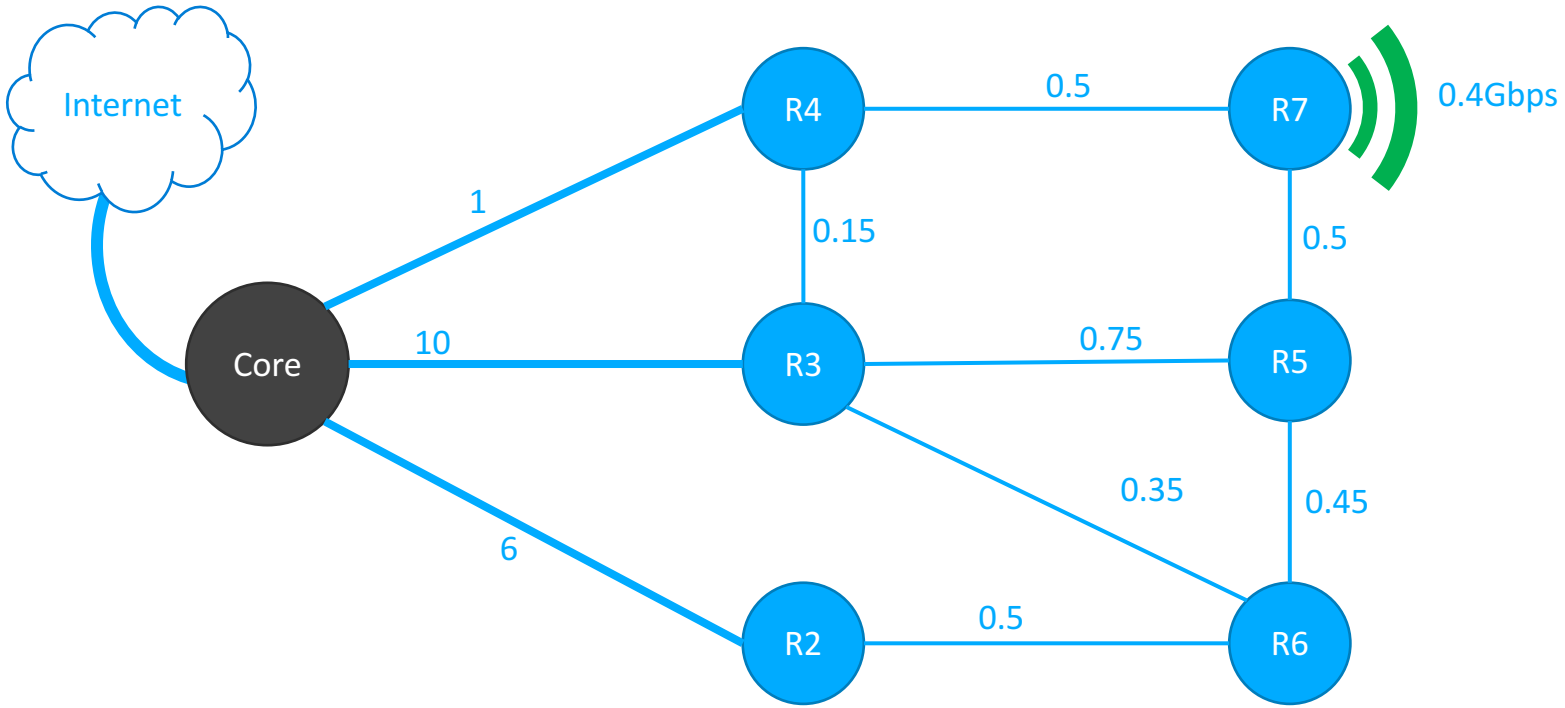

The BLU project
Phase2

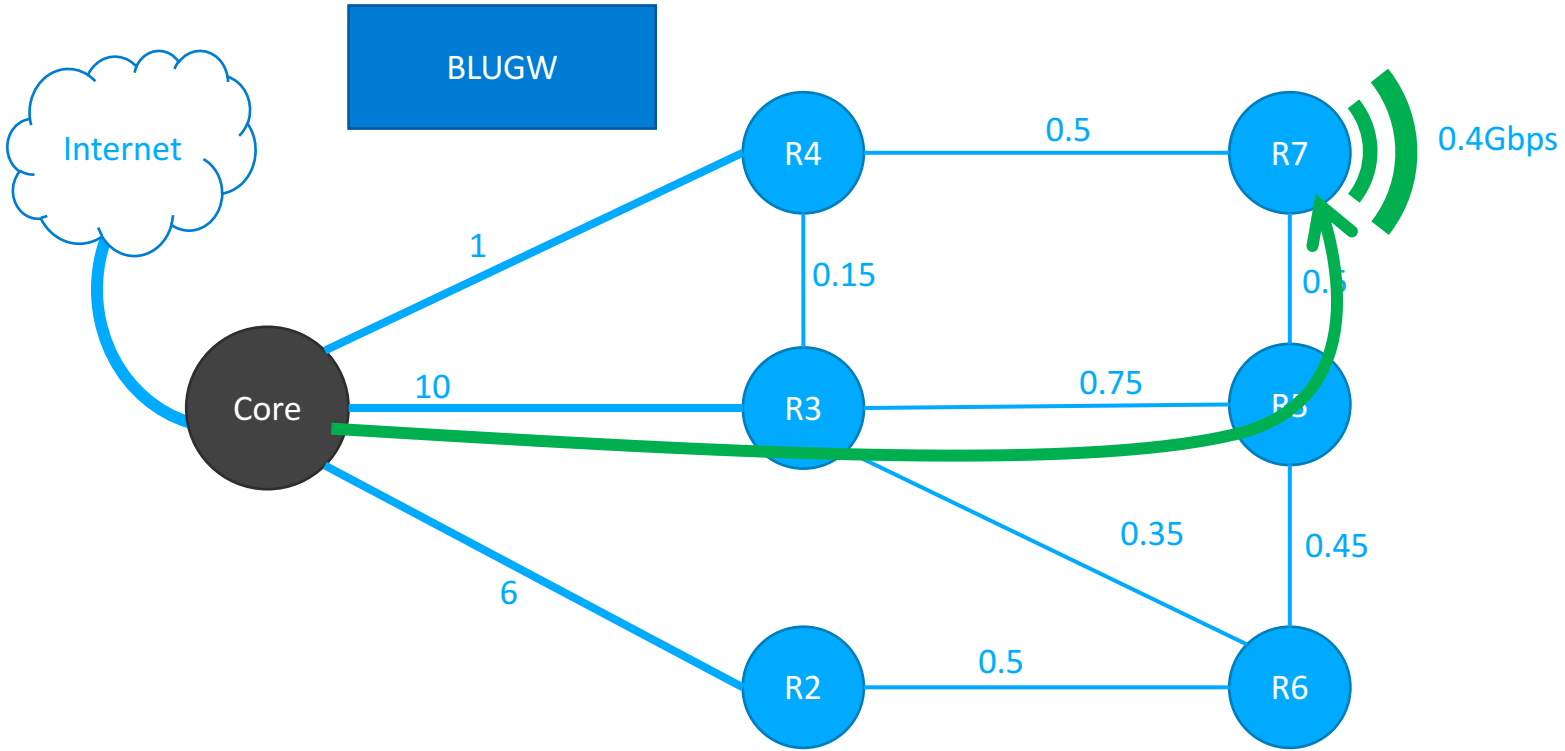

The BLU project
Phase2

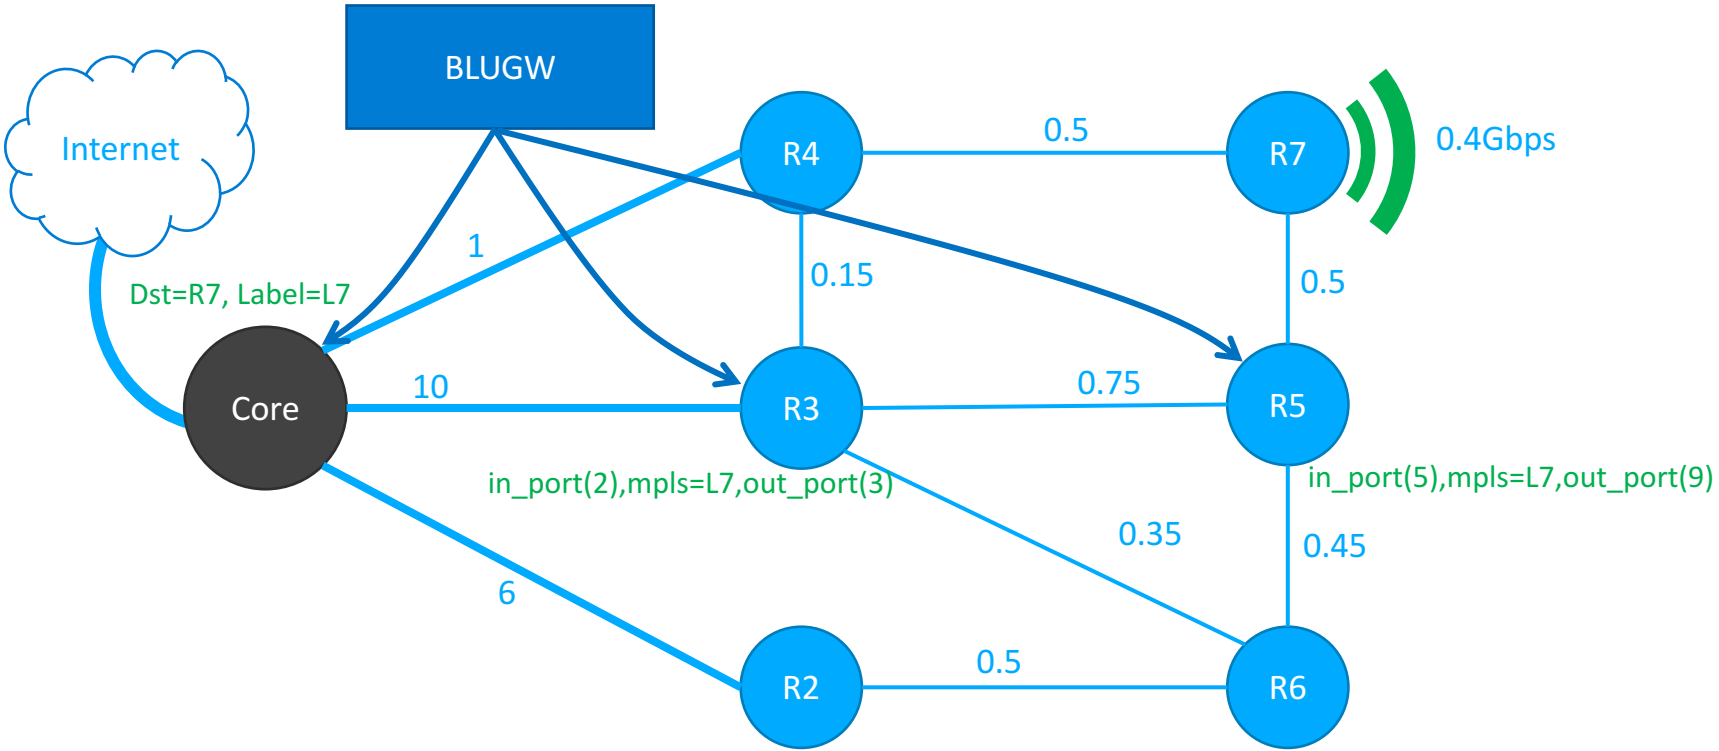

The BLU project
Phase2

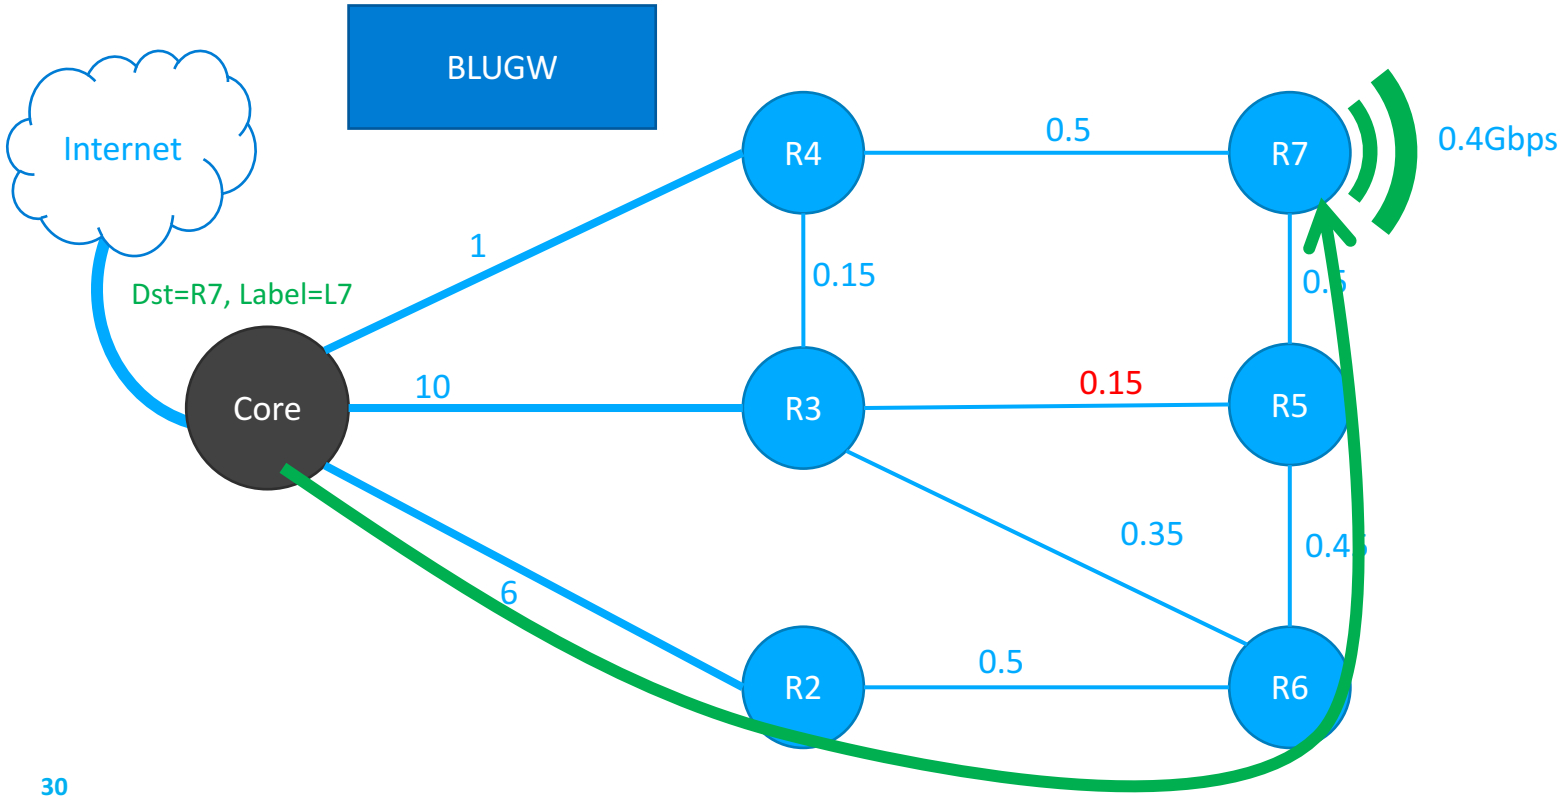

The BLU project
Phase2

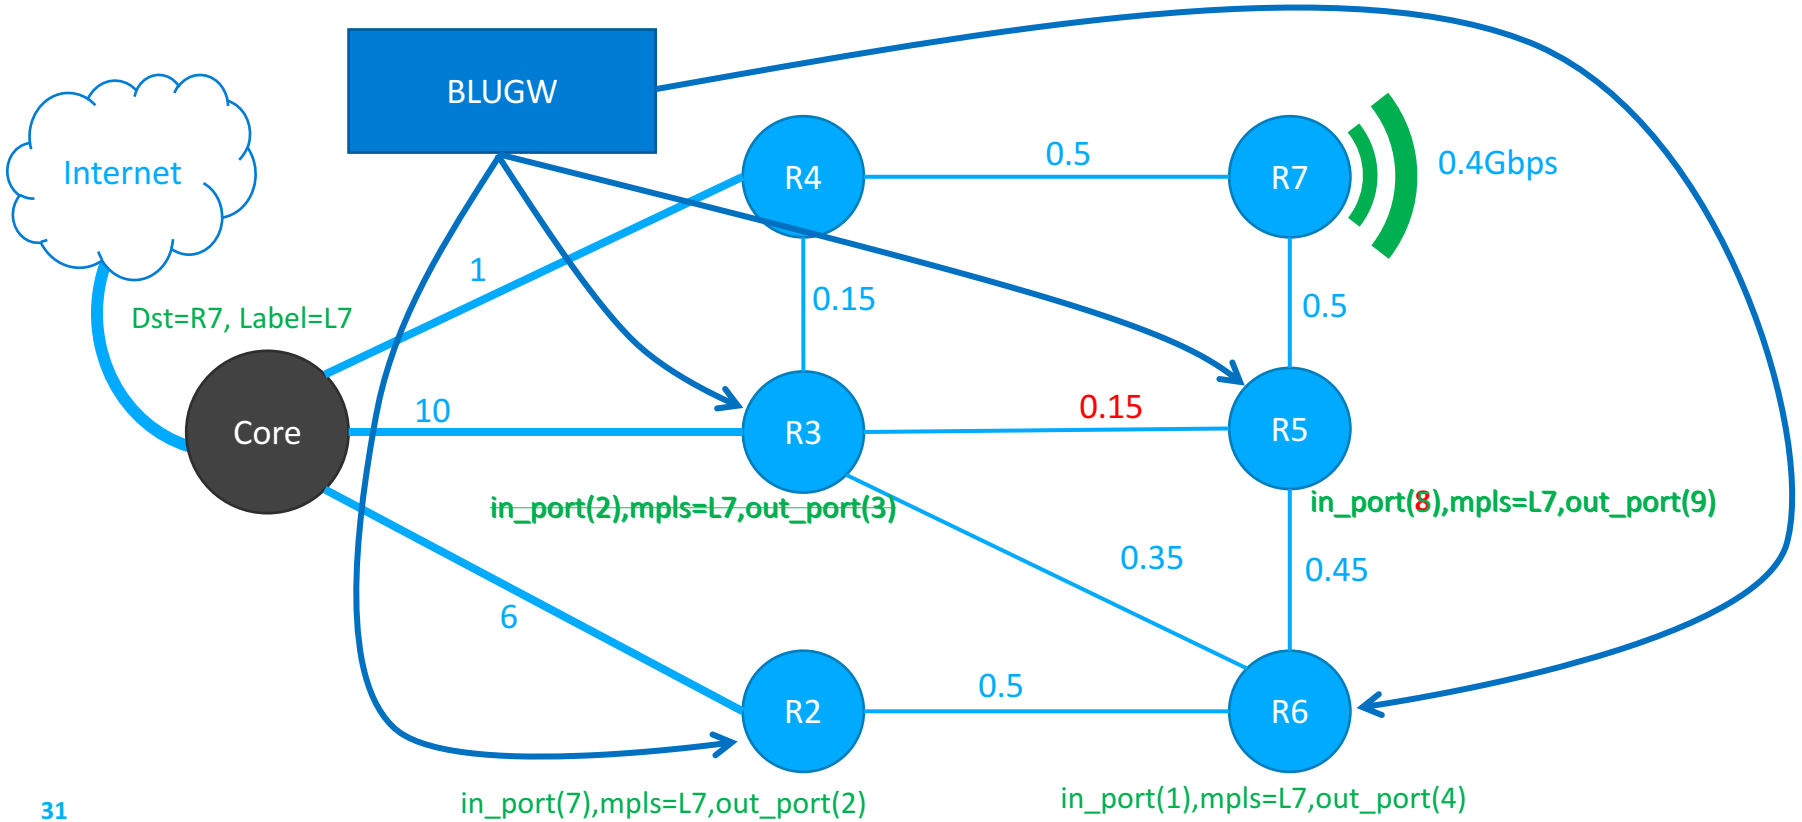

The BLU project
Phase2 - Issues

- OpenFlow granularity is rather small
  - A single flow through the network requires all the BLUs along the path to be (re)programmed
- Every time the network changes (!!!) the controller needs to re-compute the routes and upgrade the BLU
- Upgrades **need to be atomic**, high risk of inconsistent state of the network
- There is no a "BIG RED BUTTON" to be pressed in case of critical issues!
  - How to go back to a safe configuration in case of troubles?

The BLU project
Phase3

- Segment Routing philosophy
- IGP running on the BLU routers
  - eBGP, RFC3107
- RFC 3107 + SR policy for downlink traffic labeling
- Label switching (OVS)
  - eBGP governs Open vSwitch: active routes determine label switching flows

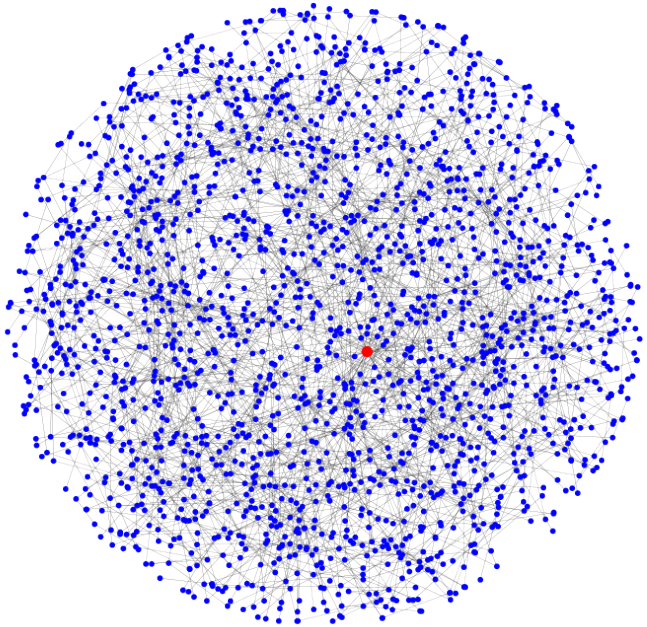

The BLU project
Phase3

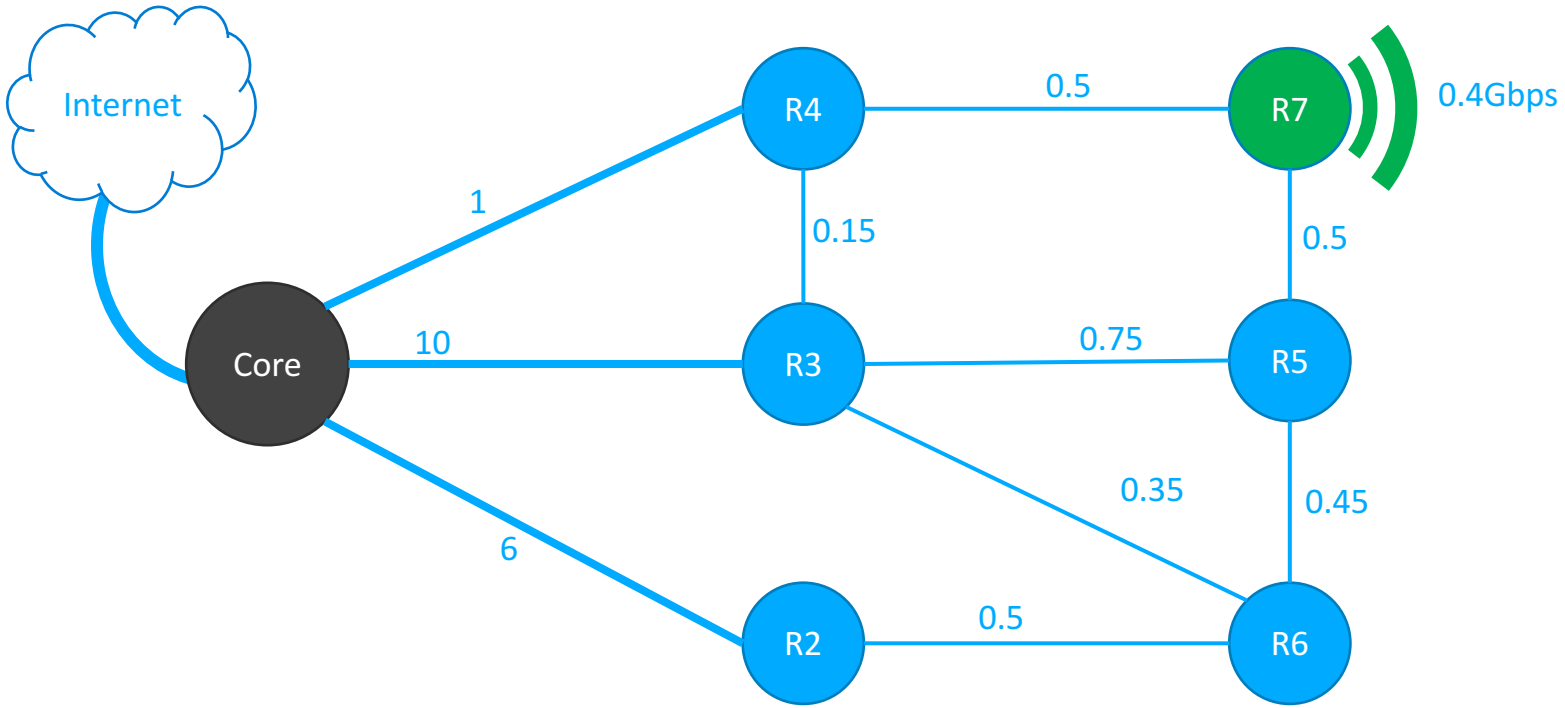

The BLU project
Phase3

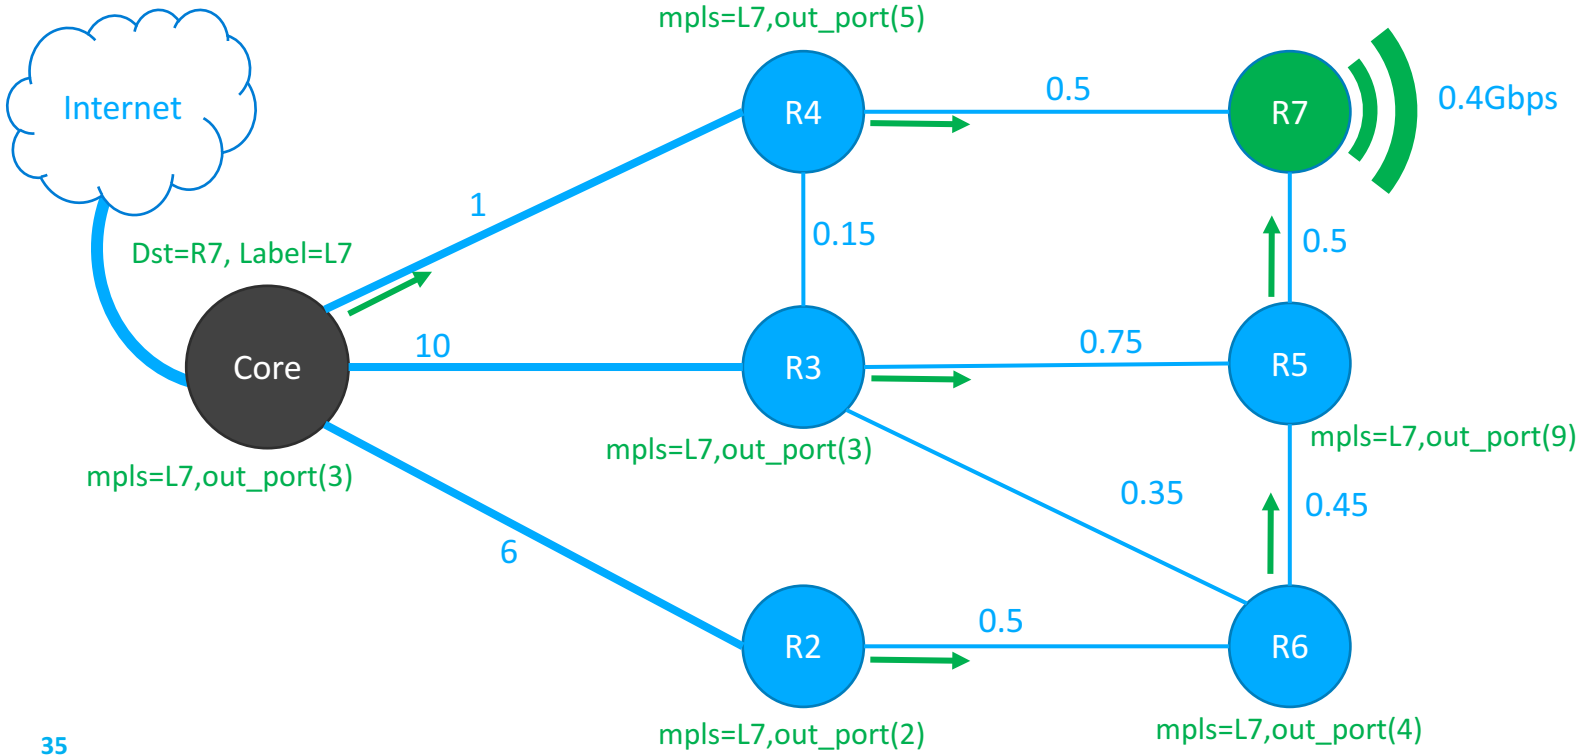

The BLU project
Phase3

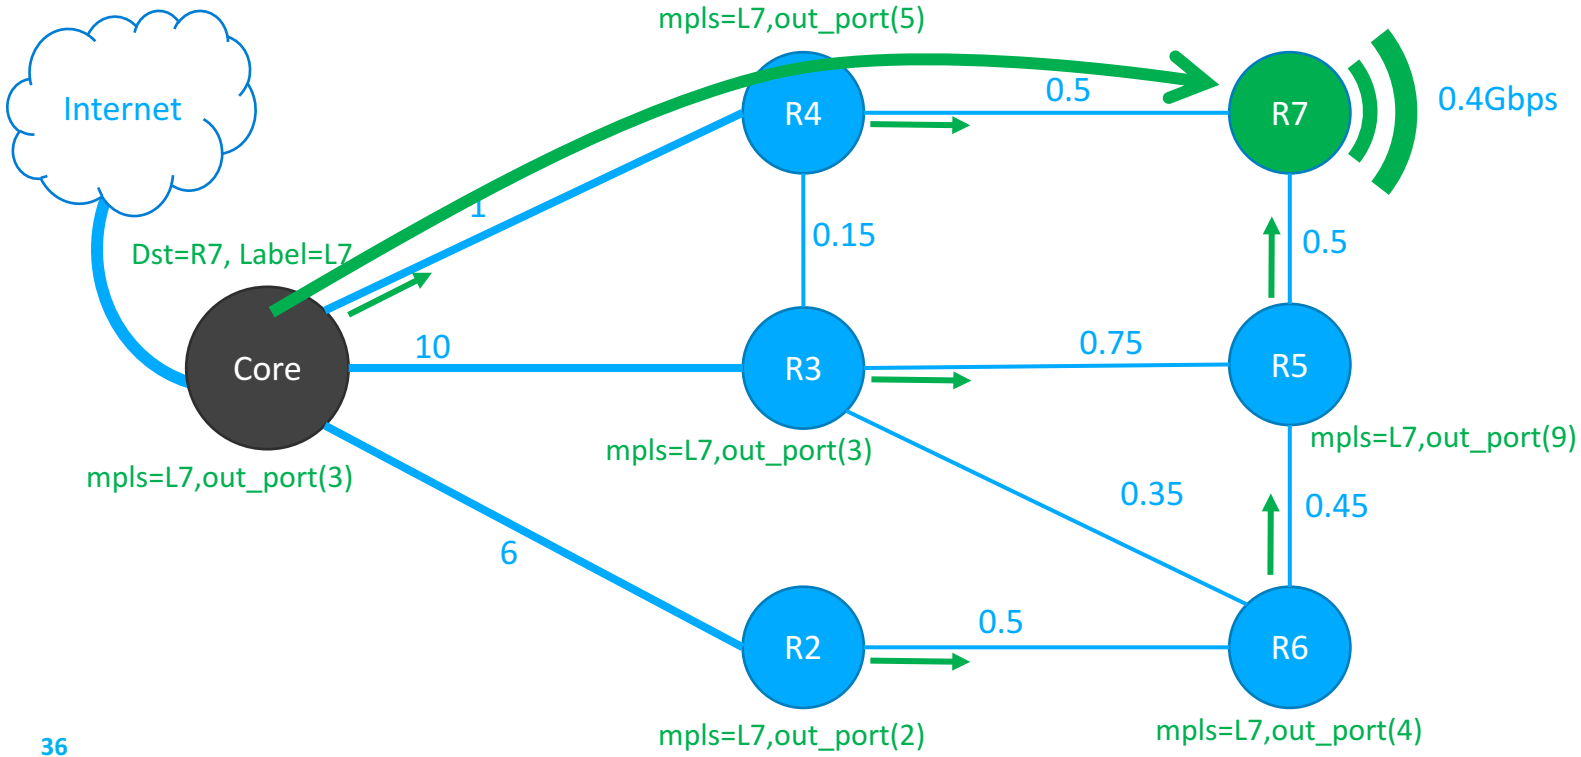

The BLU project
Phase3 – Traffic Engineering

Internet dove gli altri non arrivano.

- How to optimize the traffic flows?

- BLUGW still knows
  - (quasi-)Real time topology
  - Residual bandwidth
  - Desired access ifaces bandwidth
  - **NEW: BLUs' IGP routing tables**

- Optimization goes through these points:
  - Find a waypoint where to route the traffic to be optimized
  - Instruct the CORE on the required label stack

The BLU project
Phase3

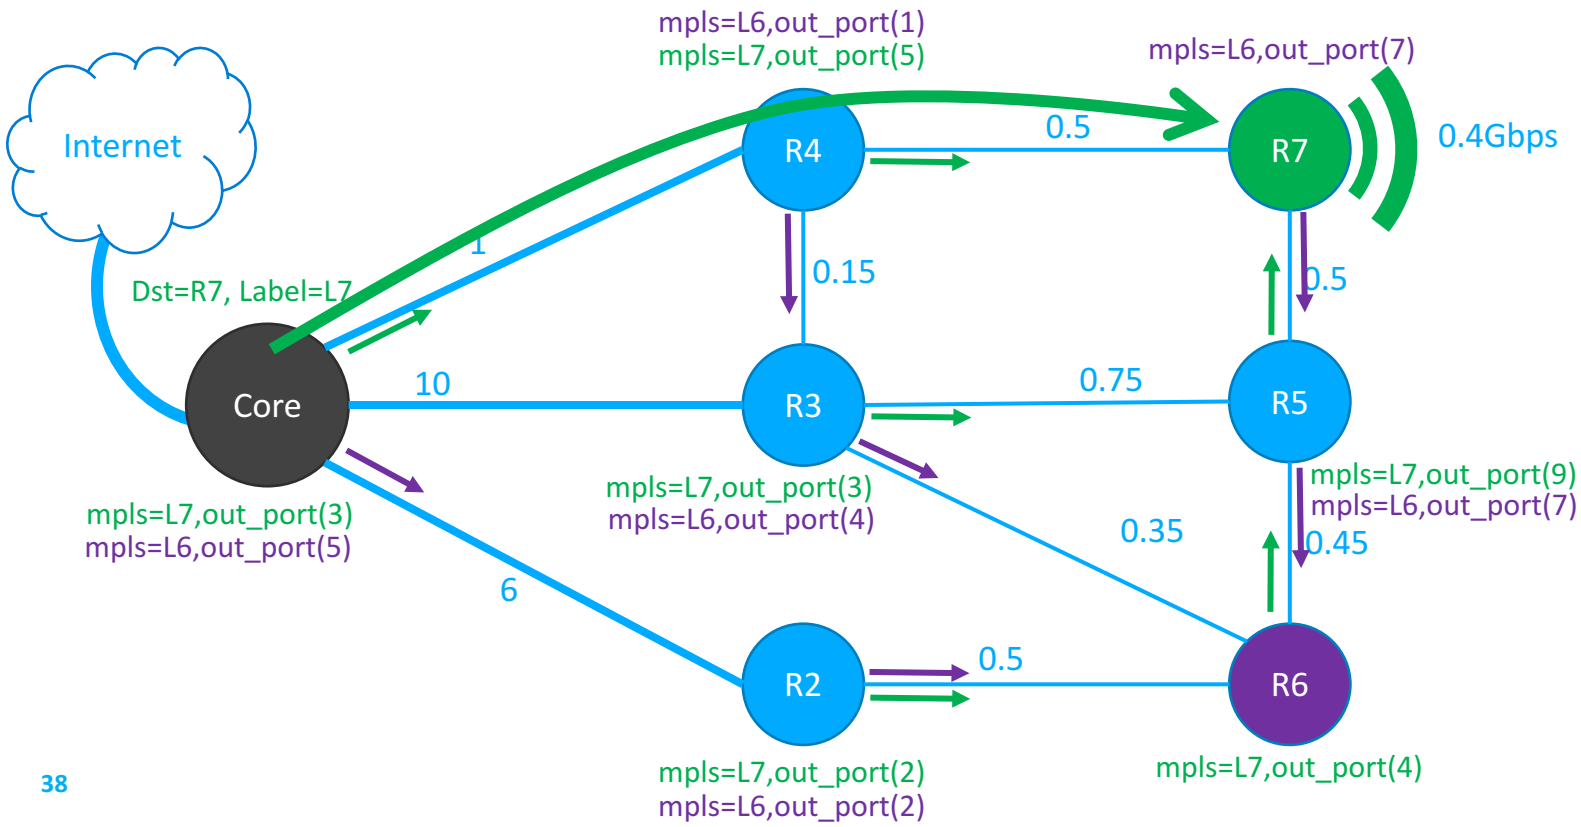

The BLU project
Phase3

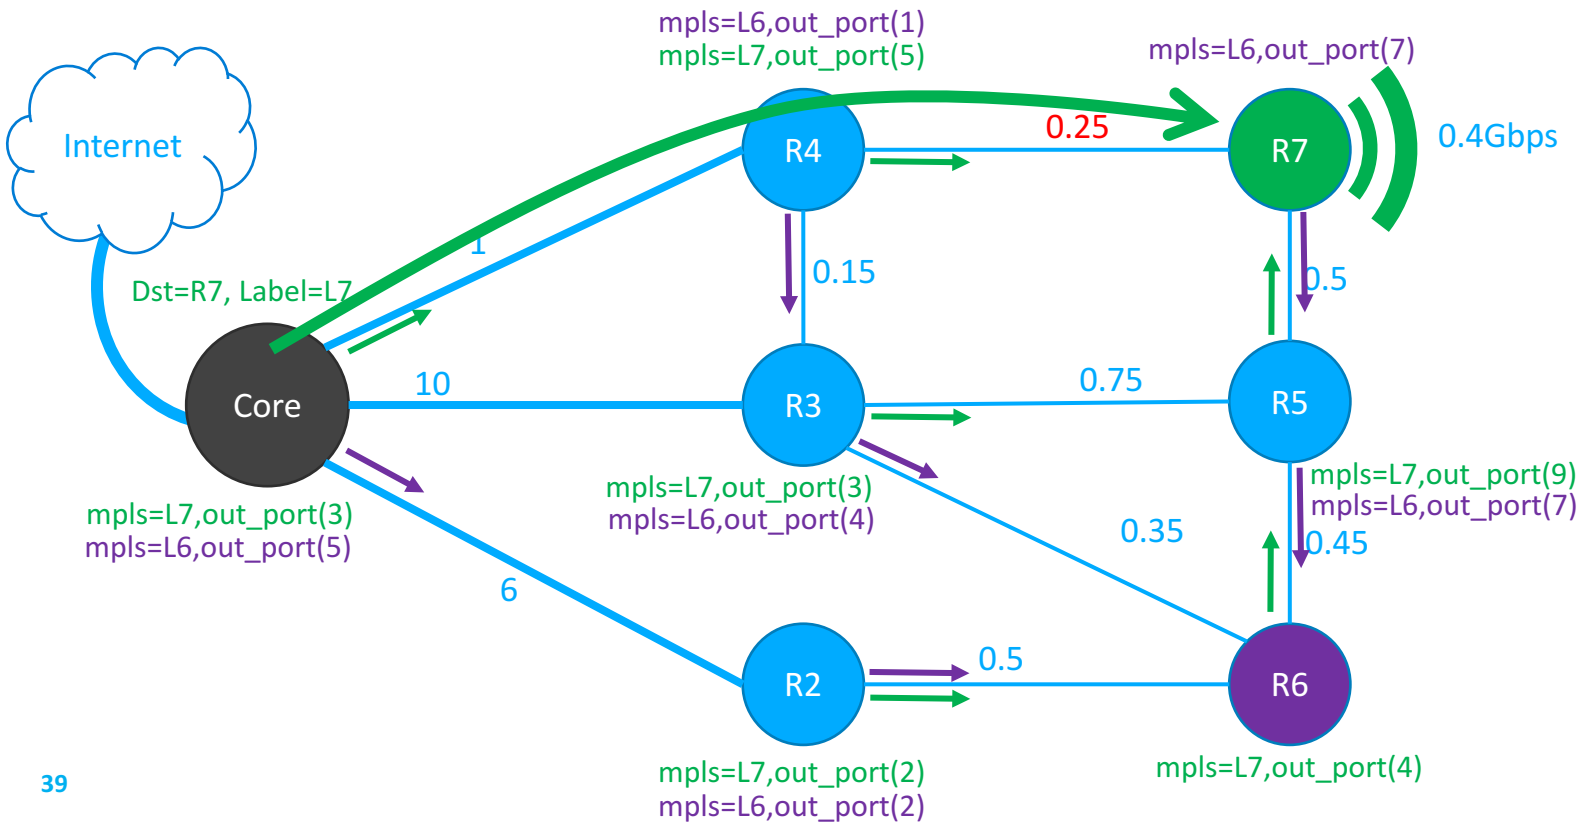

The BLU project
Phase3

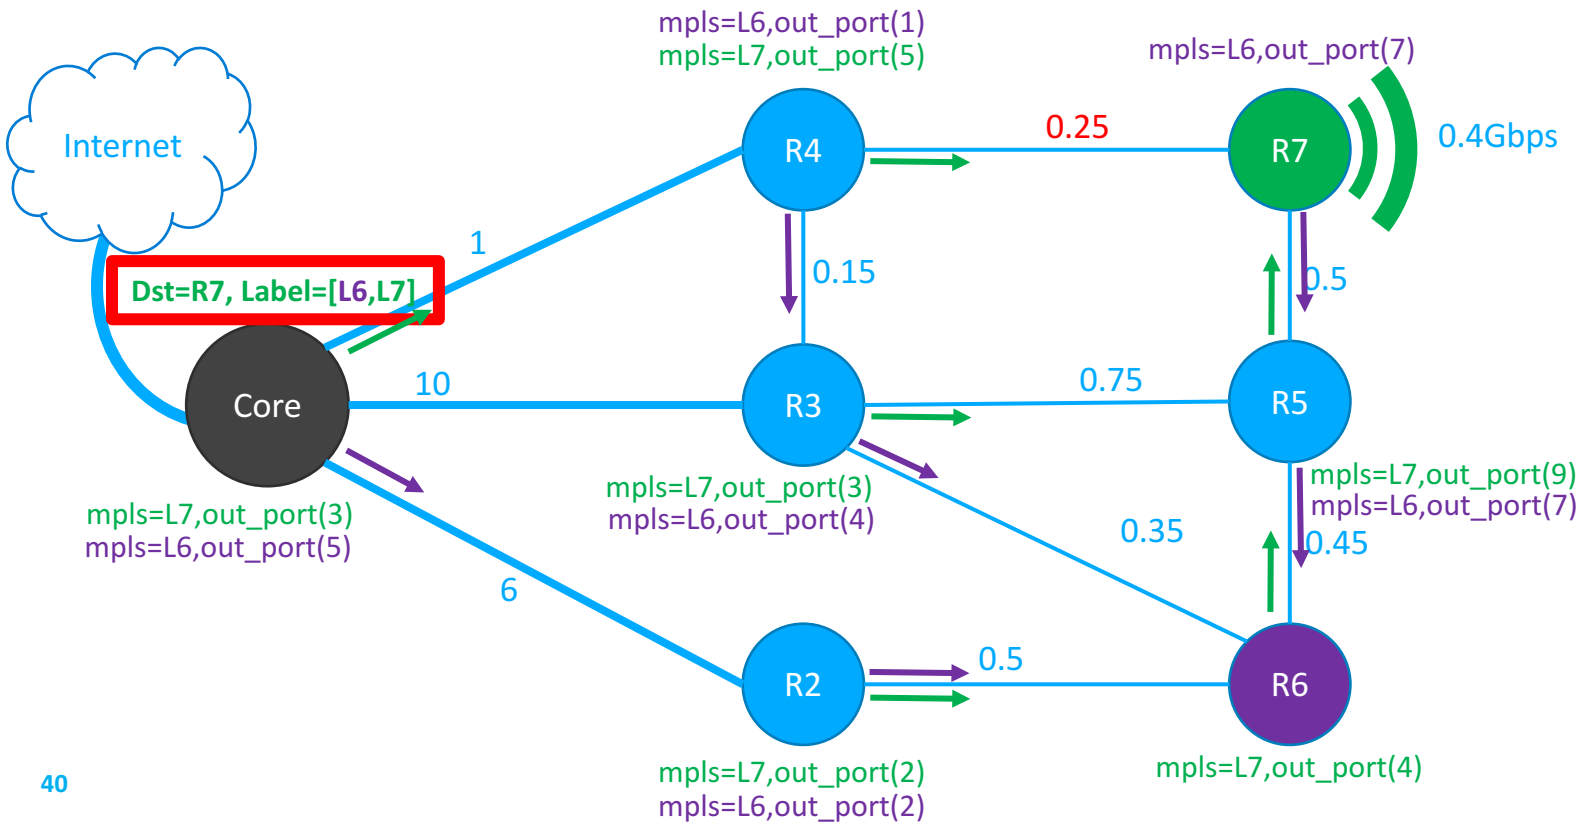

The BLU project
Phase3

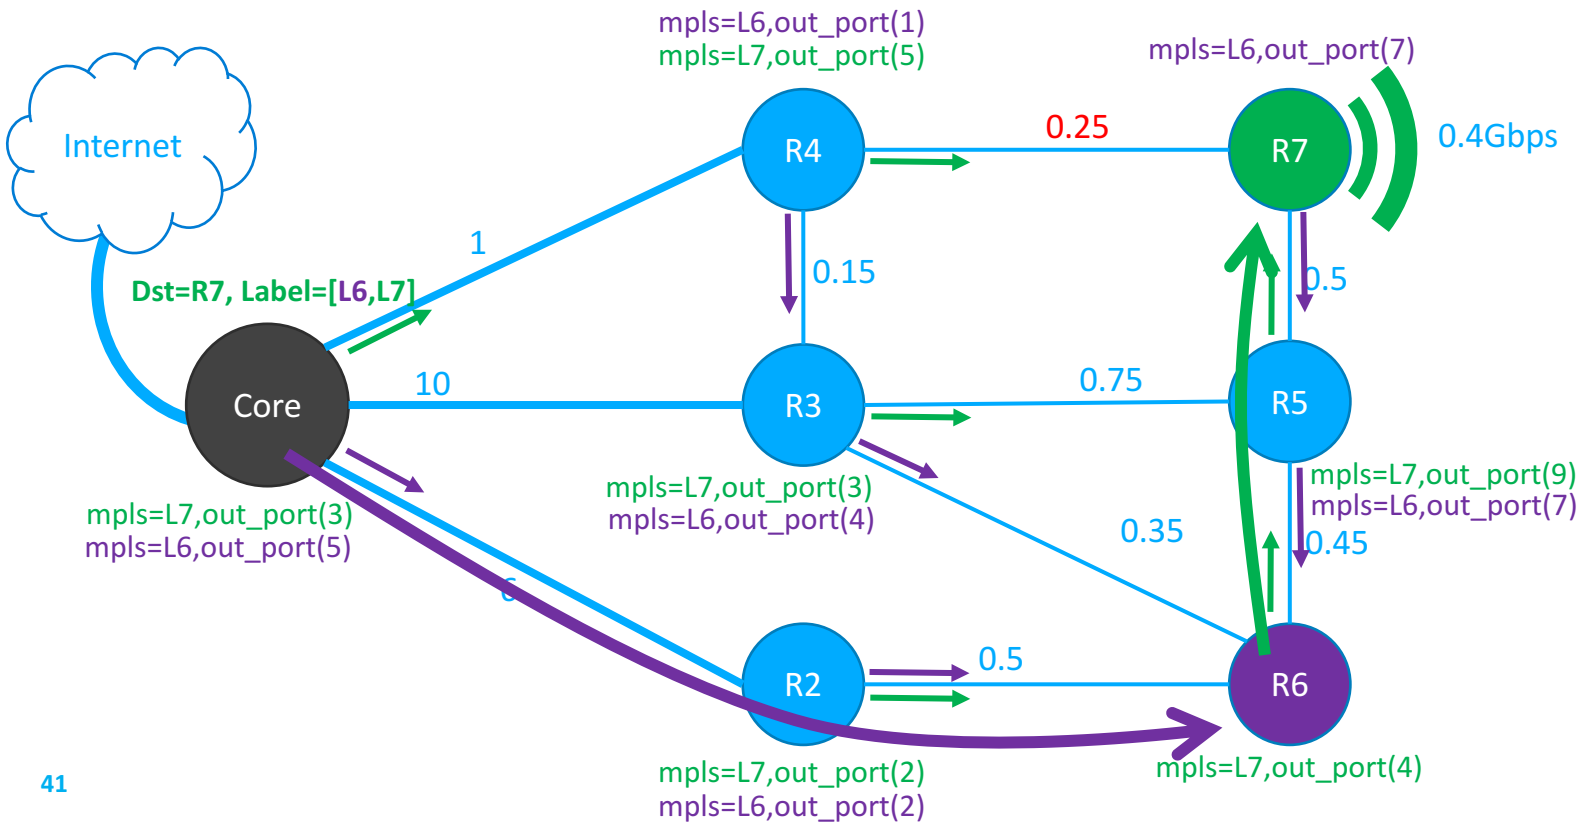

The BLU project
Phase3 - Pros

- No need to re-program the BLUs from BLUGW
  - They automatically learn how to forward MPLS labels based on IGP routes
- Optimizer:
  - Comes into play only in case of congestion
  - Many tools to tackle the optimization problem
    - Operations research, LP, SMT, etc. + a lot of heuristics
- We do have a “Big Red Button” now!
  - In case of troubles, turning off the optimizer reverts to shortest path forwarding

The BLU project
The Future

- New hardware ()
  - Motivation: need of more than 2x10G interfaces on important sites
  - Single  instead of many BLUs (power efficiency)
  - Intel based
  - BLUos compliant

- Latency-based optimizations for “delay-aware” traffic (e.g, voice, video streaming, etc.)
  - BLUs can measure latency for their BH links
  - Optimizer can take these values into consideration

- BLU2BLU (almost) for free!
  - src BLU tags traffic with the dst BLU’s label
  - …or stack of labels (depending on what the optimizer says)

Internet dove gli altri non arrivano.

Thank you!

blu@eolo.it

If you are interested in the BLU project,
do not hesitate to contact us!