LIÈGE université

November 1st, 2017

Through the Wormhole: Tracking Invisible MPLS Tunnels

Yves VANAUBEL
Pascal MÉRINDOL
Jean-Jacques PANSIOT
Benoit DONNET

Agenda

- MPLS background
- Invisible MPLS tunnels
- Measurement Campaign and Results

Agenda

- **MPLS Background**
  - Label Stack Entries
  - MPLS Network
- Invisible MPLS tunnels
- Measurement Campaign and Results

MPLS Label Stack Entries

- **L**abel **S**tack **E**ntries (LSE) :
  - 32 bits
  - Inserted between the MAC and the IP layer

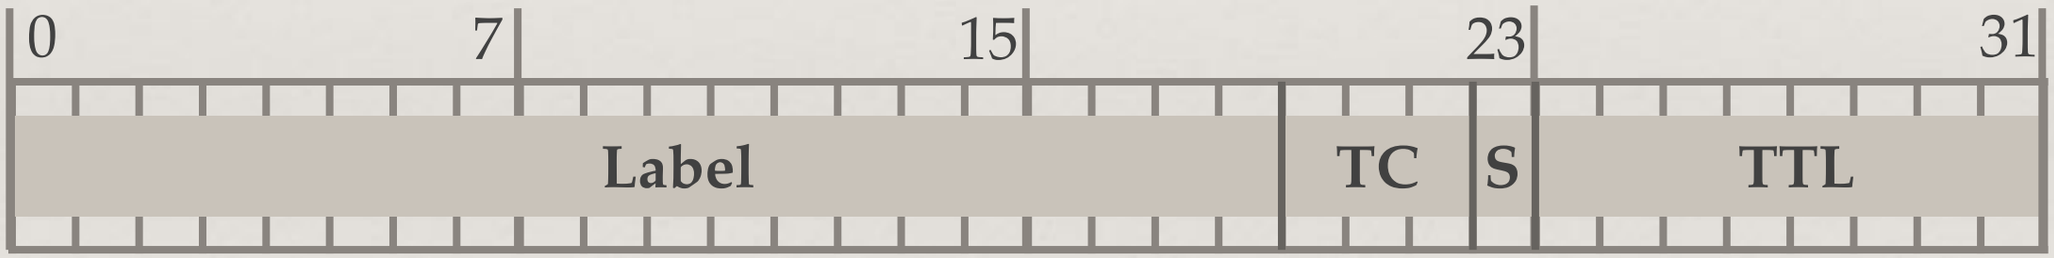

- Label : Label value, 20 bits
- TC: **T**raffic **C**lass field, 3 bits
- S: Bottom of stack, 1 bit
- TTL: **T**ime **T**o **L**ive, 8 bits

MPLS Network

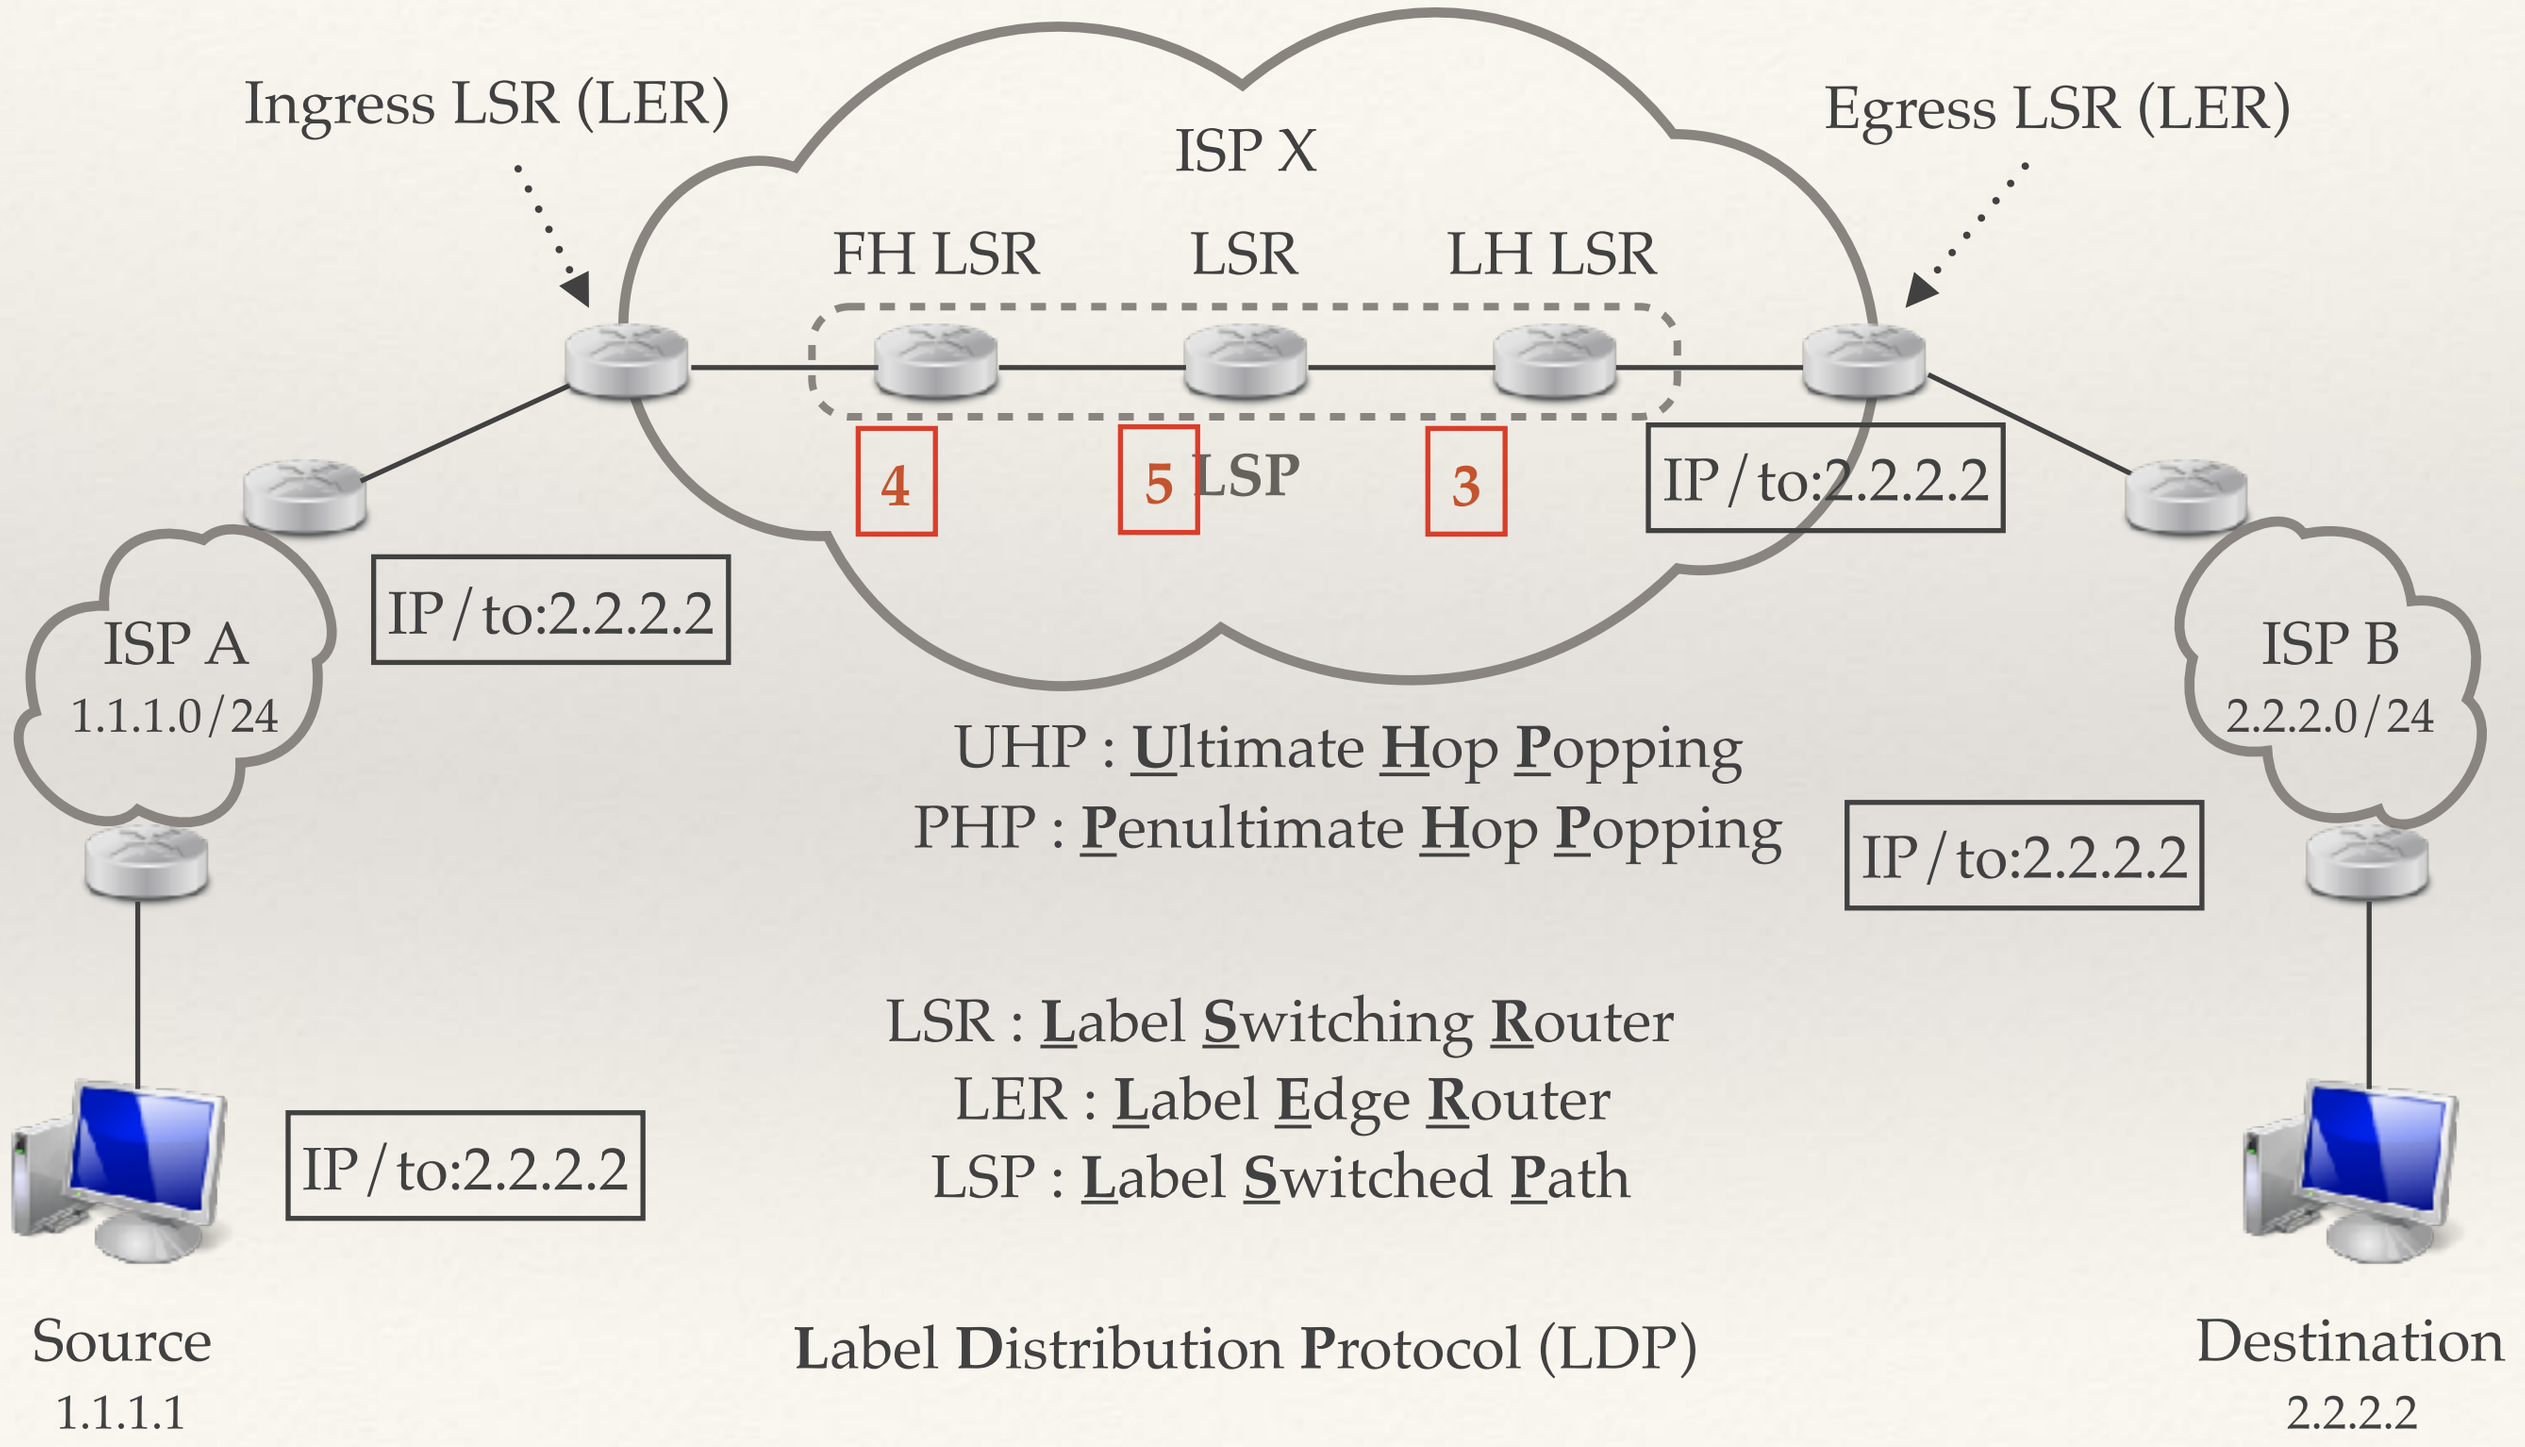

UHP : **U**ltimate **H**op **P**opping
PHP : **P**enultimate **H**op **P**opping

LSR : **L**abel **S**witching **R**outer
LER : **L**abel **E**dge **R**outer
LSP : **L**abel **S**witched **P**ath

Label **D**istribution **P**rotocol (LDP)

Agenda

- MPLS Background
- **Invisible MPLS tunnels**
  - Definition
  - Impact on the Topology Inference
  - Revelation
- Measurement Campaign and Results

MPLS Tunnel Discovery

- Classical MPLS tunnels can be revealed based on standard active measurement tools (`traceroute`)
- Two features are required:
  - **ICMP extension** ([RFC4950]):
    - ✓ If an MPLS router must forge an ICMP *time exceeded* message, it should quote the MPLS LSE into it.
  - **TTL propagation** ([RFC3443]):
    - ✓ The ingress router of an MPLS tunnel should initialize the LSE-TTL with the value inside the IP-TTL field.
    - ✓ The opposite operation is done by the egress LER.

Explicit Tunnels

- The two options are enabled
- This kind of tunnel is perfectly visible with `traceroute`

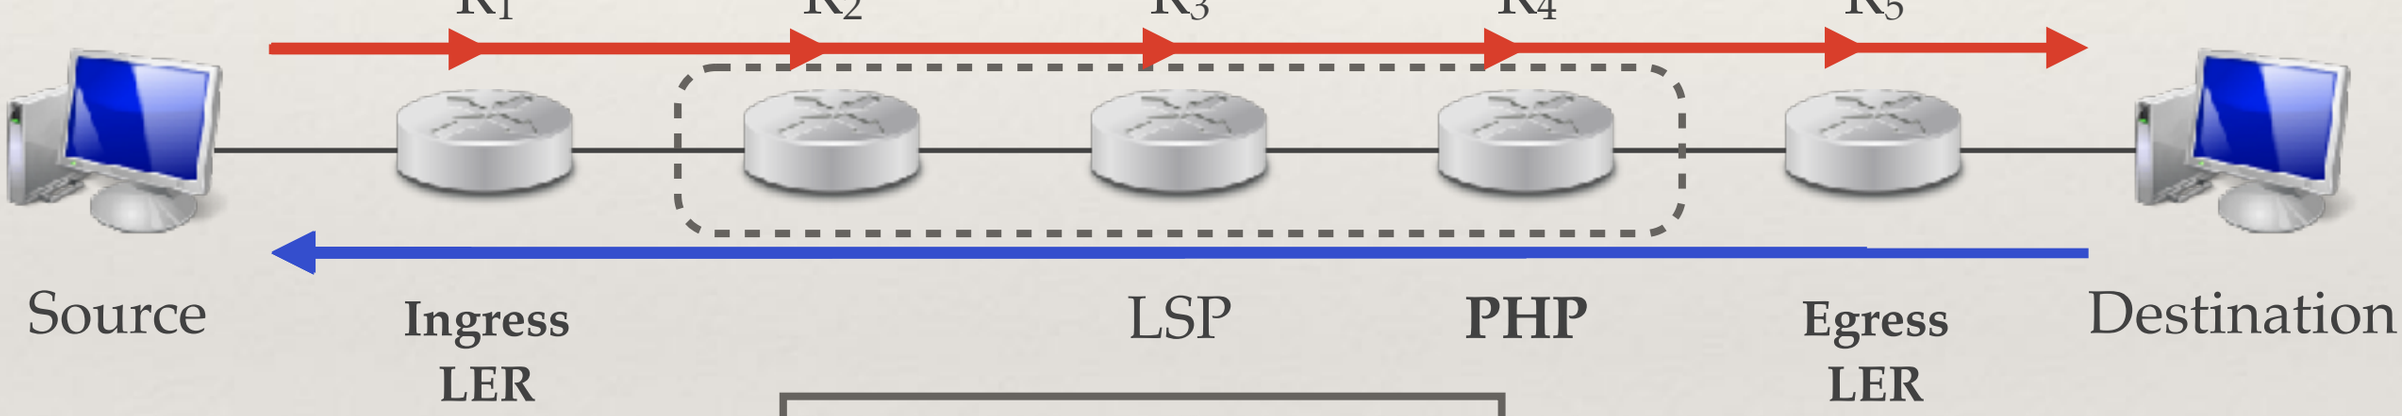

Traceroute output:

1. R_1
2. R_2 - *MPLS tag*
3. R_3 - *MPLS tag*
4. R_4 - *MPLS tag*
5. R_5
6. Destination

Invisible Tunnels

- With invisible tunnels, the TTL propagation is disabled
- Only ingress/egress LERs visible

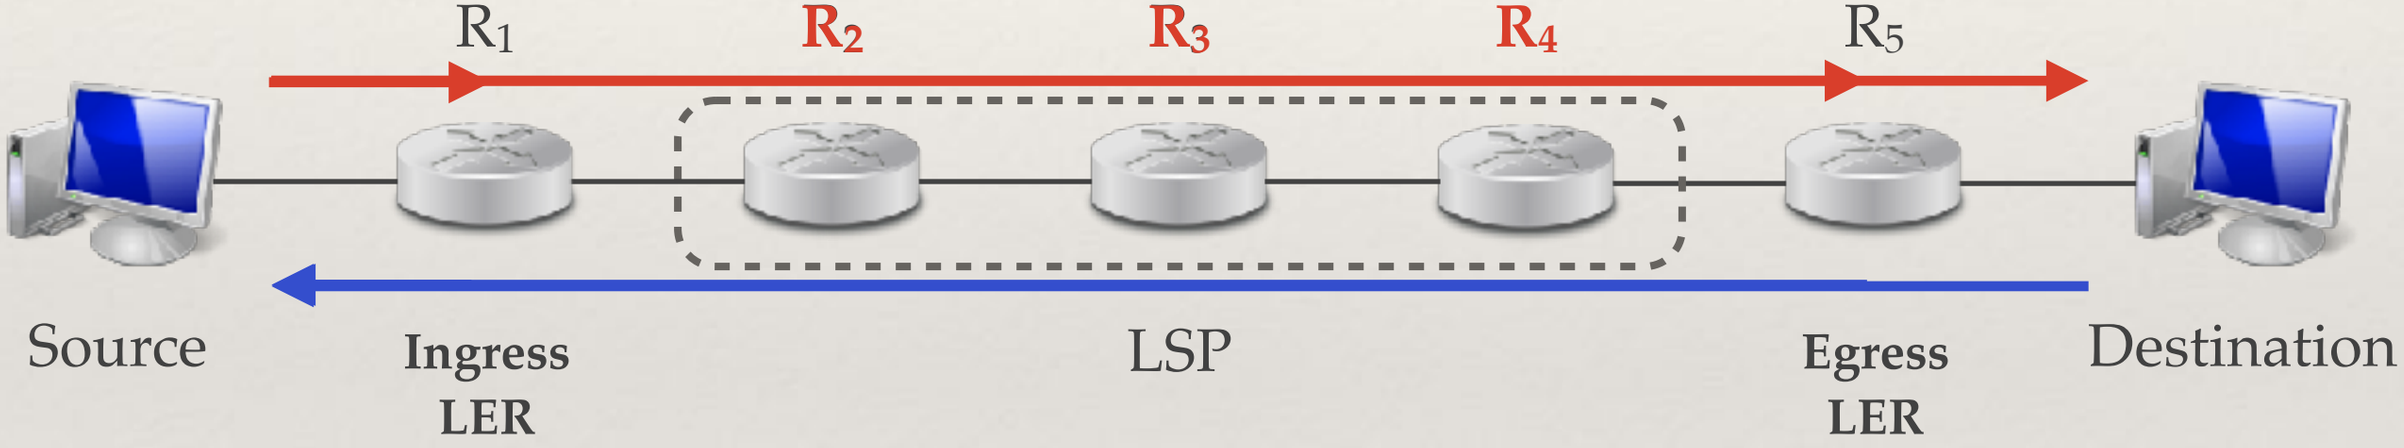

Traceroute output:

1. R_1
2. R_5
3. Destination

False IP link (R1 → R5) inference!

Impact on the Topology Inference

- Internal MPLS routers are hidden from `traceroute`
- An entry point of an MPLS network appears as the neighbor of all exit points
- The whole layer-3 network turns into a dense mesh of **H**igh **D**egree **N**odes (HDN)

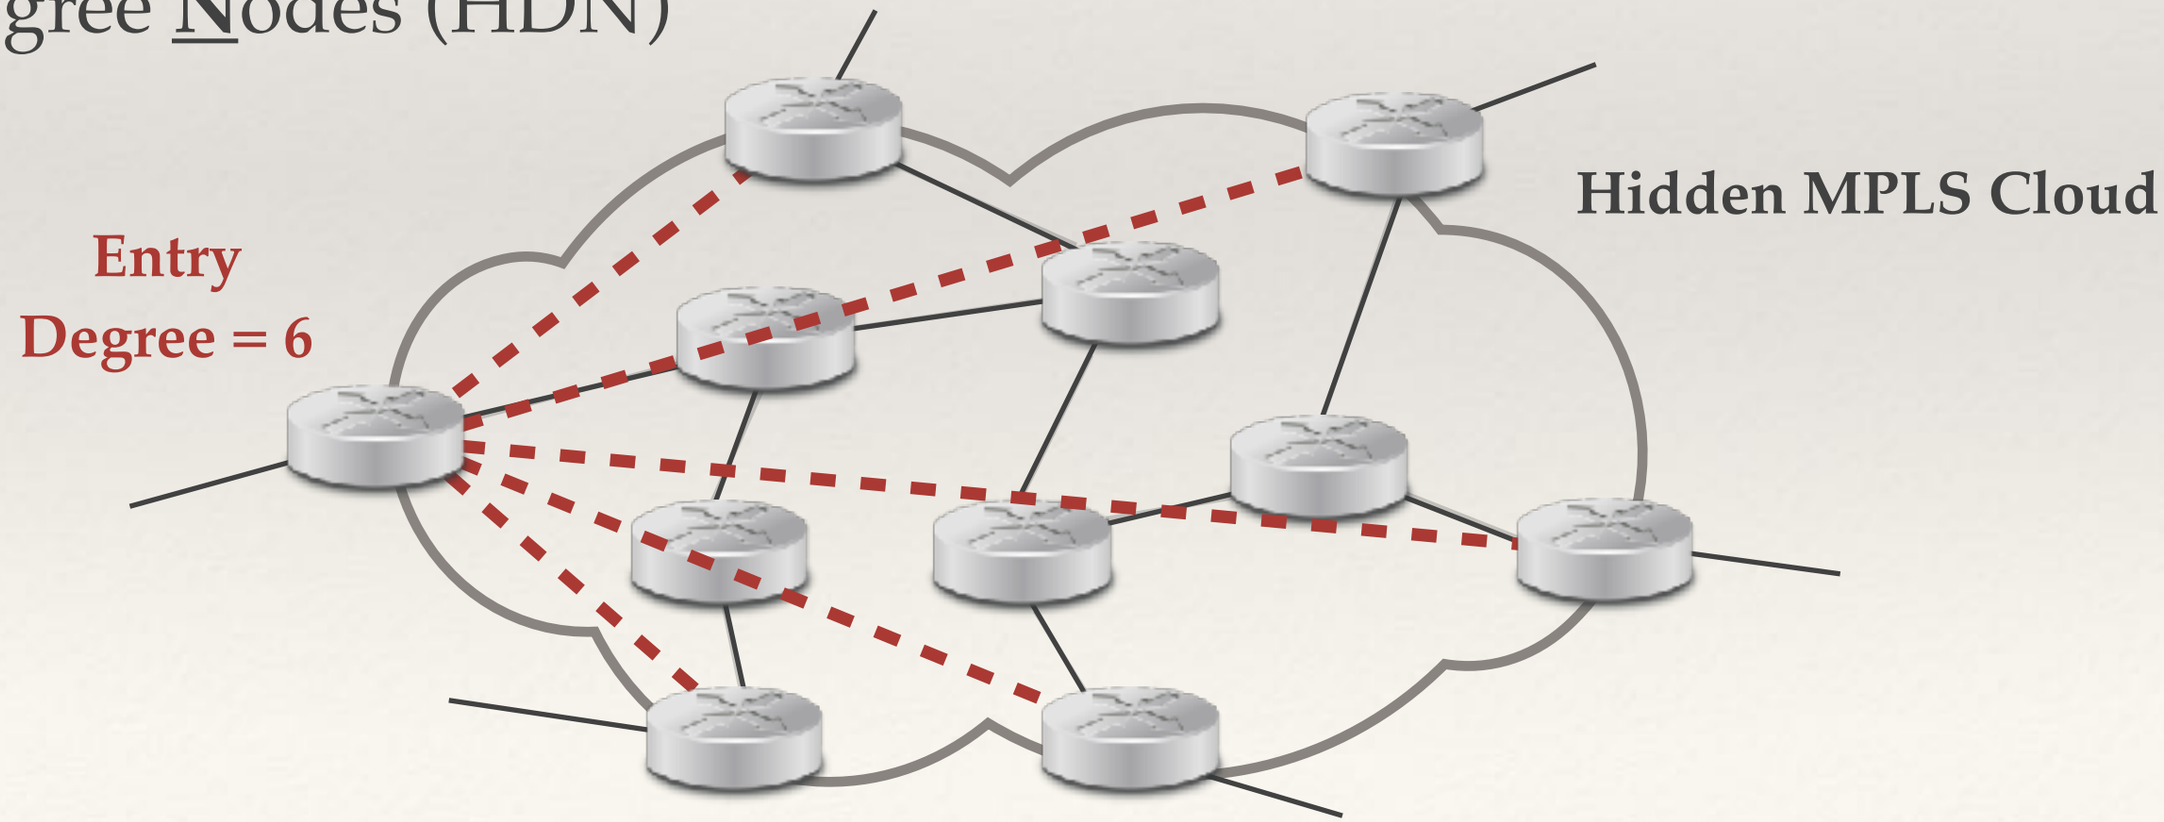

High Degree Node

- A node is a HDN if it has at least 128 neighbors
  - 128 is a lower bound relative to well-known physical provider edge hardware
  - Reasonable balance between the volume of probes sent and the amount of interesting data collected

Invisible Tunnels - Revelation

- **D**irect **P**ath **R**evelation (DPR)
  - For networks not using MPLS for internal routing
  - Mostly Juniper devices (default behavior)
- **B**ackward **R**ecursive **P**ath **R**evelation (BRPR)
  - For networks using MPLS for all prefixes (internal and external)
  - Mostly CISCO routers (default behavior)

Direct Path Revelation (DPR)

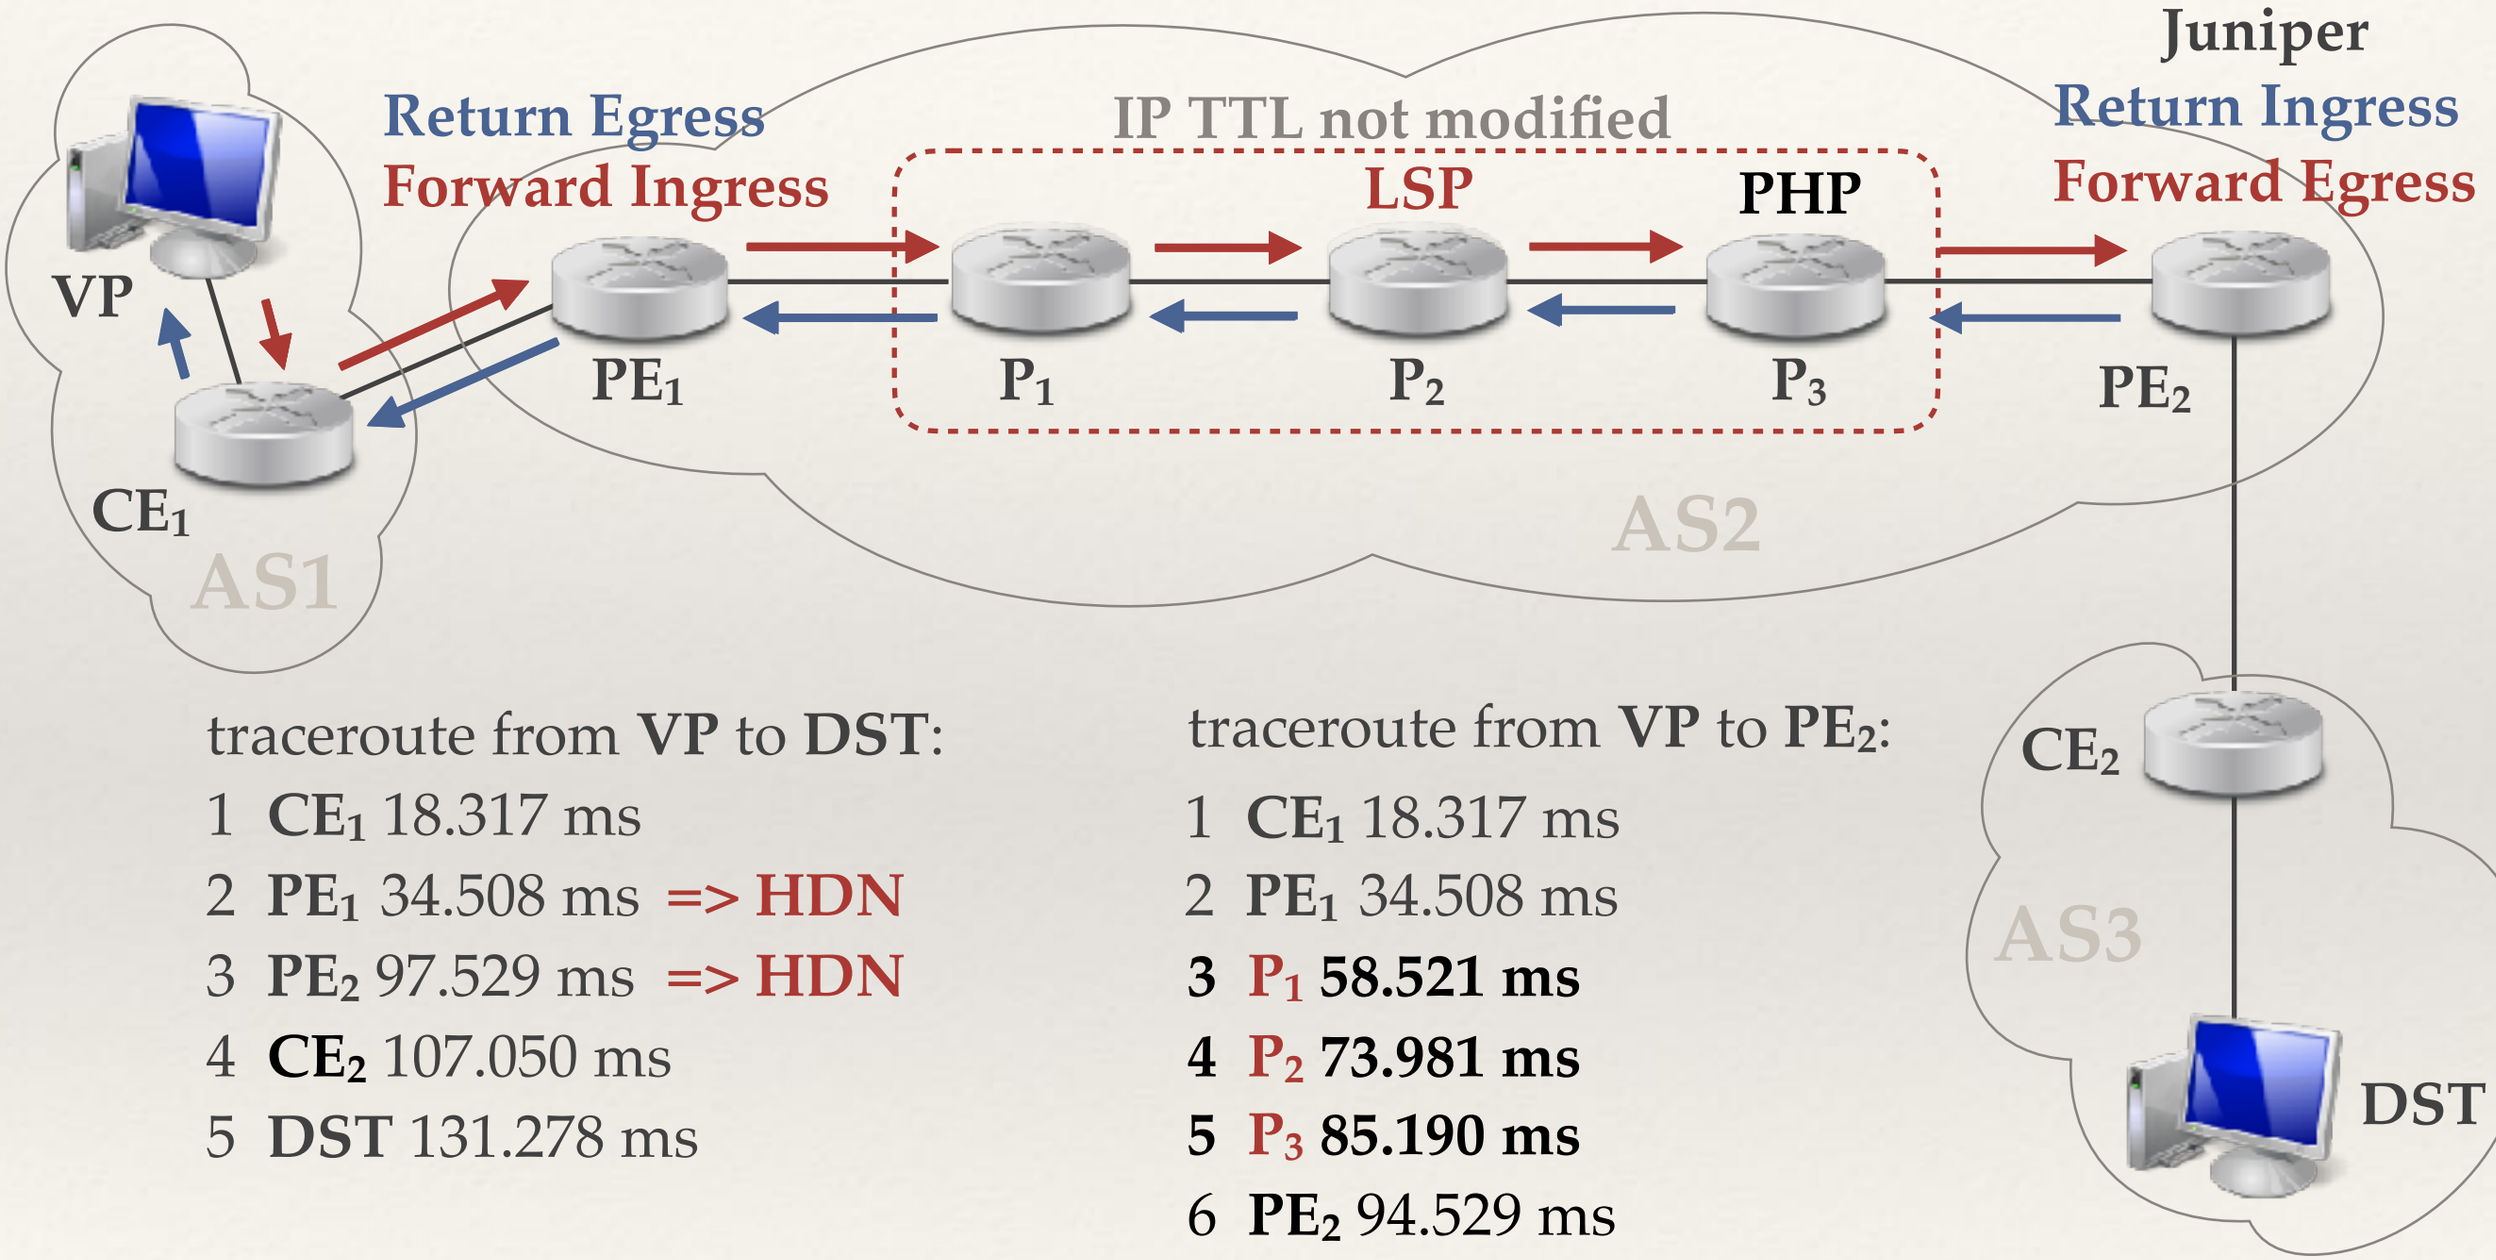

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms
2 **PE$_1$** 34.508 ms **=> HDN**
3 **PE$_2$** 97.529 ms **=> HDN**
4 **CE$_2$** 107.050 ms
5 **DST** 131.278 ms

traceroute from **VP** to **PE$_2$**:

1 **CE$_1$** 18.317 ms
2 **PE$_1$** 34.508 ms
3 P$_1$ 58.521 ms
4 P$_2$ 73.981 ms
5 P$_3$ 85.190 ms
6 **PE$_2$** 94.529 ms

Simple IP forwarding if MPLS not used for internal traffic

=> Try to run a trace to an internal prefix and see if routers reveal themselves

Backward Recursive Path Revelation (BRPR)

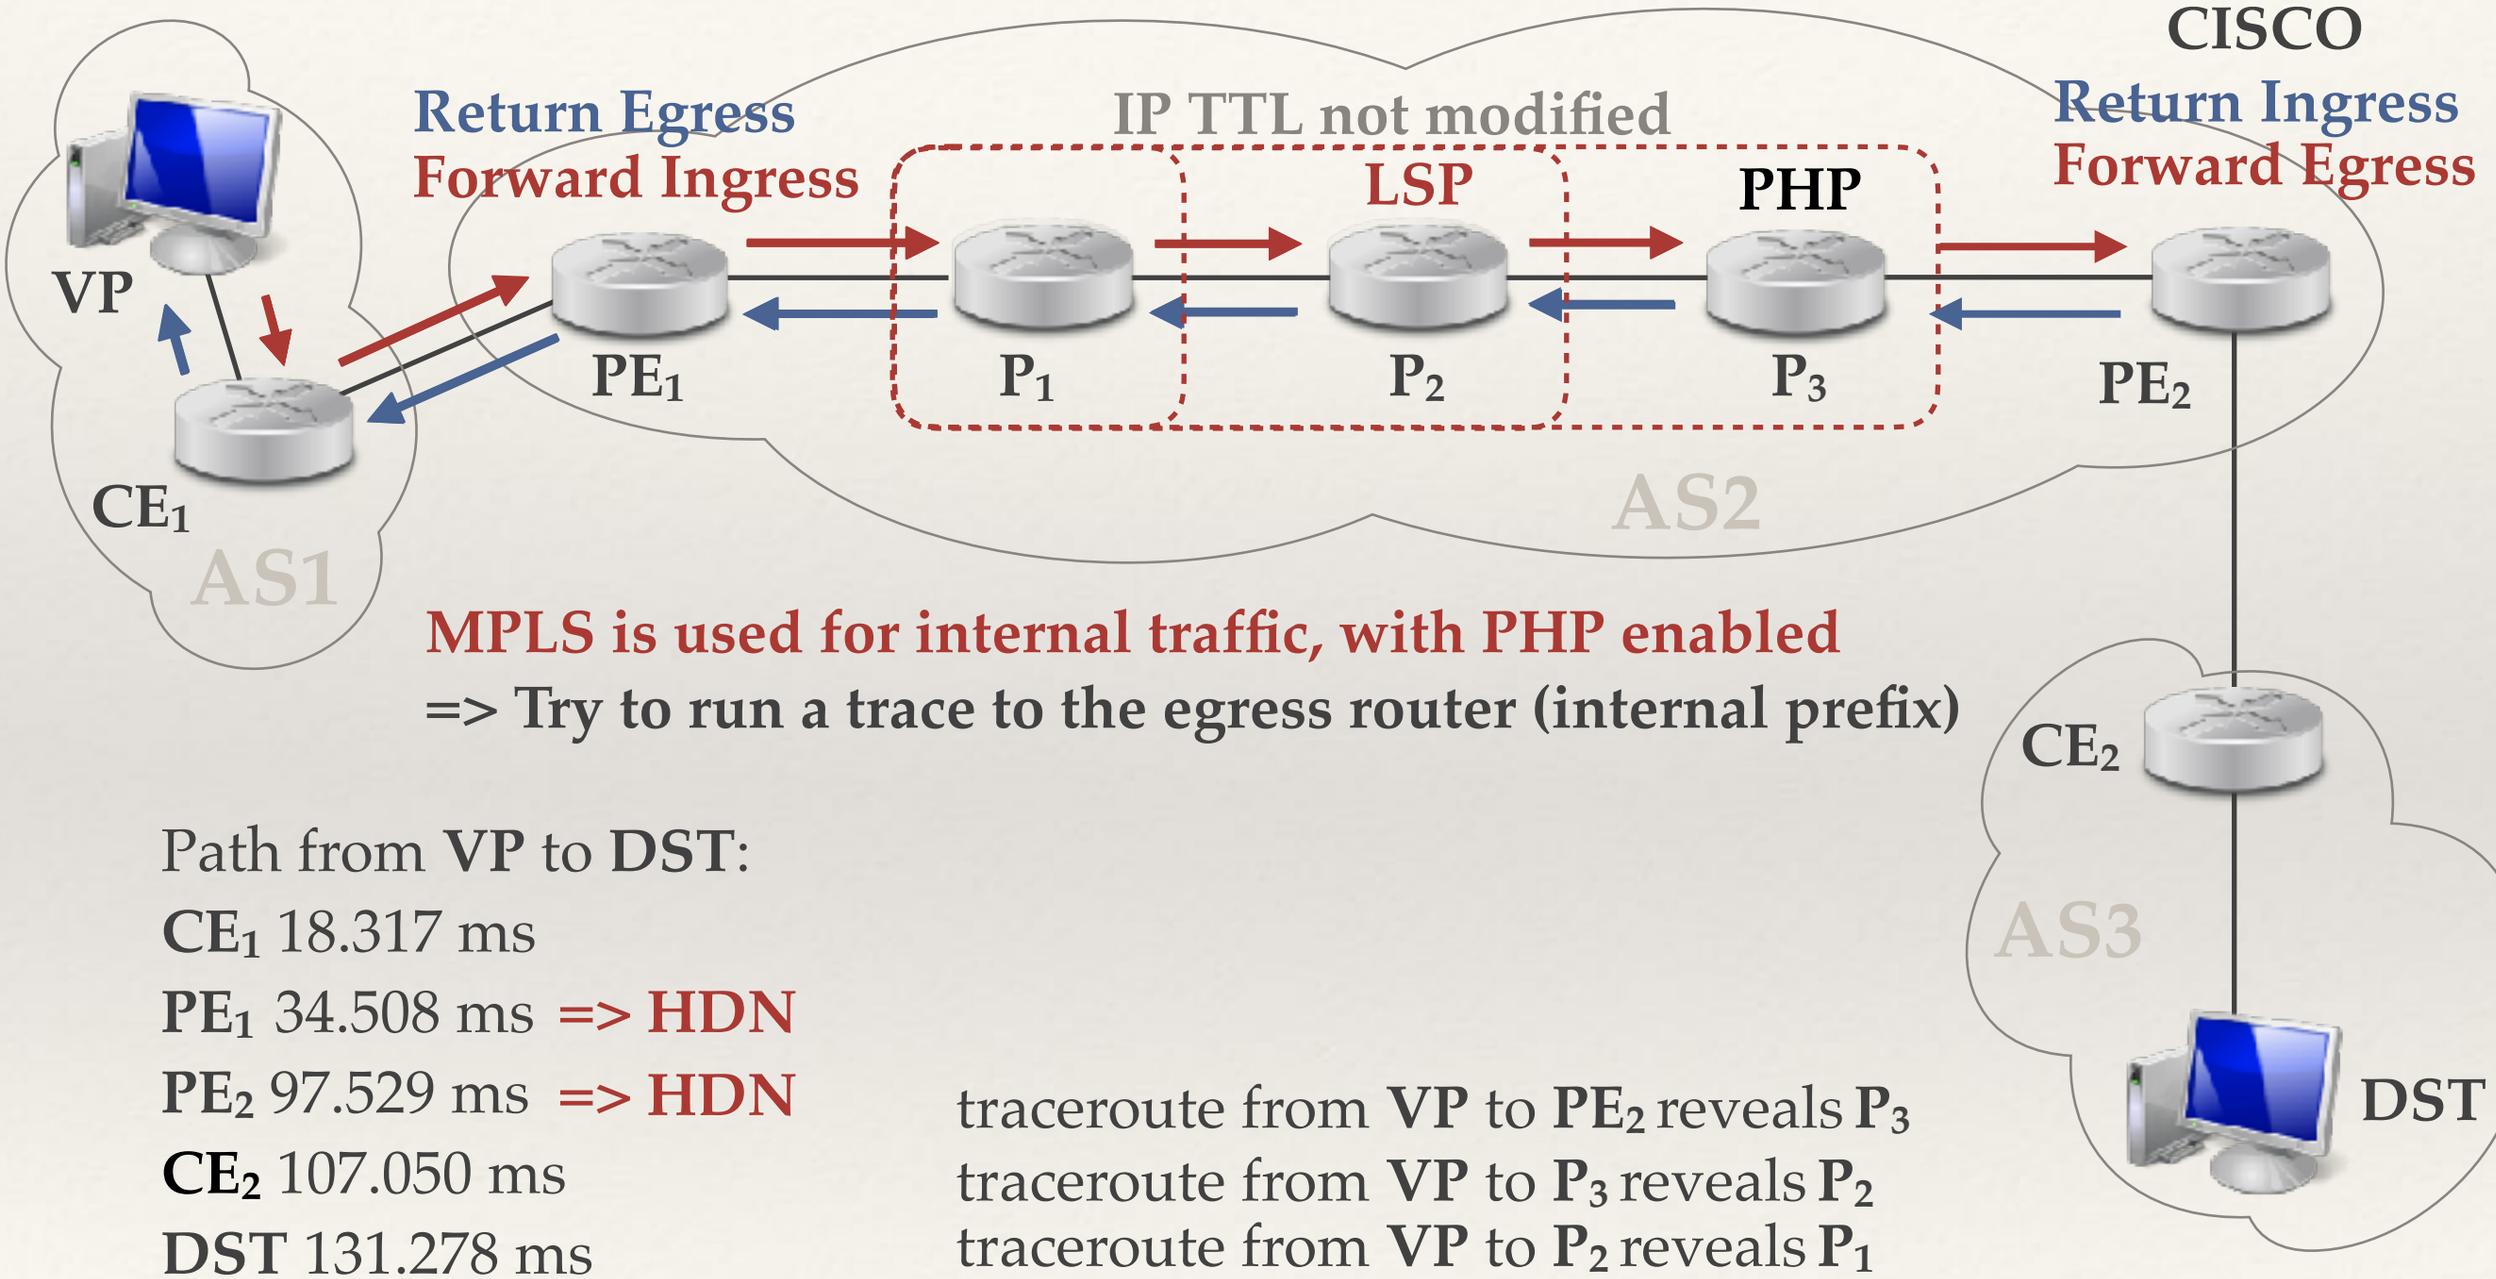

MPLS is used for internal traffic, with PHP enabled

=> Try to run a trace to the egress router (internal prefix)

Path from **VP** to **DST**:

CE$_1$ 18.317 ms

PE$_1$ 34.508 ms **=> HDN**

PE$_2$ 97.529 ms **=> HDN**

CE$_2$ 107.050 ms

DST 131.278 ms

traceroute from **VP** to **PE$_2$** reveals **P$_3$**

traceroute from **VP** to **P$_3$** reveals **P$_2$**

traceroute from **VP** to **P$_2$** reveals **P$_1$**

traceroute from **VP** to **P$_1$** does not reveal any new node

=> STOP

Agenda

- MPLS background
- Invisible MPLS tunnels
- **Measurement Campaign and Results**

Measurement Campaign

- PlanetLab network
- 91 vantage points equally divided in 5 groups
- Selection of HDNs in CAIDA ITDK dataset
- Destinations set: HDNs and their neighbors, i.e. about 1.3M IP addresses
- Destinations distributed amongst the 5 groups
- Scamper with `paris-traceroute`
- Each IP address in the traces pinged for fingerprinting
- About 19 days of measurement

Measurement Results

- 13,771 revealed invisible tunnels
  - 61% with DPR
  - 16% with BRPR
  - 23% with DPR/BRPR (1 hop, impossible to discriminate between the two techniques)
- 5193 revealed public IP addresses

Invisible Tunnels Length

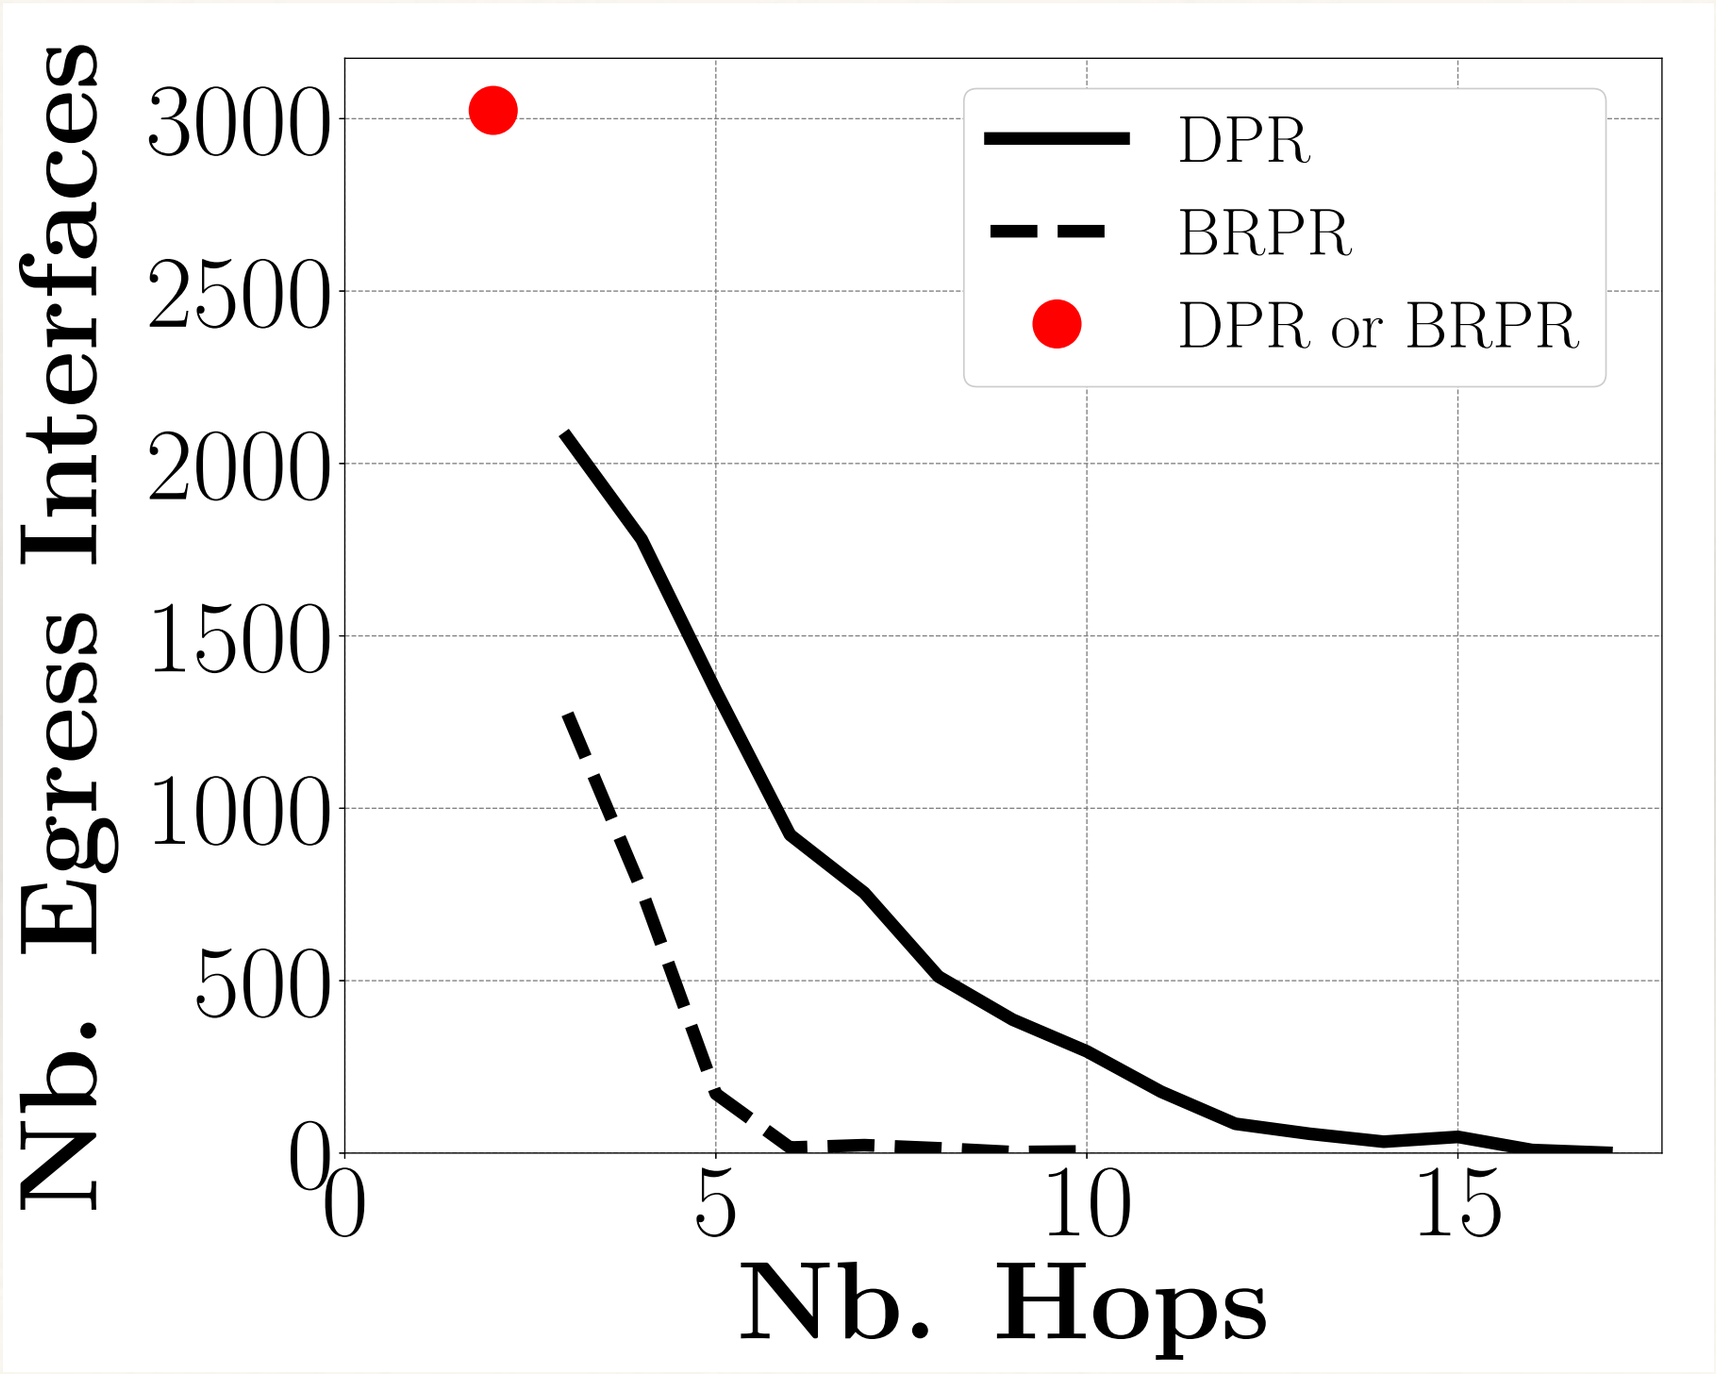

Impact of Invisible Tunnel on Internet Models

- Degree distribution

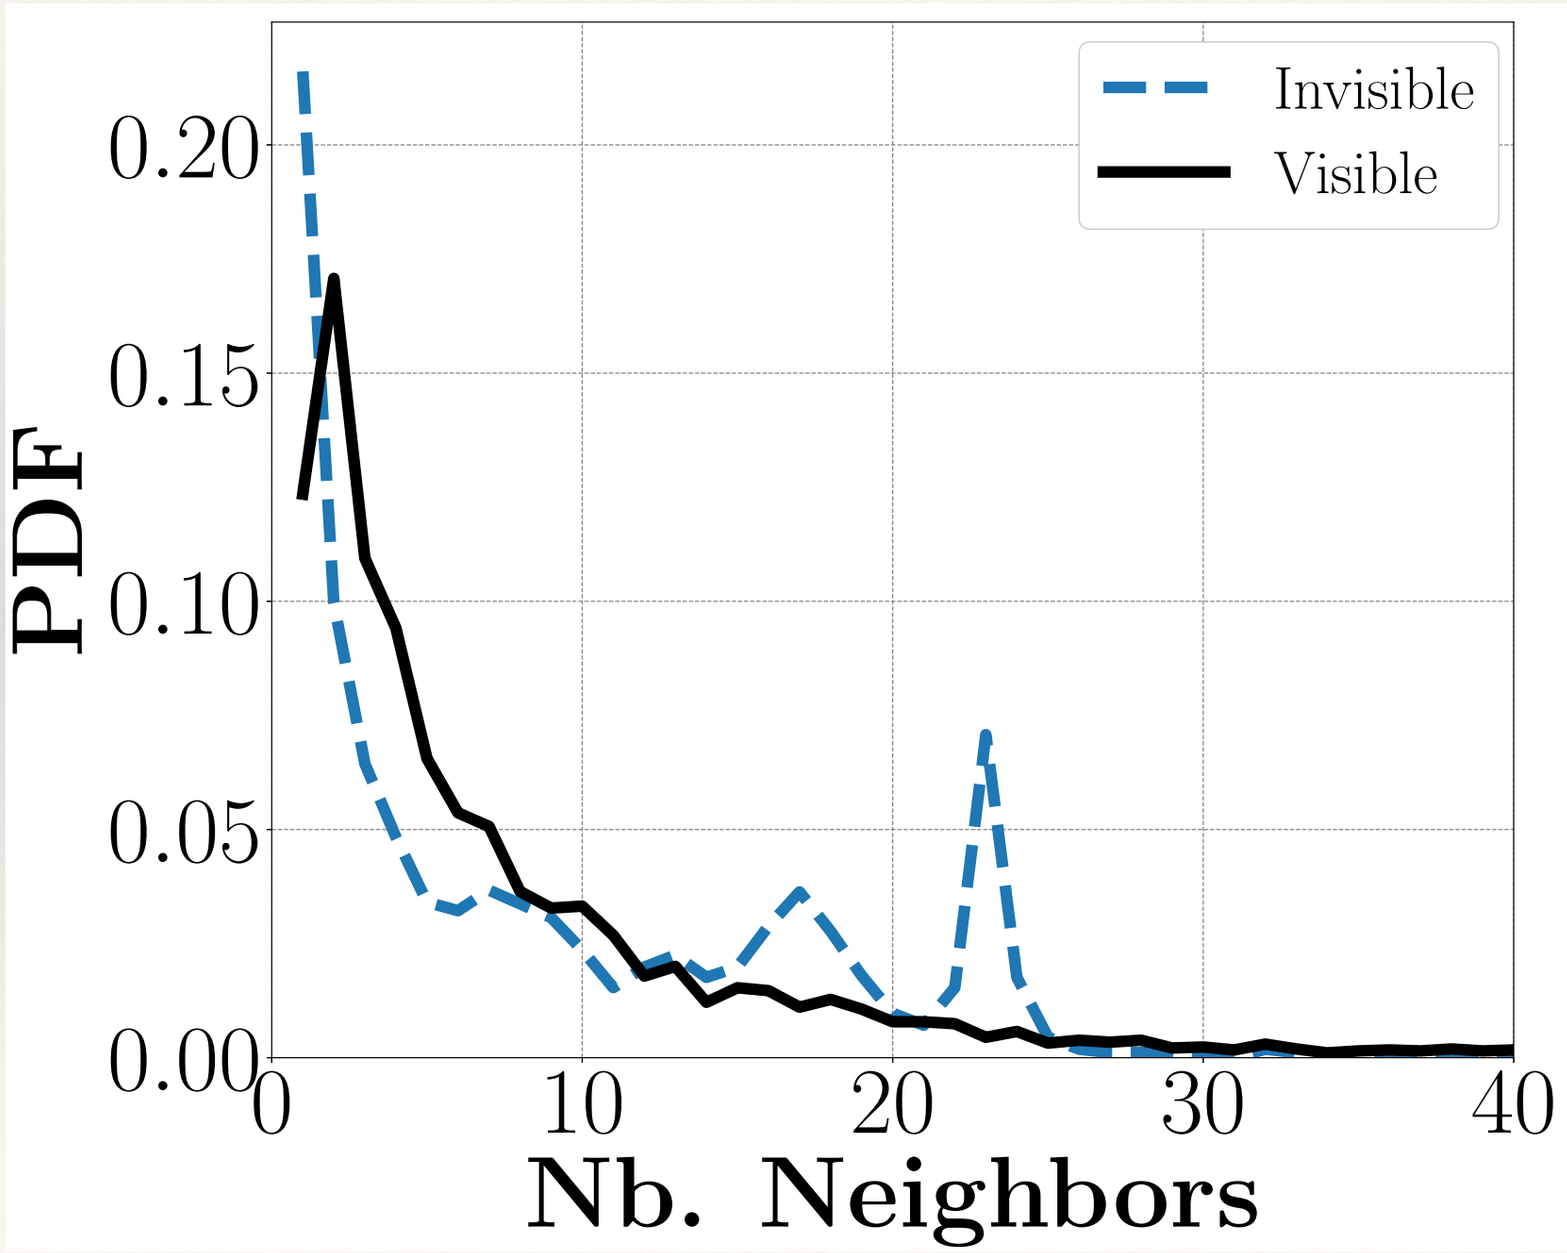

Impact of Invisible Tunnel on Internet Models

- Path lengths

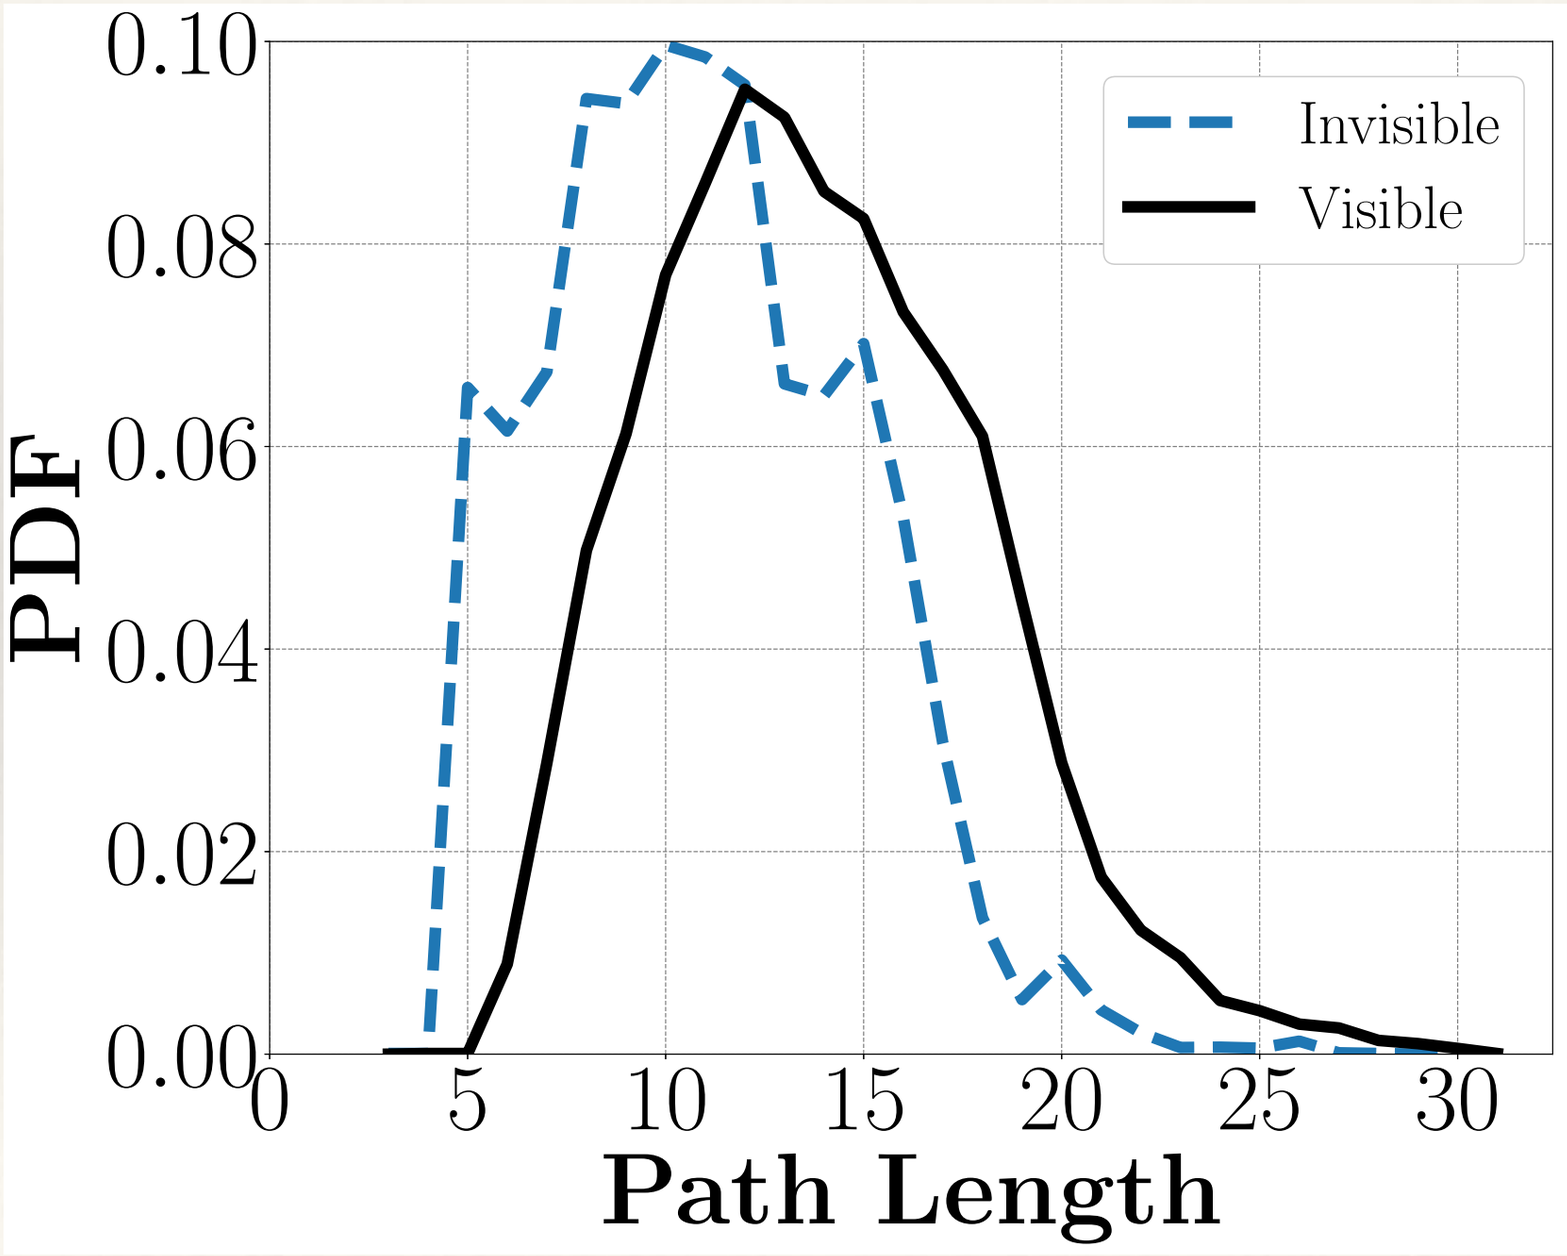

Conclusions

- New techniques to infer the presence and reveal invisible MPLS tunnels
- Validation based on GNS3 emulations
- Gain knowledge on the internal architecture of opaque MPLS ASes
- Help improving Internet models

Conclusions

- Other techniques allow to infer the length of invisible tunnels without revealing the content
  - Can be used as triggers before applying the revelation methods
  - Allow a modification of `traceroute` to run hidden MPLS tunnel revelations based on the triggers
- Dataset and GNS3 validation models publicly available:

http://www.montefiore.ulg.ac.be/~bdonnet/mpls