TNT: Revealing All MPLS Tunnels

Y. Vanaubel, Jean-Romain Luttringer,
P. Mérindol, B. Donnet, J.-J. Pansiot
RIPE Meeting 77, Amsterdam, October 2018

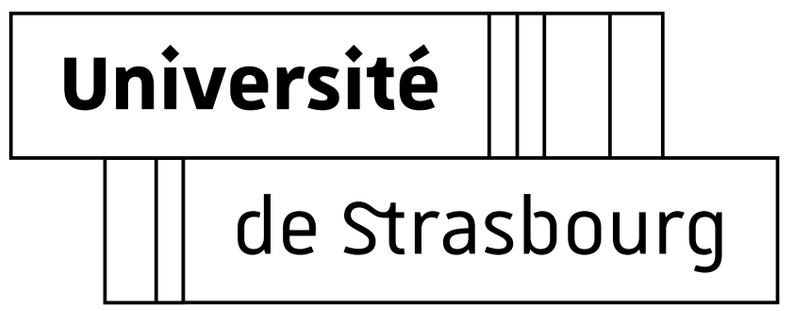

Agenda

- Motivations
- Network Fingerprinting
- MPLS Background
- TNT and Invisible Tunnels
- Conclusion

Agenda

- Motivations
- Network Fingerprinting
- MPLS Background
- TNT and MPLS Invisible Tunnels
- Conclusion

Motivations

- General knowledge
  - are MPLS tunnels prevalent in the Internet?
  - is the technology widely used by operators?
- Network design
  - how MPLS is used by operators?
  - are there any particular design associated to MPLS networks?
- Network modeling
  - what are the interactions between MPLS tunnels and standard network models?

Agenda

- Motivations
- Network Fingerprinting
- MPLS Background
- TNT and MPLS Invisible Tunnels
- Conclusion

Network Fingerprinting

Network Fingerprinting

- **<u>Network fingerprinting</u>**
  - Y. Vanaubel, J.-J. Pansiot, P. Mérindol, B. Donnet. *Network Fingerprinting: TTL-Based Router Signatures*. In Proc. ACM Internet Measurement Conference (IMC). November 2013

Network Fingerprinting

- **<u>Network fingerprinting</u>**
  - Y. Vanaubel, J.-J. Pansiot, P. Mérindol, B. Donnet. *Network Fingerprinting: TTL-Based Router Signatures*. In Proc. ACM Internet Measurement Conference (IMC). November 2013
- **<u>Fingerprinting</u>**
  - action of grouping network devices into disjoint classes

Network Fingerprinting

- **Network fingerprinting**
  - Y. Vanaubel, J.-J. Pansiot, P. Mérindol, B. Donnet. *Network Fingerprinting: TTL-Based Router Signatures*. In Proc. ACM Internet Measurement Conference (IMC). November 2013
- **Fingerprinting**
  - action of grouping network devices into disjoint classes
- **Signature**
  - set of information collected thanks to fingerprinting

Network Fingerprinting (2)

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])
- In practice, iTTL may depend on

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])
- In practice, iTTL may depend on
  - hardware
    - ✓ CISCO, Juniper

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])
- In practice, iTTL may depend on
  - hardware
    - ✓ CISCO, Juniper
  - operating system
    - ✓ JunOS, JunOSE, IOS, ...

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])
- In practice, iTTL may depend on
  - hardware
    - ✓ CISCO, Juniper
  - operating system
    - ✓ JunOS, JunOSE, IOS, ...
  - protocol used to transmit the message
    - ✓ ICMP, UDP, TCP, ...

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])
- In practice, iTTL may depend on
  - hardware
    - ✓ CISCO, Juniper
  - operating system
    - ✓ JunOS, JunOSE, IOS, ...
  - protocol used to transmit the message
    - ✓ ICMP, UDP, TCP, ...
  - the type of message
    - ✓ information of error message

Network Fingerprinting (2)

- Fingerprinting is based on *initial TTL* (iTTL) value
  - when forging packets, devices should initialize the IP-TTL to 64 ([RFC1700])
- In practice, iTTL may depend on
  - hardware
    - ✓ CISCO, Juniper
  - operating system
    - ✓ JunOS, JunOSE, IOS, ...
  - protocol used to transmit the message
    - ✓ ICMP, UDP, TCP, ...
  - the type of message
    - ✓ information of error message
- iTTL used by devices: ***32***, ***64***, ***128***, ***255***

Network Fingerprinting (3)

- Basic idea of network fingerprinting
  - soliciting routers with different probes to receive different types of (ICMP) replies
  - infer their iTTL value
    - ✓ smallest integer in {***32***, ***64***, ***128***, ***255***} larger than the received value
  - derive a signature of the type

$$< iTTL_1, iTTL_2, iTTL_3, \ldots, iTTL_n >$$

Network Fingerprinting (3)

- Basic idea of network fingerprinting
  - soliciting routers with different probes to receive different types of (ICMP) replies
  - infer their iTTL value
    - ✓ smallest integer in {***32***, ***64***, ***128***, ***255***} larger than the received value
  - derive a signature of the type

$$< iTTL_1, iTTL_2, iTTL_3, \ldots, iTTL_n >$$

- Distribution of signatures already valuable with 2 iTTLs
  - ICMP `time_exceeded` (TE)
  - ICMP `echo_reply` (ER)

Network Fingerprinting (4)

- Signatures for major manufacturers

Manufacturer	<TE, ER>
Cisco	**<255, 255>**
Juniper (JunOS)	**<255, 64>**
Juniper (JunOSE)	**<128, 128>**
Brocade, Alcatel, and Linux Boxes	**<64, 64>**

Agenda

- Motivations
- Network Fingerprinting
- MPLS Background
  - Label Stack Entry
  - MPLS Network
- TNT and MPLS Invisible Tunnels
- Conclusion

Label Stack Entry

- **<u>Label Stack Entry</u>** (LSE)
  - 32 bits

Label Stack Entry

- **<u>Label Stack Entry</u>** (LSE)
  - 32 bits

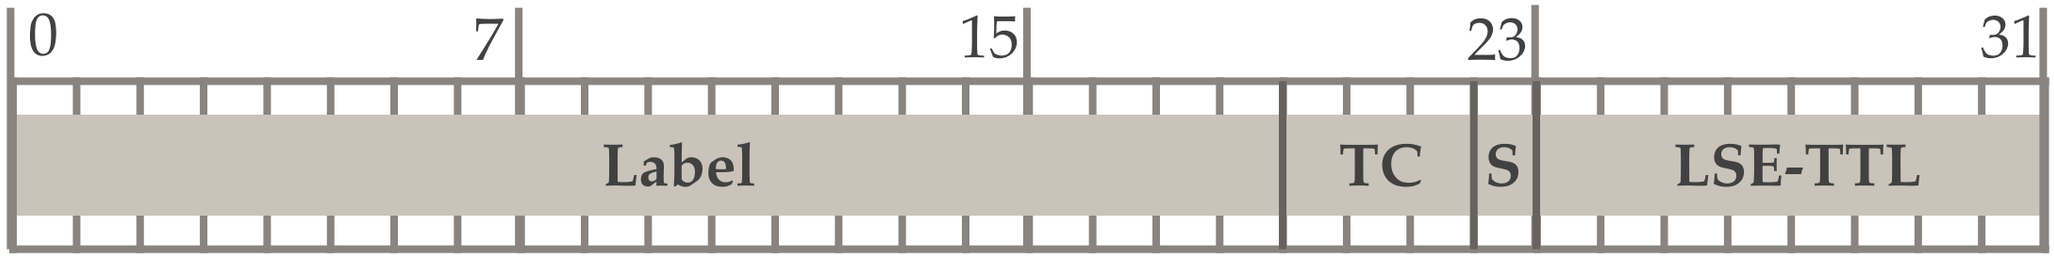

- Label : Label value, 20 bits
  - 0 - 1048575
  - 0 - 15 reserved by IETF
- TC: <u>T</u>raffic <u>C</u>lass field, 3 bits
- S: Bottom of stack, 1 bit
- LSE-TTL: <u>T</u>ime <u>T</u>o <u>L</u>ive, 8 bits

MPLS Network

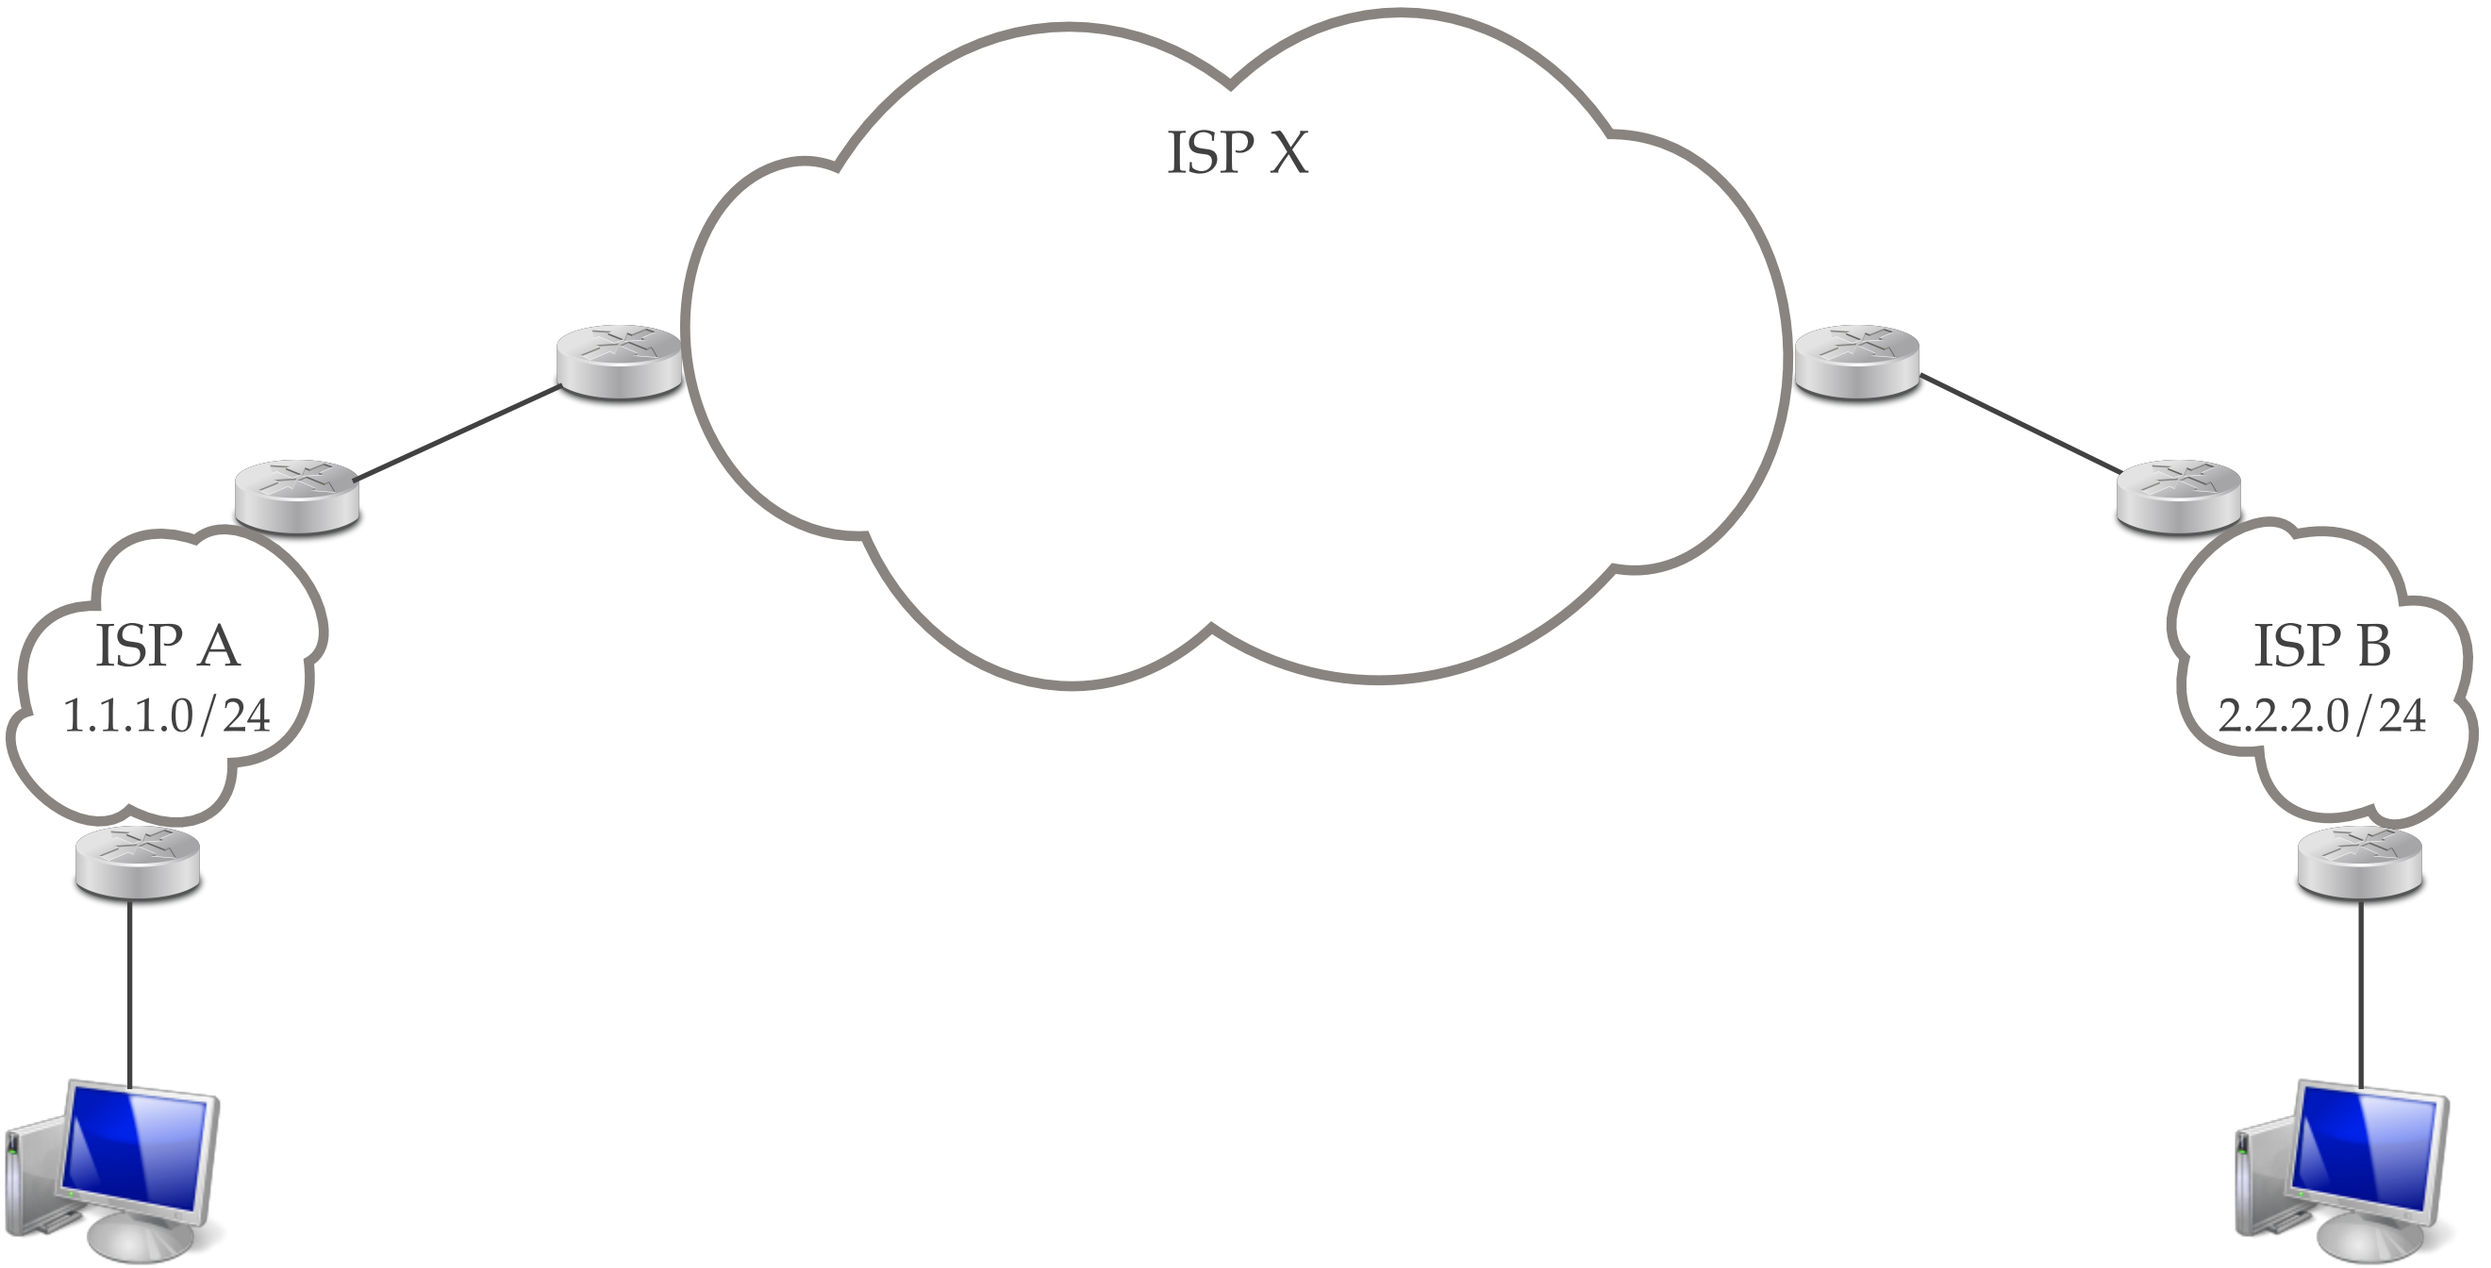

MPLS Network

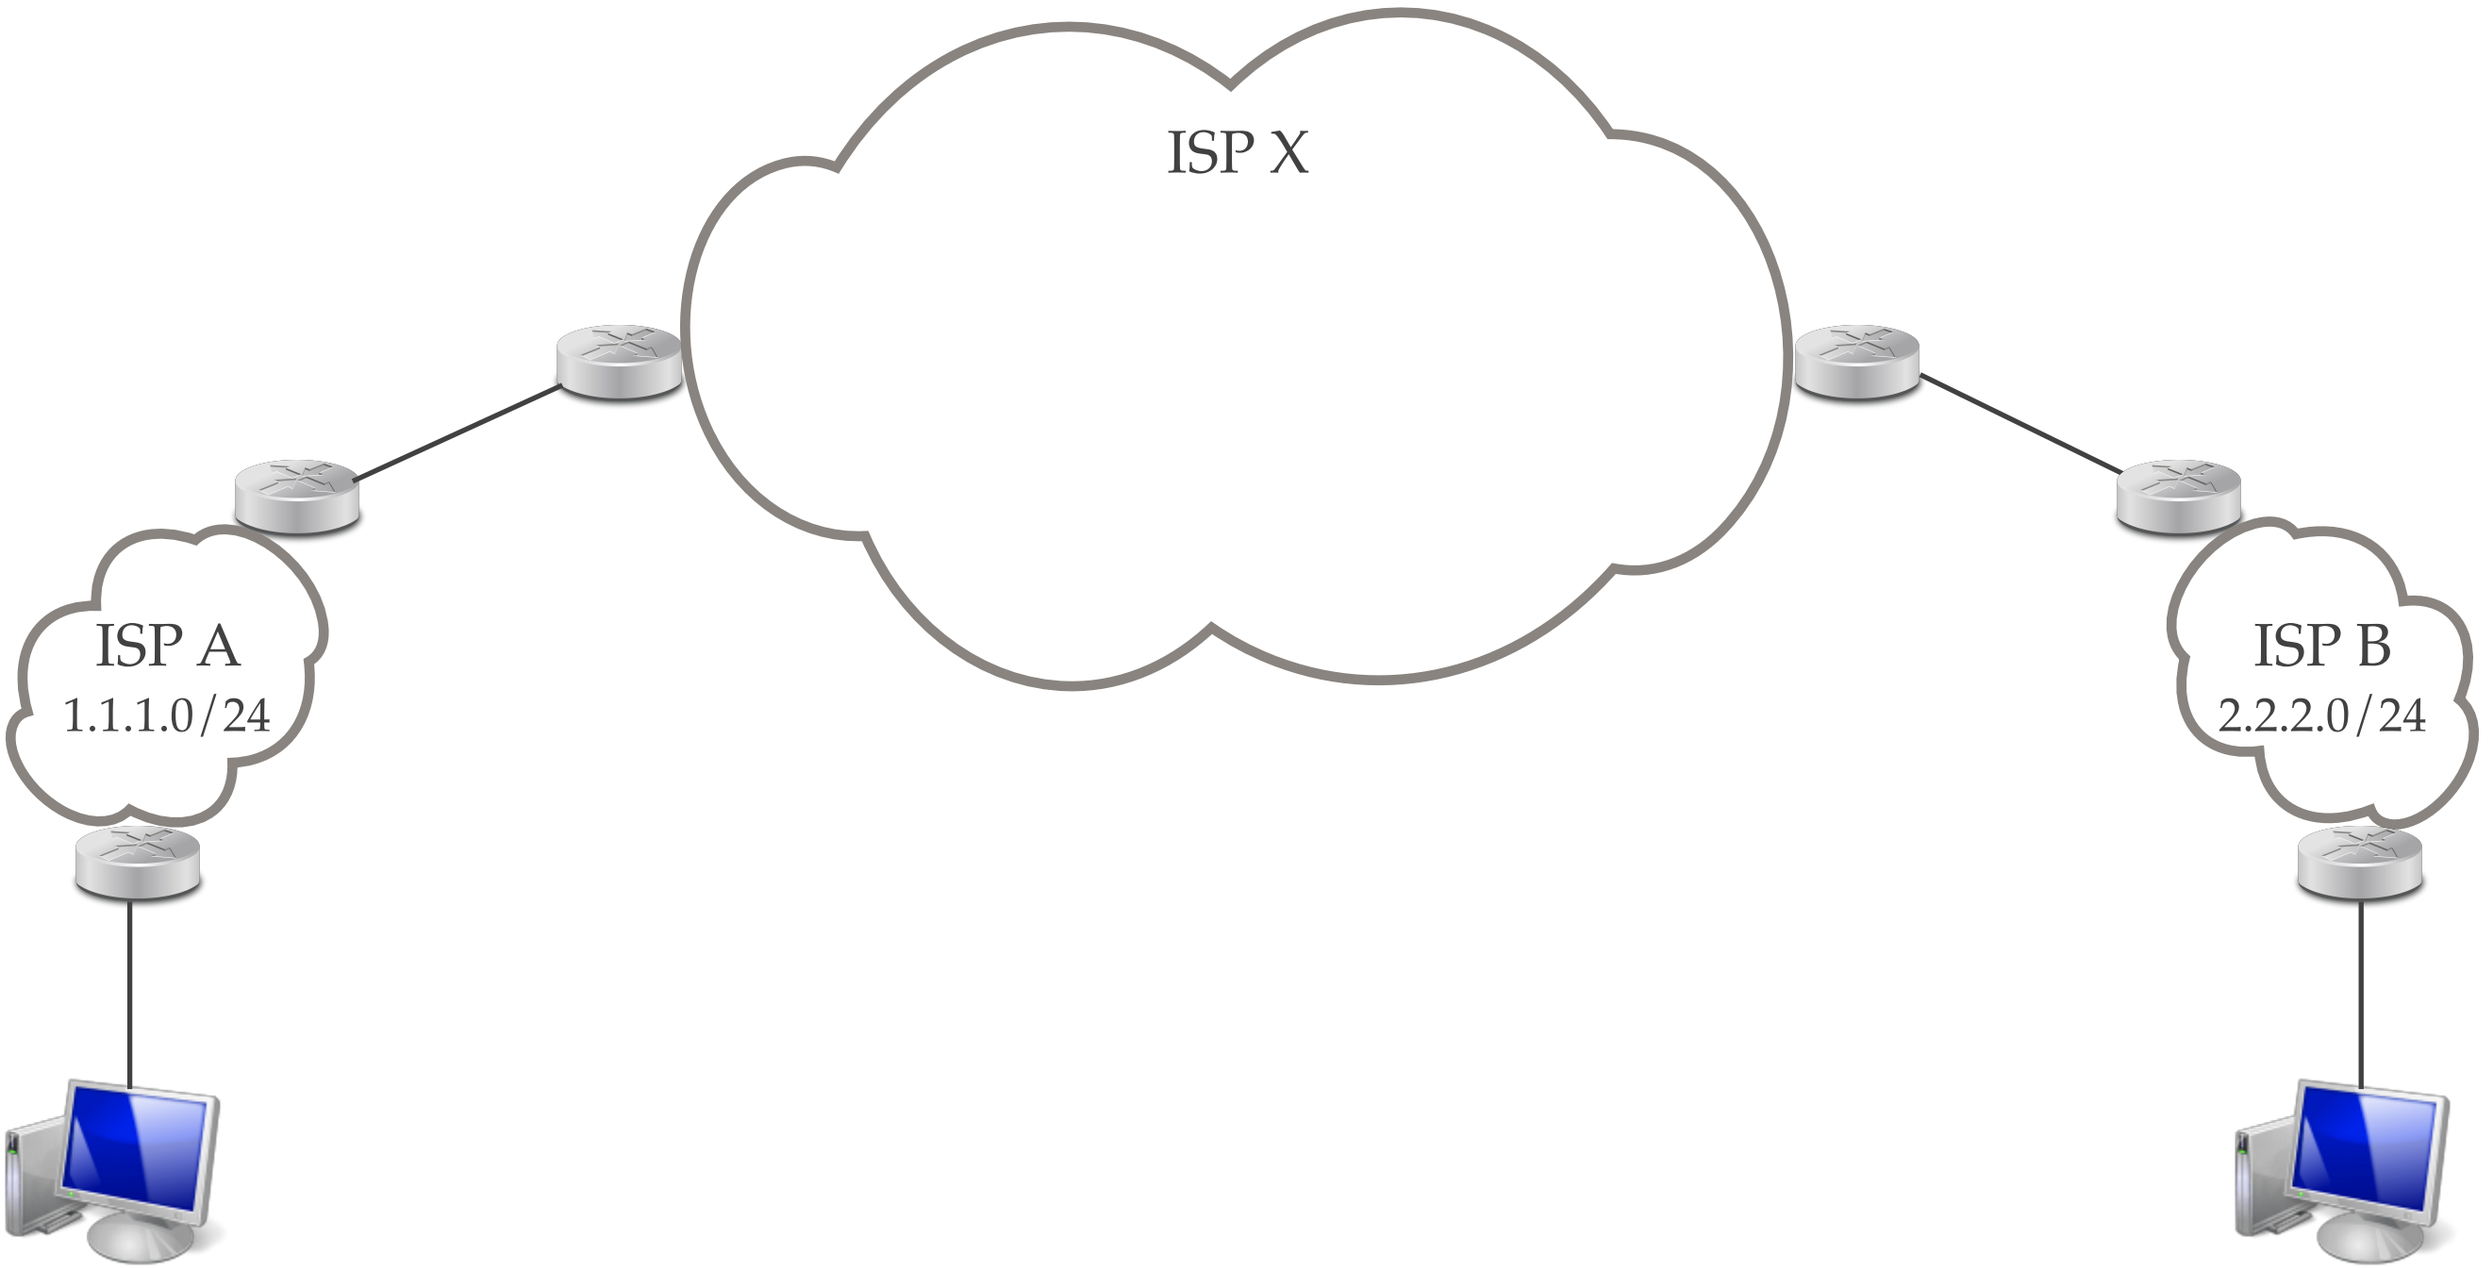

MPLS Network

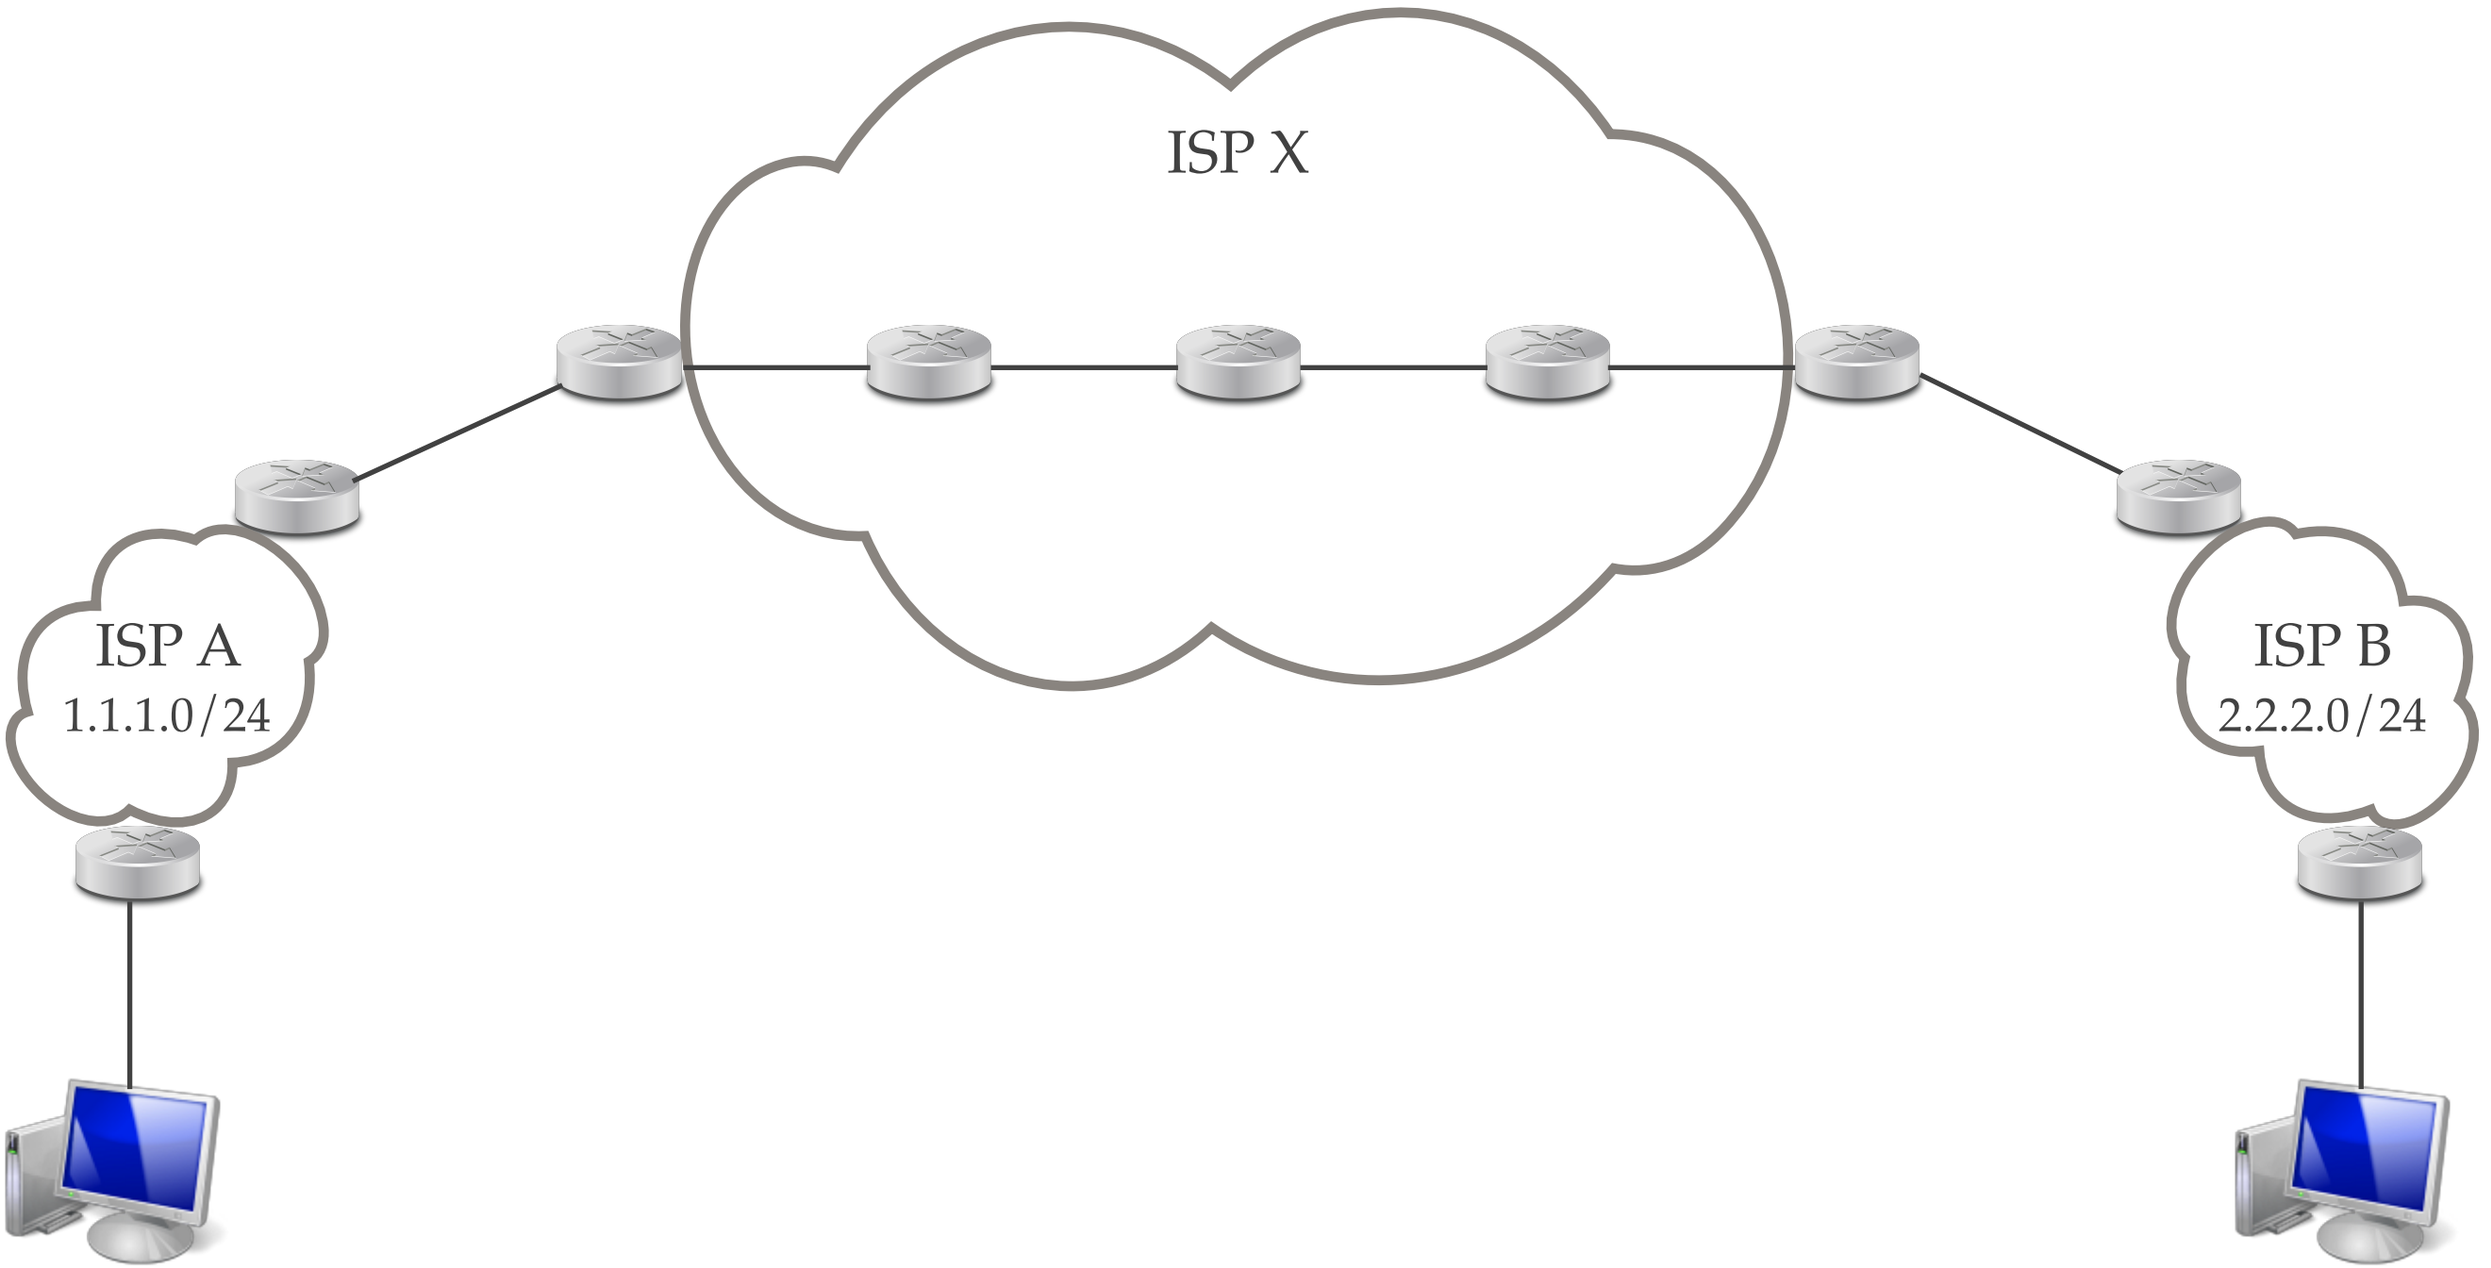

MPLS Network

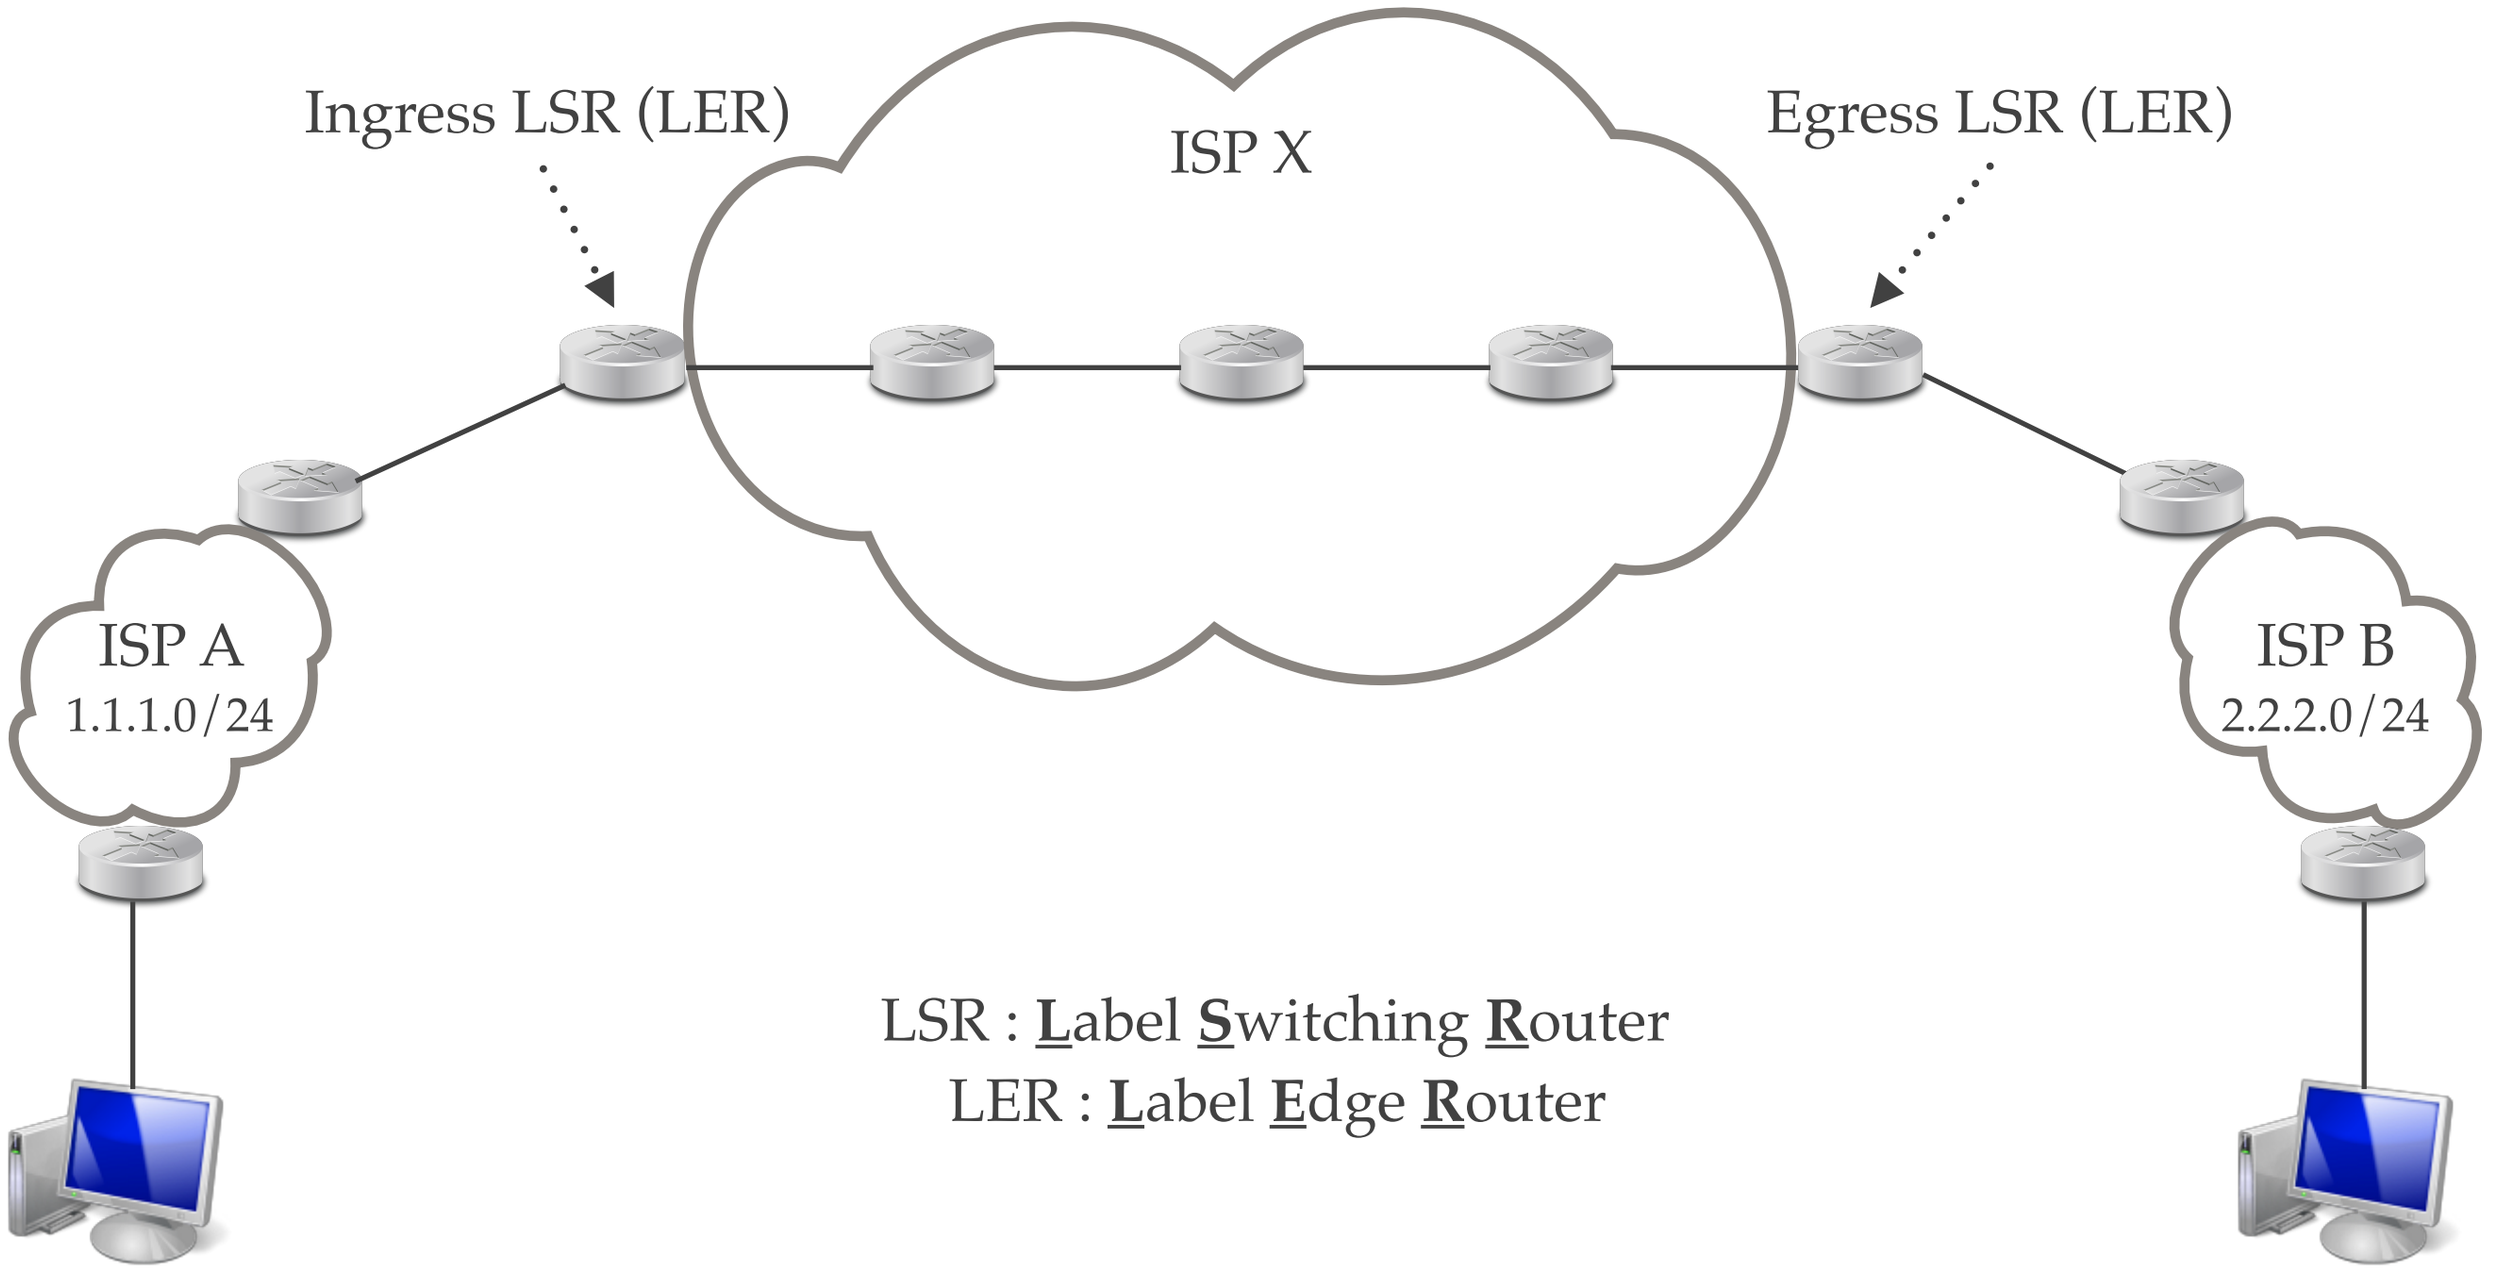

MPLS Network

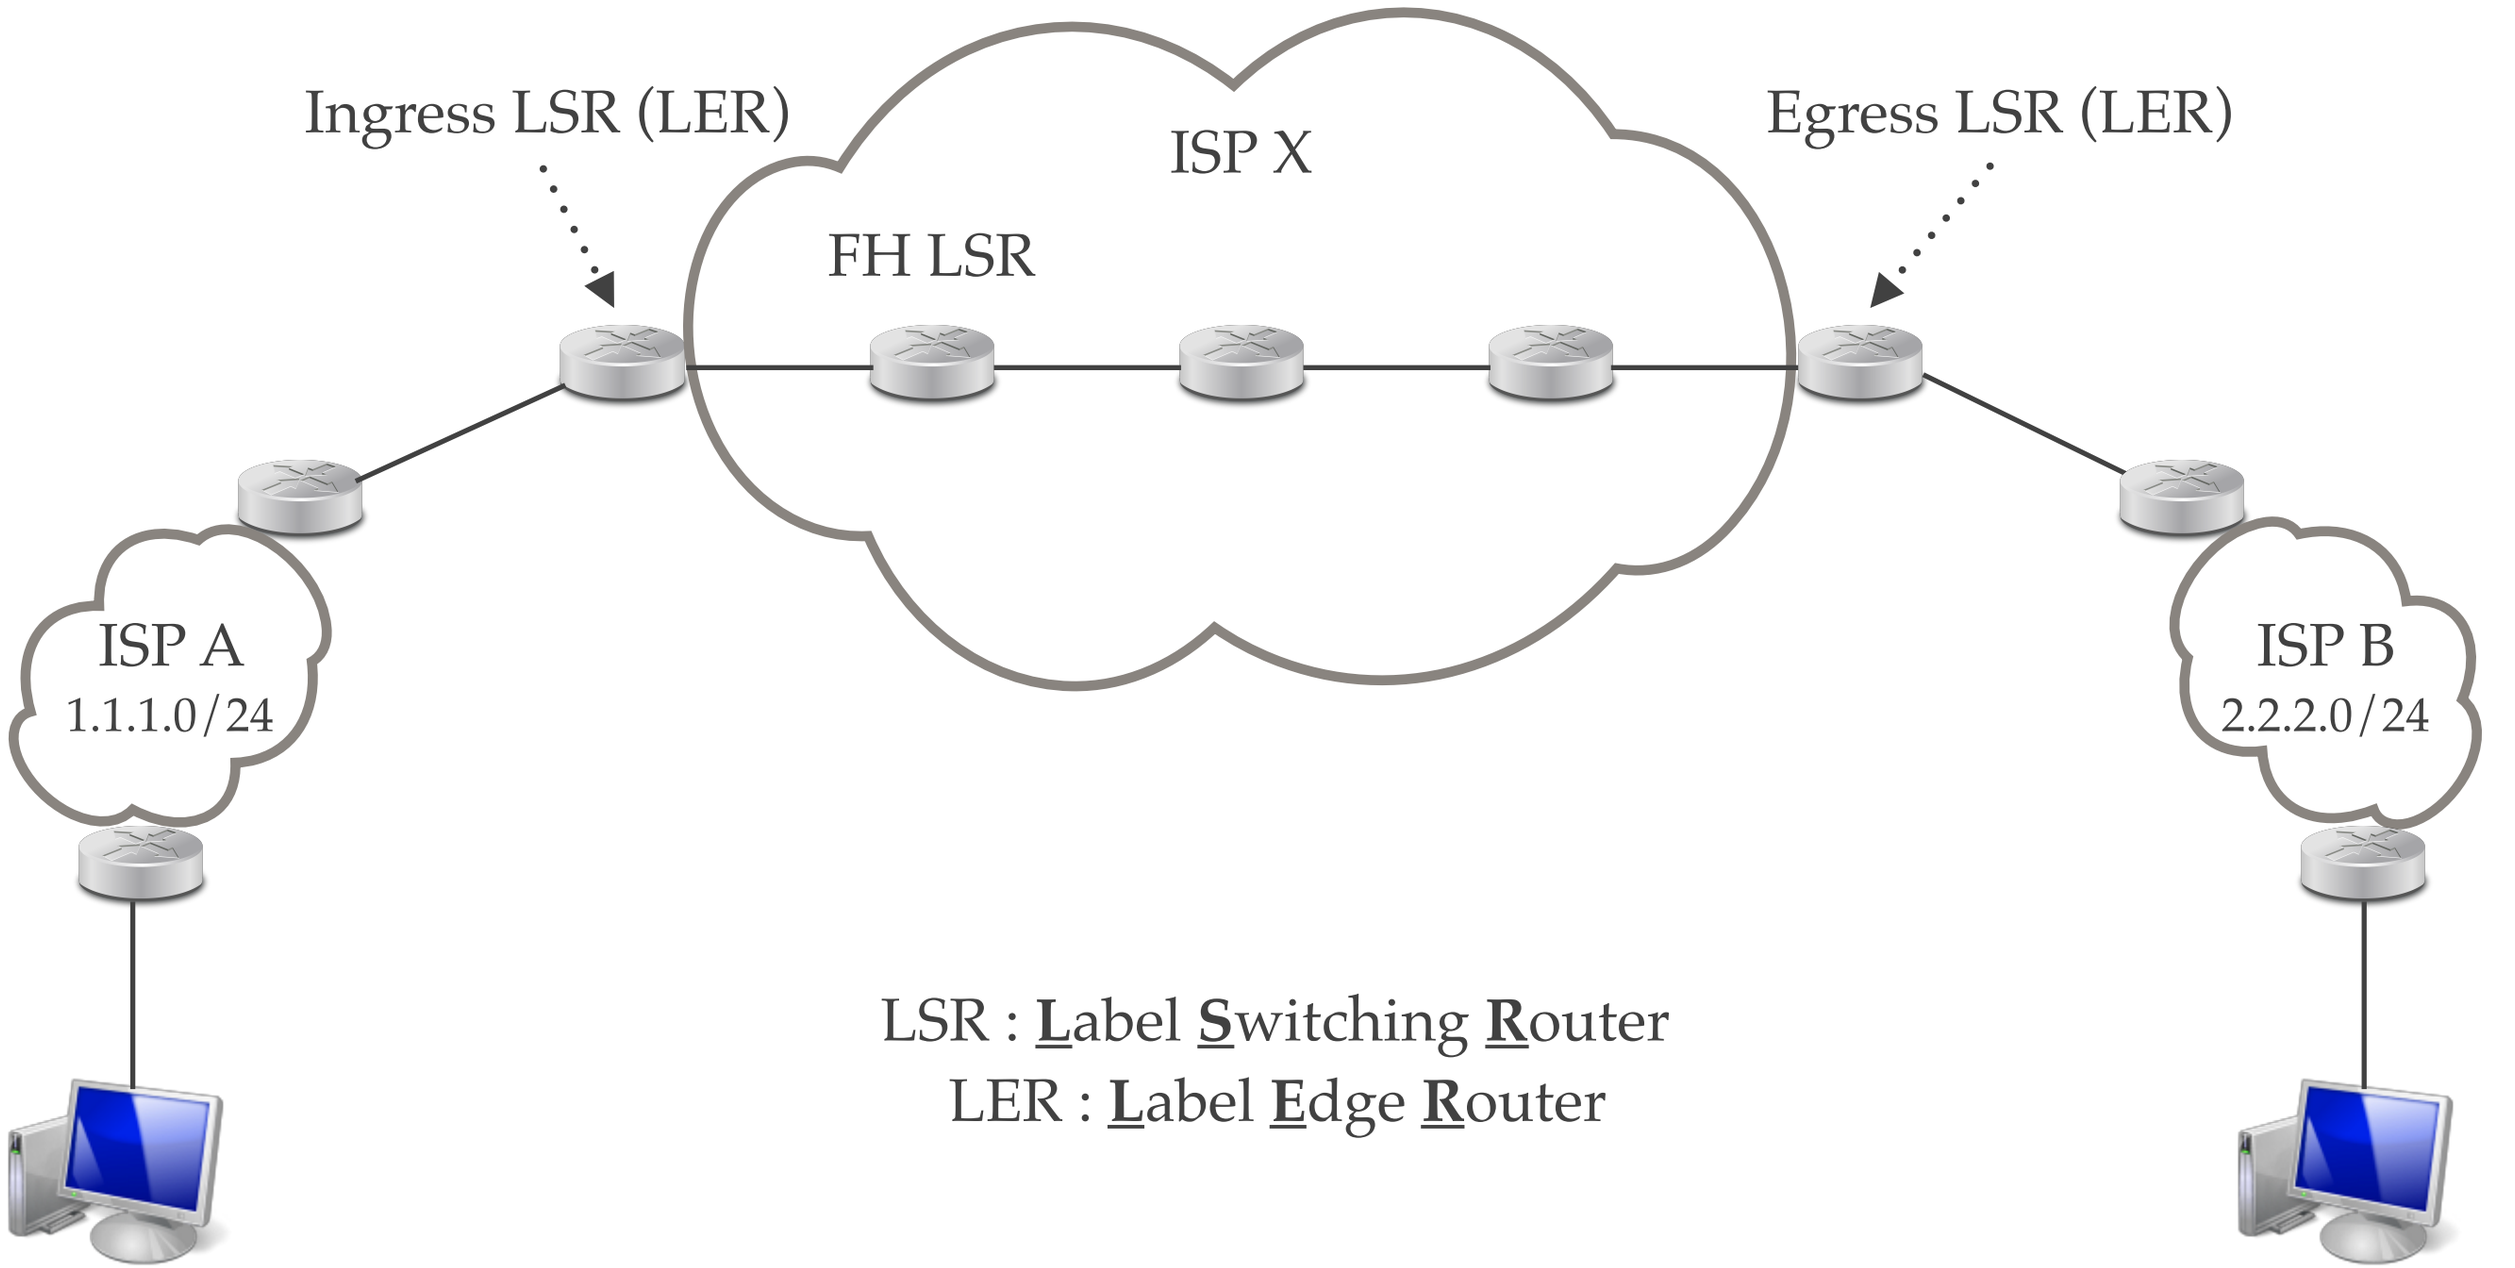

MPLS Network

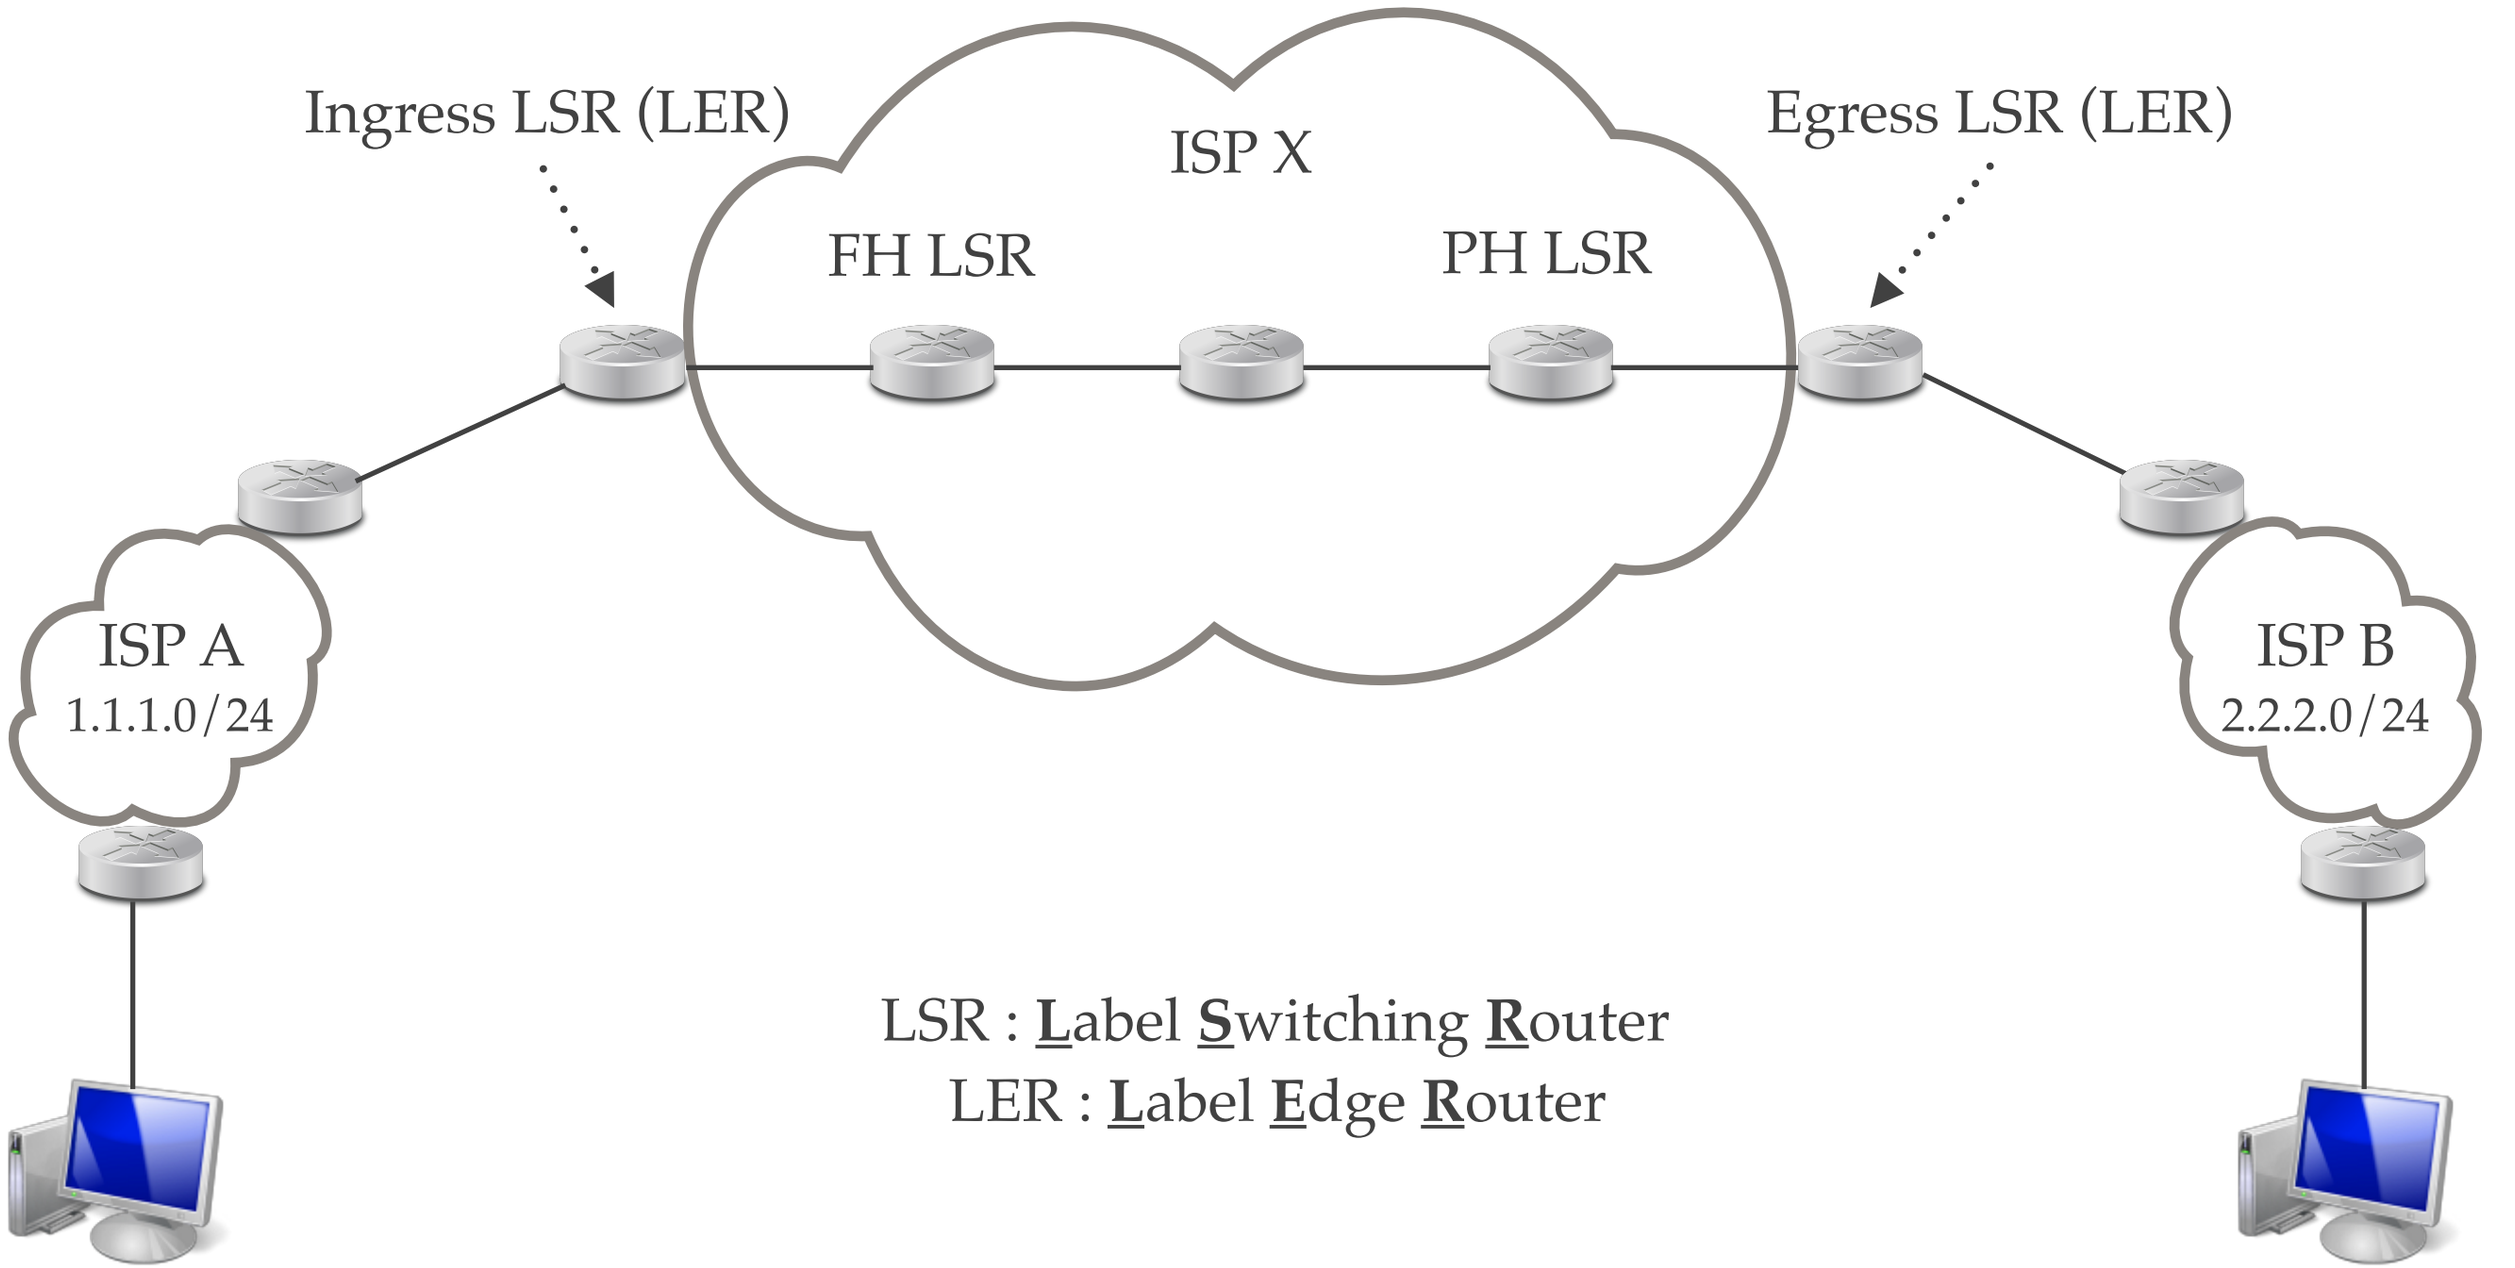

MPLS Network

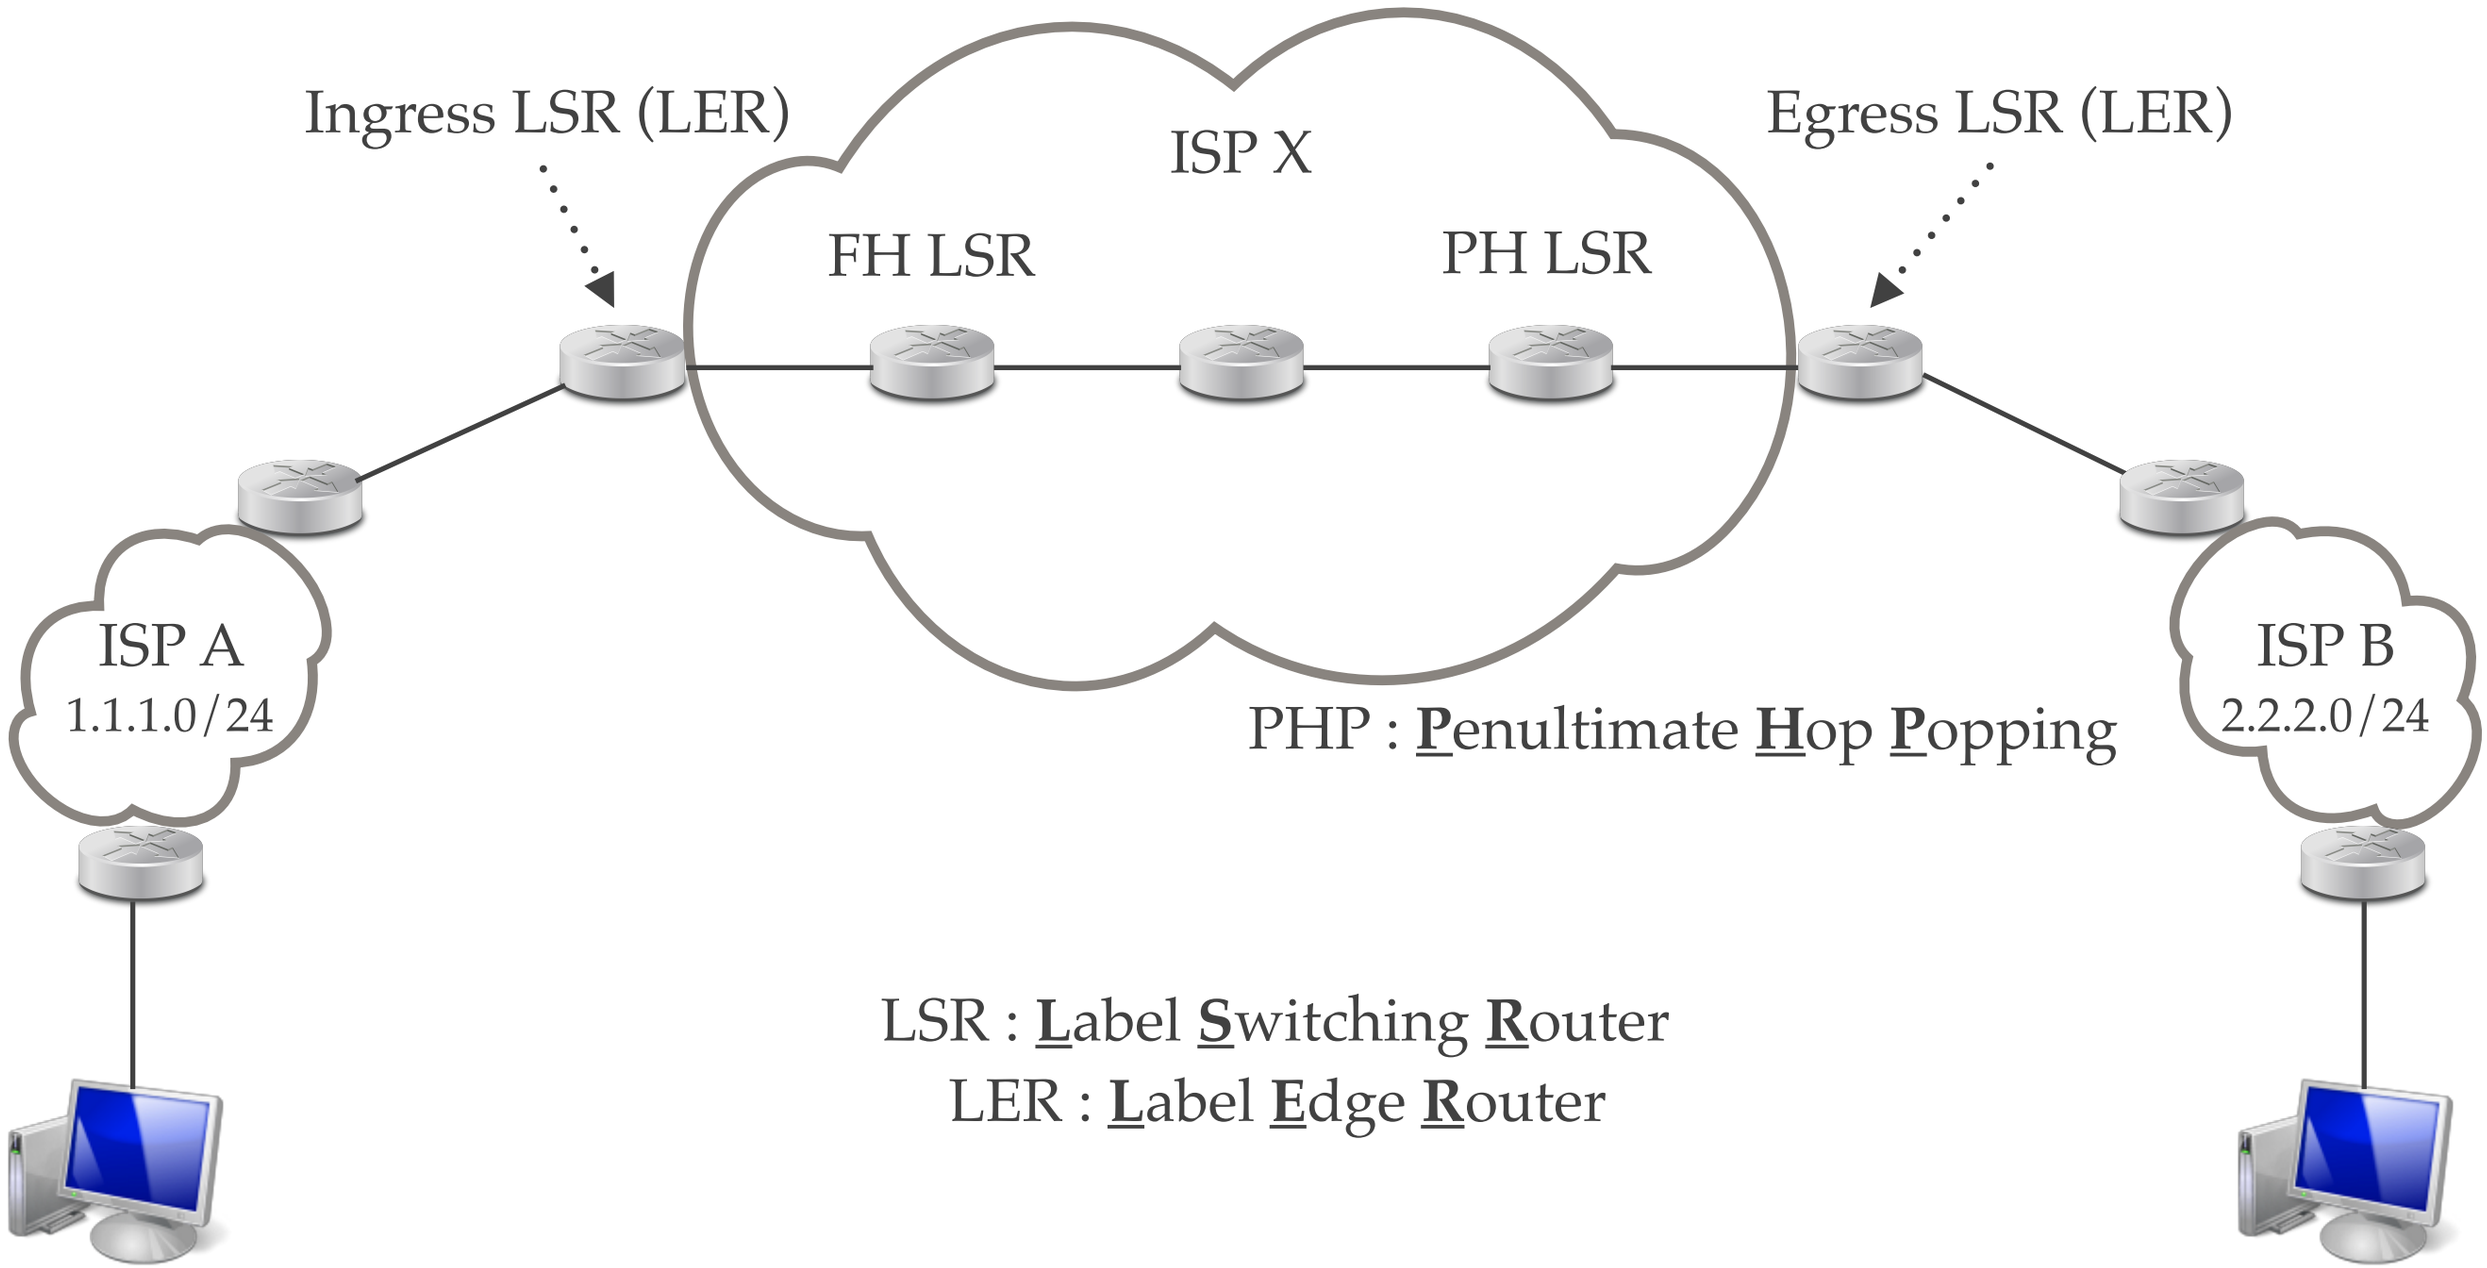

MPLS Network

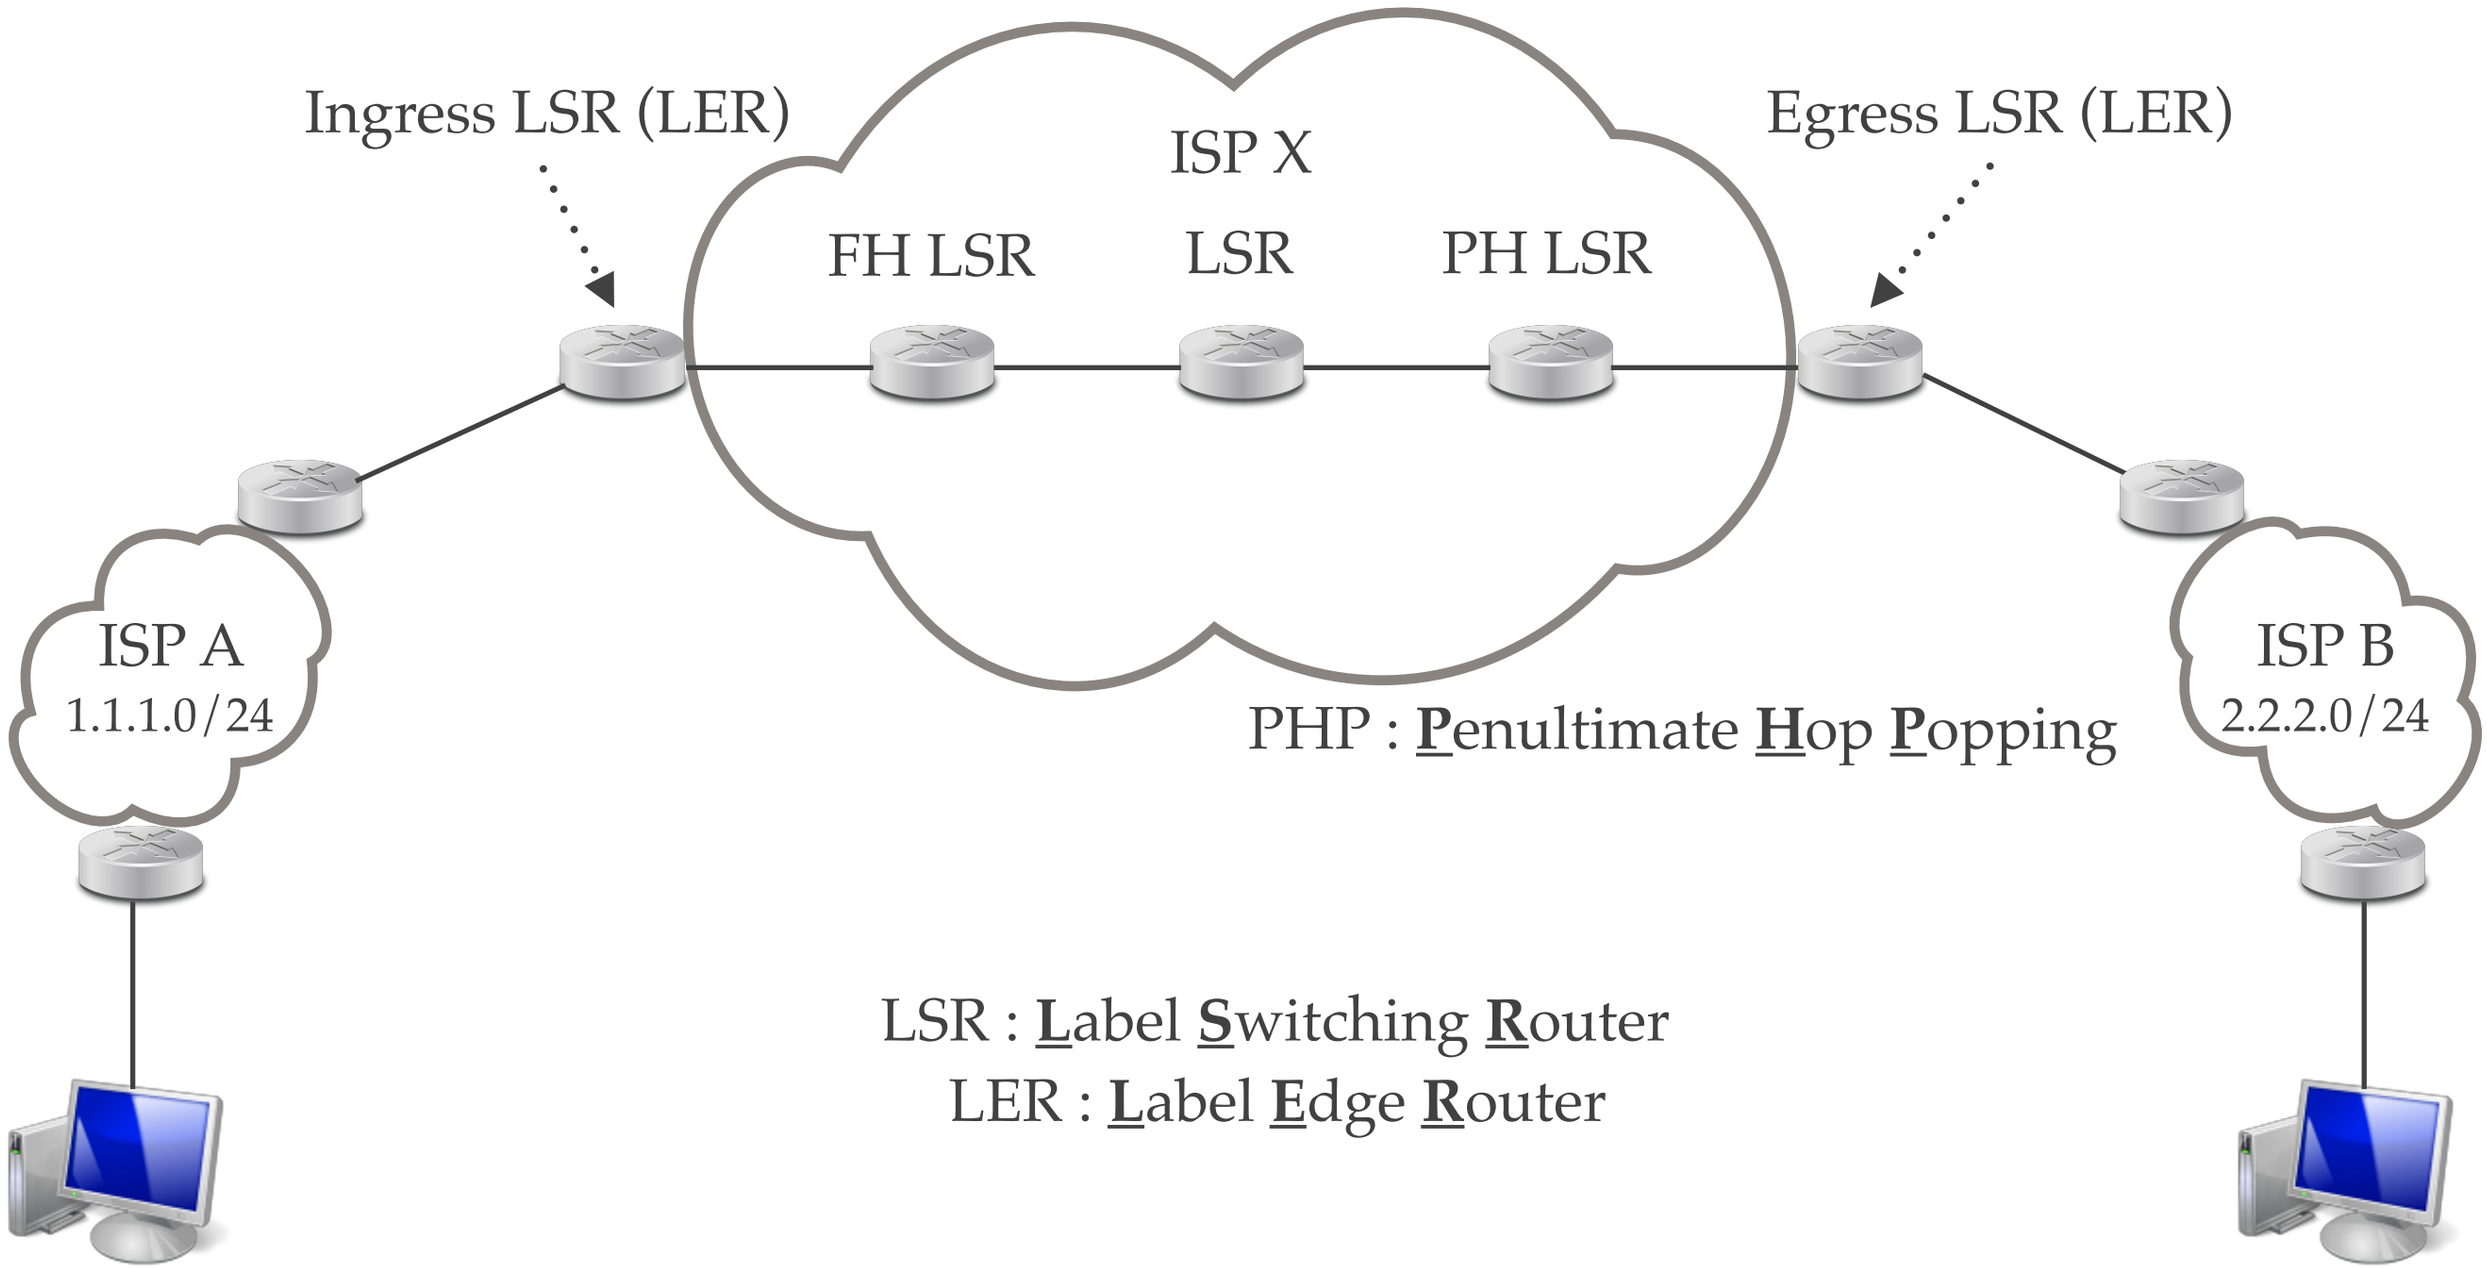

MPLS Network

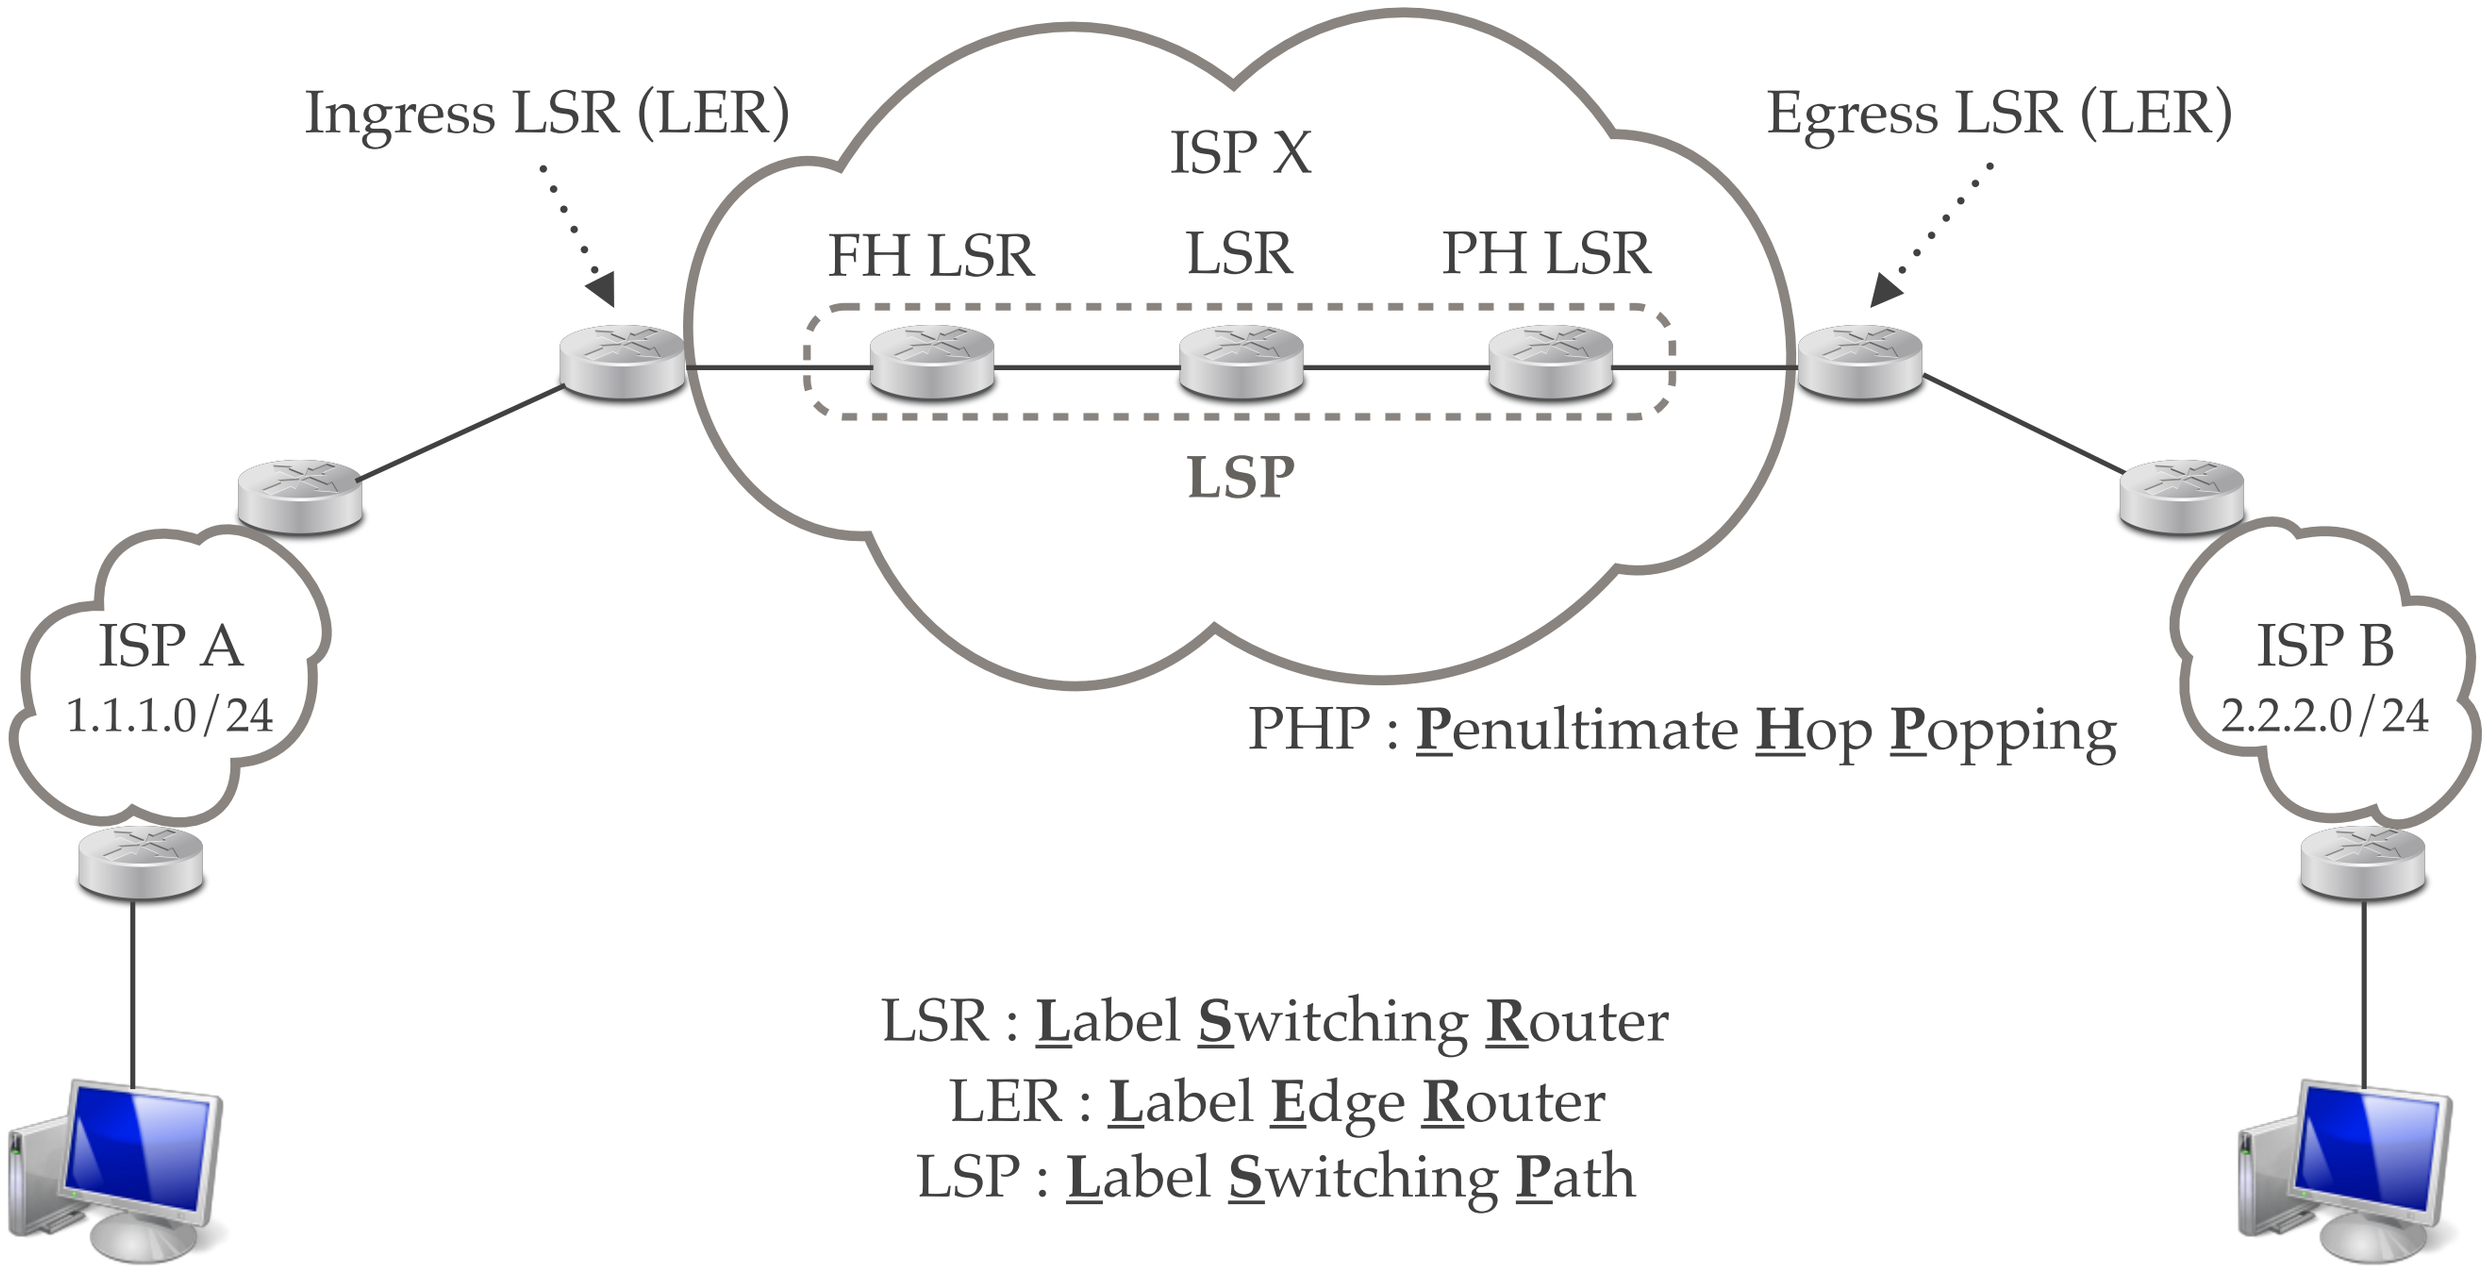

MPLS Network

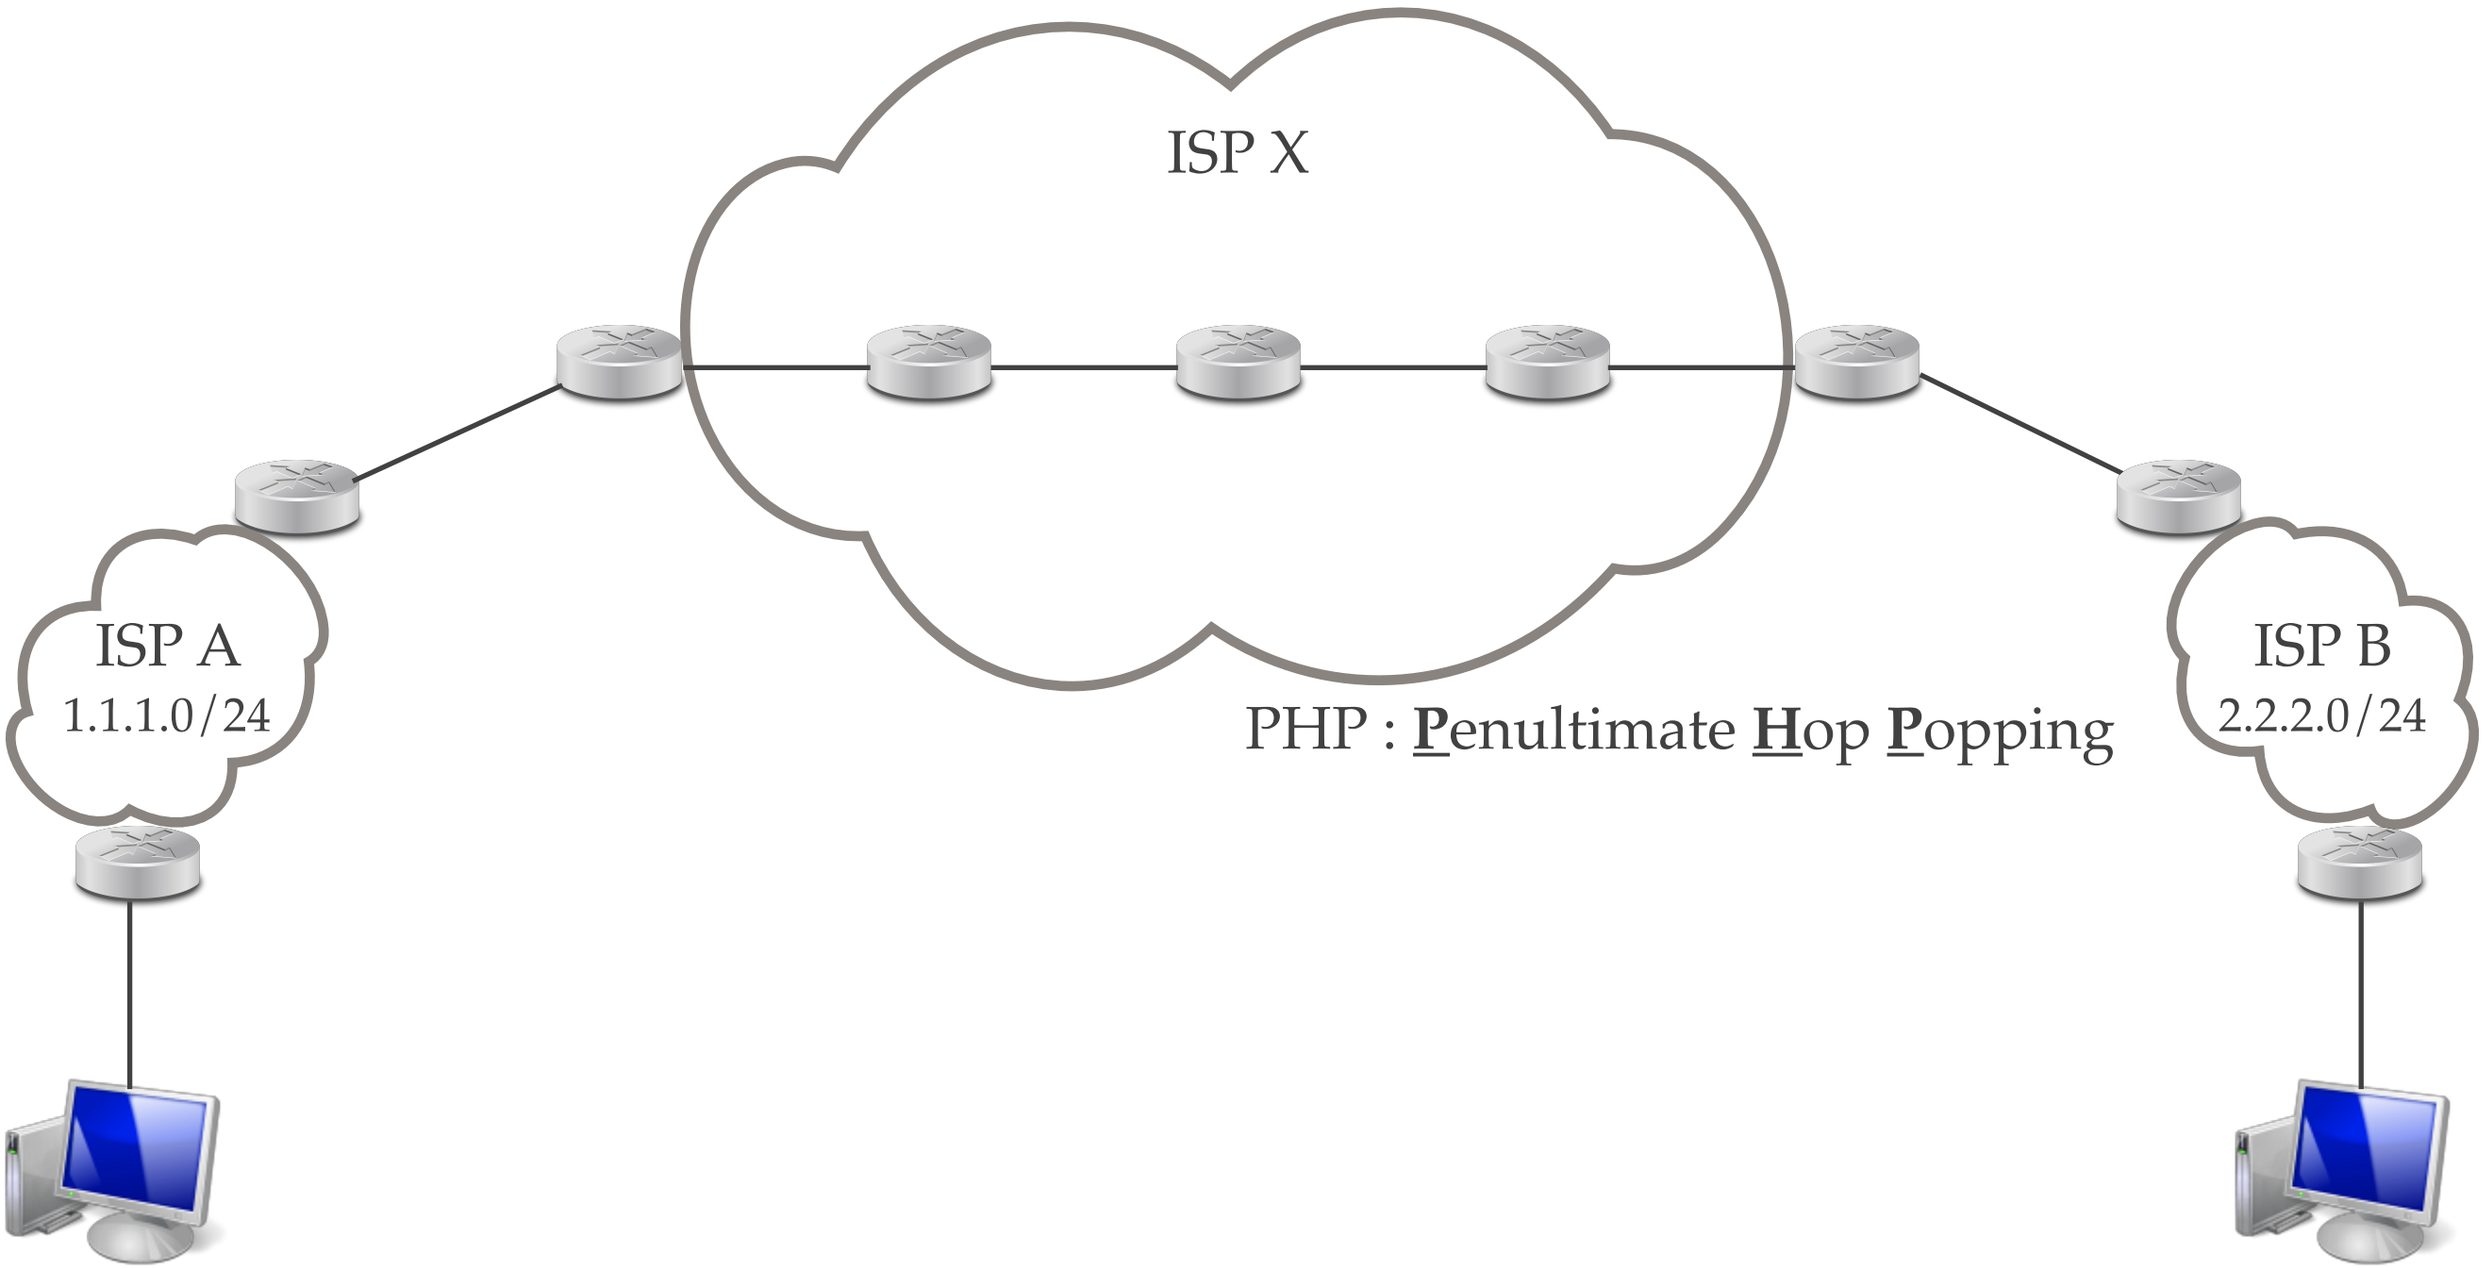

MPLS Network

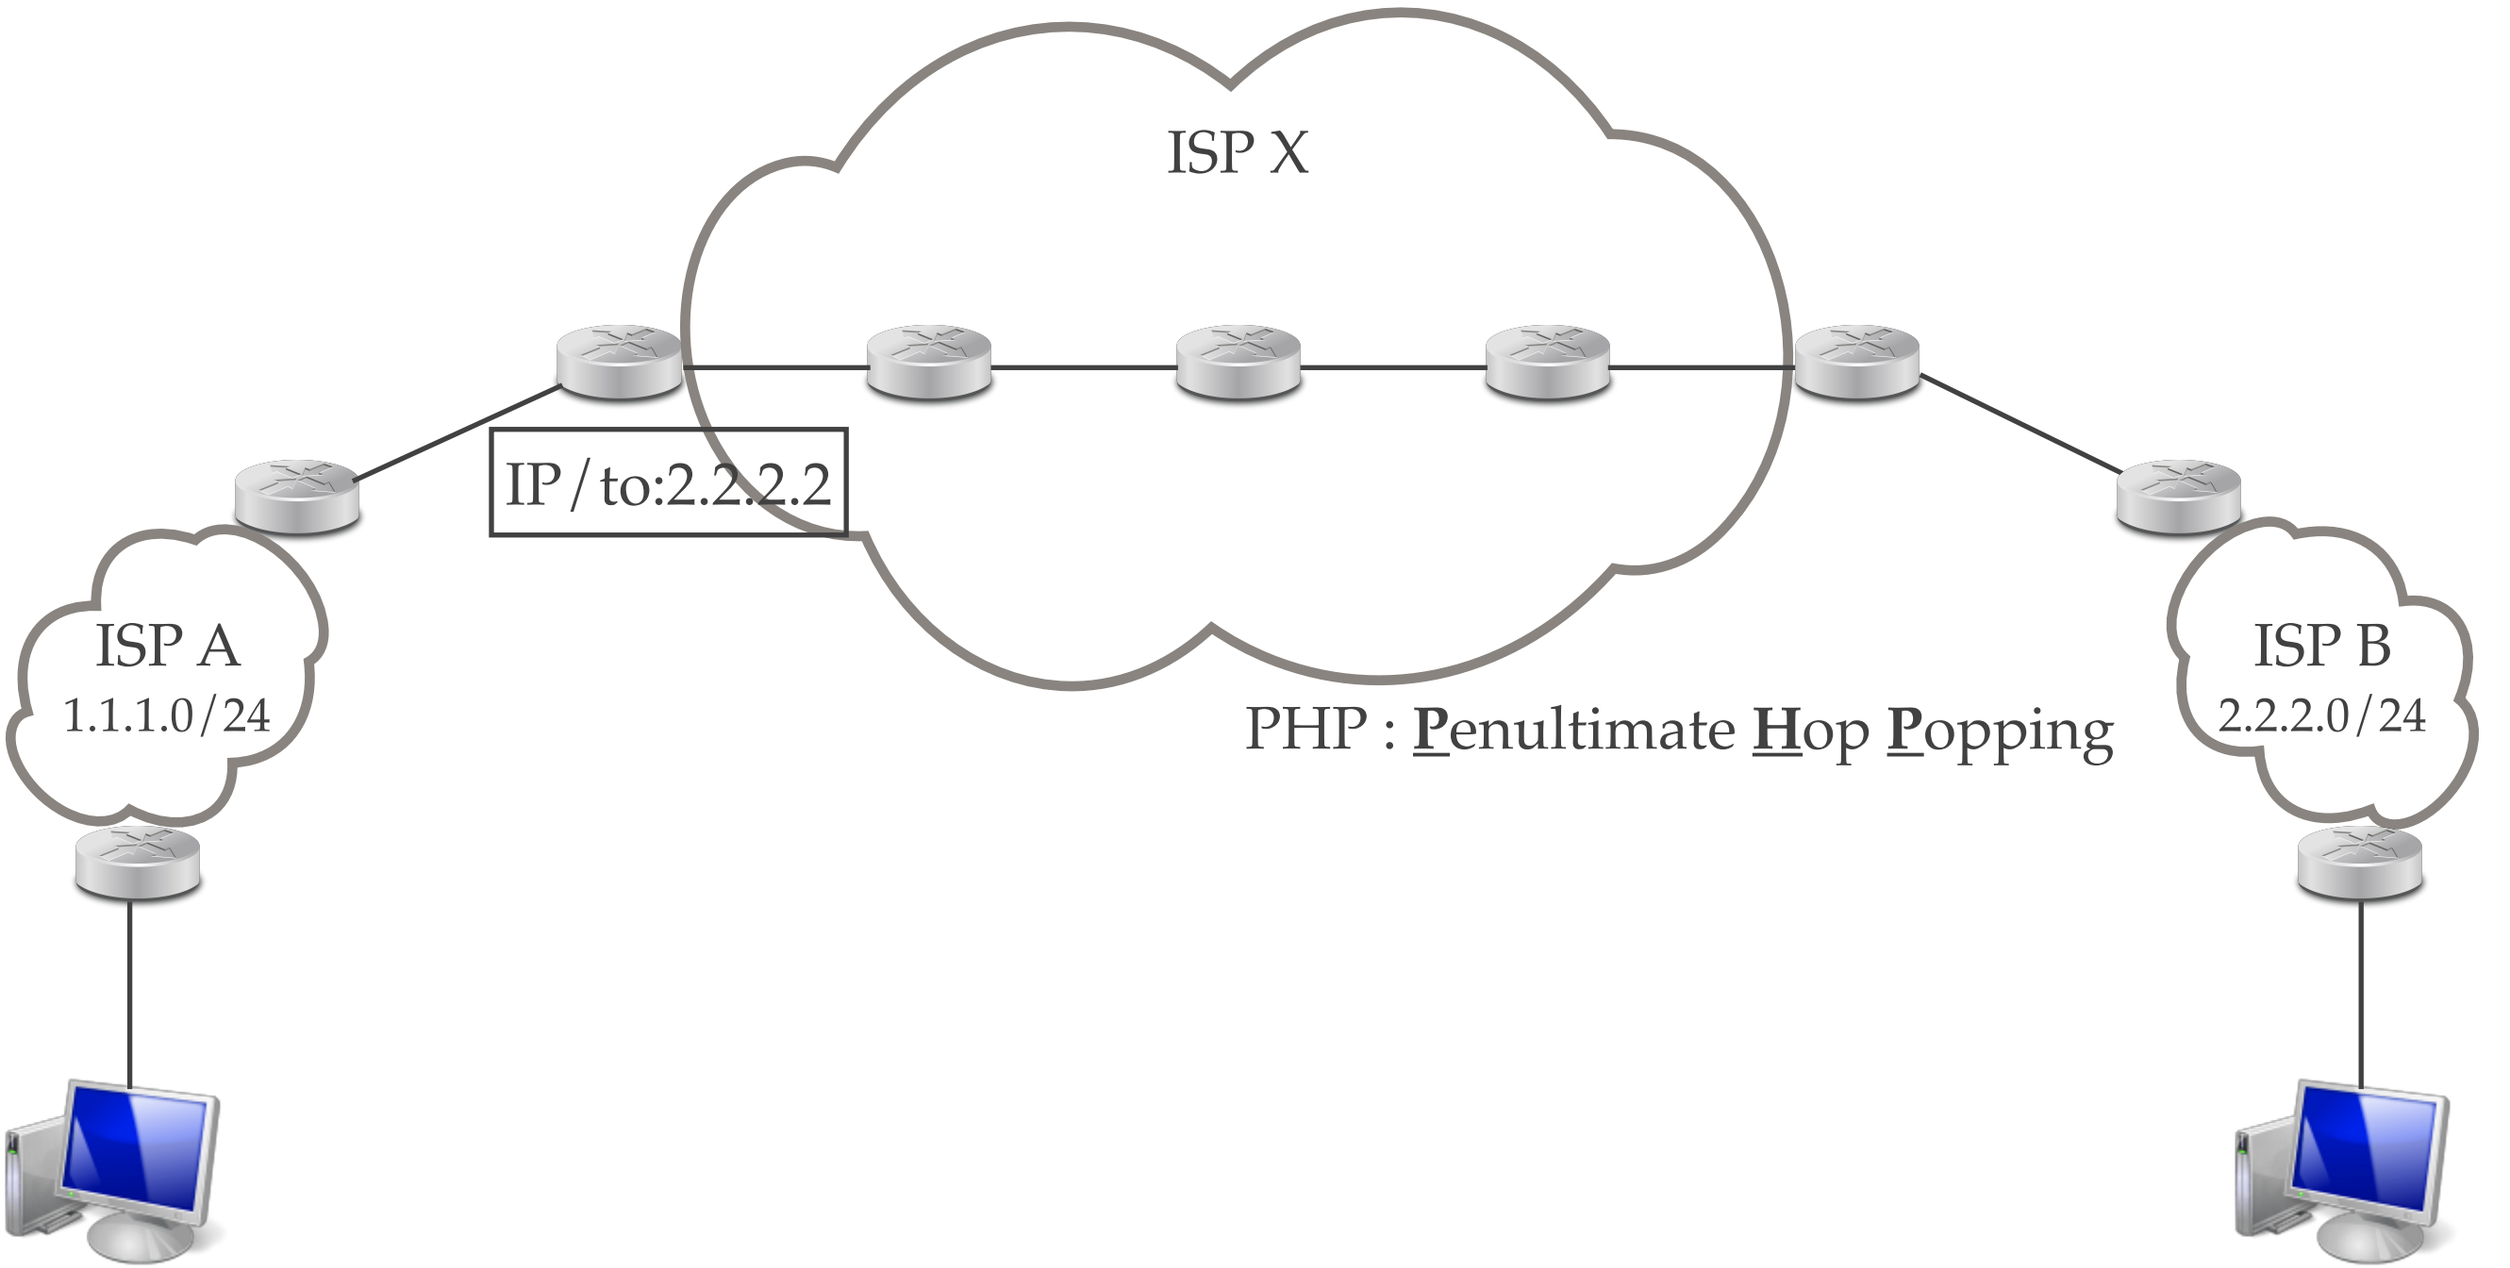

MPLS Network

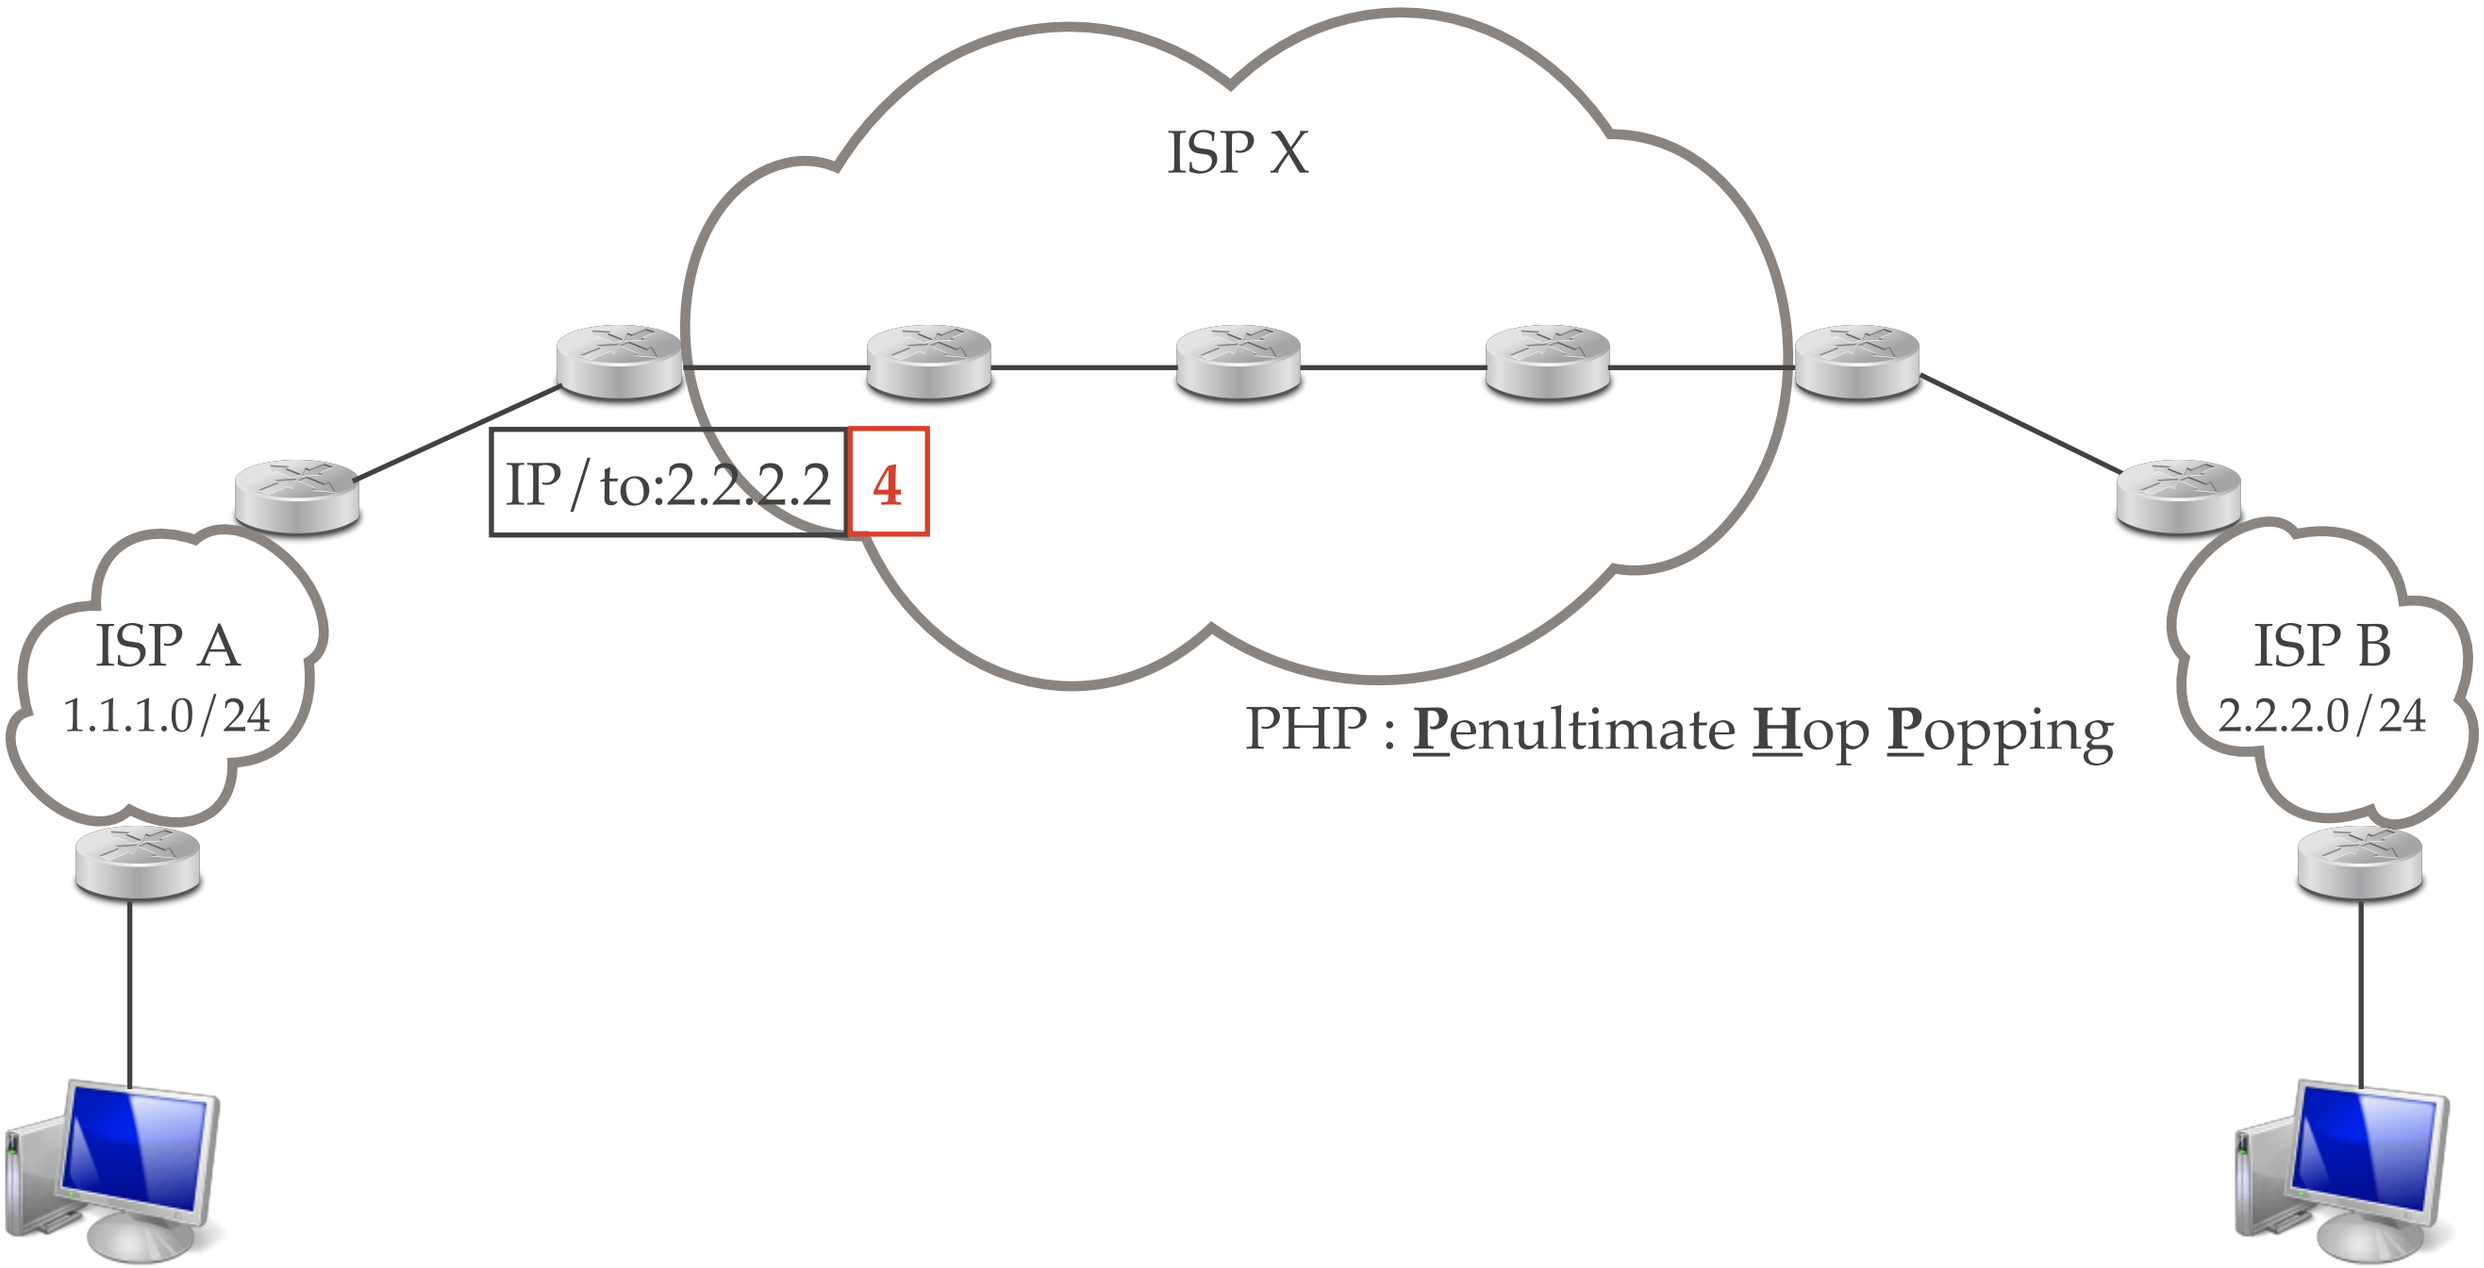

MPLS Network

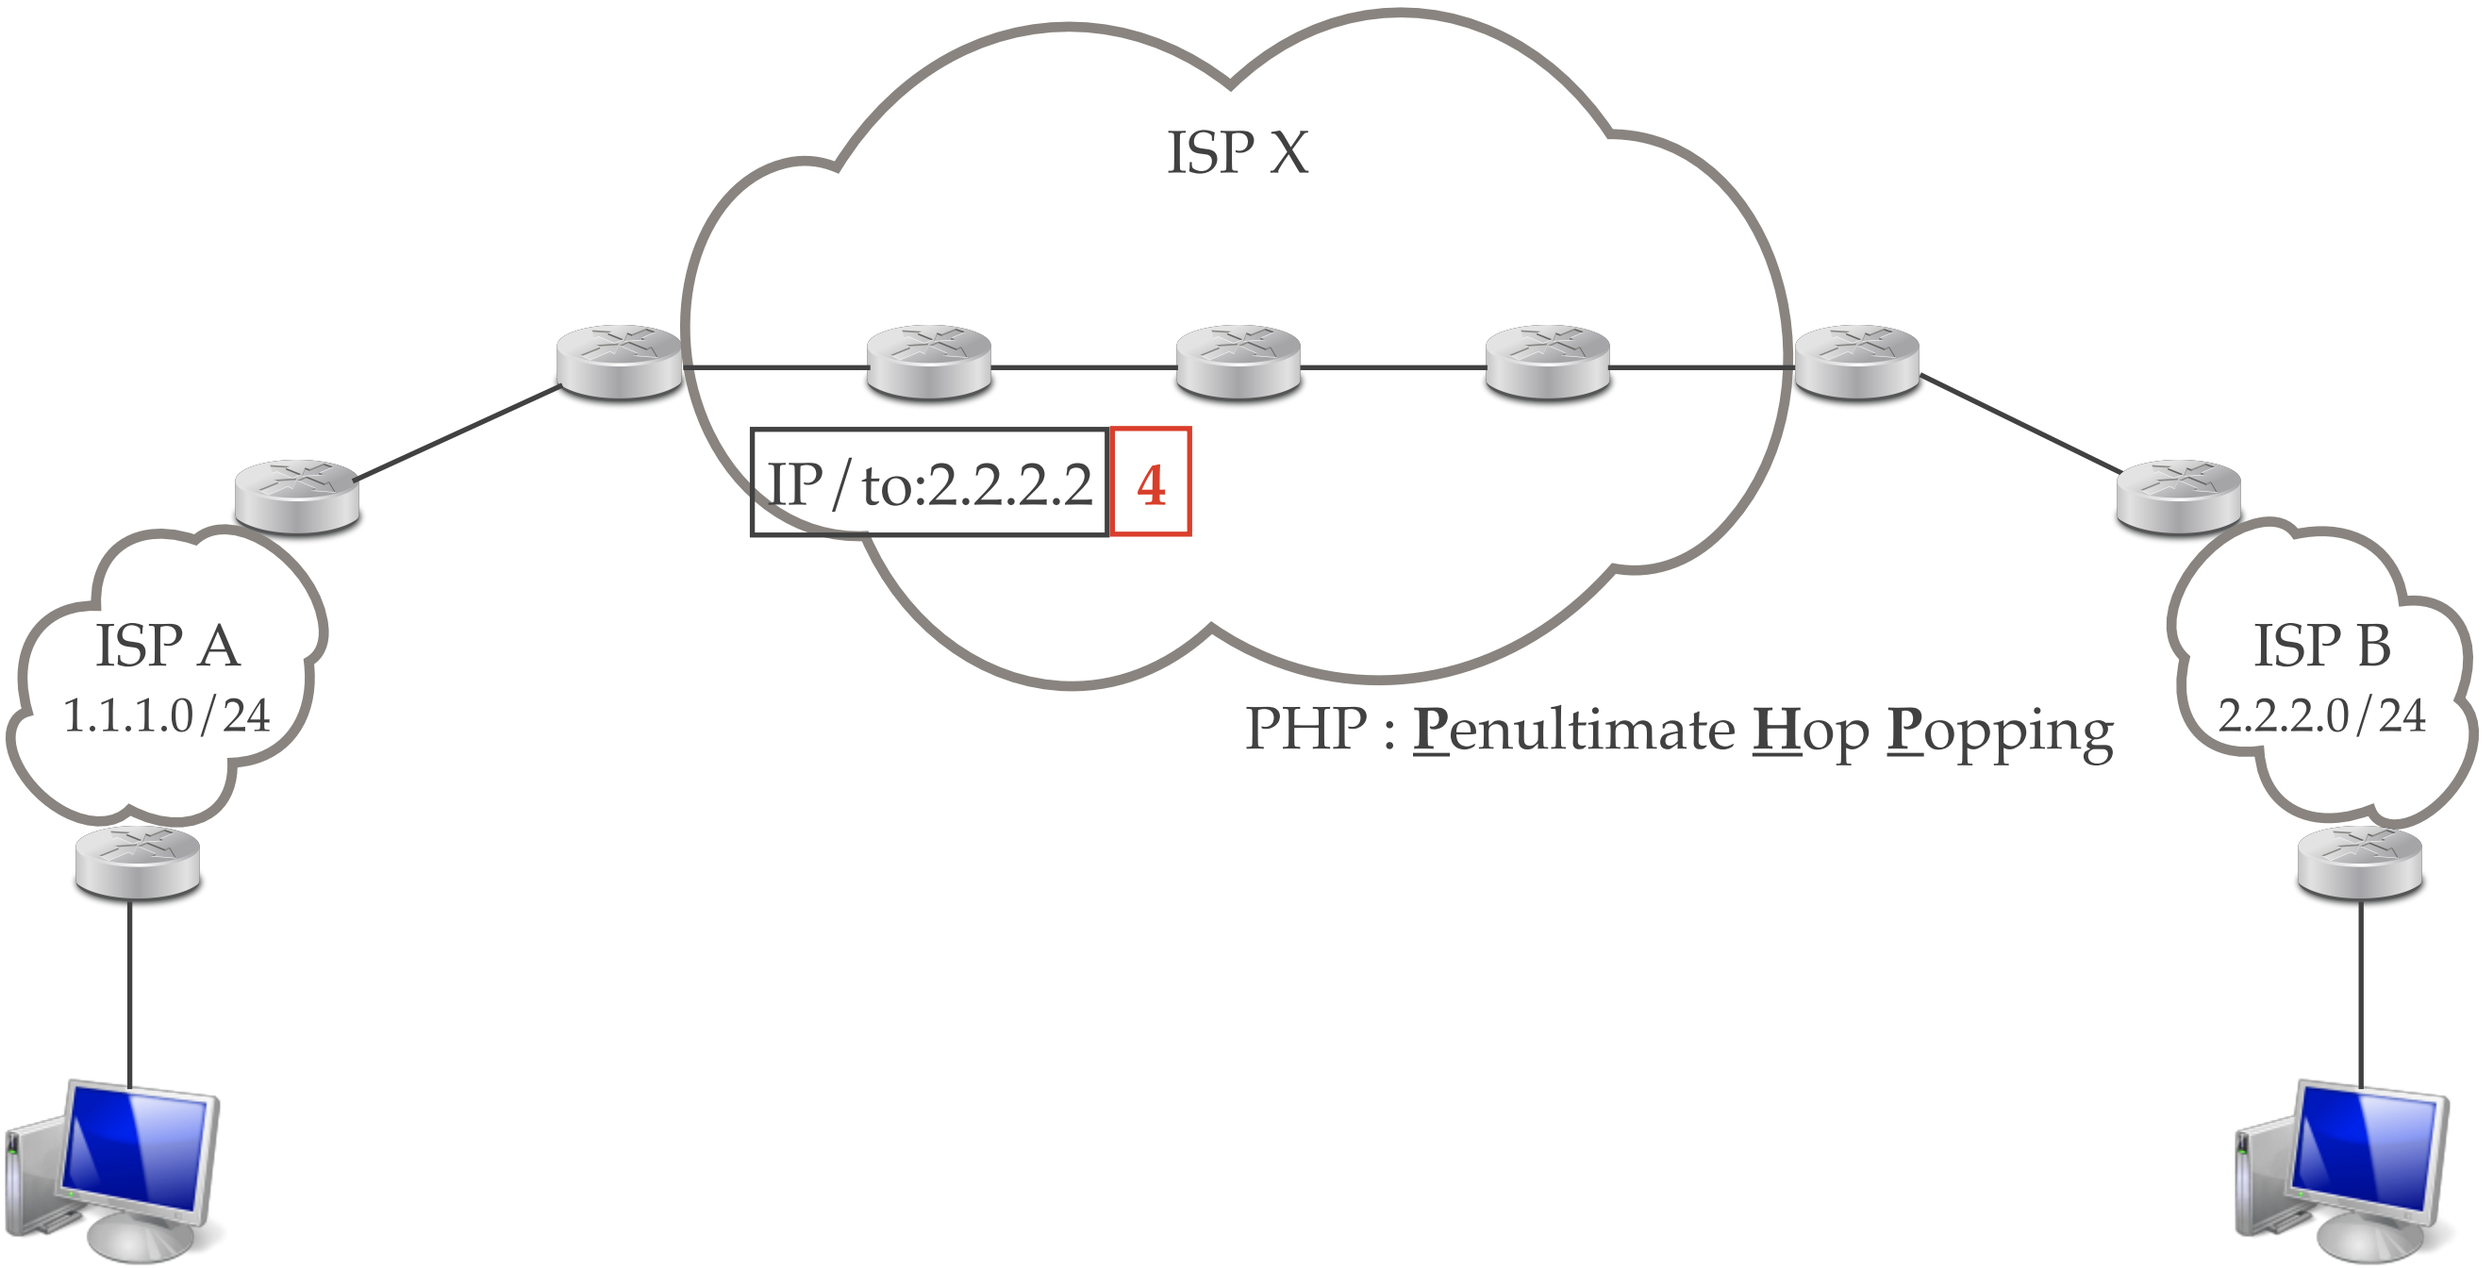

MPLS Network

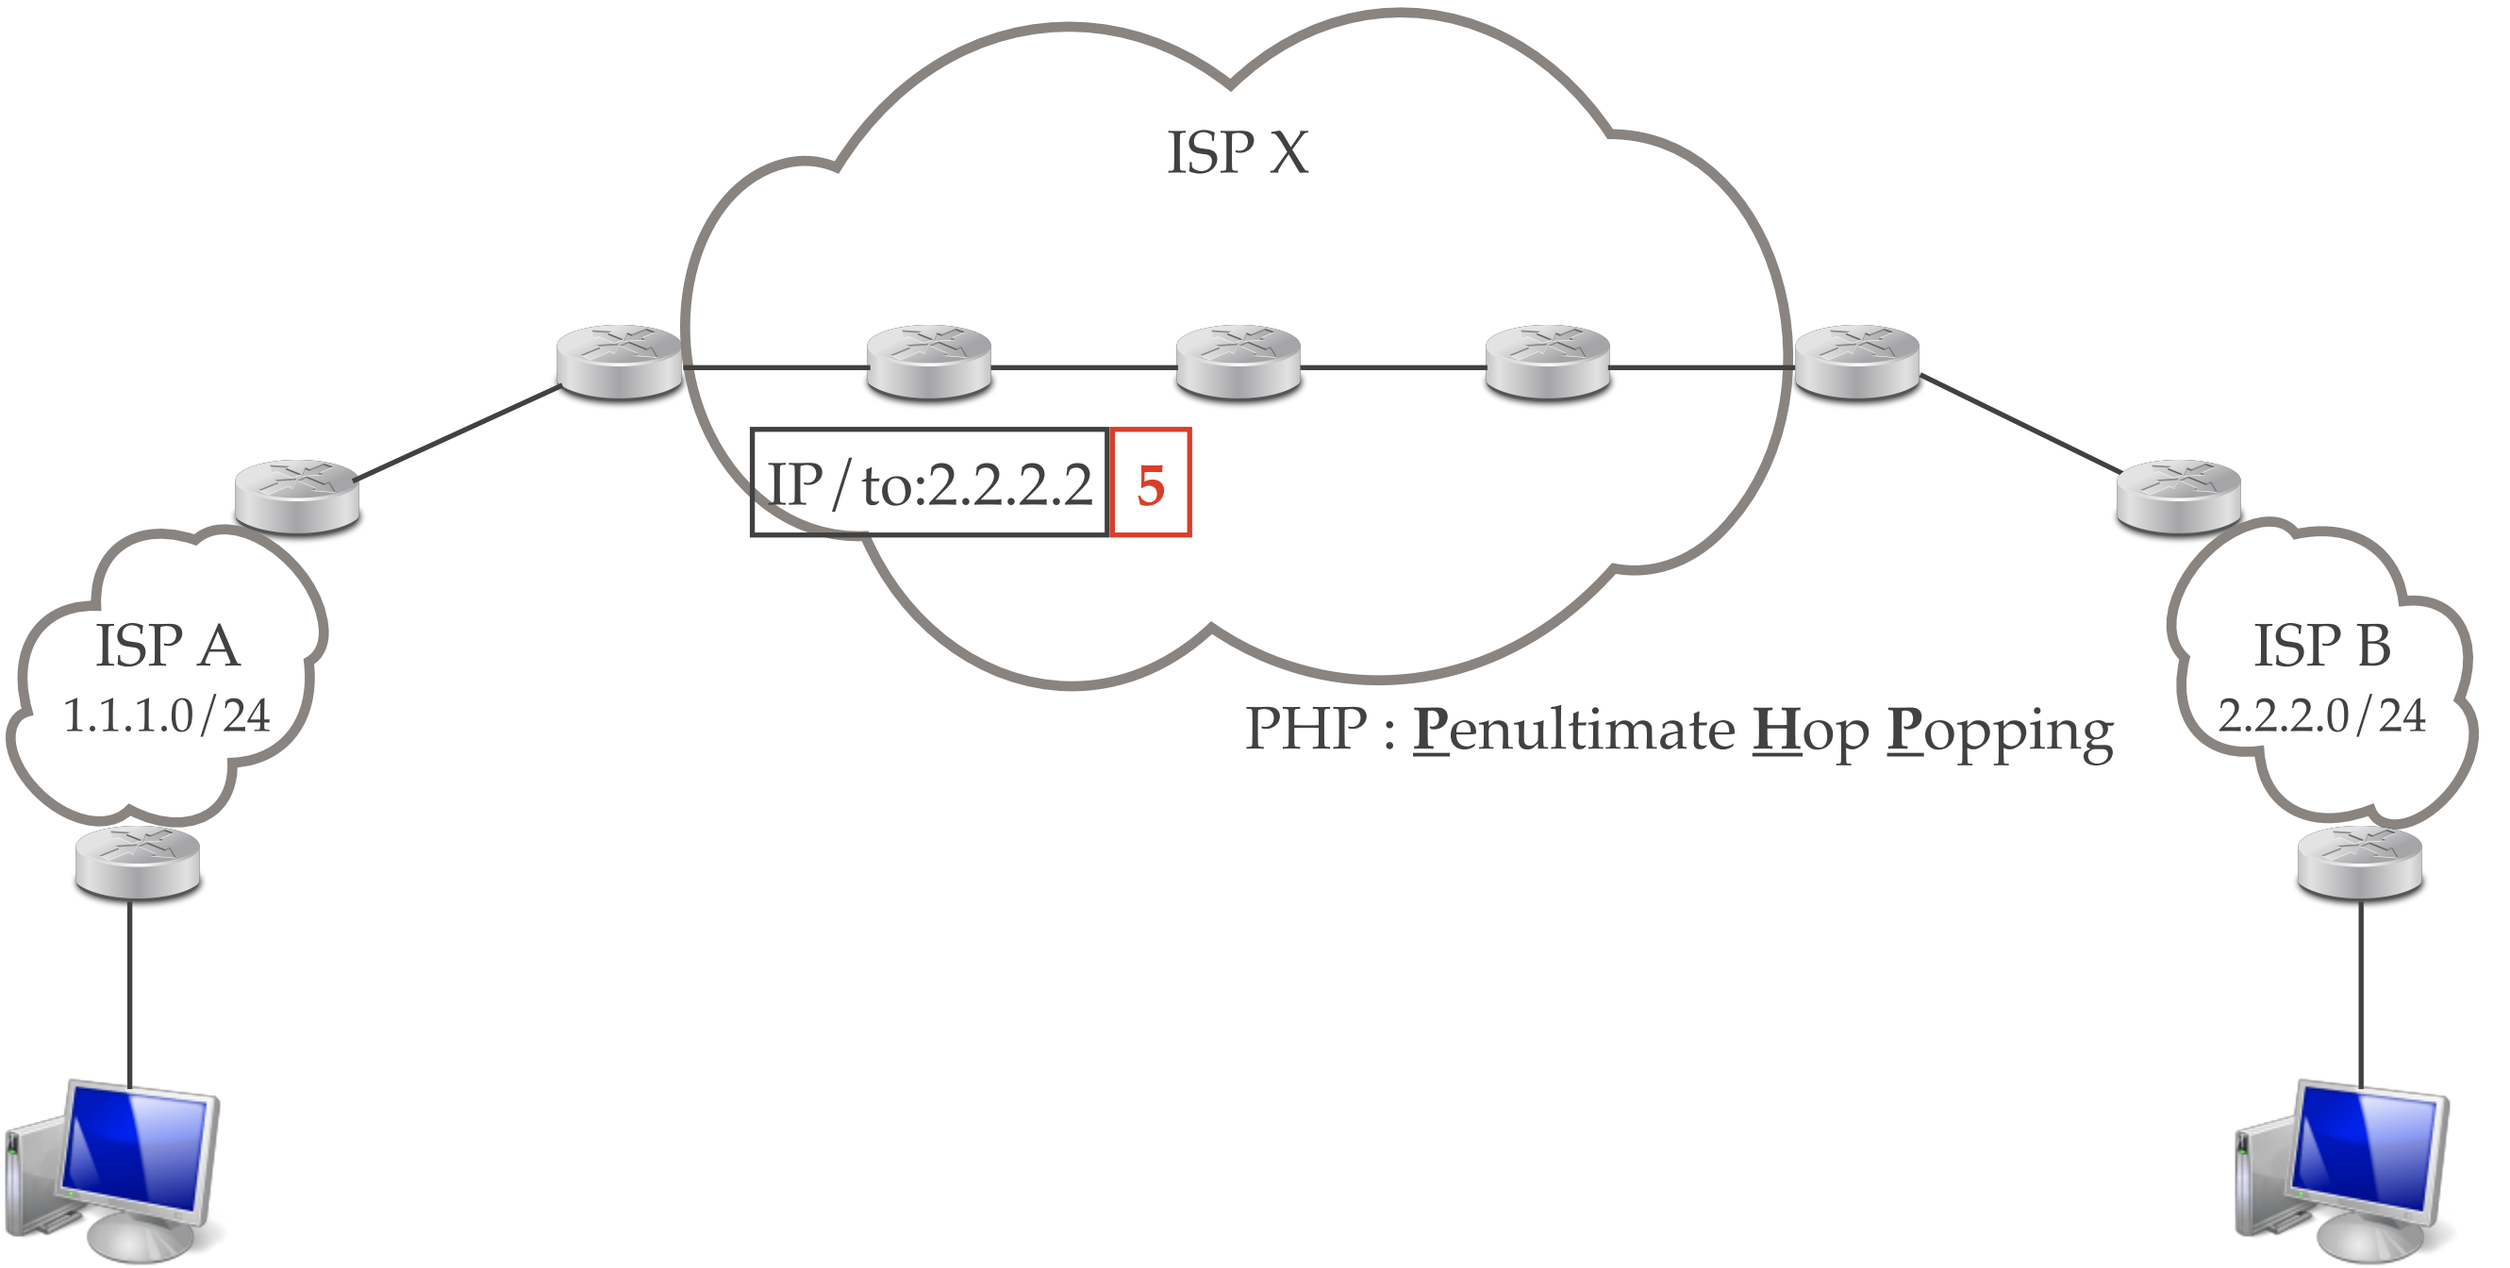

MPLS Network

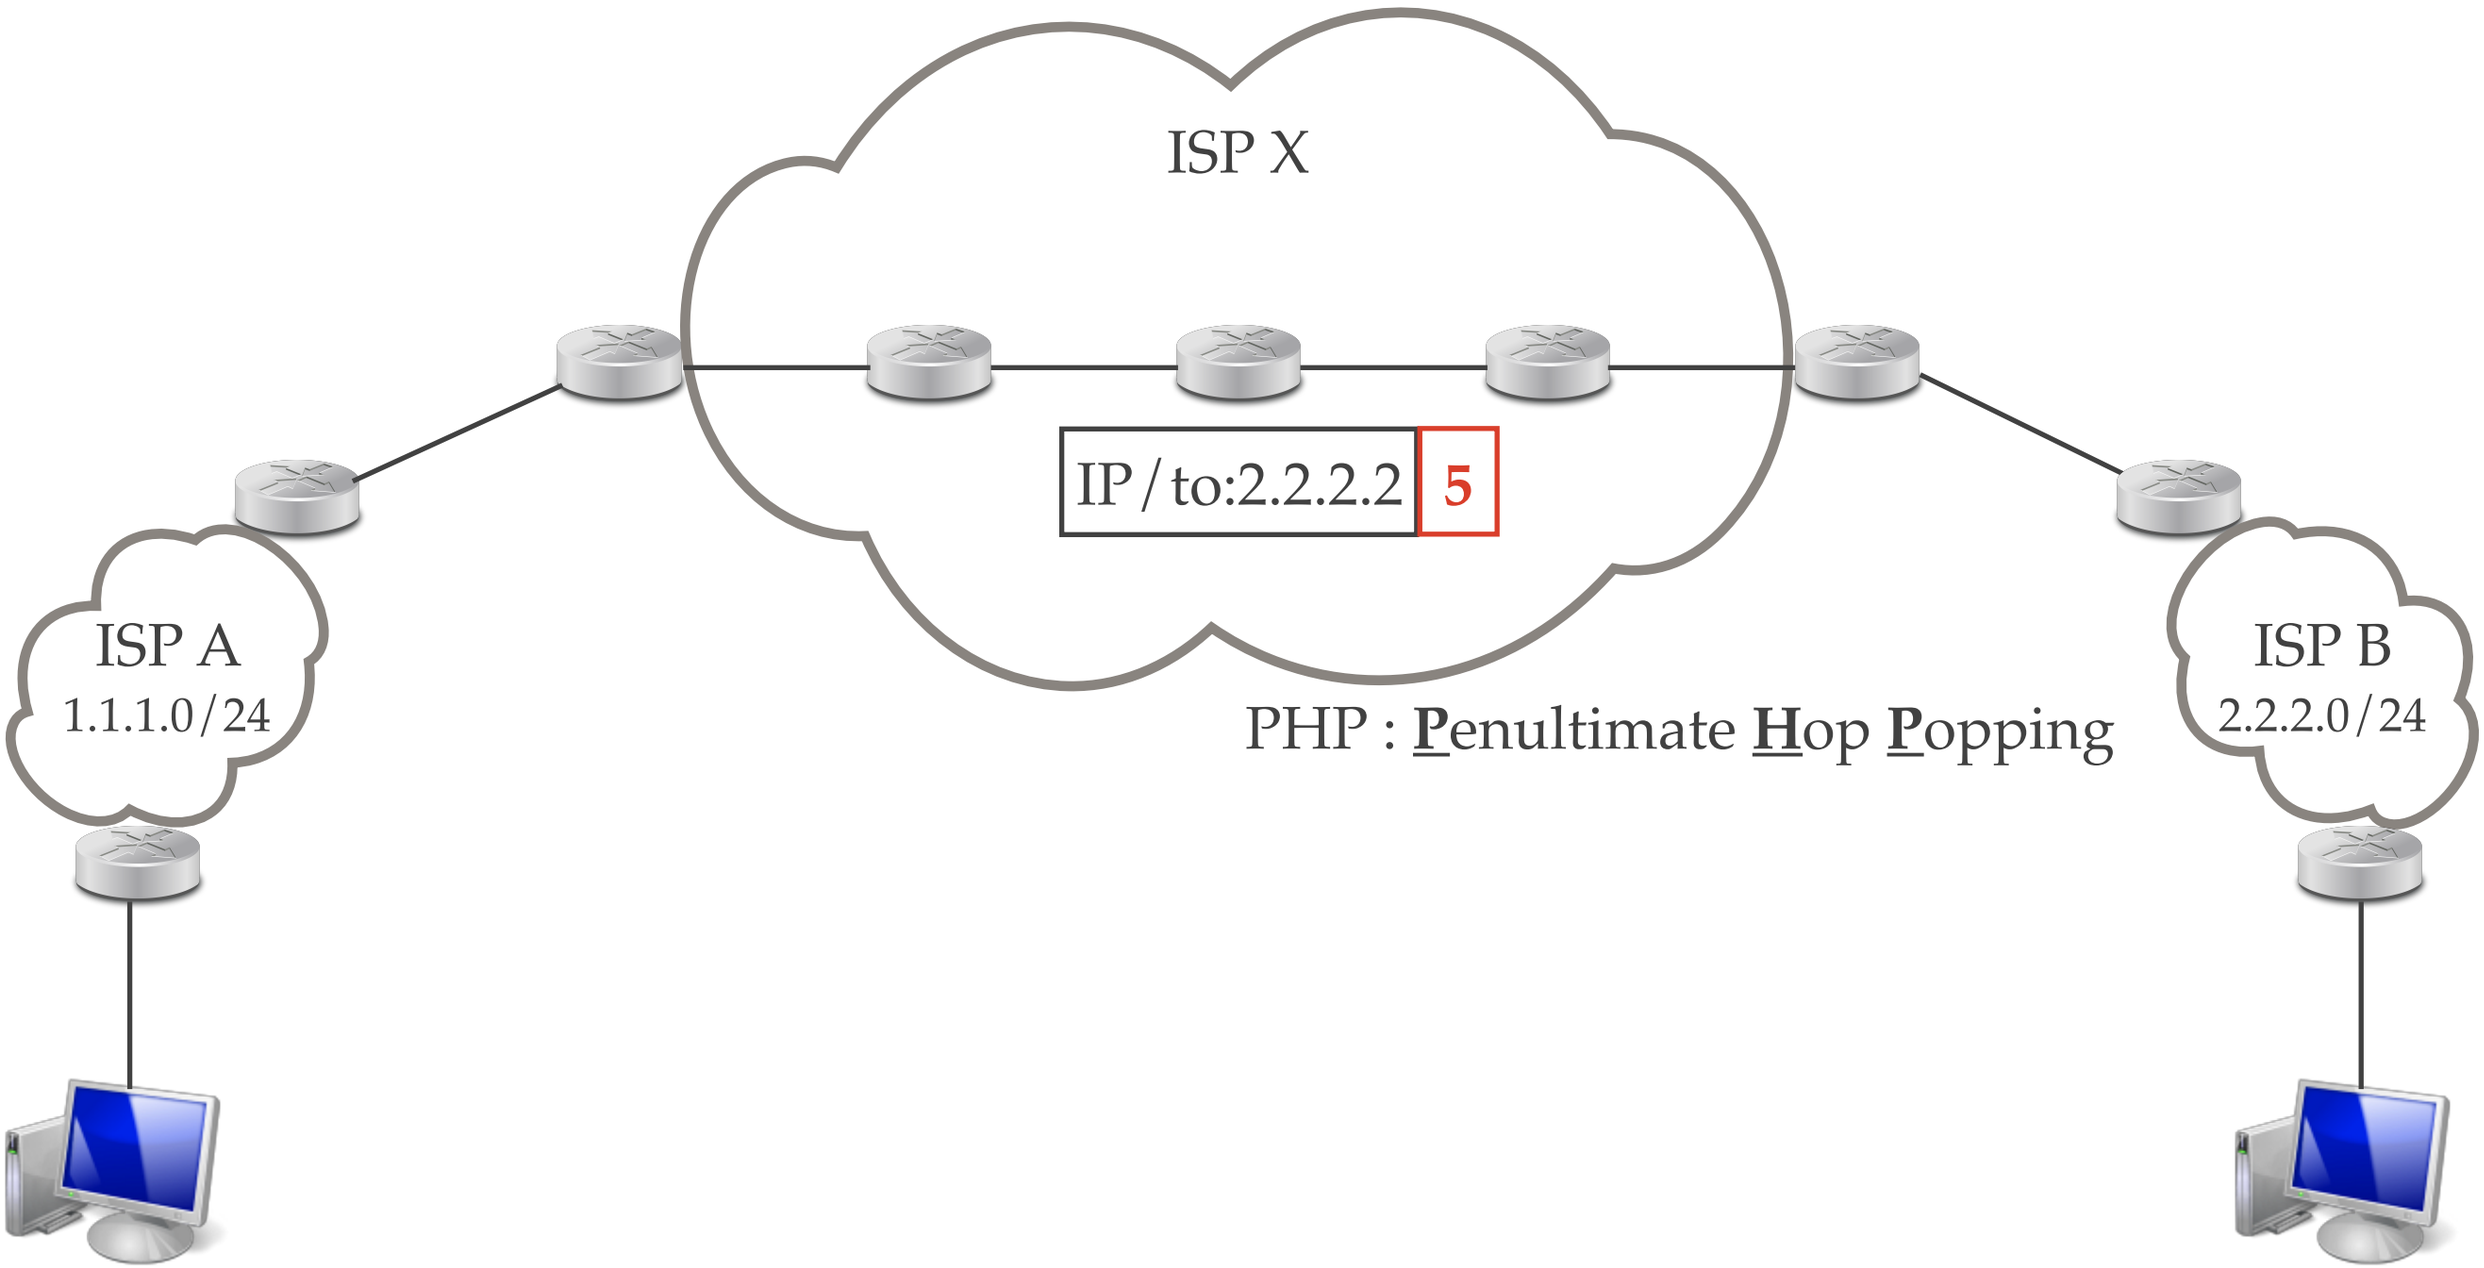

MPLS Network

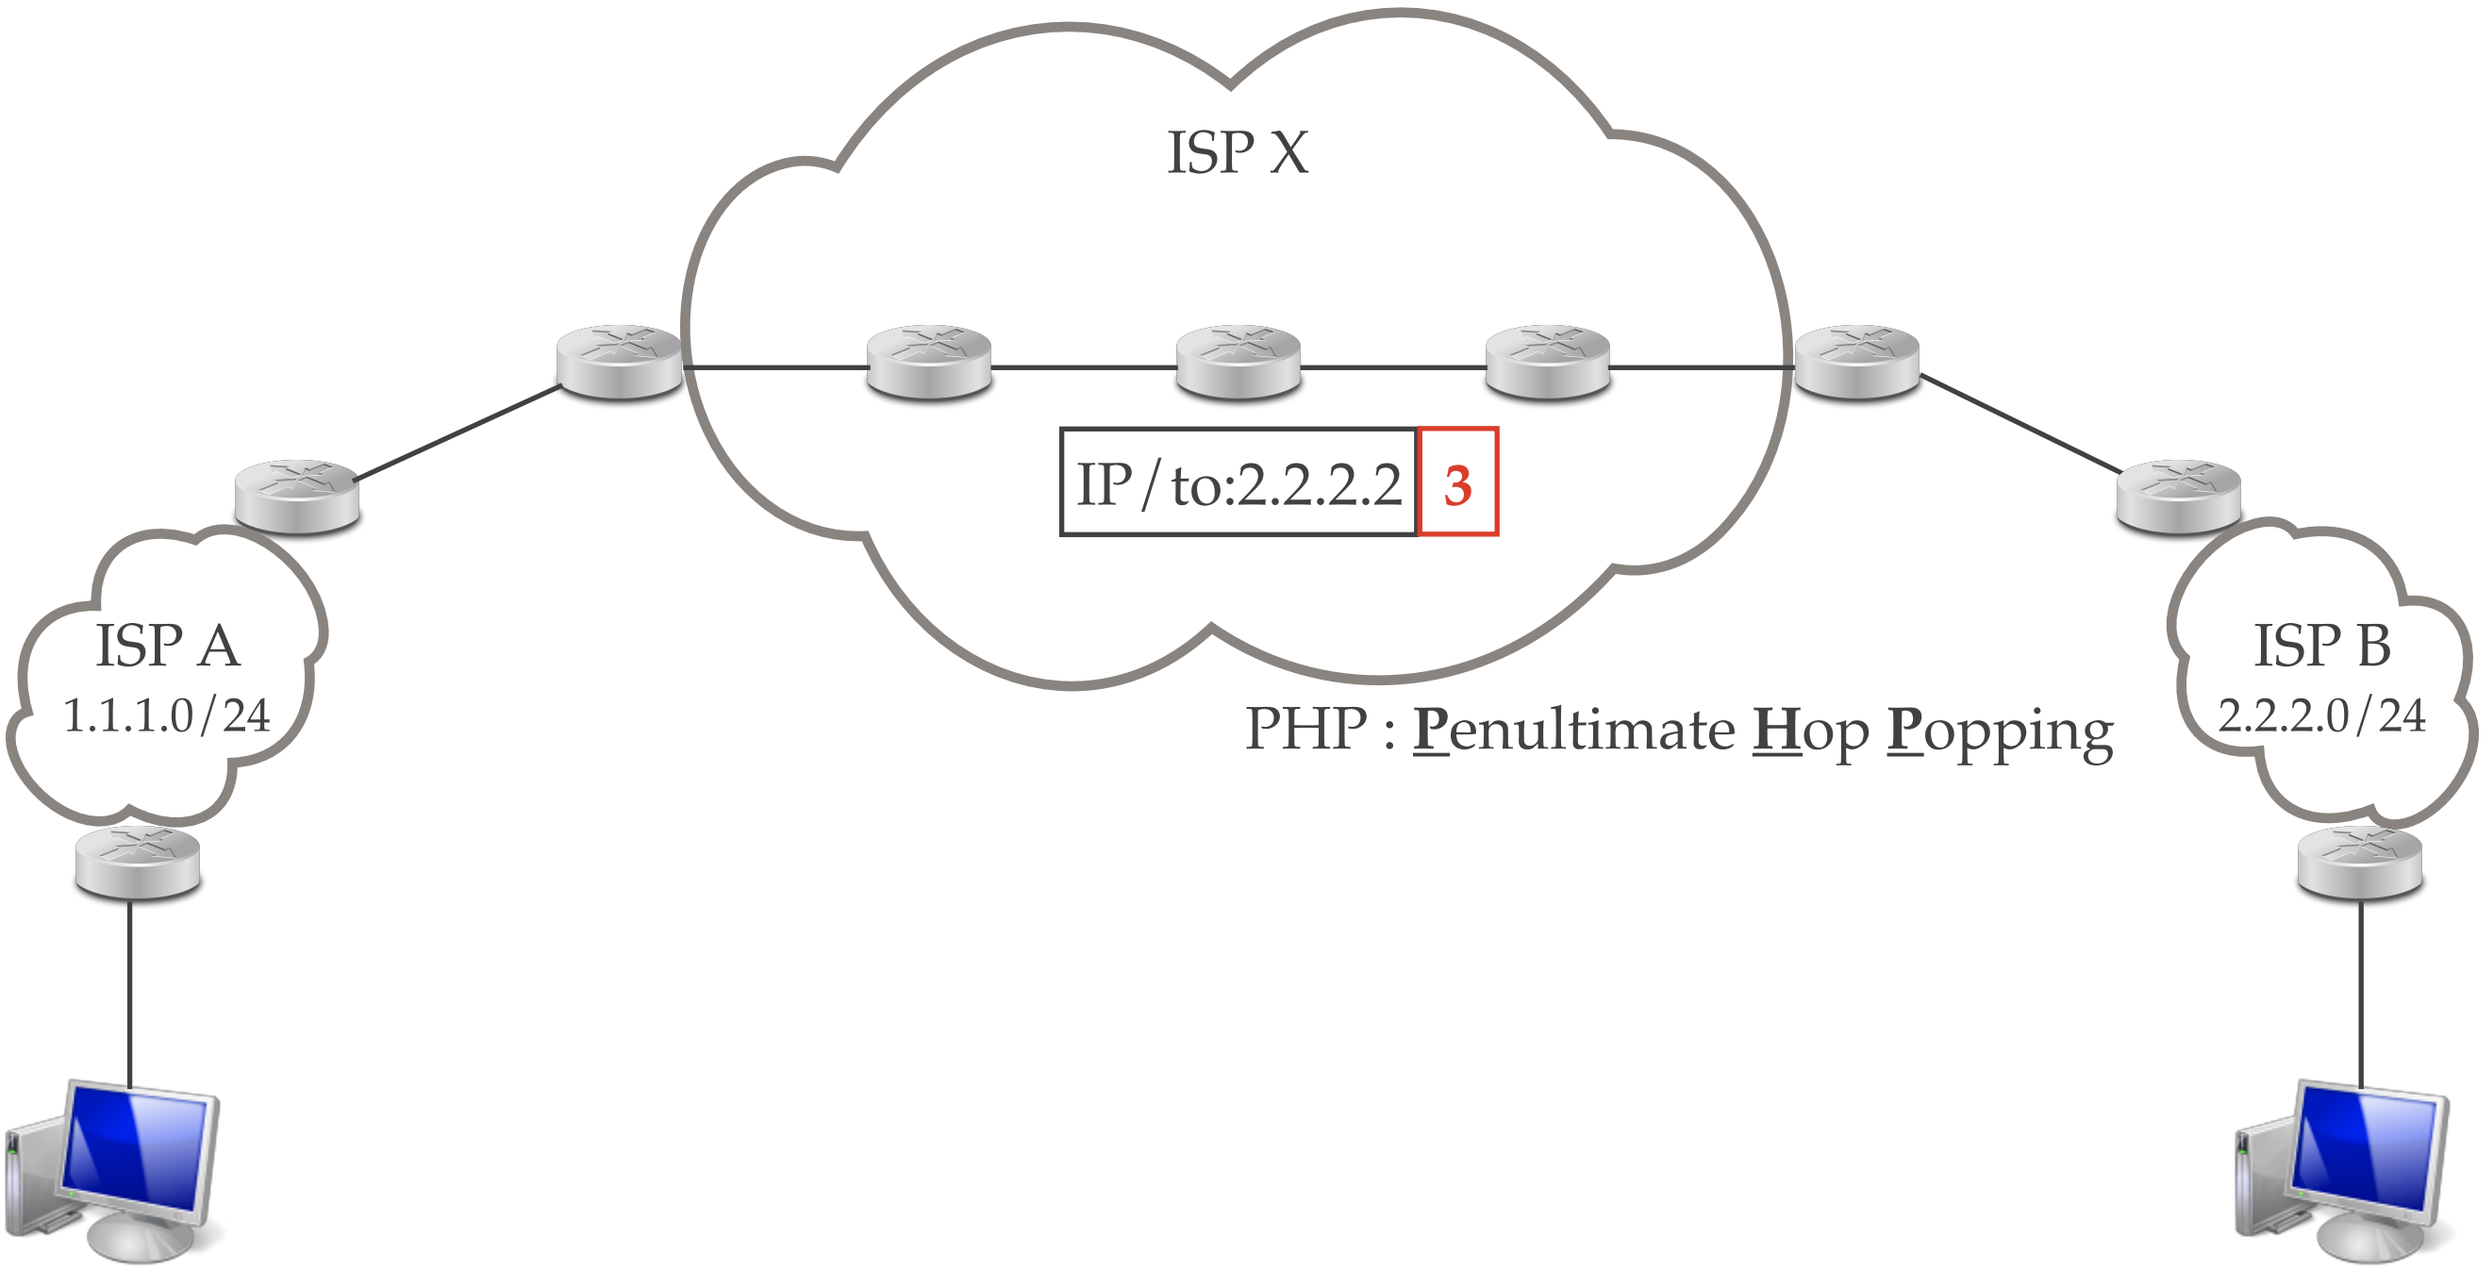

MPLS Network

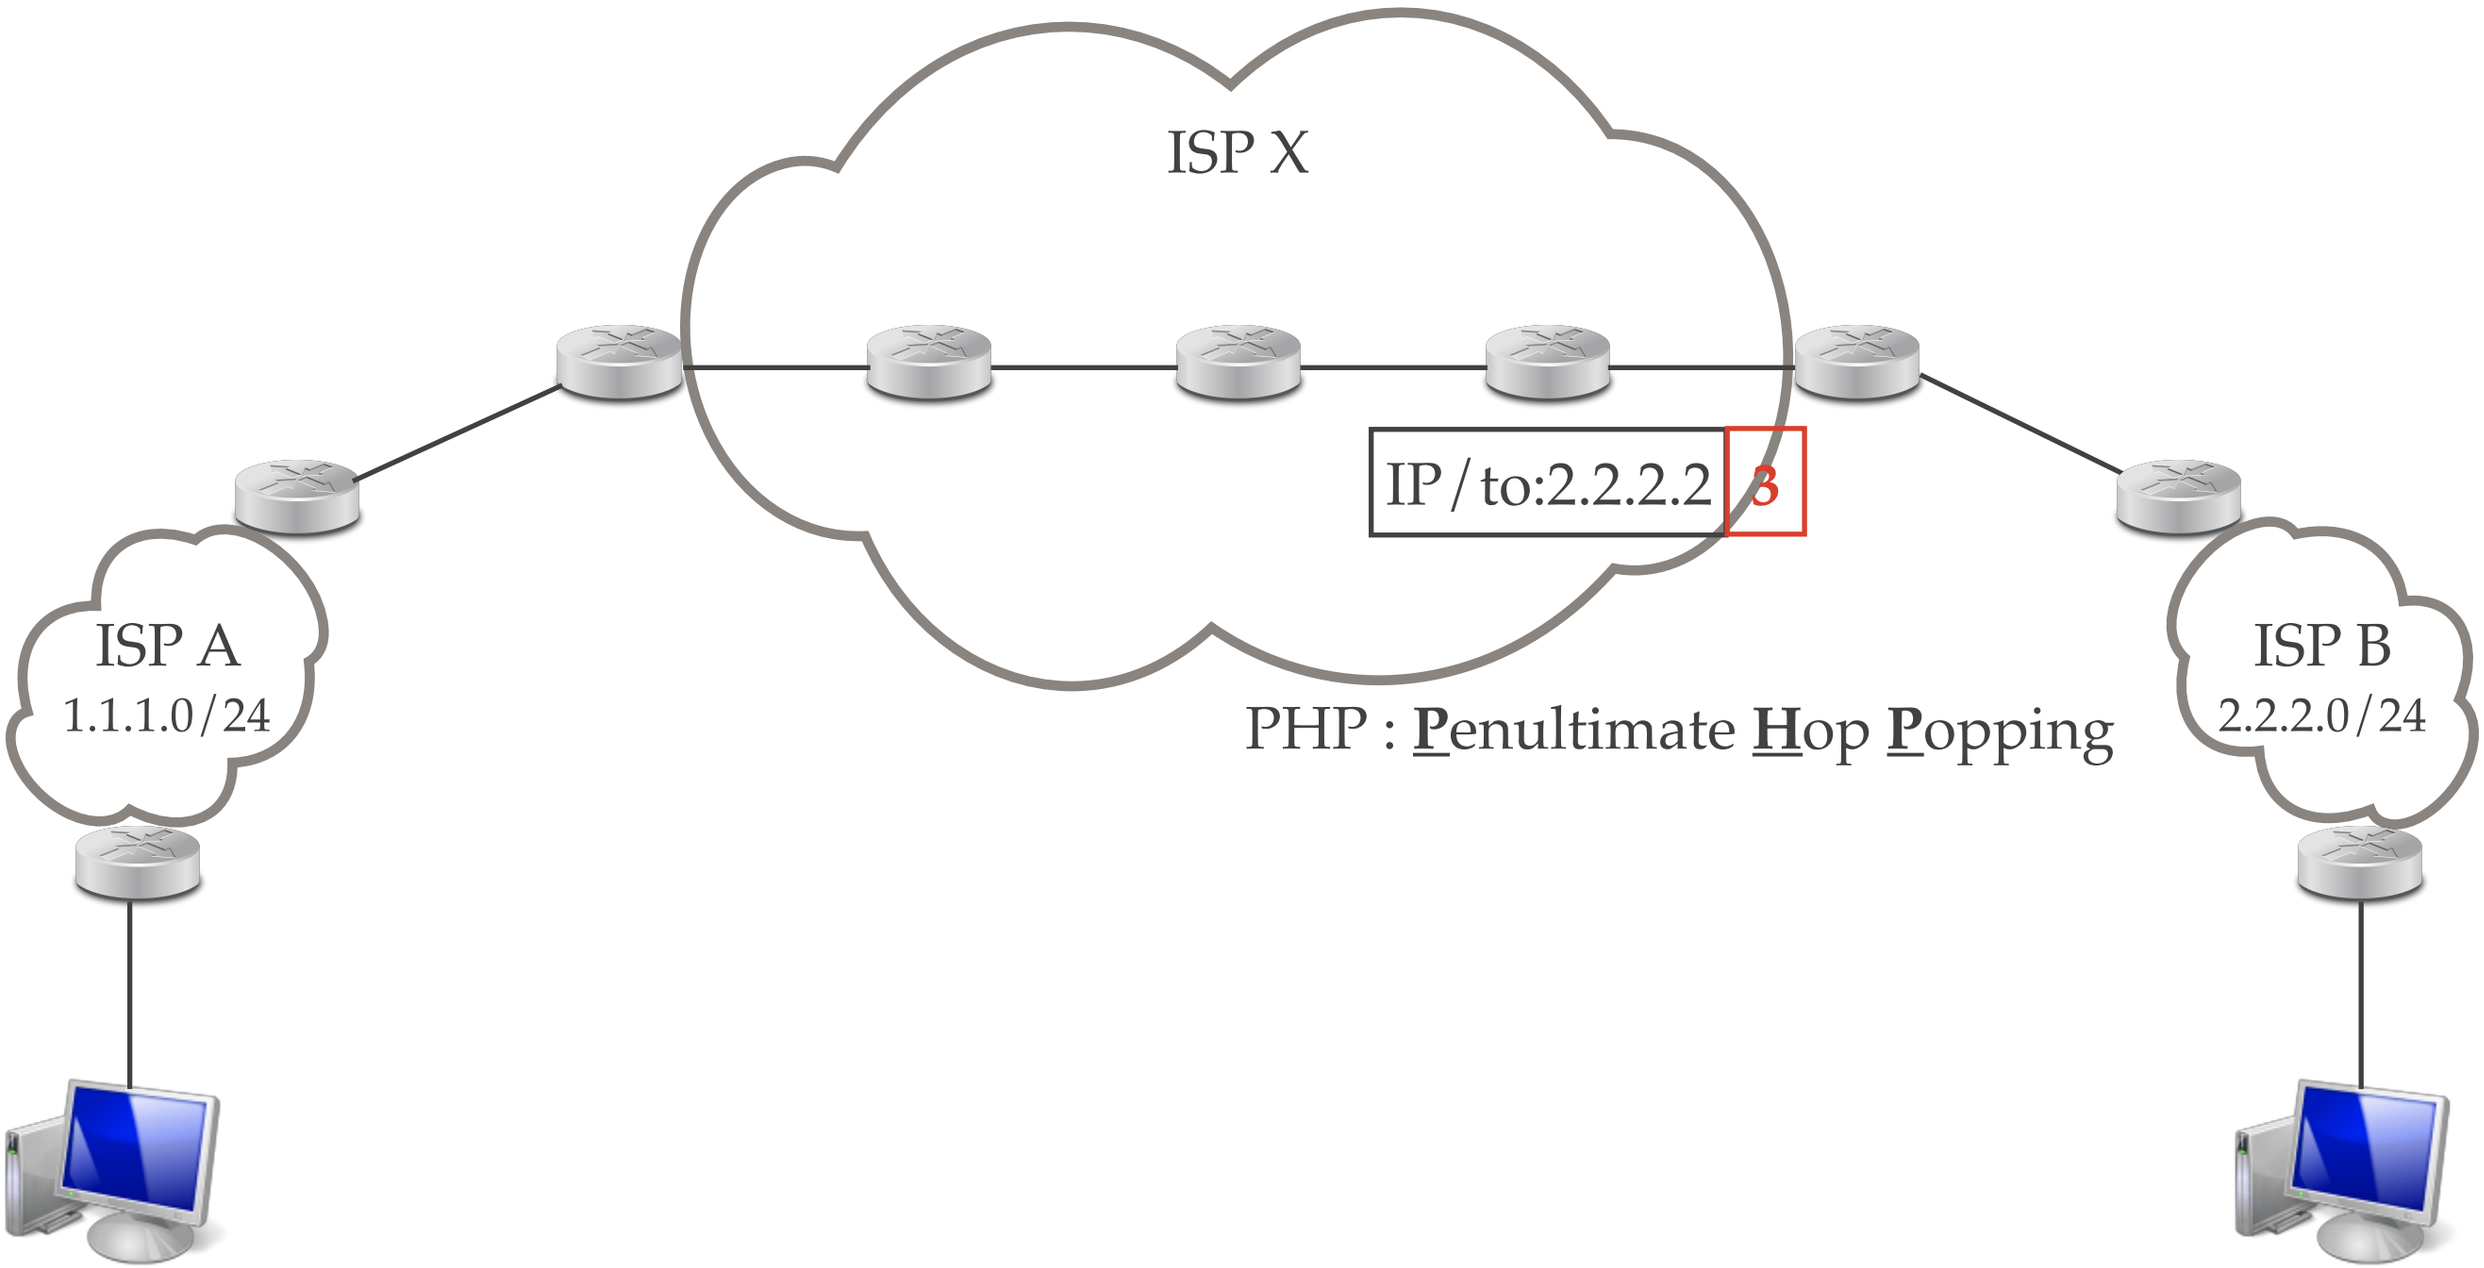

MPLS Network

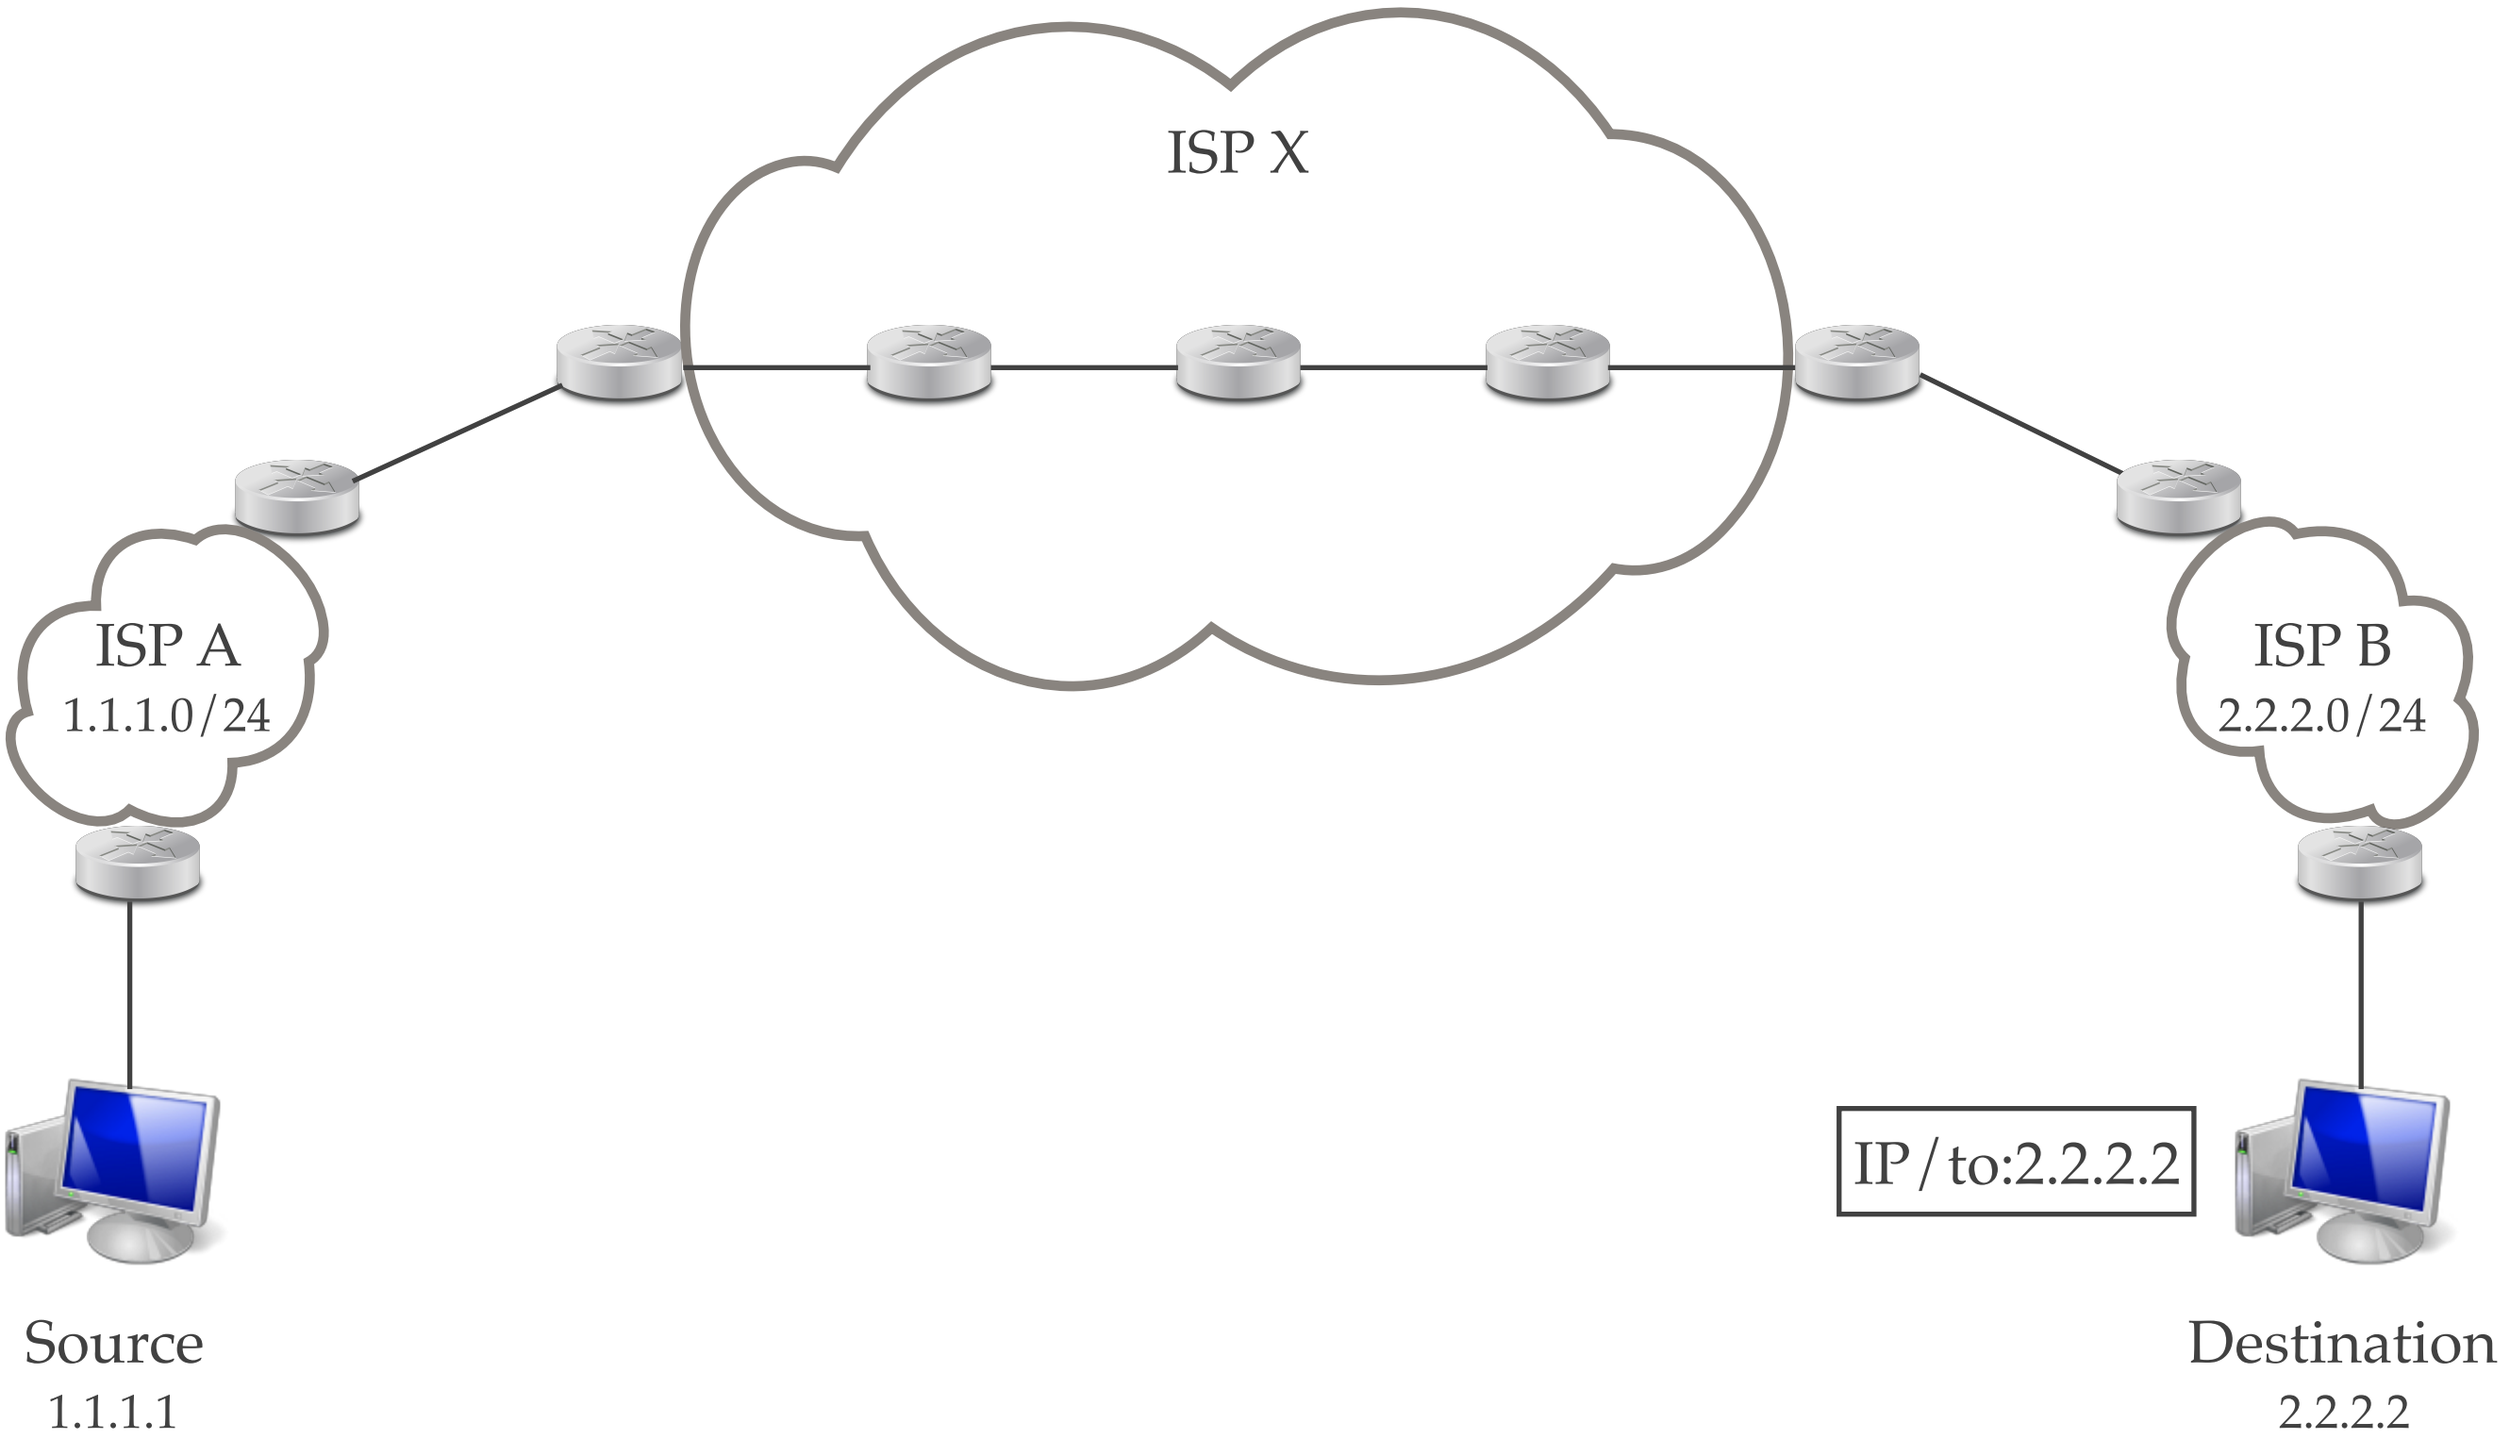

Agenda

- Motivations
- Network Fingerprinting
- MPLS Background
- TNT and MPLS Invisible Tunnels
  - Measuring MPLS
  - MPLS Tunnels Taxonomy
  - Revealing Invisible Tunnels
  - Results
- Conclusion

Measuring MPLS

Measuring MPLS

- The discovery of MPLS can be based on standard active measurement tools
  - B. Donnet, M. Luckie, P. Mérindol, J.-J. Pansiot. *Revealing MPLS Tunnels Obscured from Traceroute*. In ACM SIGCOMM Computer Communication Review. 42(2), pp. 87-93. April 2012.

Measuring MPLS

- The discovery of MPLS can be based on standard active measurement tools
  - B. Donnet, M. Luckie, P. Mérindol, J.-J. Pansiot. *Revealing MPLS Tunnels Obscured from Traceroute*. In ACM SIGCOMM Computer Communication Review. 42(2), pp. 87-93. April 2012.
- Two options are required

Measuring MPLS

- The discovery of MPLS can be based on standard active measurement tools
  - B. Donnet, M. Luckie, P. Mérindol, J.-J. Pansiot. *Revealing MPLS Tunnels Obscured from Traceroute*. In ACM SIGCOMM Computer Communication Review. 42(2), pp. 87-93. April 2012.
- Two options are required
  1. **ICMP extension** ([RFC4950])
     - ✓ if an MPLS router must forge an ICMP `time_exceeded` message, it should quote the MPLS LSE stack in it

Measuring MPLS

- The discovery of MPLS can be based on standard active measurement tools
  - B. Donnet, M. Luckie, P. Mérindol, J.-J. Pansiot. *Revealing MPLS Tunnels Obscured from Traceroute*. In ACM SIGCOMM Computer Communication Review. 42(2), pp. 87-93. April 2012.
- Two options are required
  1. **ICMP extension** ([RFC4950])
     - ✓ if an MPLS router must forge an ICMP `time_exceeded` message, it should quote the MPLS LSE stack in it
  2. **TTL propagate** ([RFC3443])
     - ✓ the ingress LER of an MPLS should initialize the LSE-TTL with the value inside the IP-TTL field (iTTL)
     - ✓ the opposite operation is done by the egress LER (oTTL)

Measuring MPLS (2)

- Effect of TTL propagate
  - with PHP

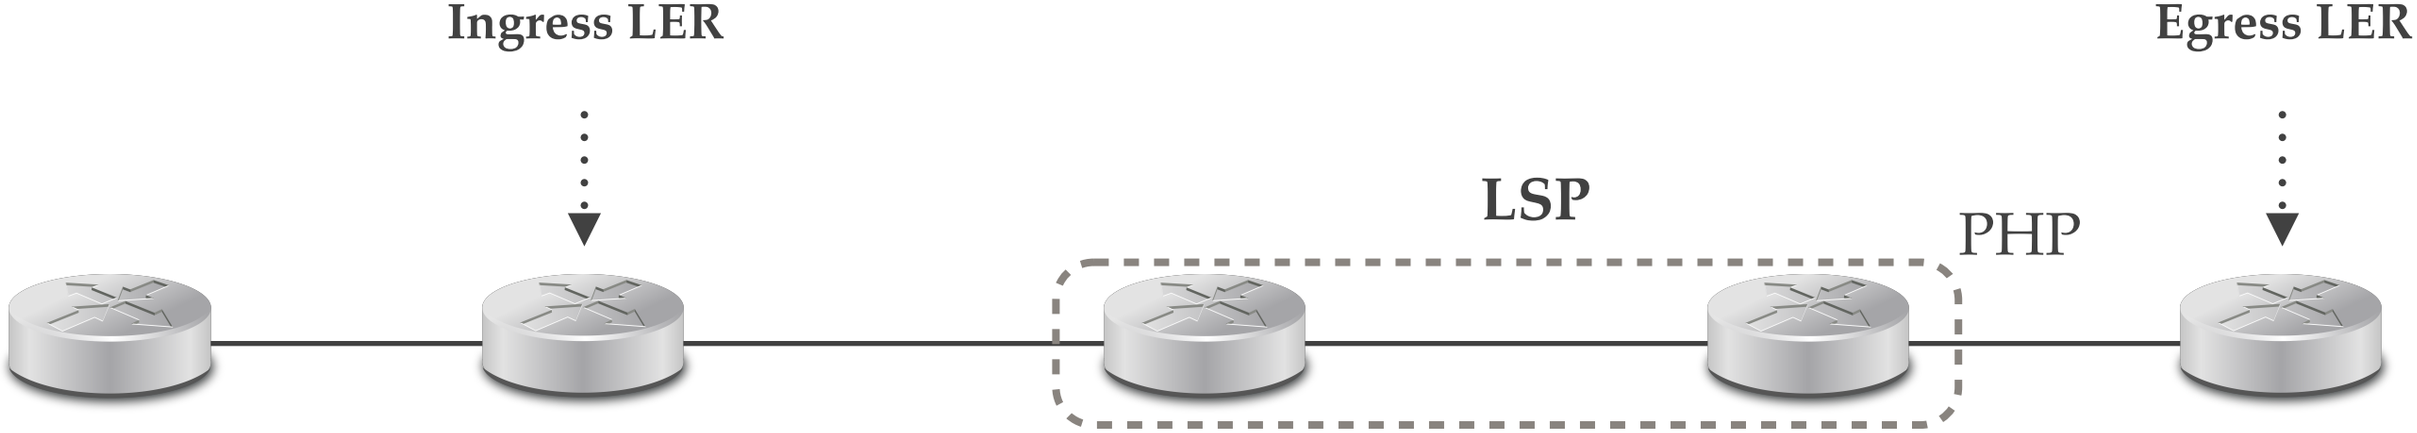

Measuring MPLS (2)

- Effect of TTL propagate
  - with PHP

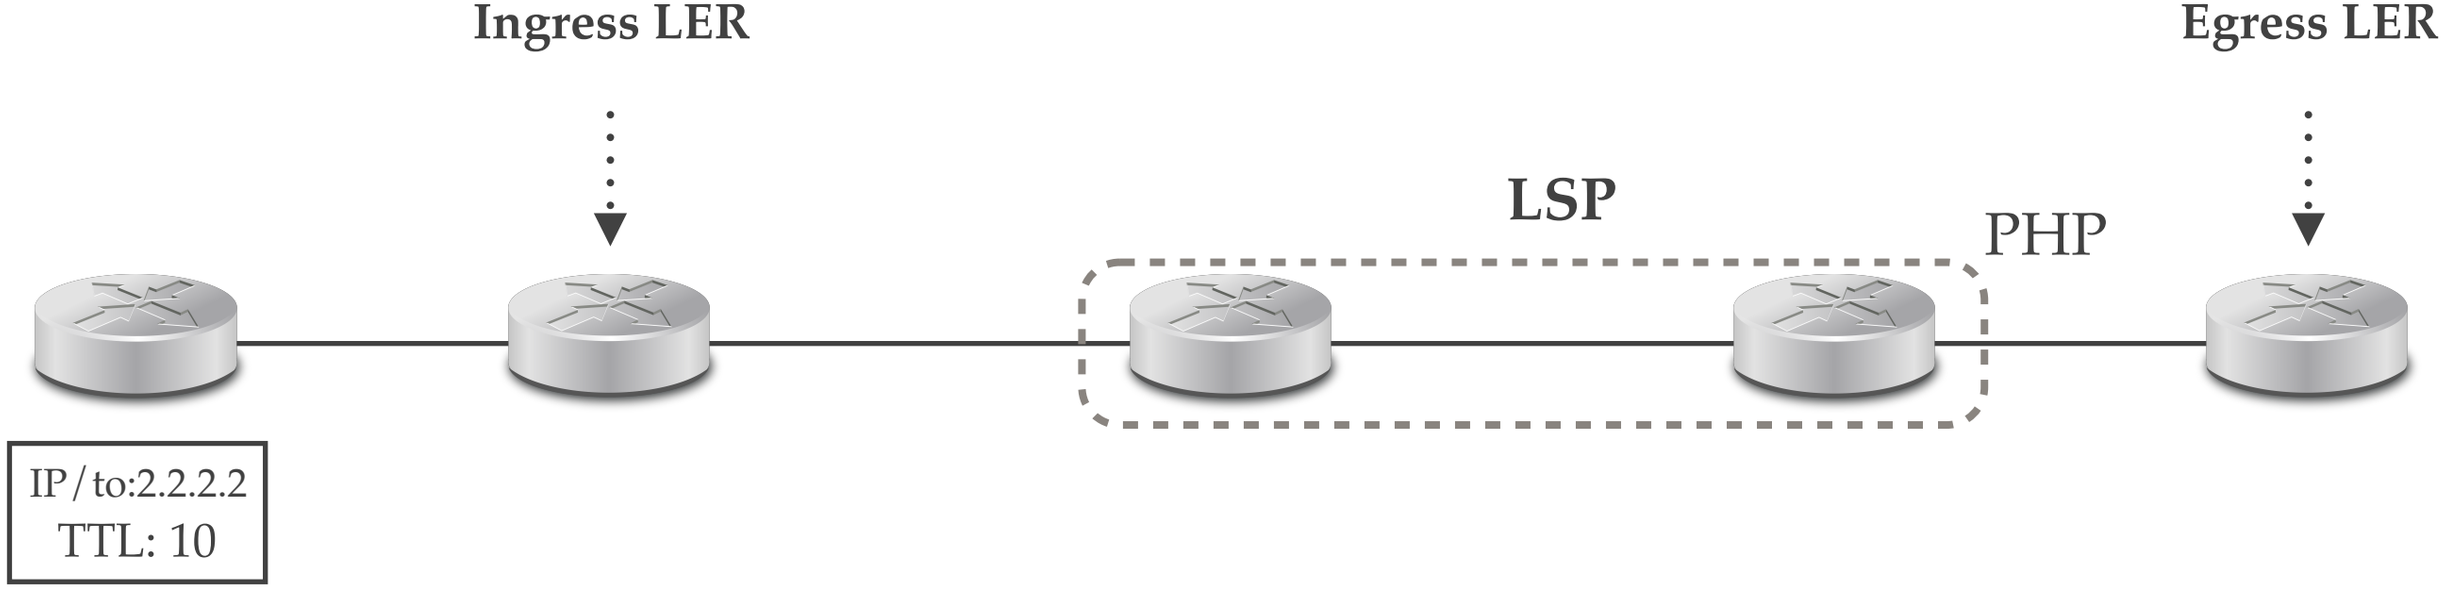

Measuring MPLS (2)

- Effect of TTL propagate
  - with PHP

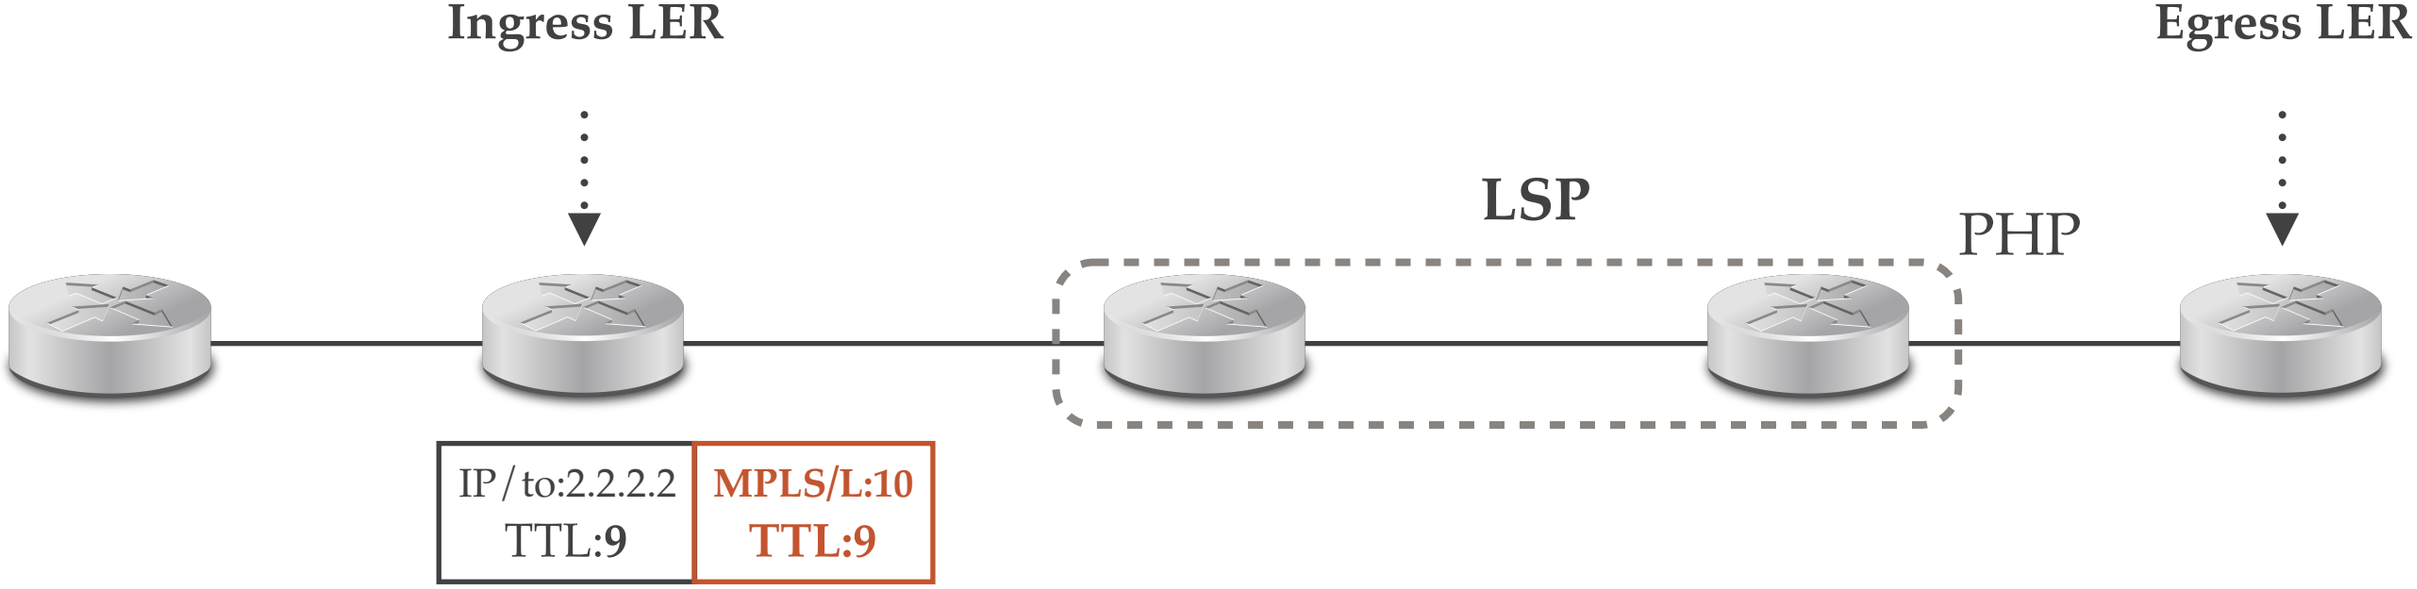

Measuring MPLS (2)

- Effect of TTL propagate
  - with PHP

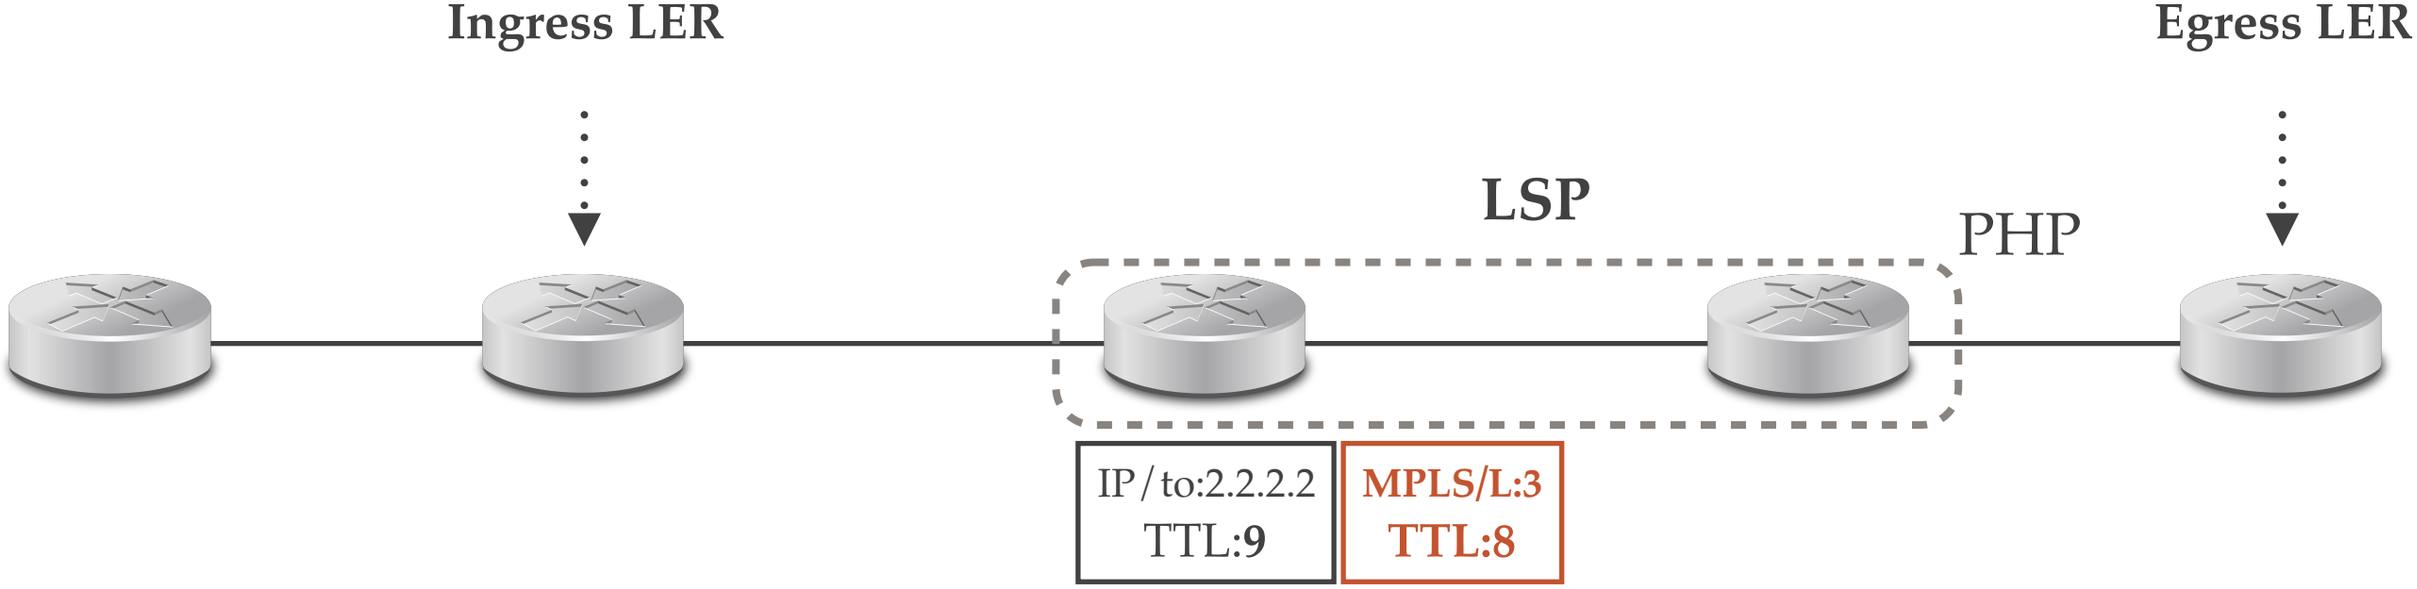

Measuring MPLS (2)

- Effect of TTL propagate
  - with PHP

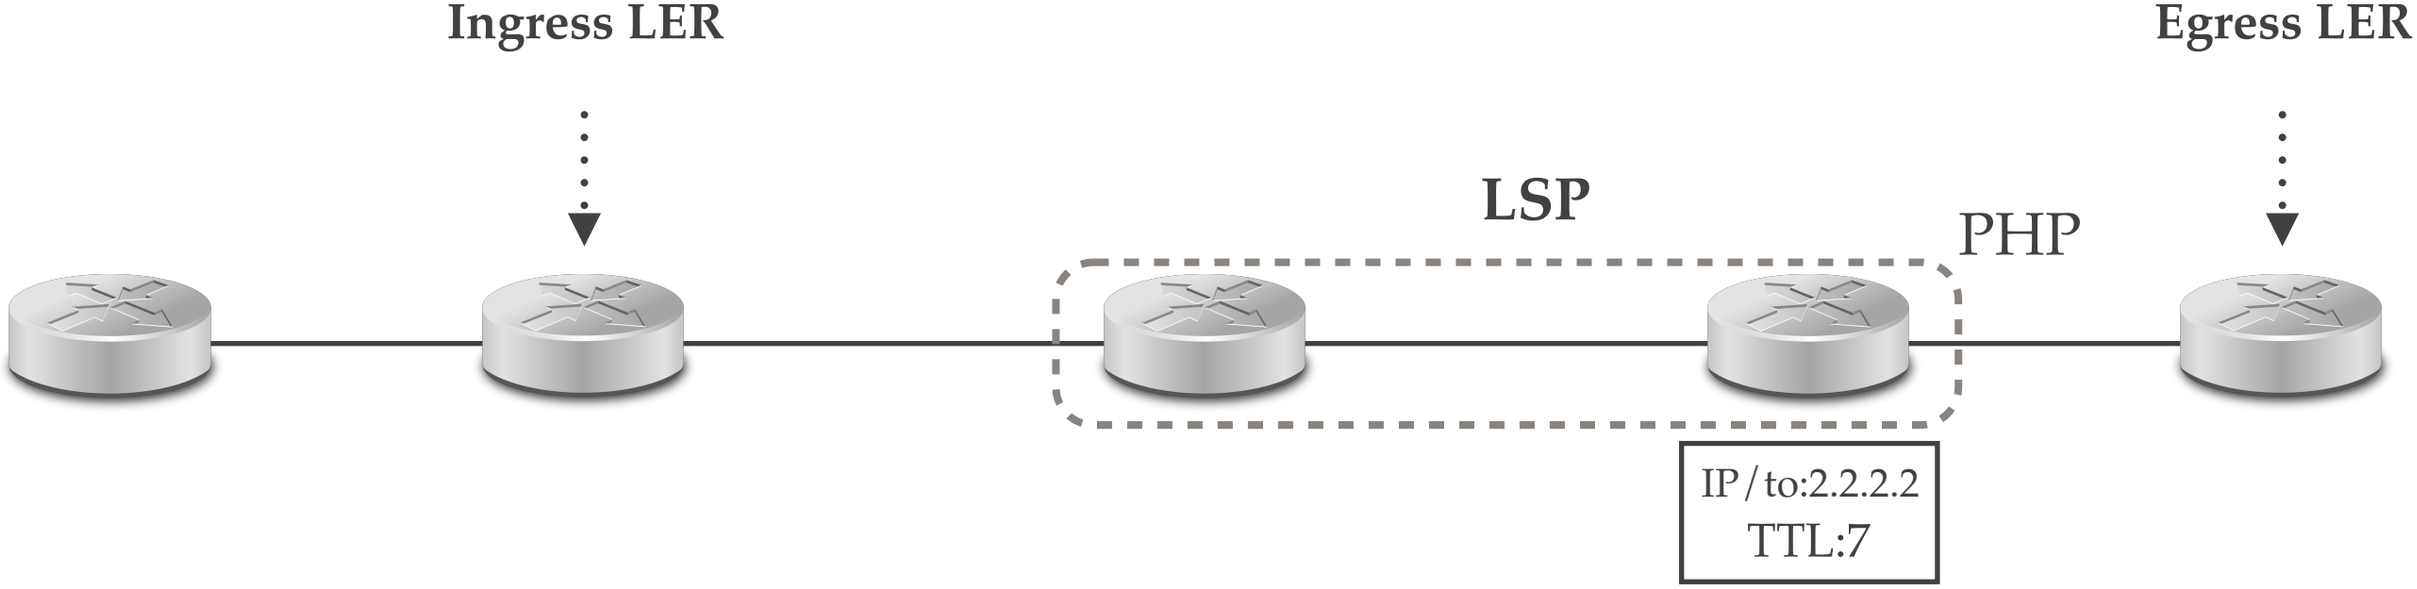

Measuring MPLS (2)

- Effect of TTL propagate
  - with PHP

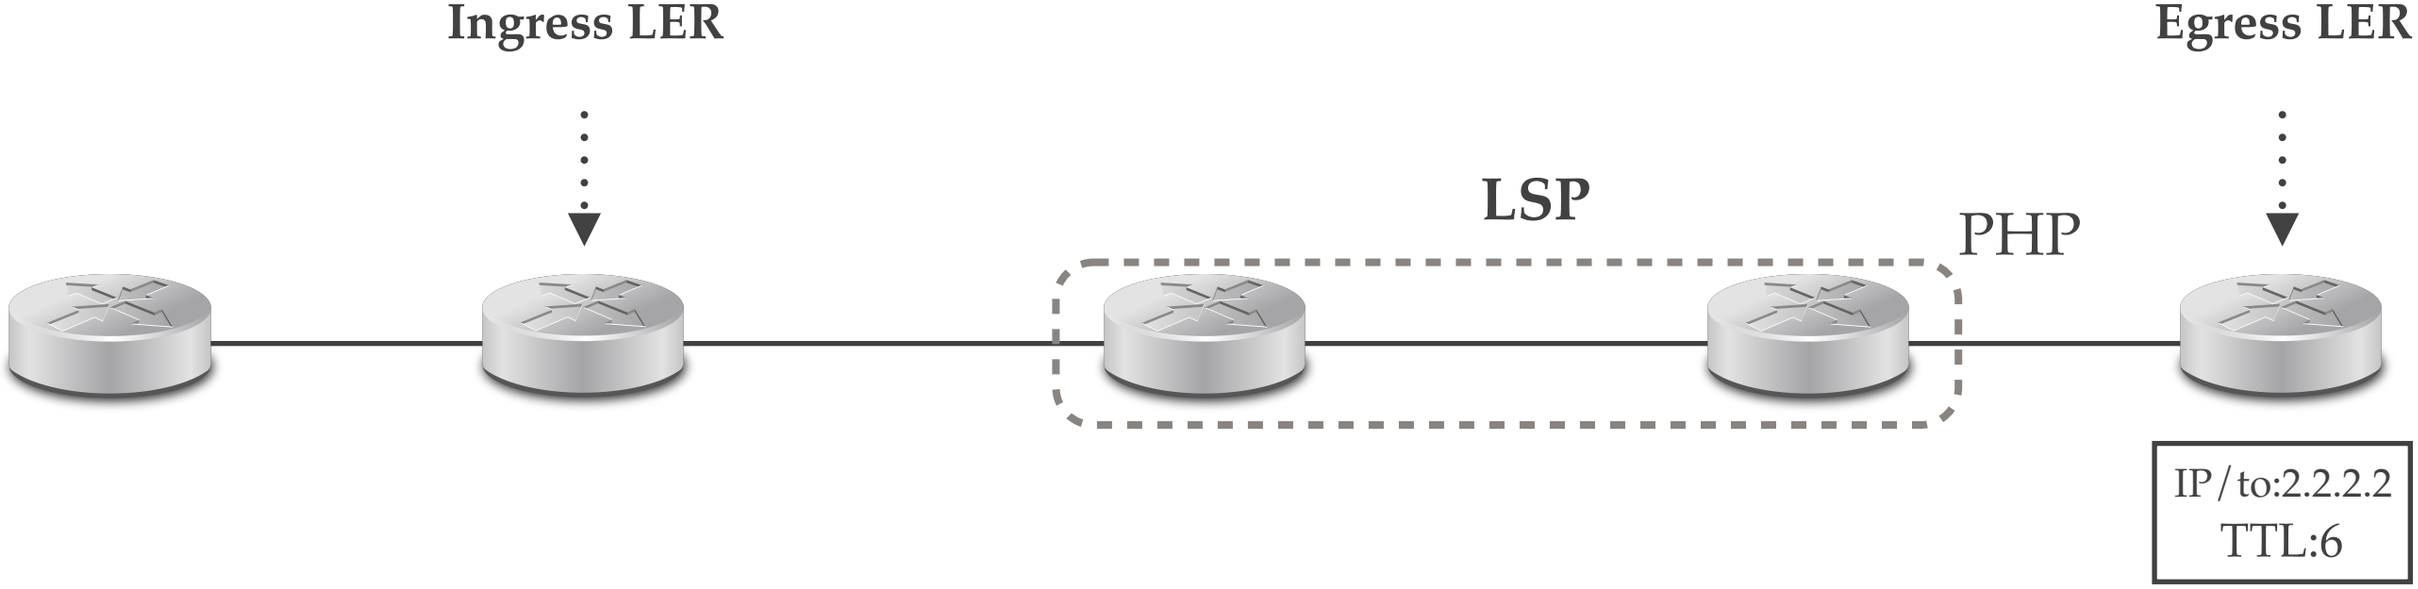

Measuring MPLS (4)

- What happens if the tunnel does not use TTL propagate
  - with PHP

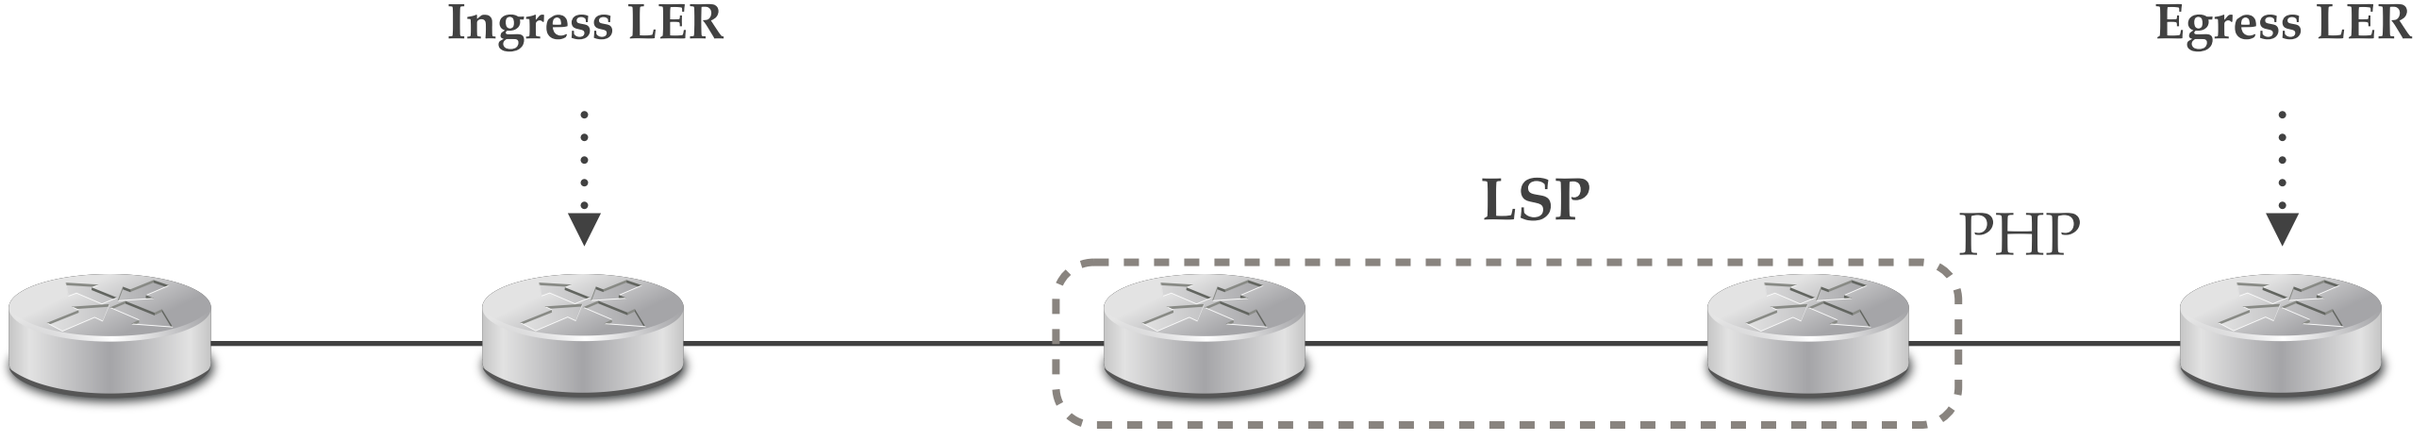

Measuring MPLS (4)

- What happens if the tunnel does not use TTL propagate
  - with PHP

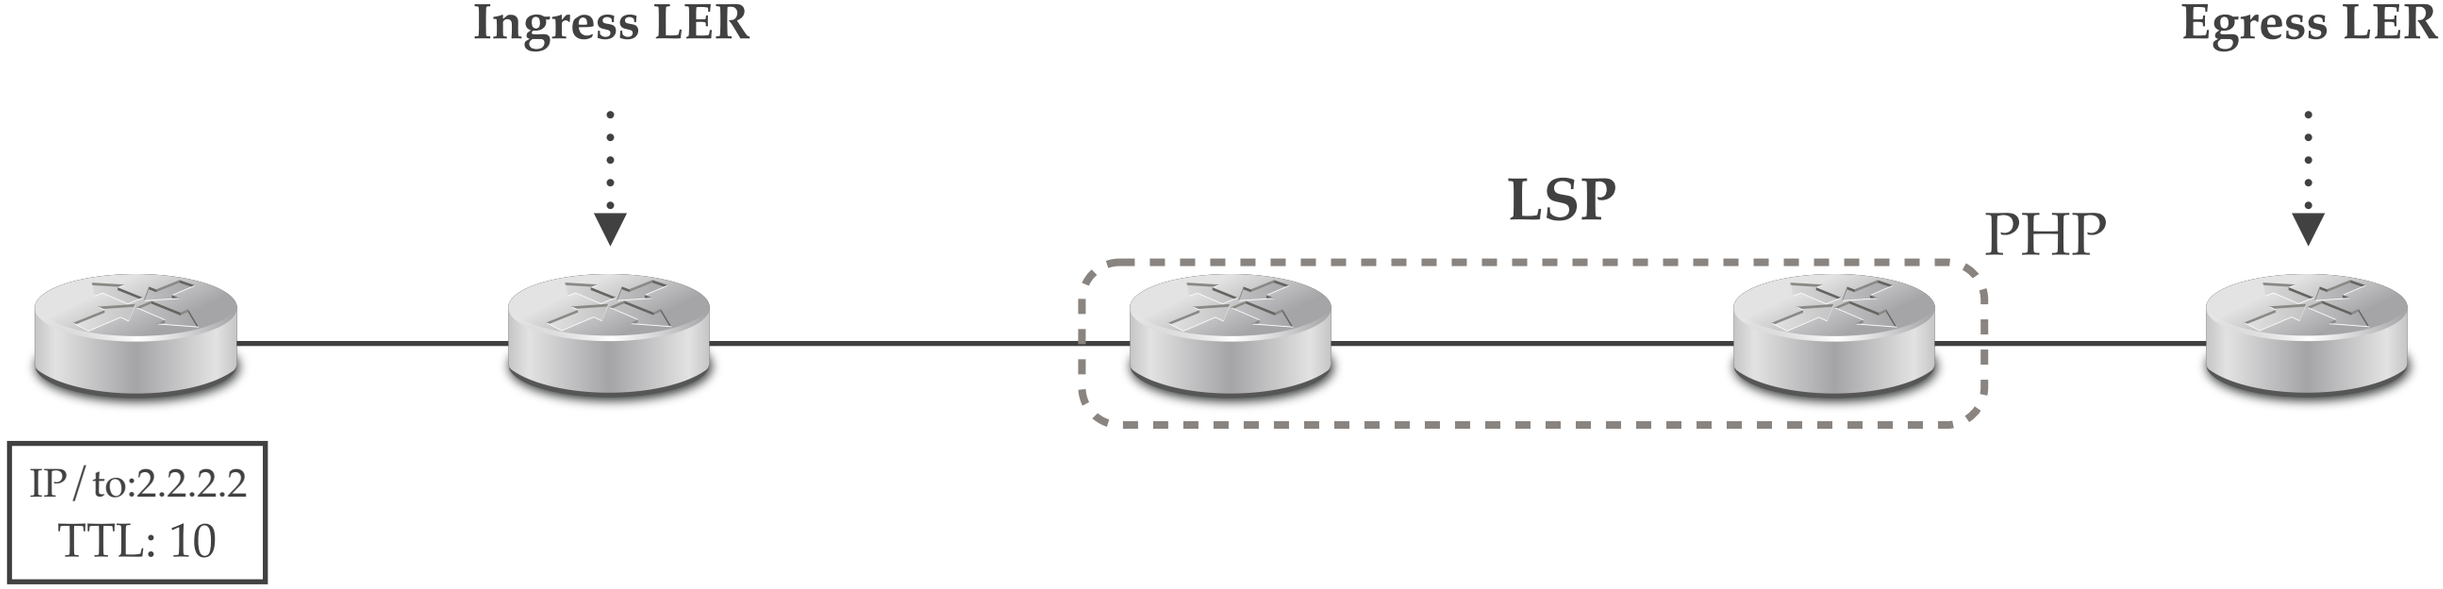

Measuring MPLS (4)

- What happens if the tunnel does not use TTL propagate
  - with PHP

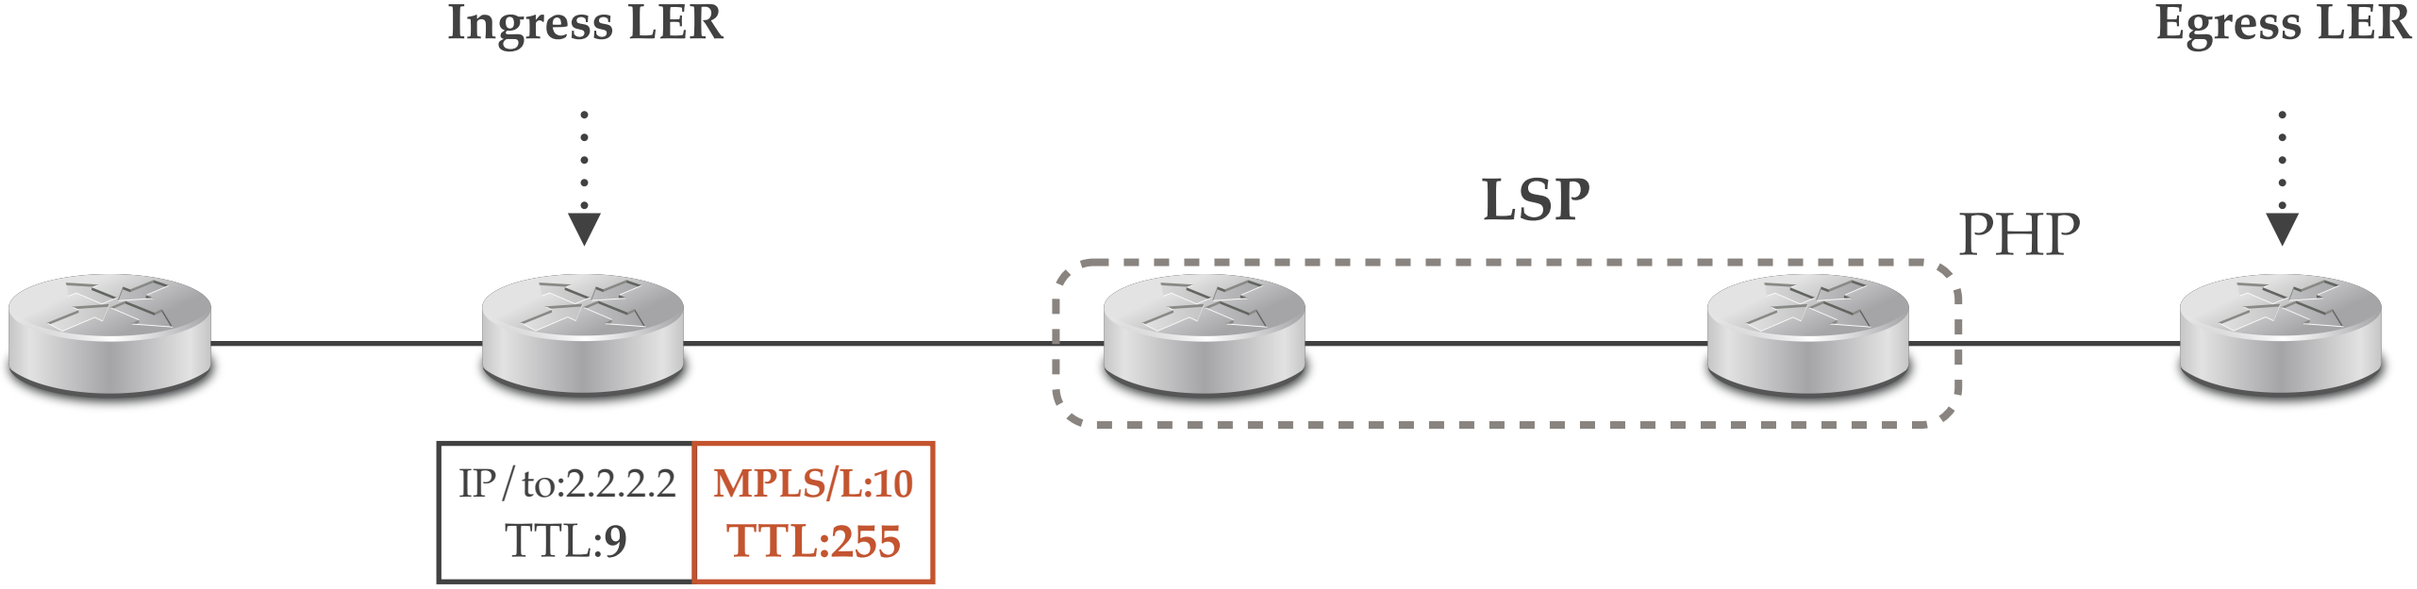

Measuring MPLS (4)

- What happens if the tunnel does not use TTL propagate
  - with PHP

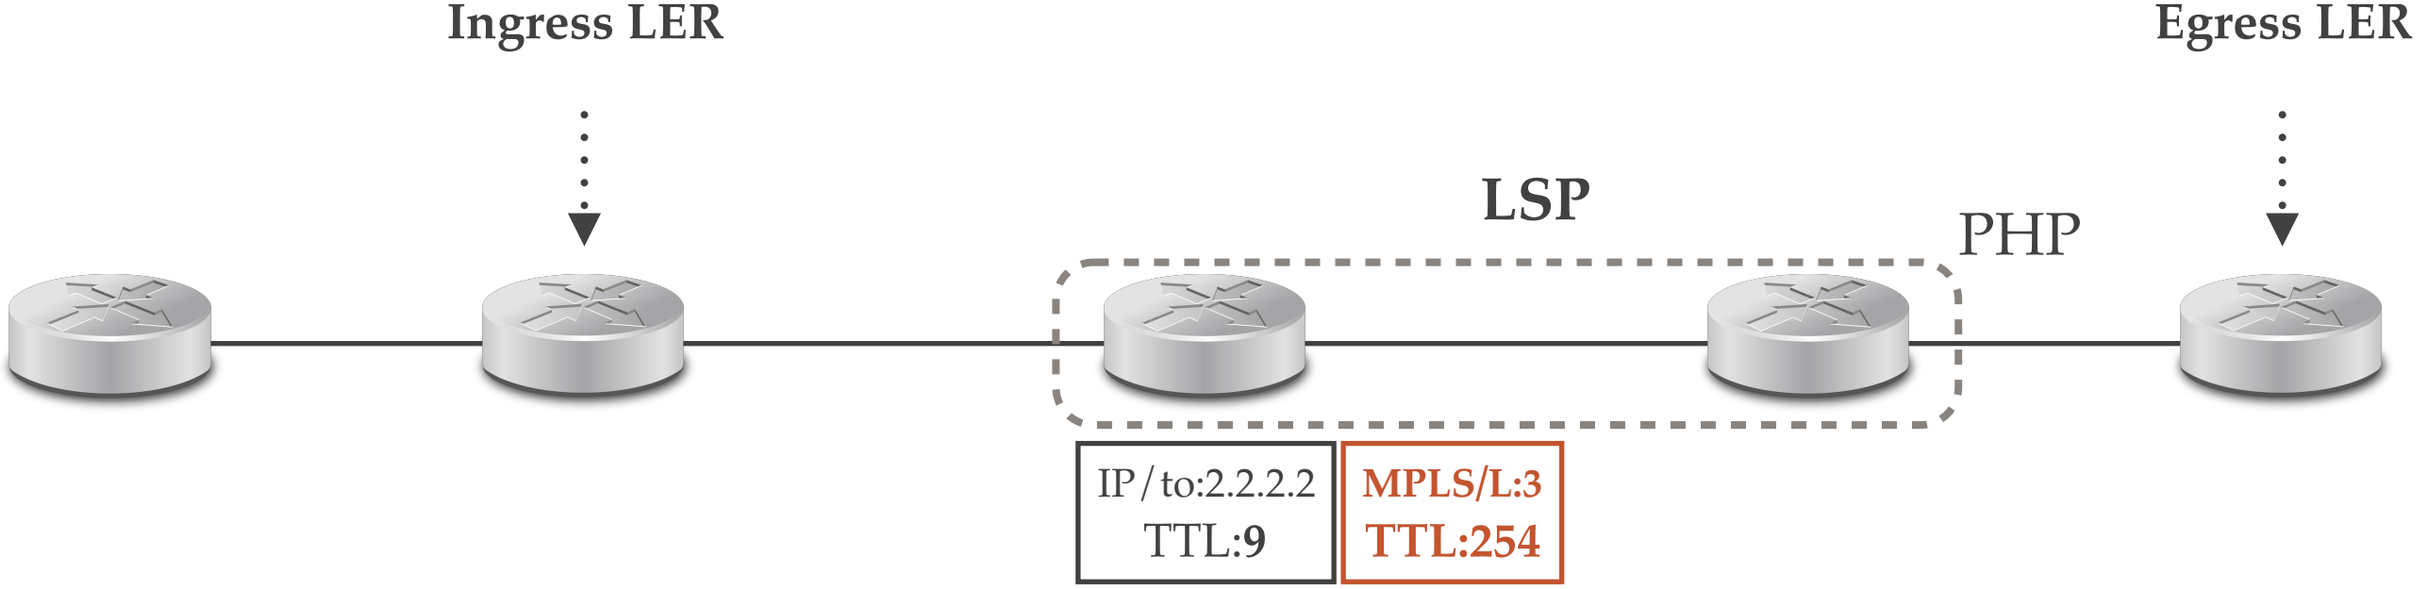

Measuring MPLS (4)

- What happens if the tunnel does not use TTL propagate
  - with PHP

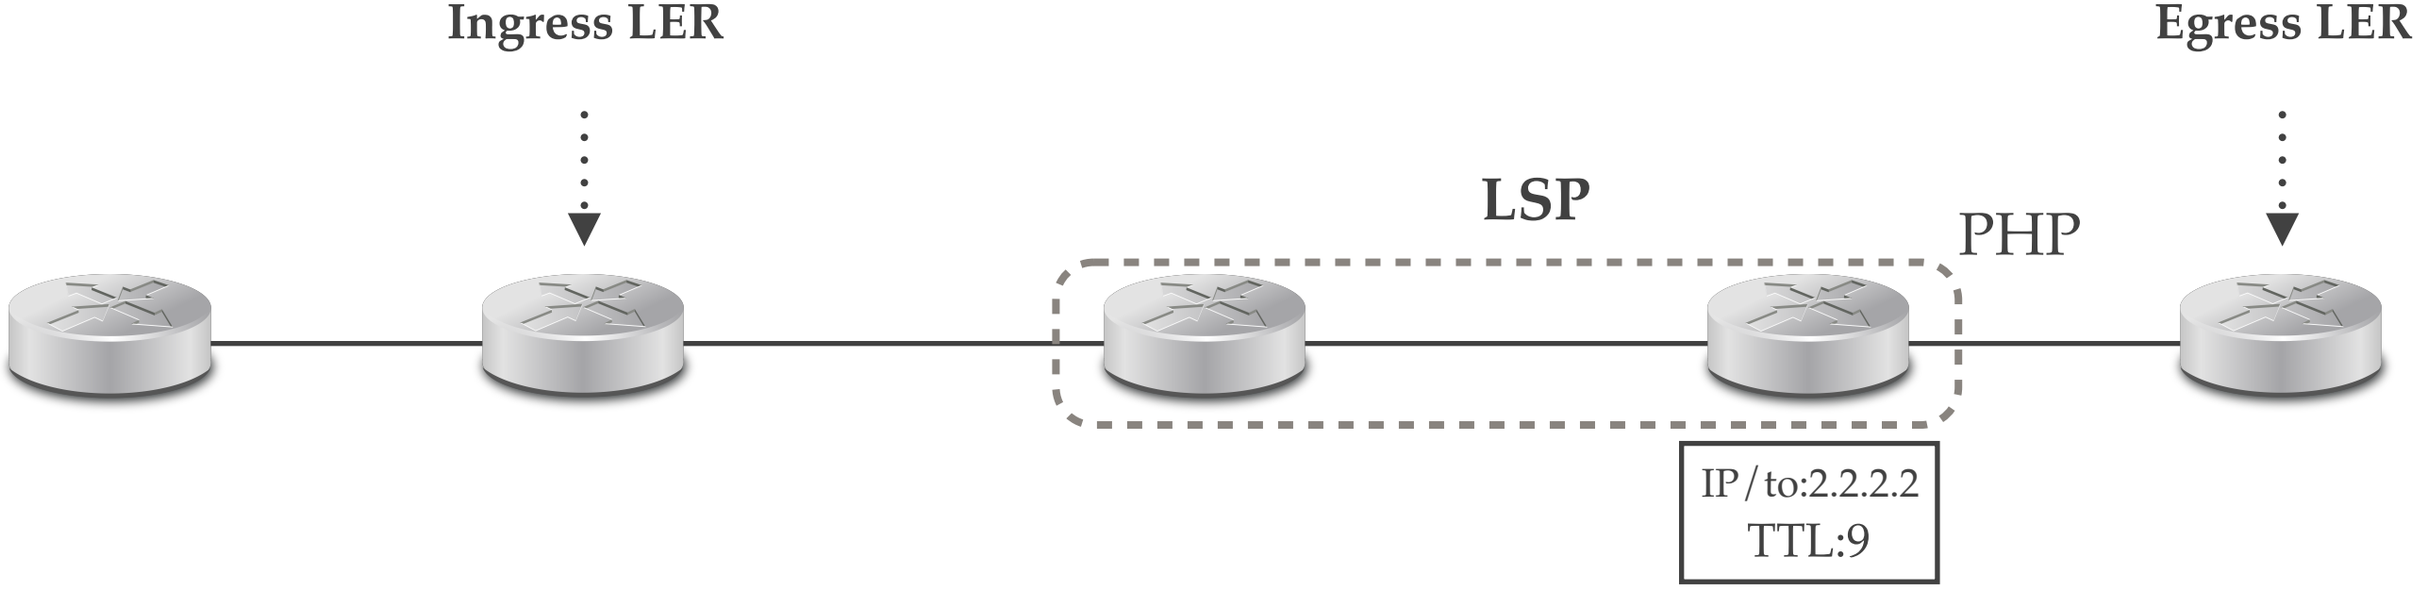

Measuring MPLS (4)

- What happens if the tunnel does not use TTL propagate
  - with PHP

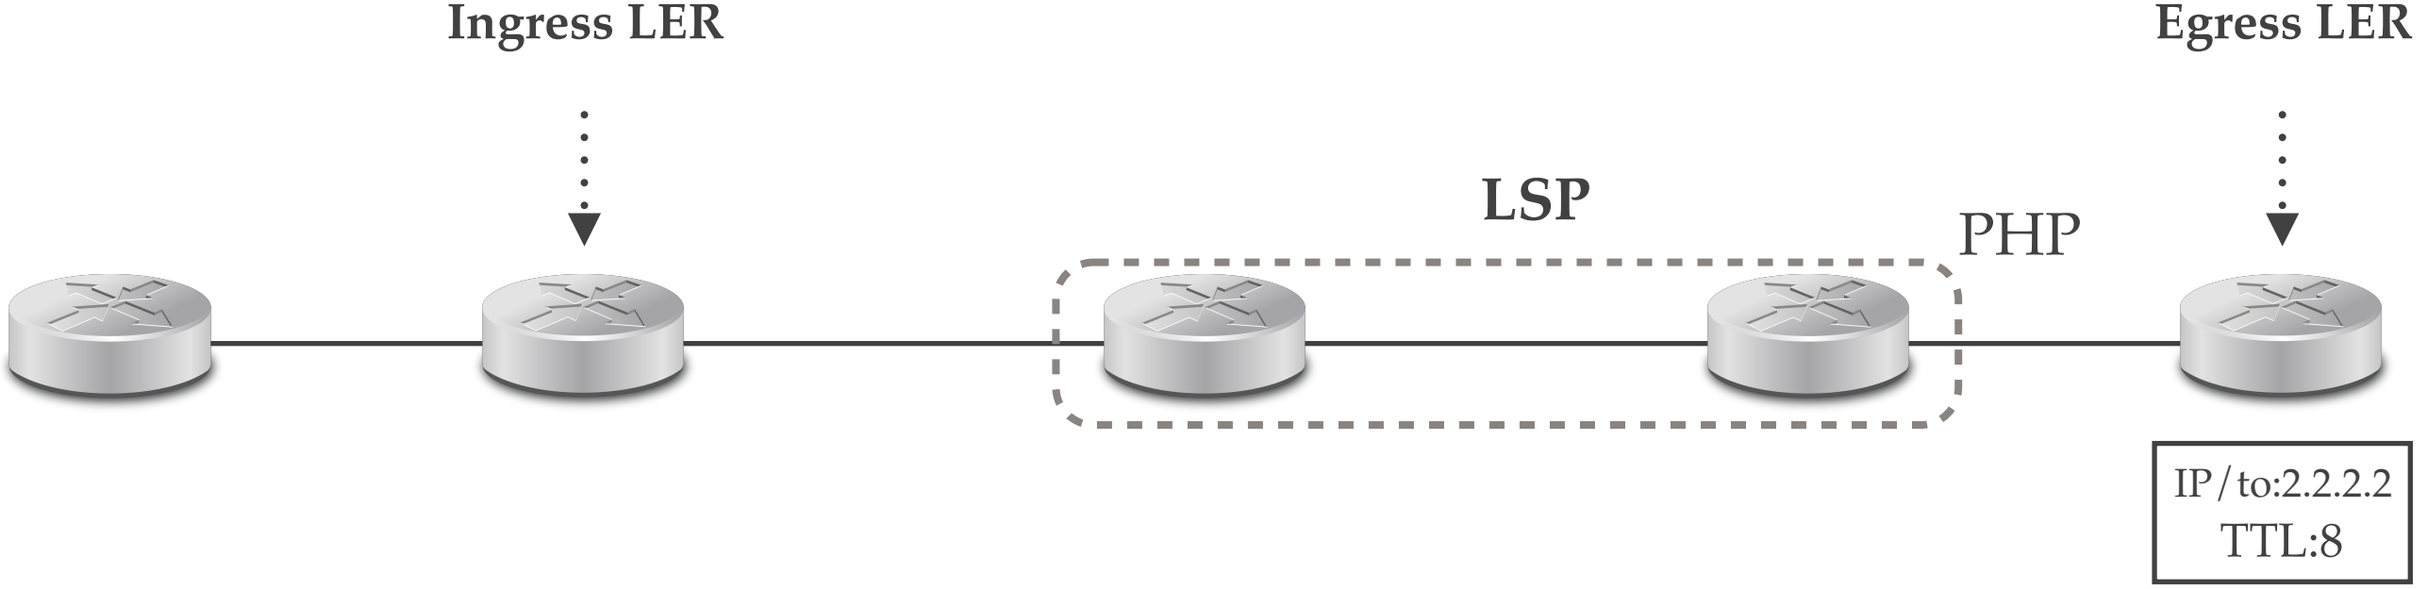

Taxonomy

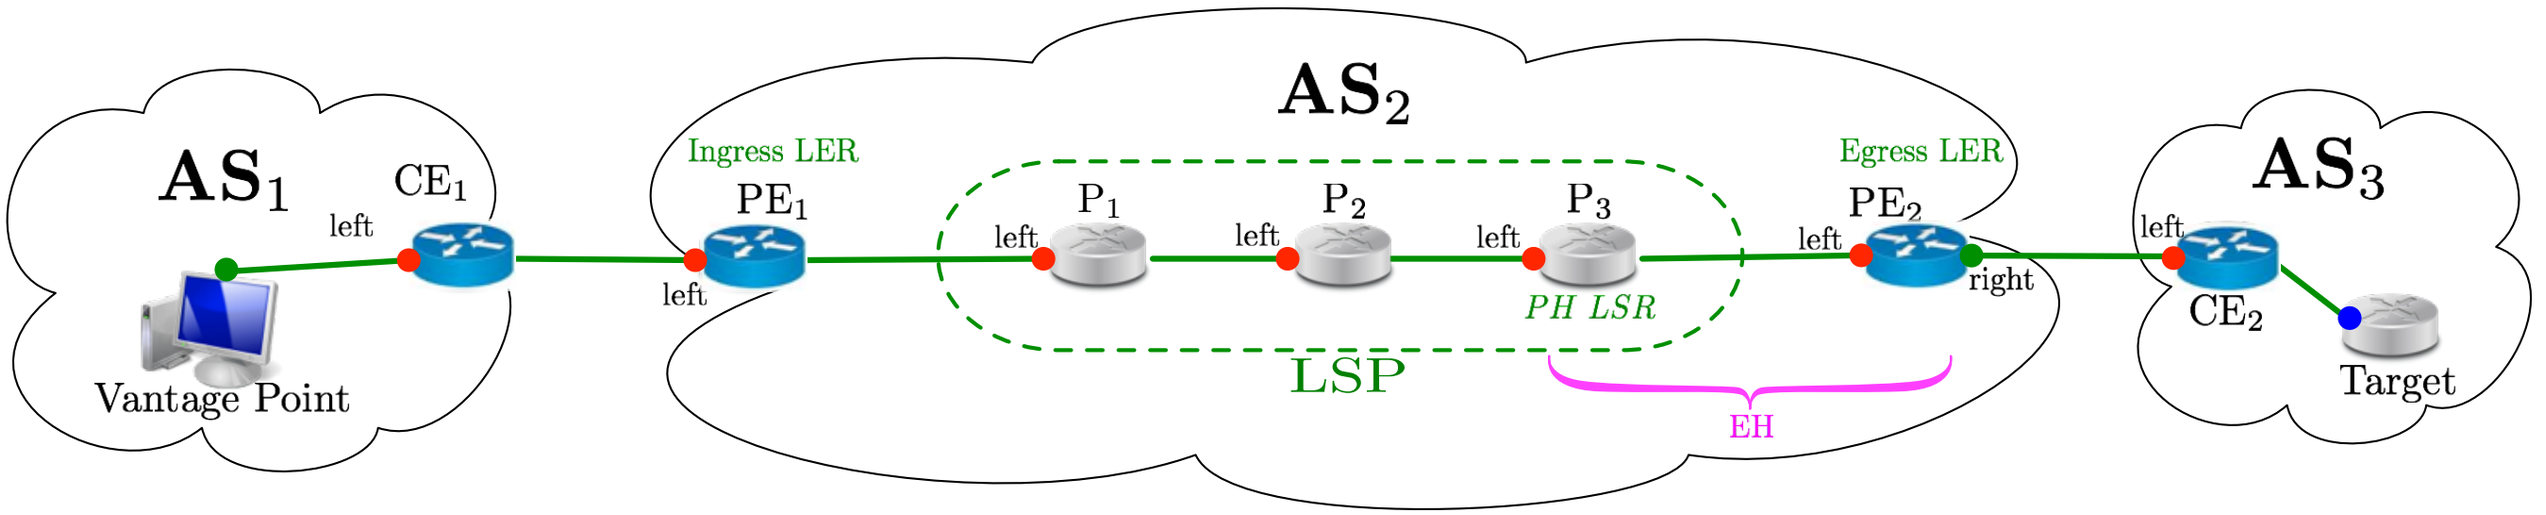

Explicit

L^T	Hop
1.	CE_1.left
2.	PE_1.left
3.	P_1.left – MPLS
4.	P_2.left – MPLS
5.	P_3.left – MPLS
6.	PE_2.left
7.	CE_2.left
8.	Target

Implicit

L^T	(L_R^{TE},L_R^{ER},qTTL)	Hop
1.	(1,1,1)	CE_1.left
2.	(2,2,1)	PE_1.left
3.	(9,3,1)	P_1.left
4.	(8,4,2)	P_2.left
5.	(5,5,3)	P_3.left
6.	(6,6,1)	PE_2.left
7.	(7,7,1)	CE_2.left
8.	(8,8,1)	Target

Opaque

L^T	(L_R^{TE},L_R^{ER})	Hop	LSE-TTL
1.	(1,1)	CE_1.left	
2.	(2,2)	PE_1.left	
3.	(6,6)	PE_2.left – MPLS	252
4.	(6,6)	CE_2.left	
5.	(6,6)	Target	

Invisible PHP

L^T	(L_R^{TE},L_R^{ER} ($\|LJ_R^{ER}$))	Hop
1.	(1,1)	CE_1.left
2.	(2,2)	PE_1.left
3.	(6,6|3)	PE_2.left
4.	(6,6|4)	CE_2.left
5.	(6,6|5)	Target

Invisible UHP

L^T	(L_R^{TE},L_R^{ER})	Hop
1.	(1,1)	CE_1.left
2.	(2,2)	PE_1.left
3.	(4,4)	CE_2.left
4.	(4,4)	CE_2.left
5.	(5,5)	Target

LSE headers — qTTL — Uturn — 1 << LSE-TTL < 255 — Rtla — Frpla — Dup_Ip

TNT indicators and triggers

Tunnels visibile to `traceroute` — Invisible tunnels revealed with TNT

Legend

IP interface

Border Router

Internal Router

L^T := `traceroute` hop
L_R^{TE} := Return path `time_exceeded` length
L_R^{ER} := Return path `echo_reply` length
LJ_R^{ER} := Return path `echo_reply` length (Juniper)

Invisible Tunnels

- In case of invisible tunnels
  - LSRs do not appear in `traceroute` output
  - MPLS labels are not included in the `time_exceeded` messages sent by the LH (in case of PHP) or the Egress LER (in case of UHP)
- We need triggers to infer their presence
  - Y. Vanaubel, P. Mérindol, J.-J. Pansiot, B. Donnet. *Through the Wormhole: Tracking MPLS Invisibile Tunnels*. In Proc. ACM Internet Measurement Conference (IMC). November 2017.
  - Y. Vanaubel, P. Mérindol, J.-J. Pansiot, B. Donnet. *TNT, Watch me Explode: A Light in the Dark for Revealing All MPLS Tunnels*. Under Submission. June 2018

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

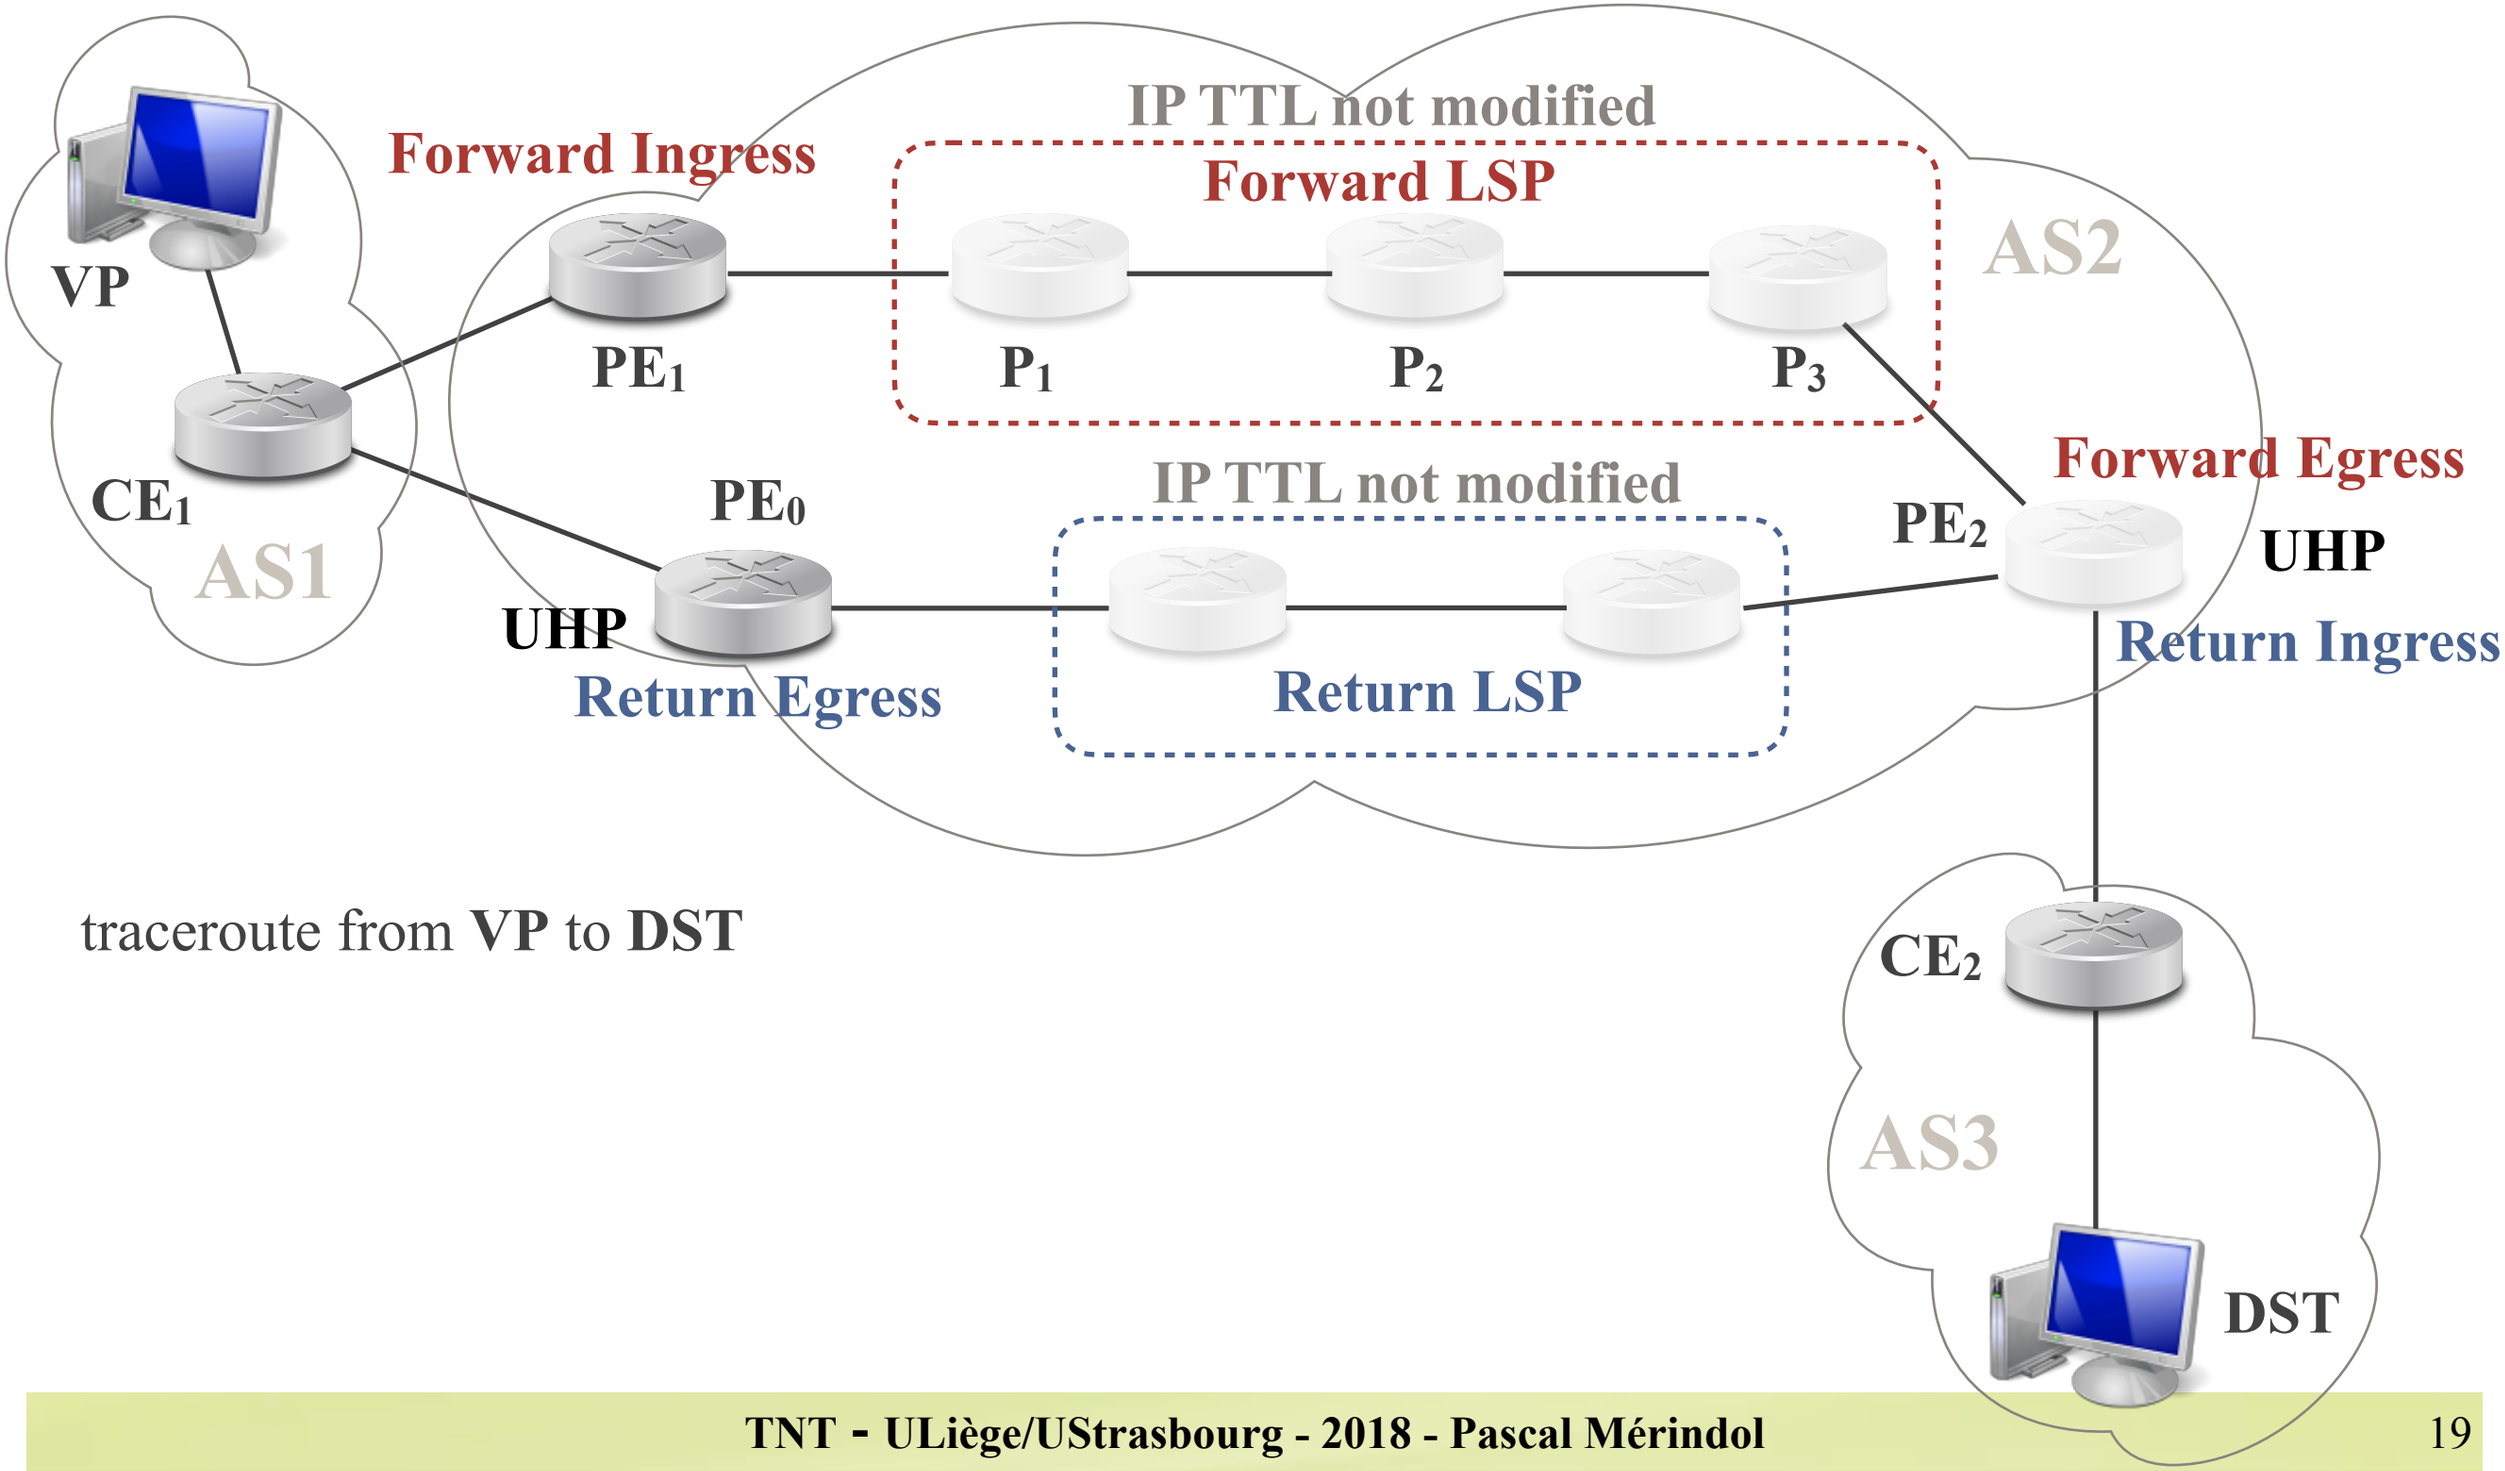

traceroute from **VP** to **DST**

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

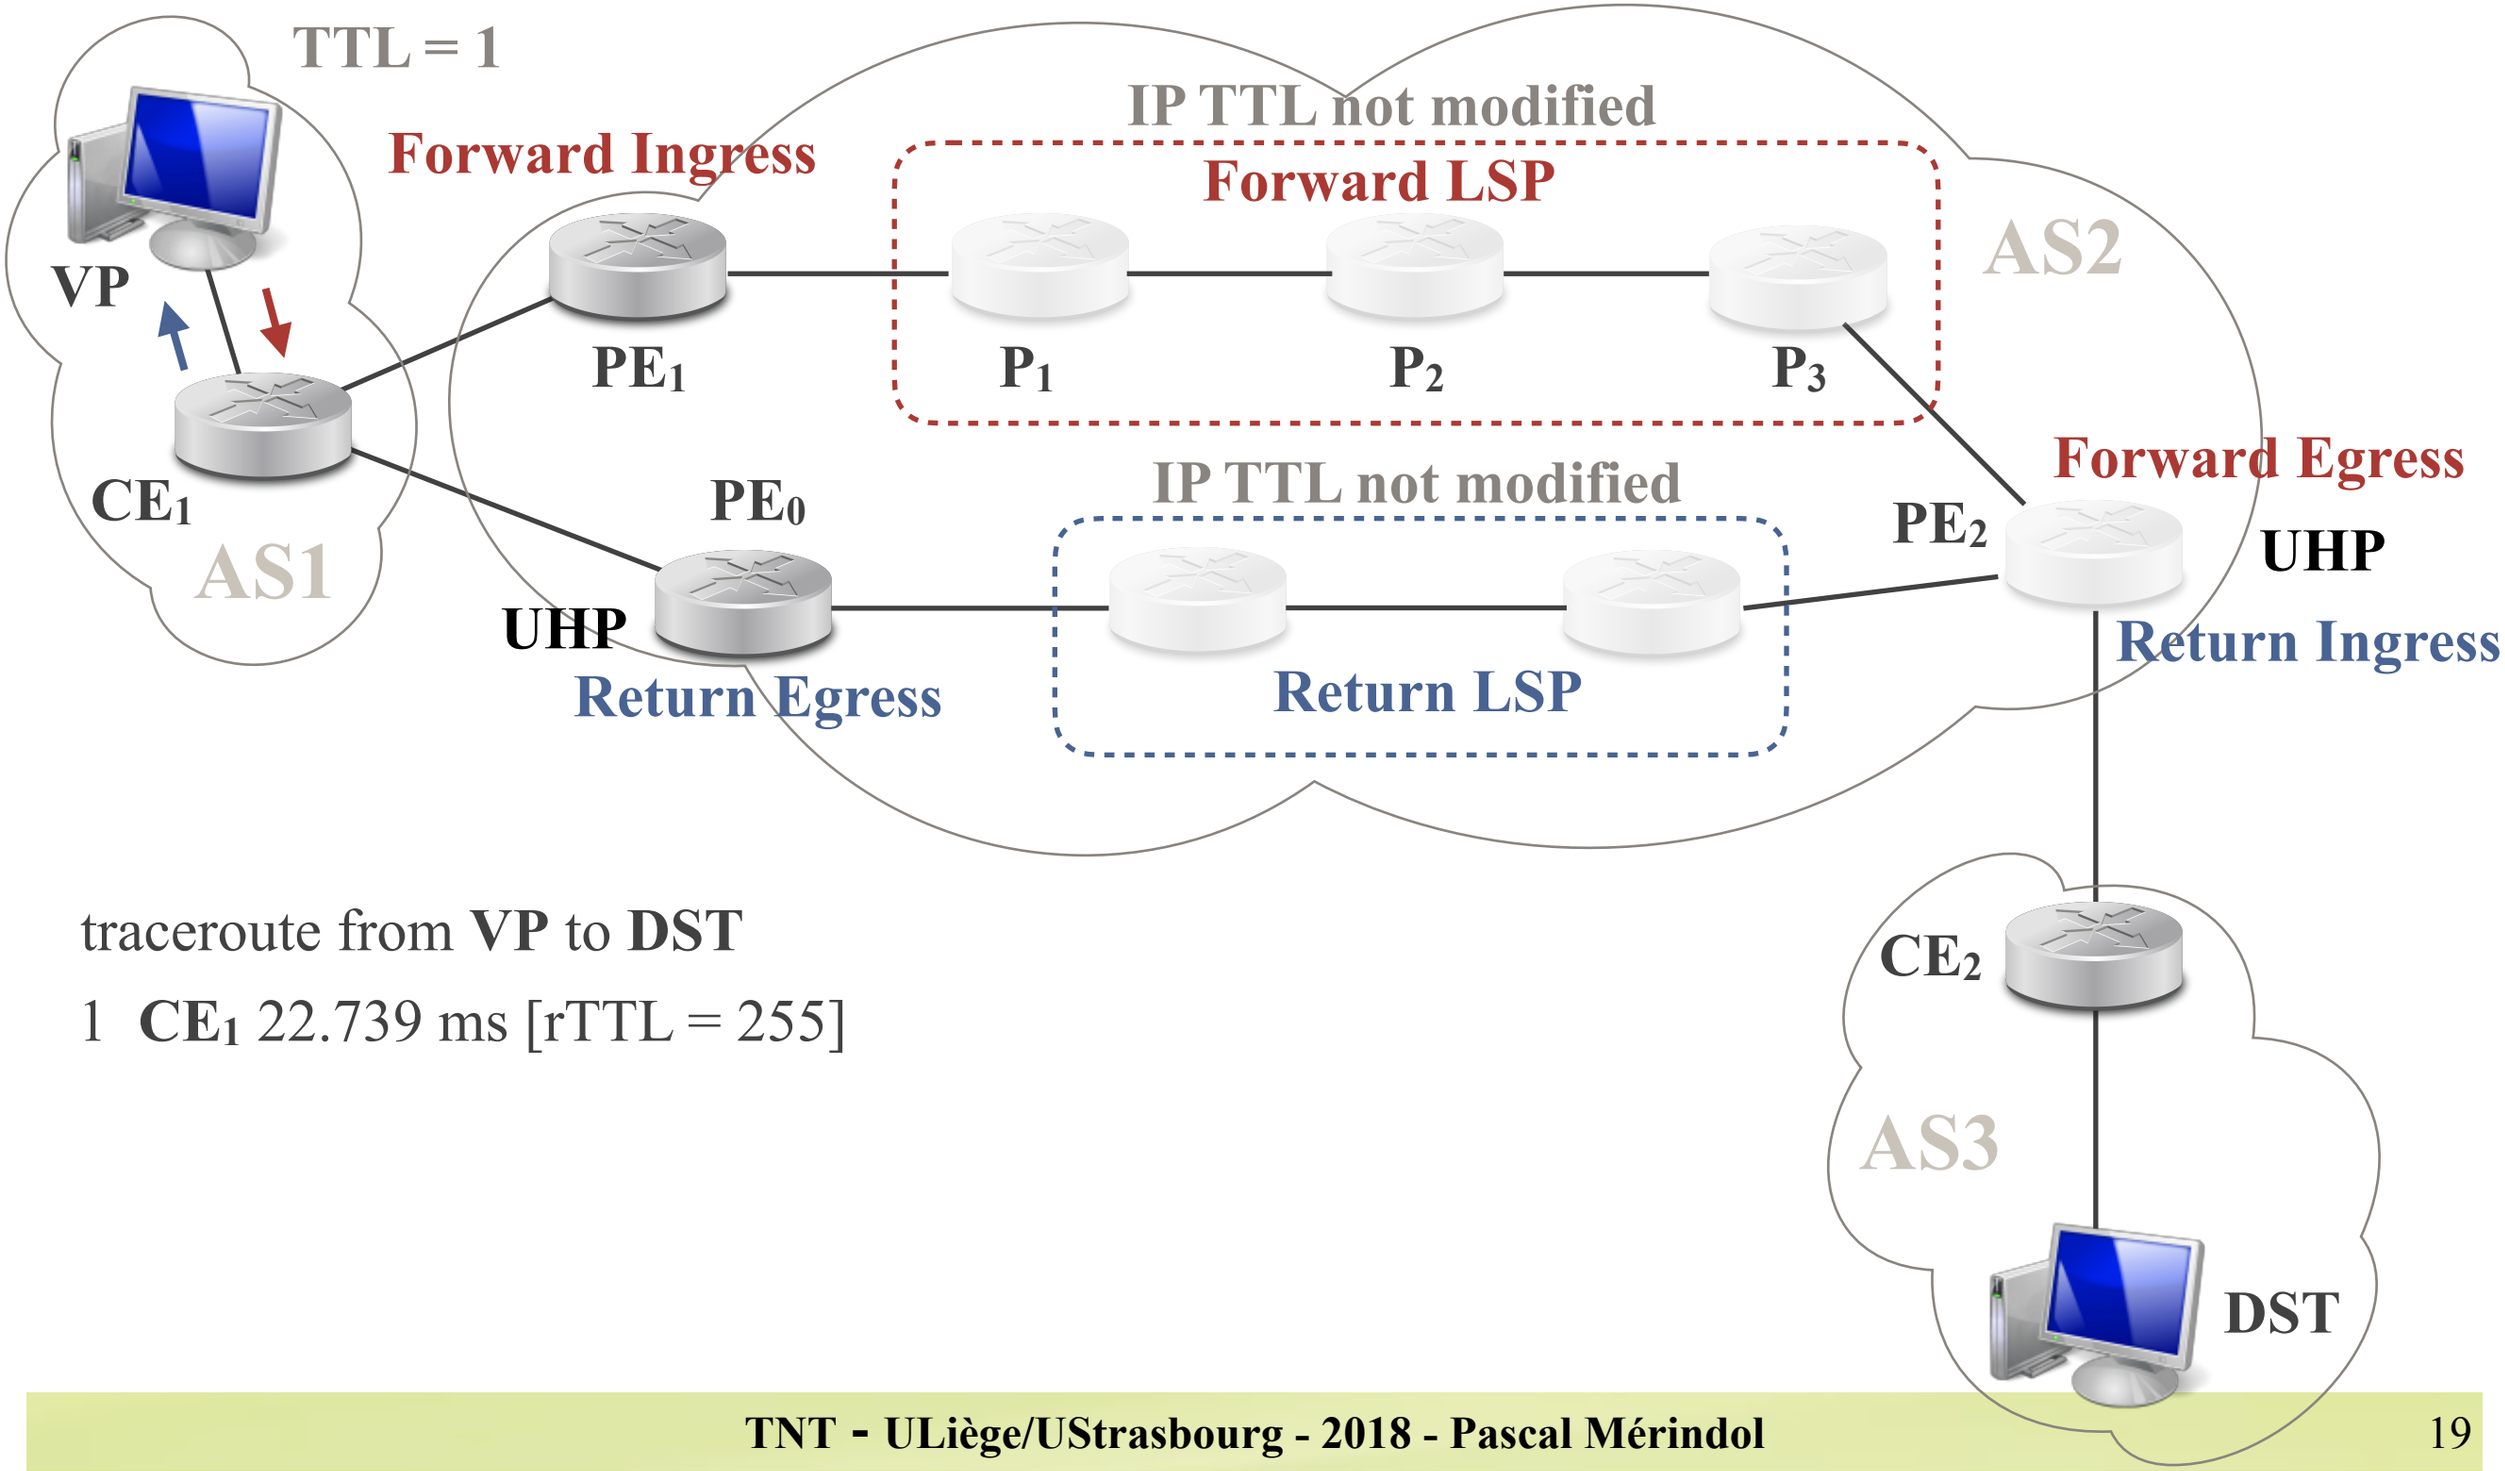

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

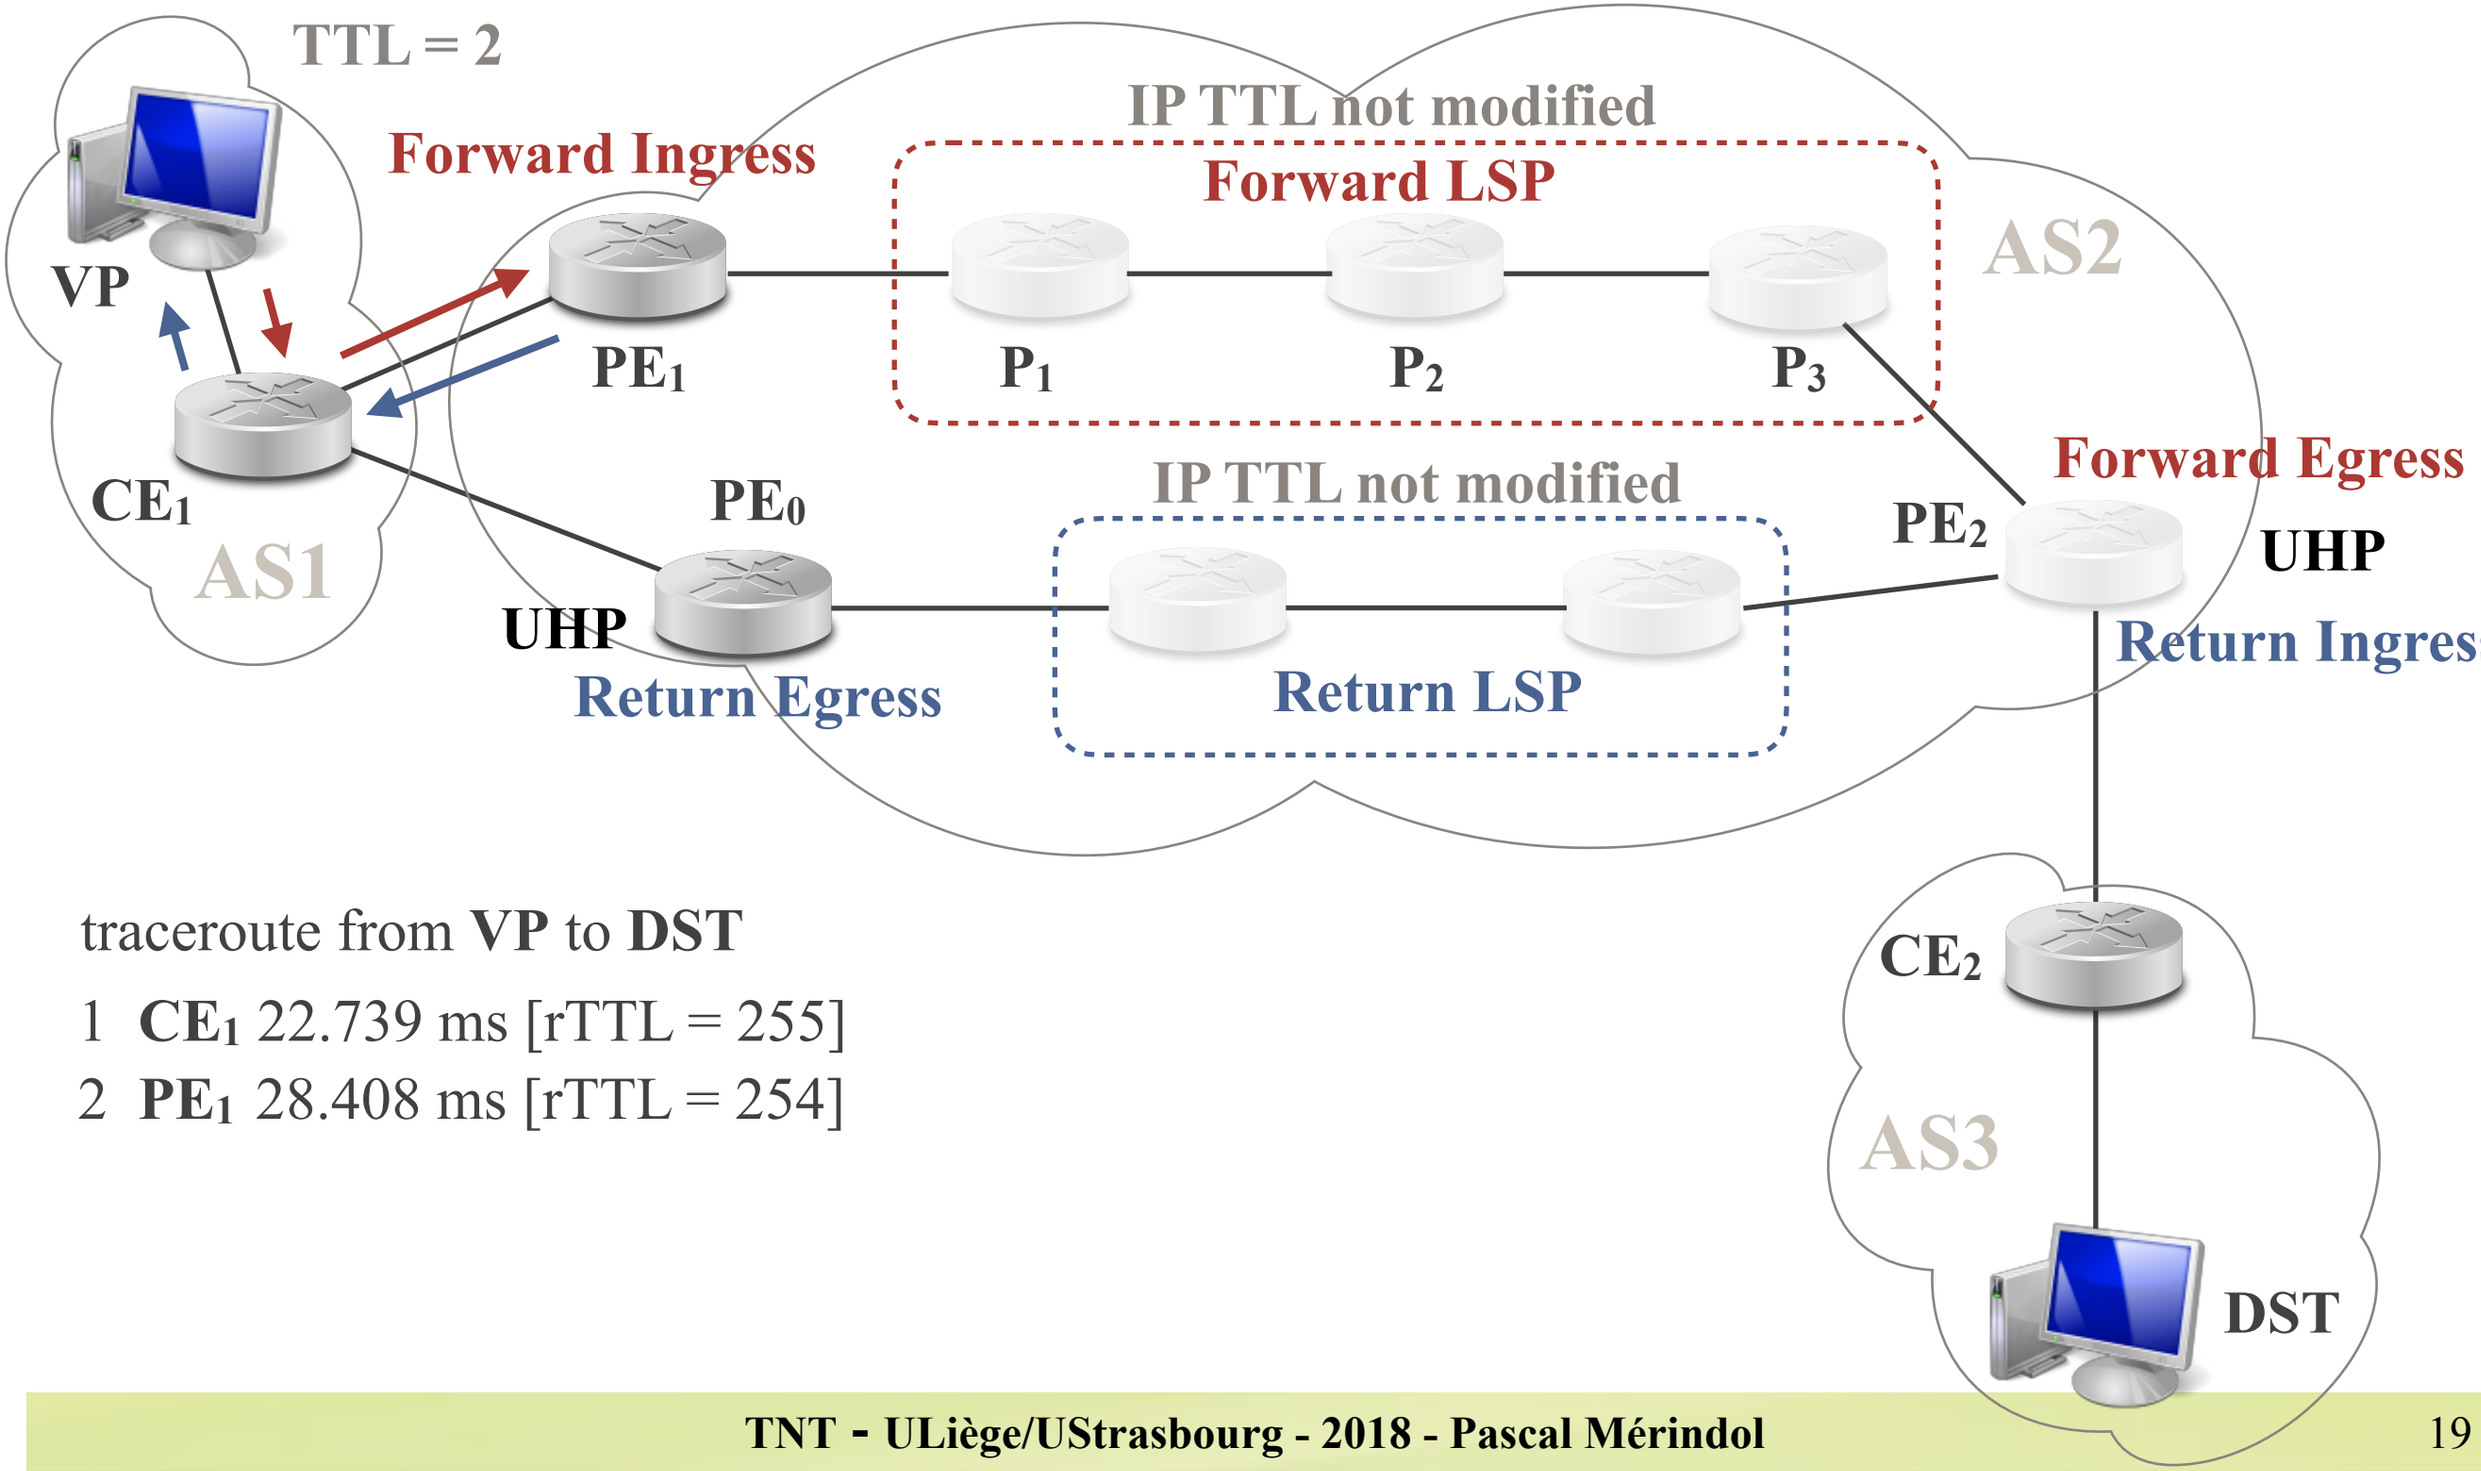

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]

2 **PE1** 28.408 ms [rTTL = 254]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

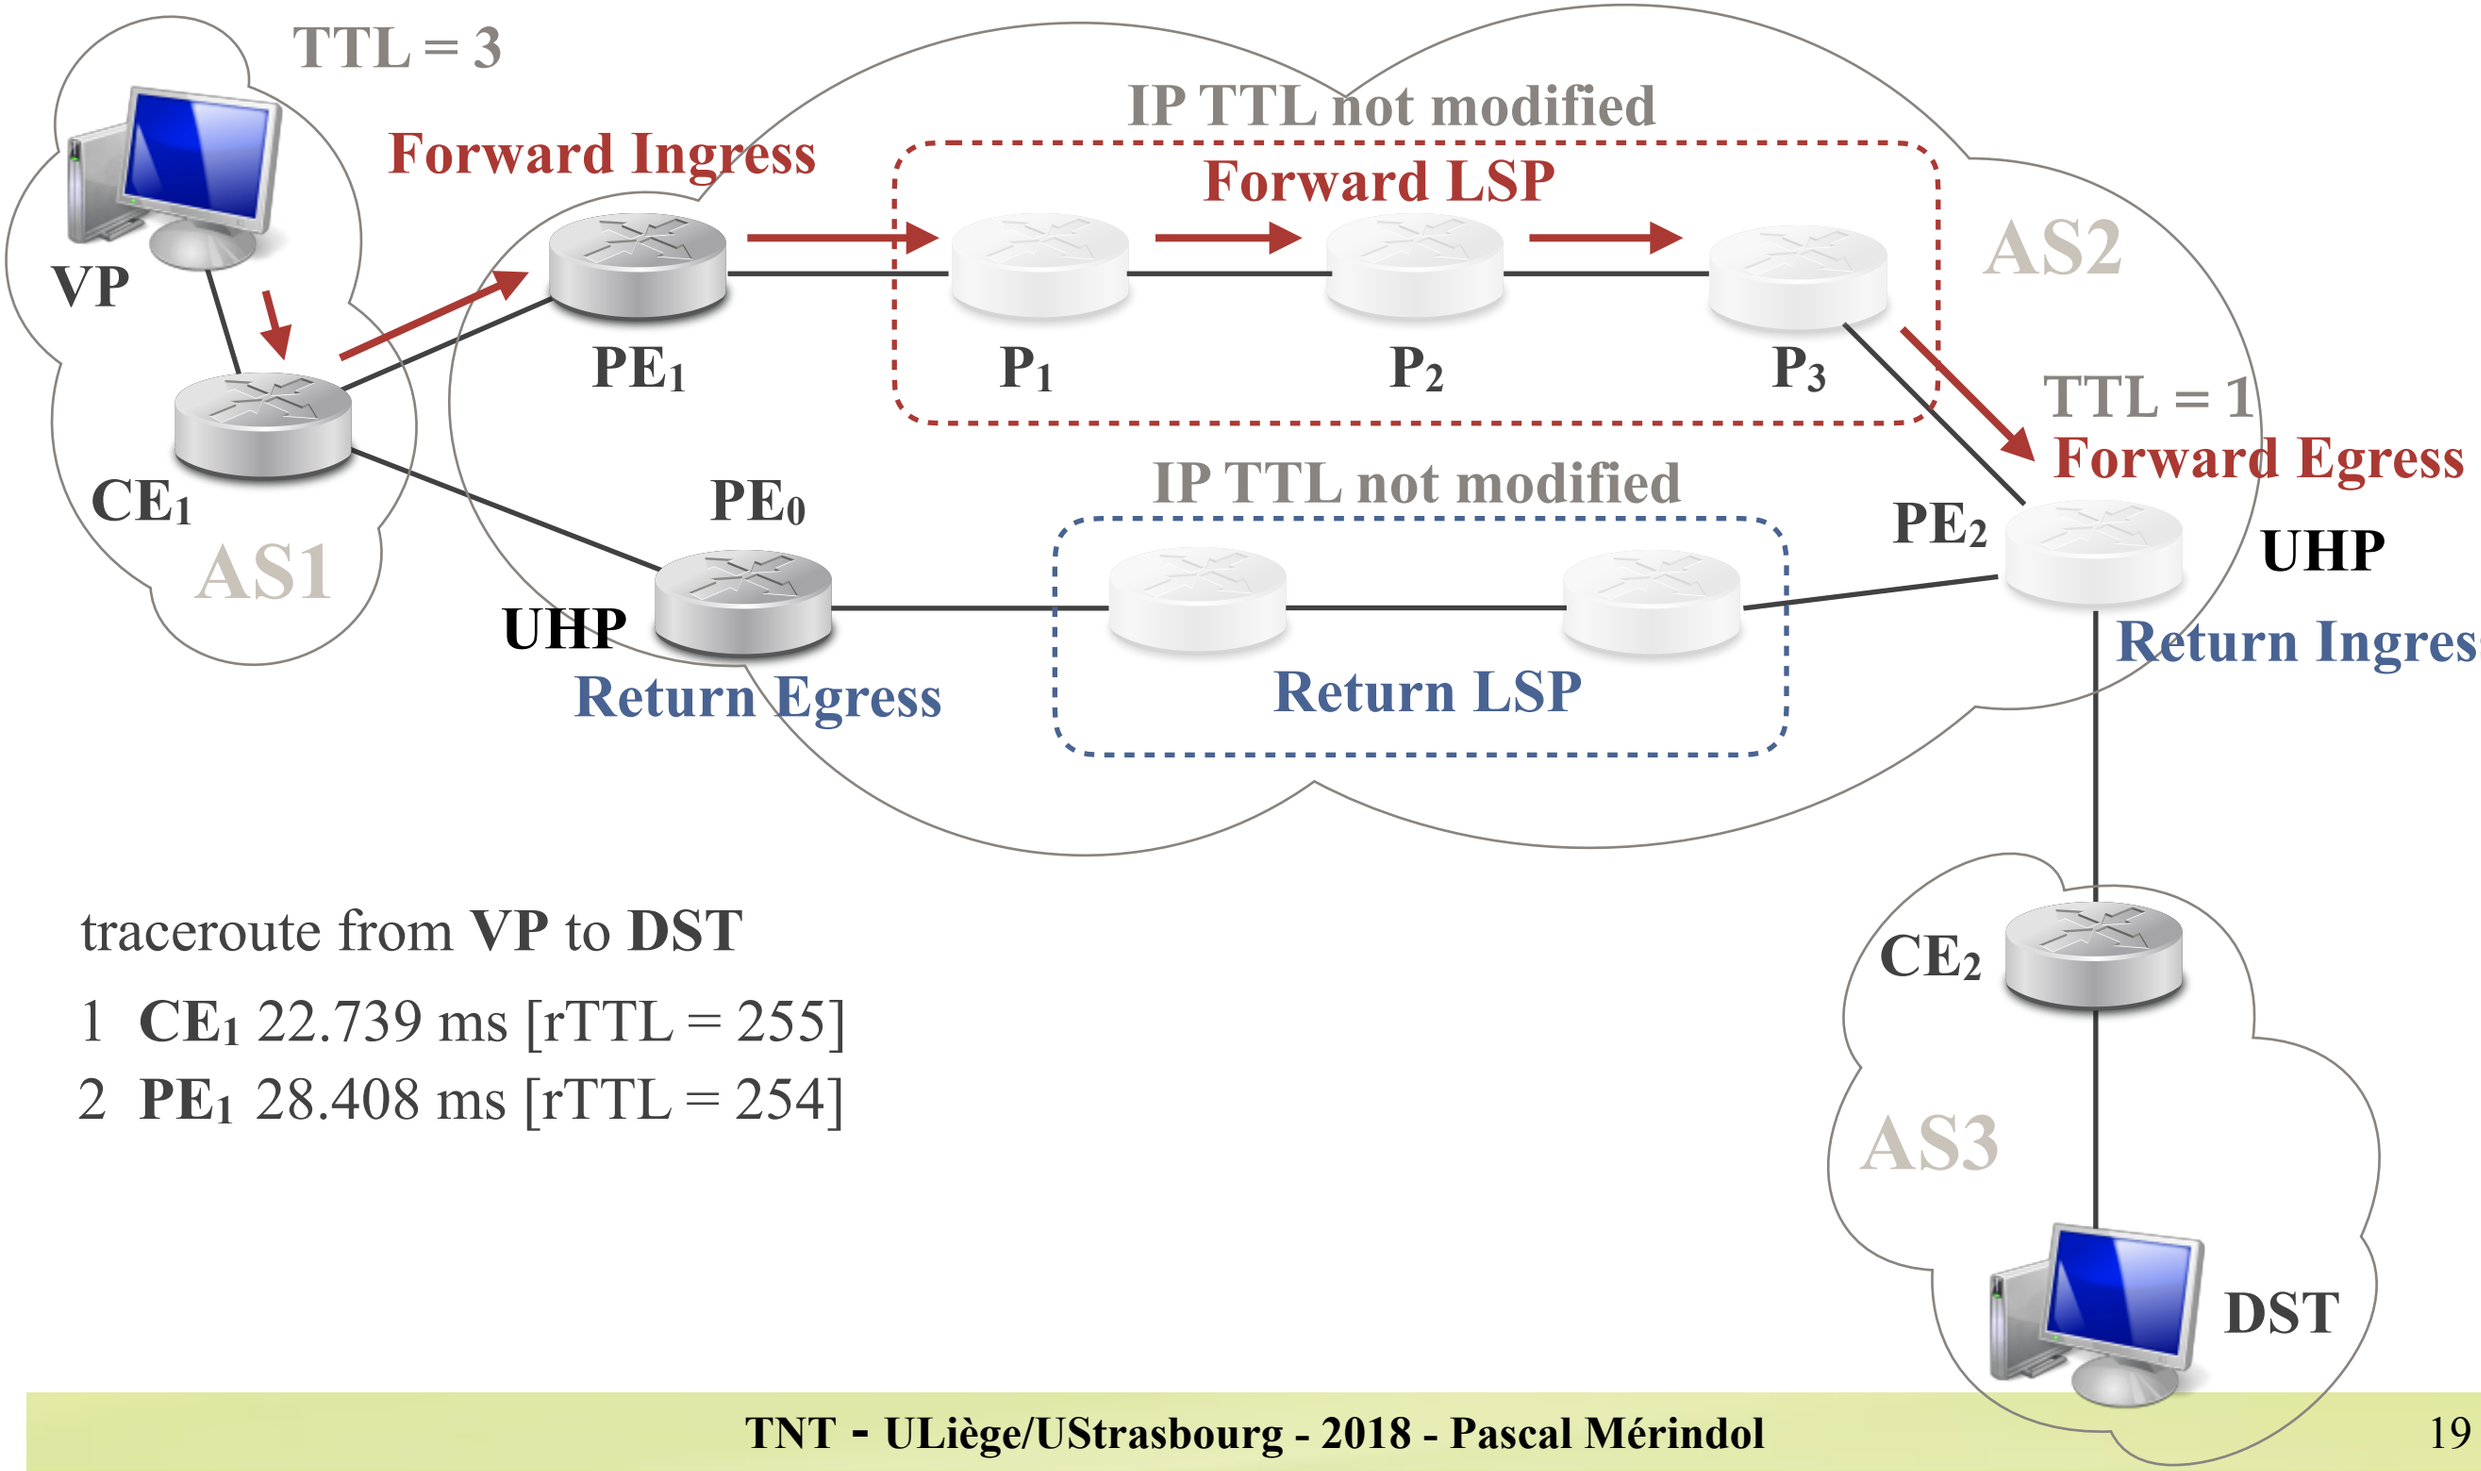

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]

2 **PE1** 28.408 ms [rTTL = 254]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

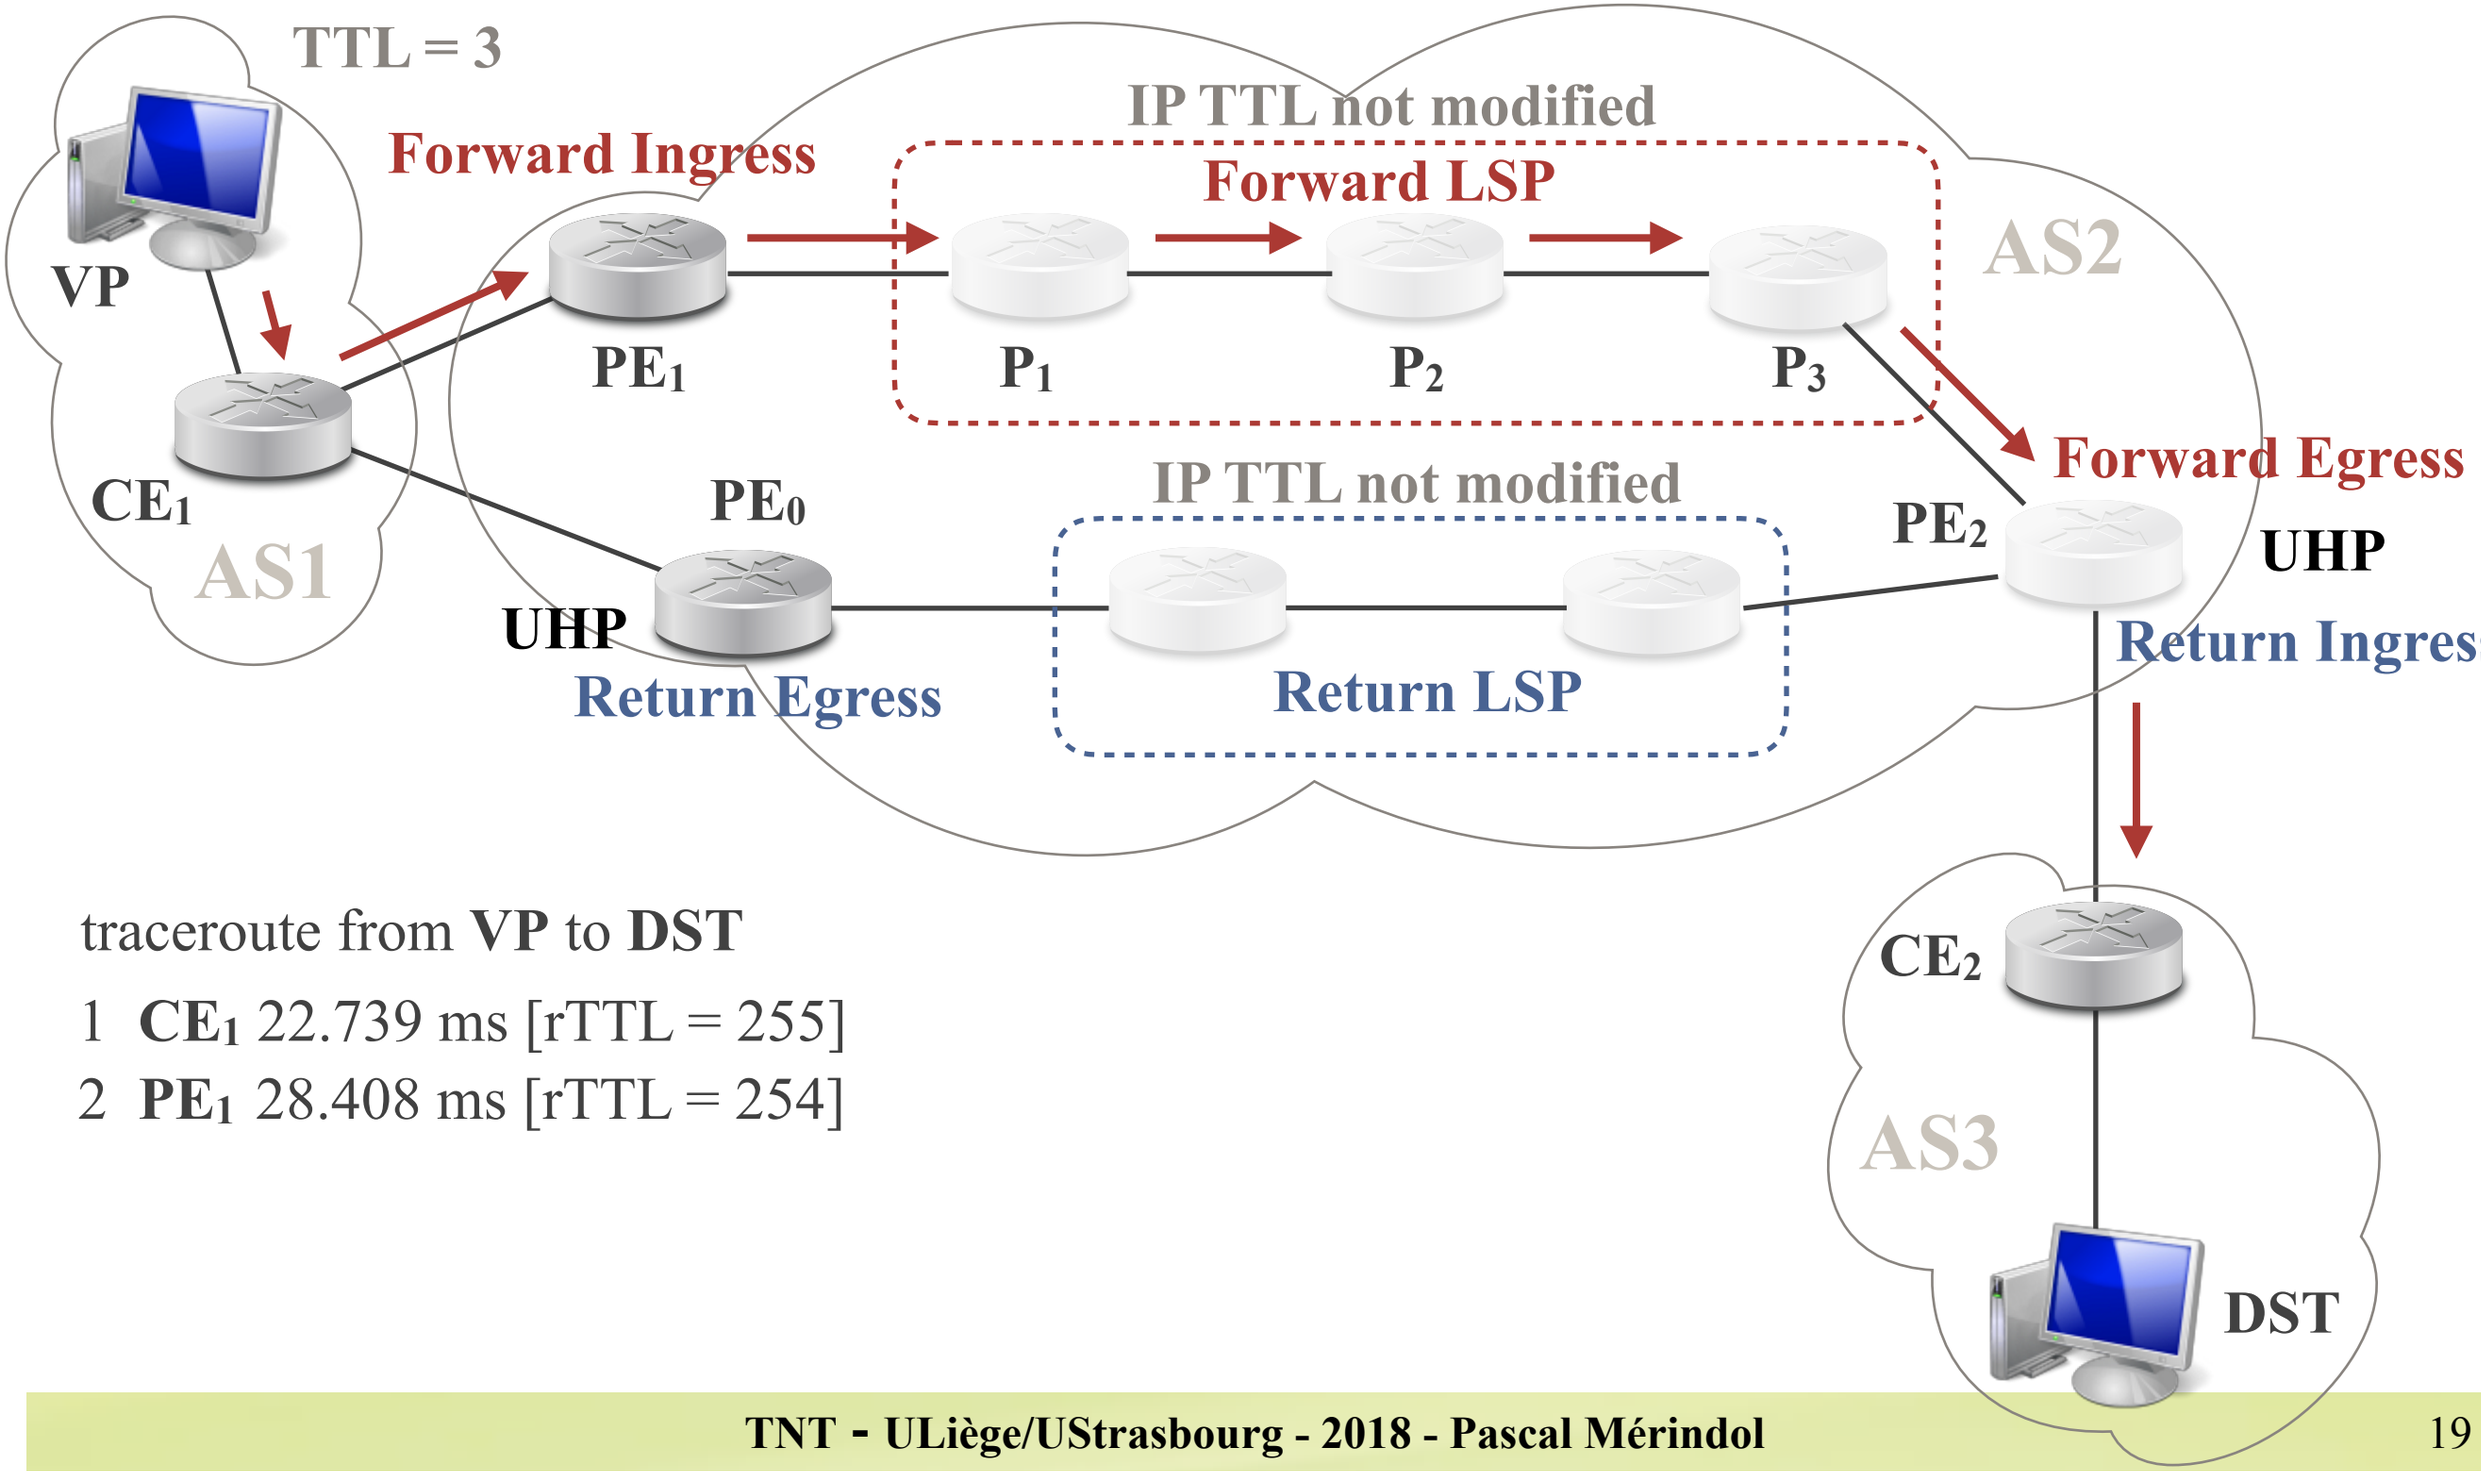

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]

2 **PE1** 28.408 ms [rTTL = 254]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

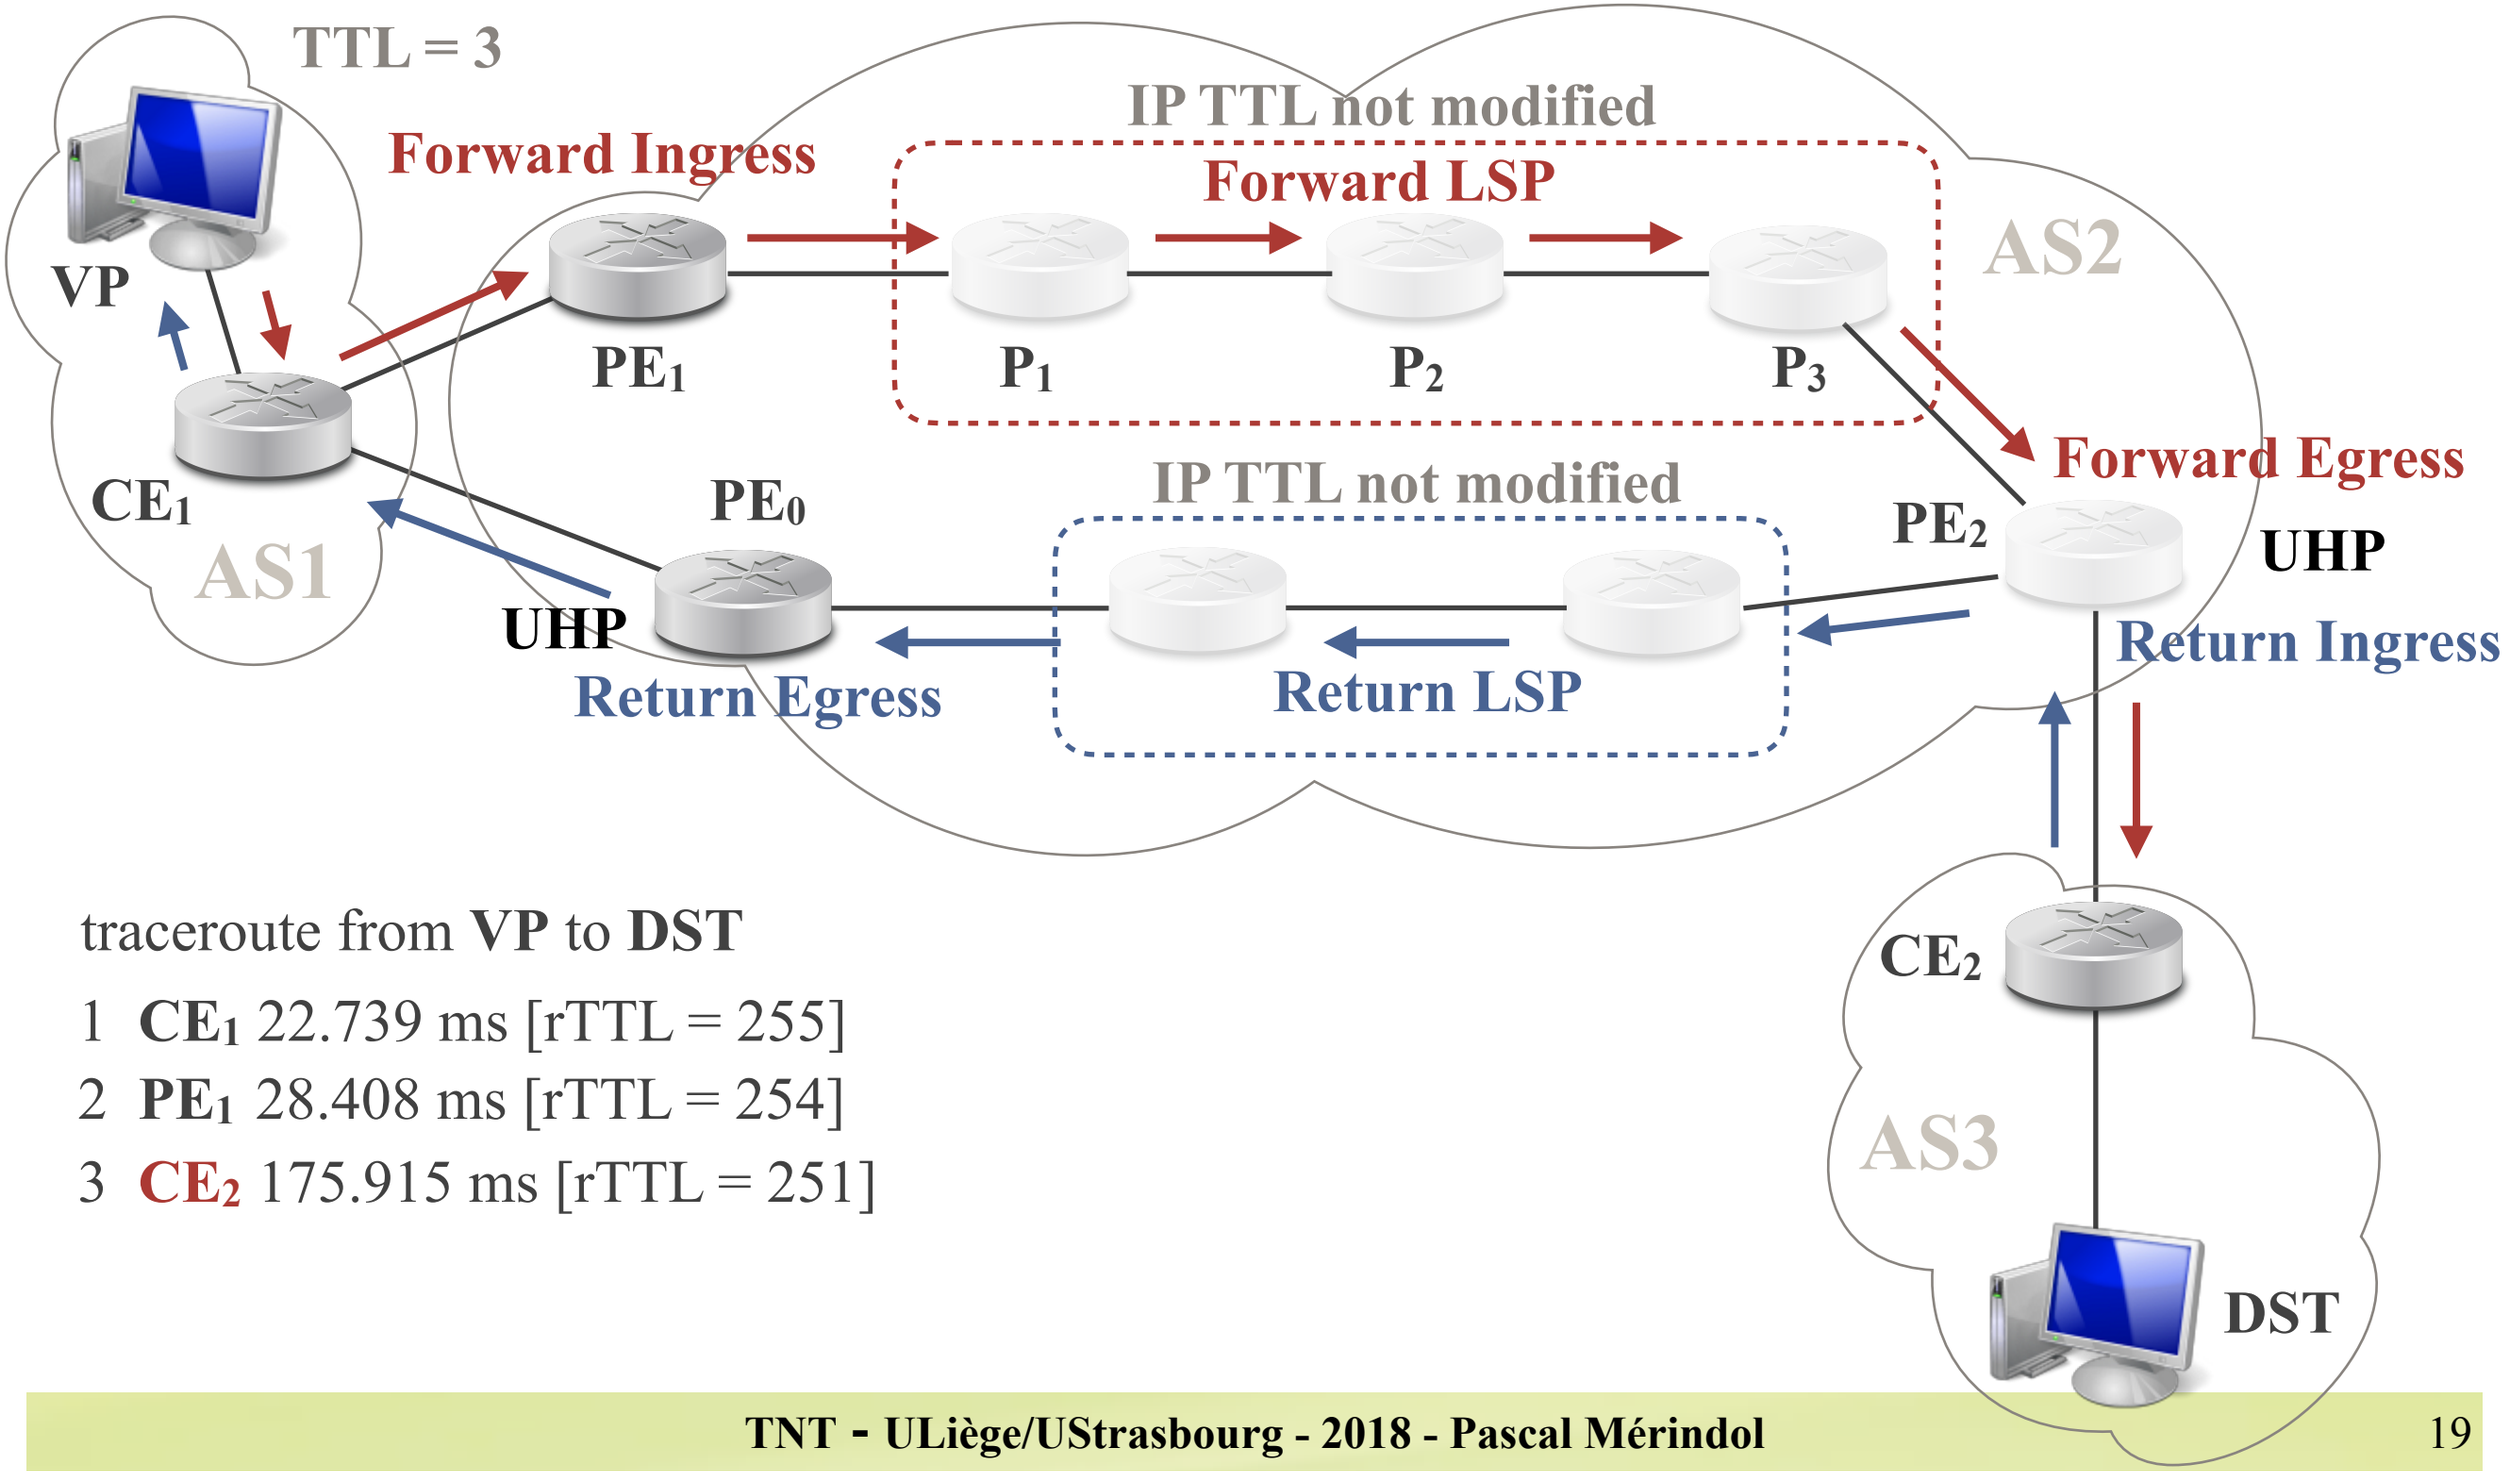

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]

2 **PE1** 28.408 ms [rTTL = 254]

3 **CE2** 175.915 ms [rTTL = 251]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

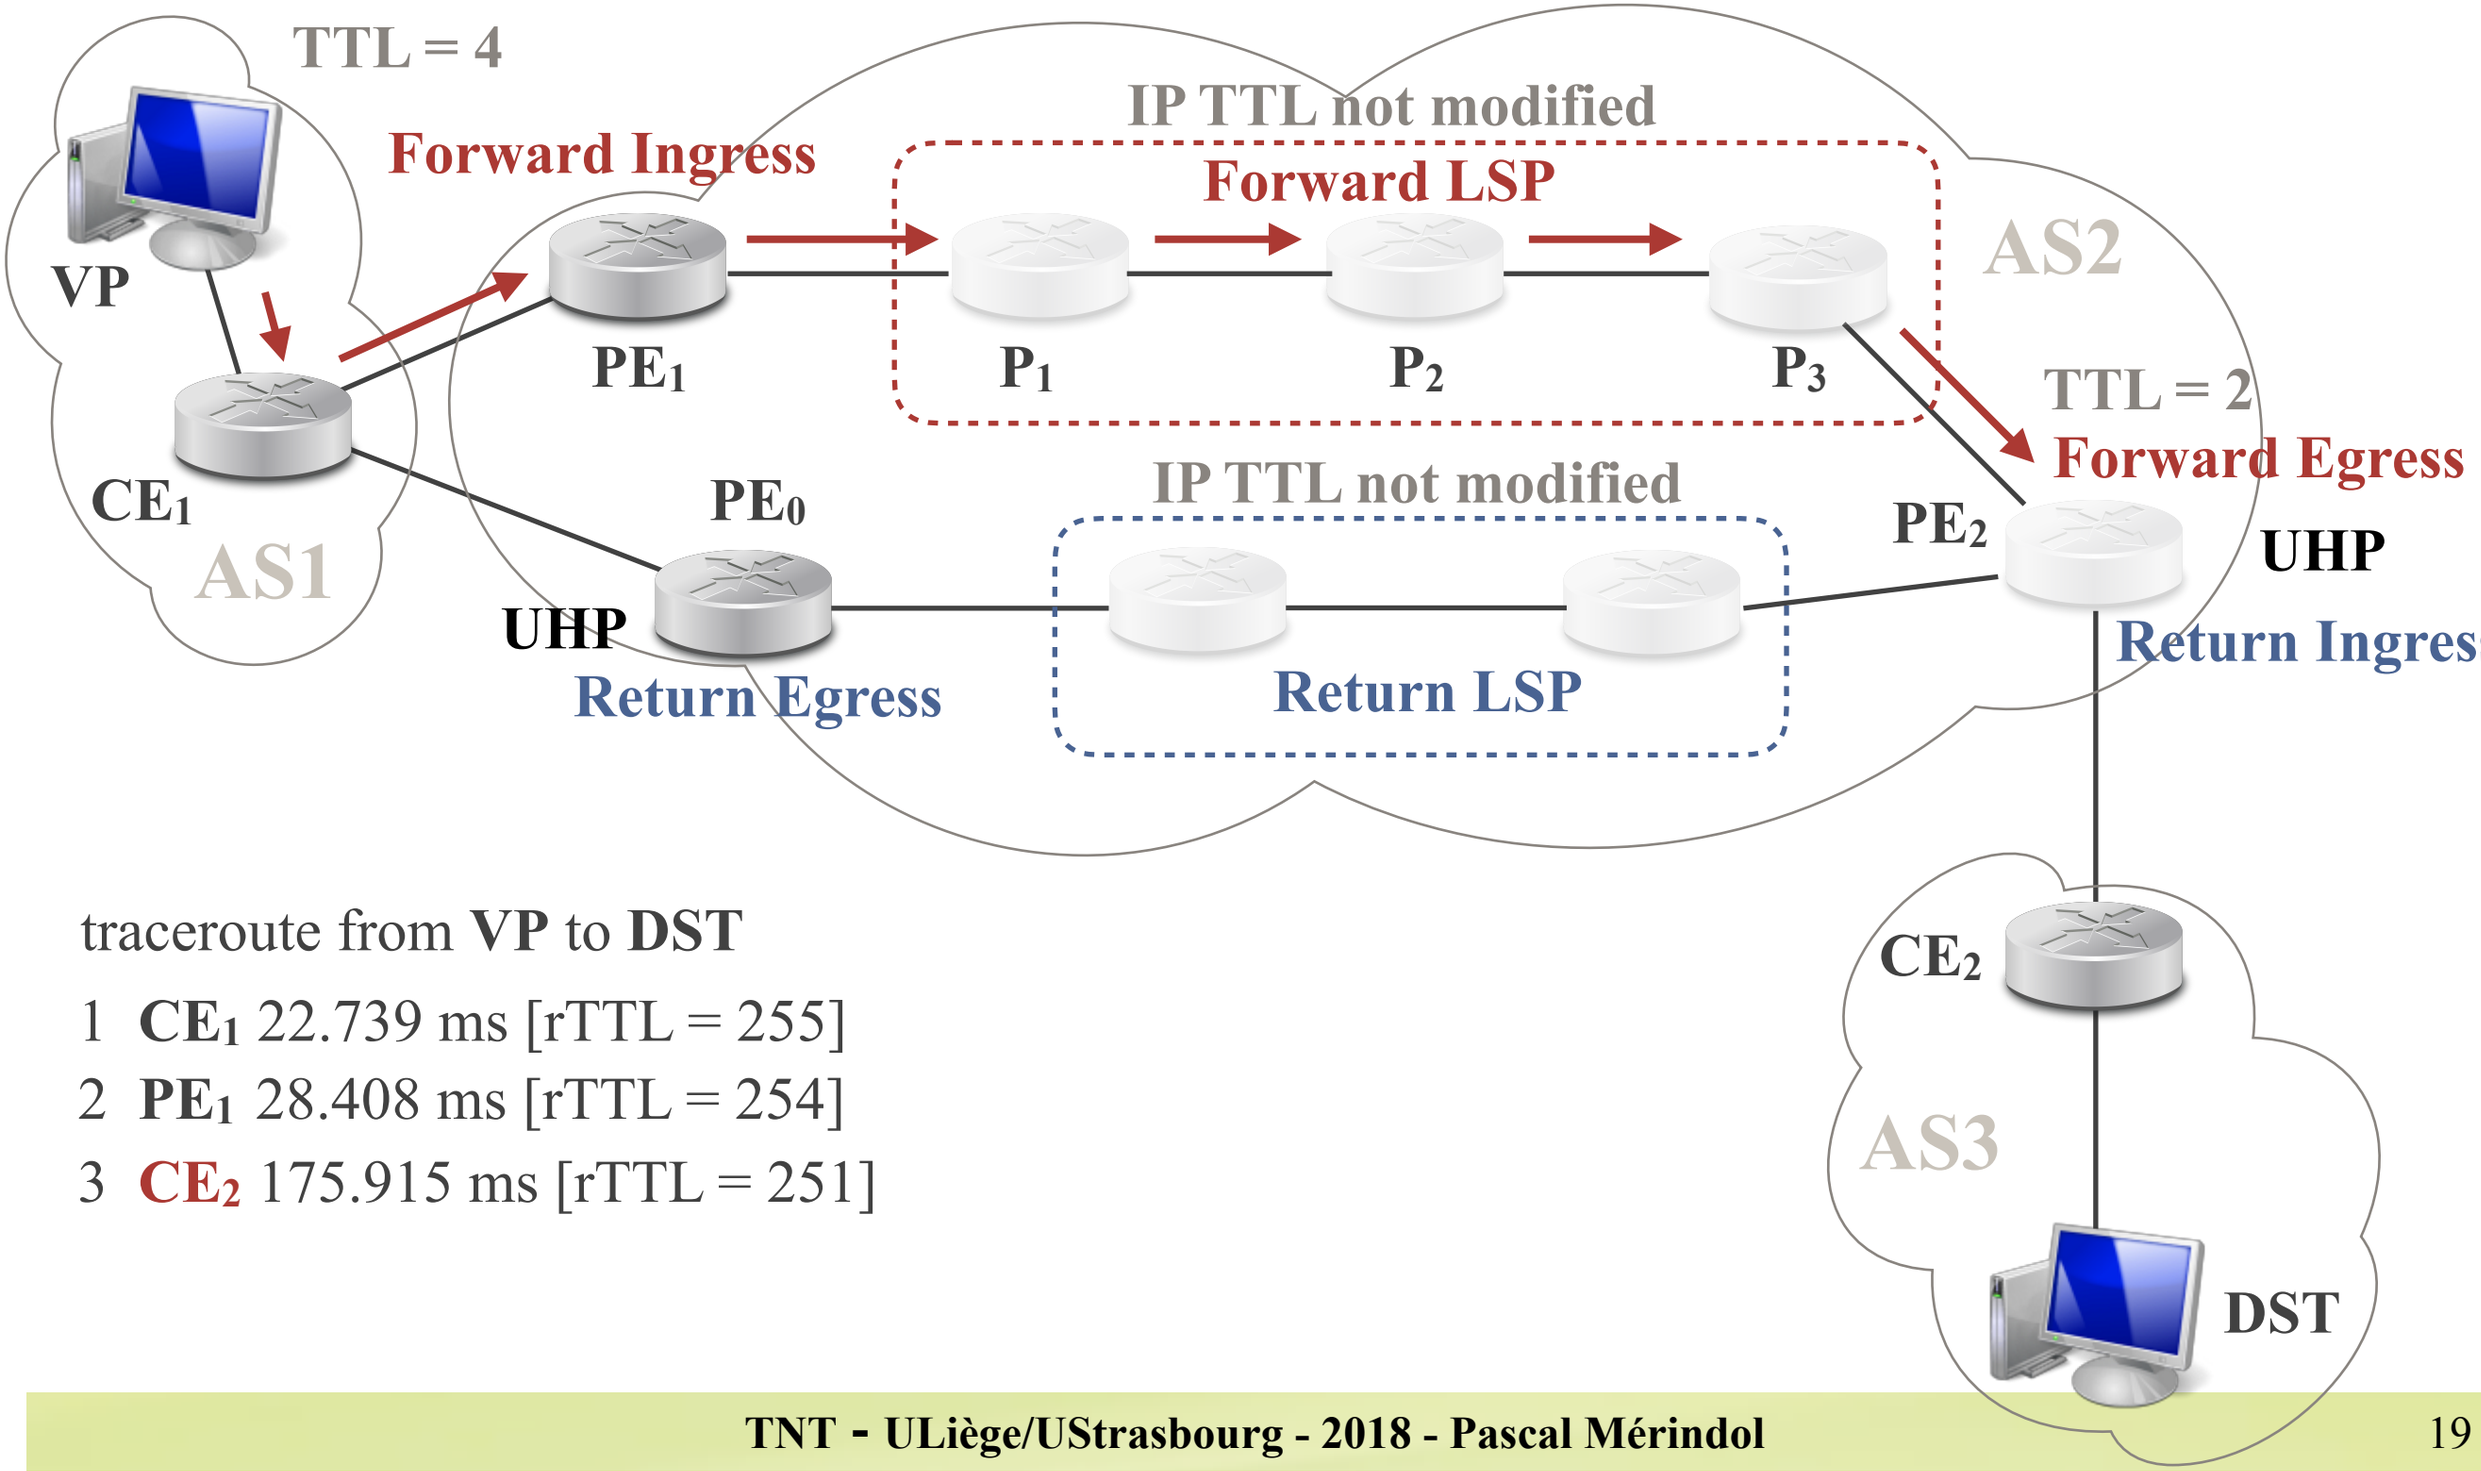

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]
2 **PE1** 28.408 ms [rTTL = 254]
3 **CE2** 175.915 ms [rTTL = 251]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

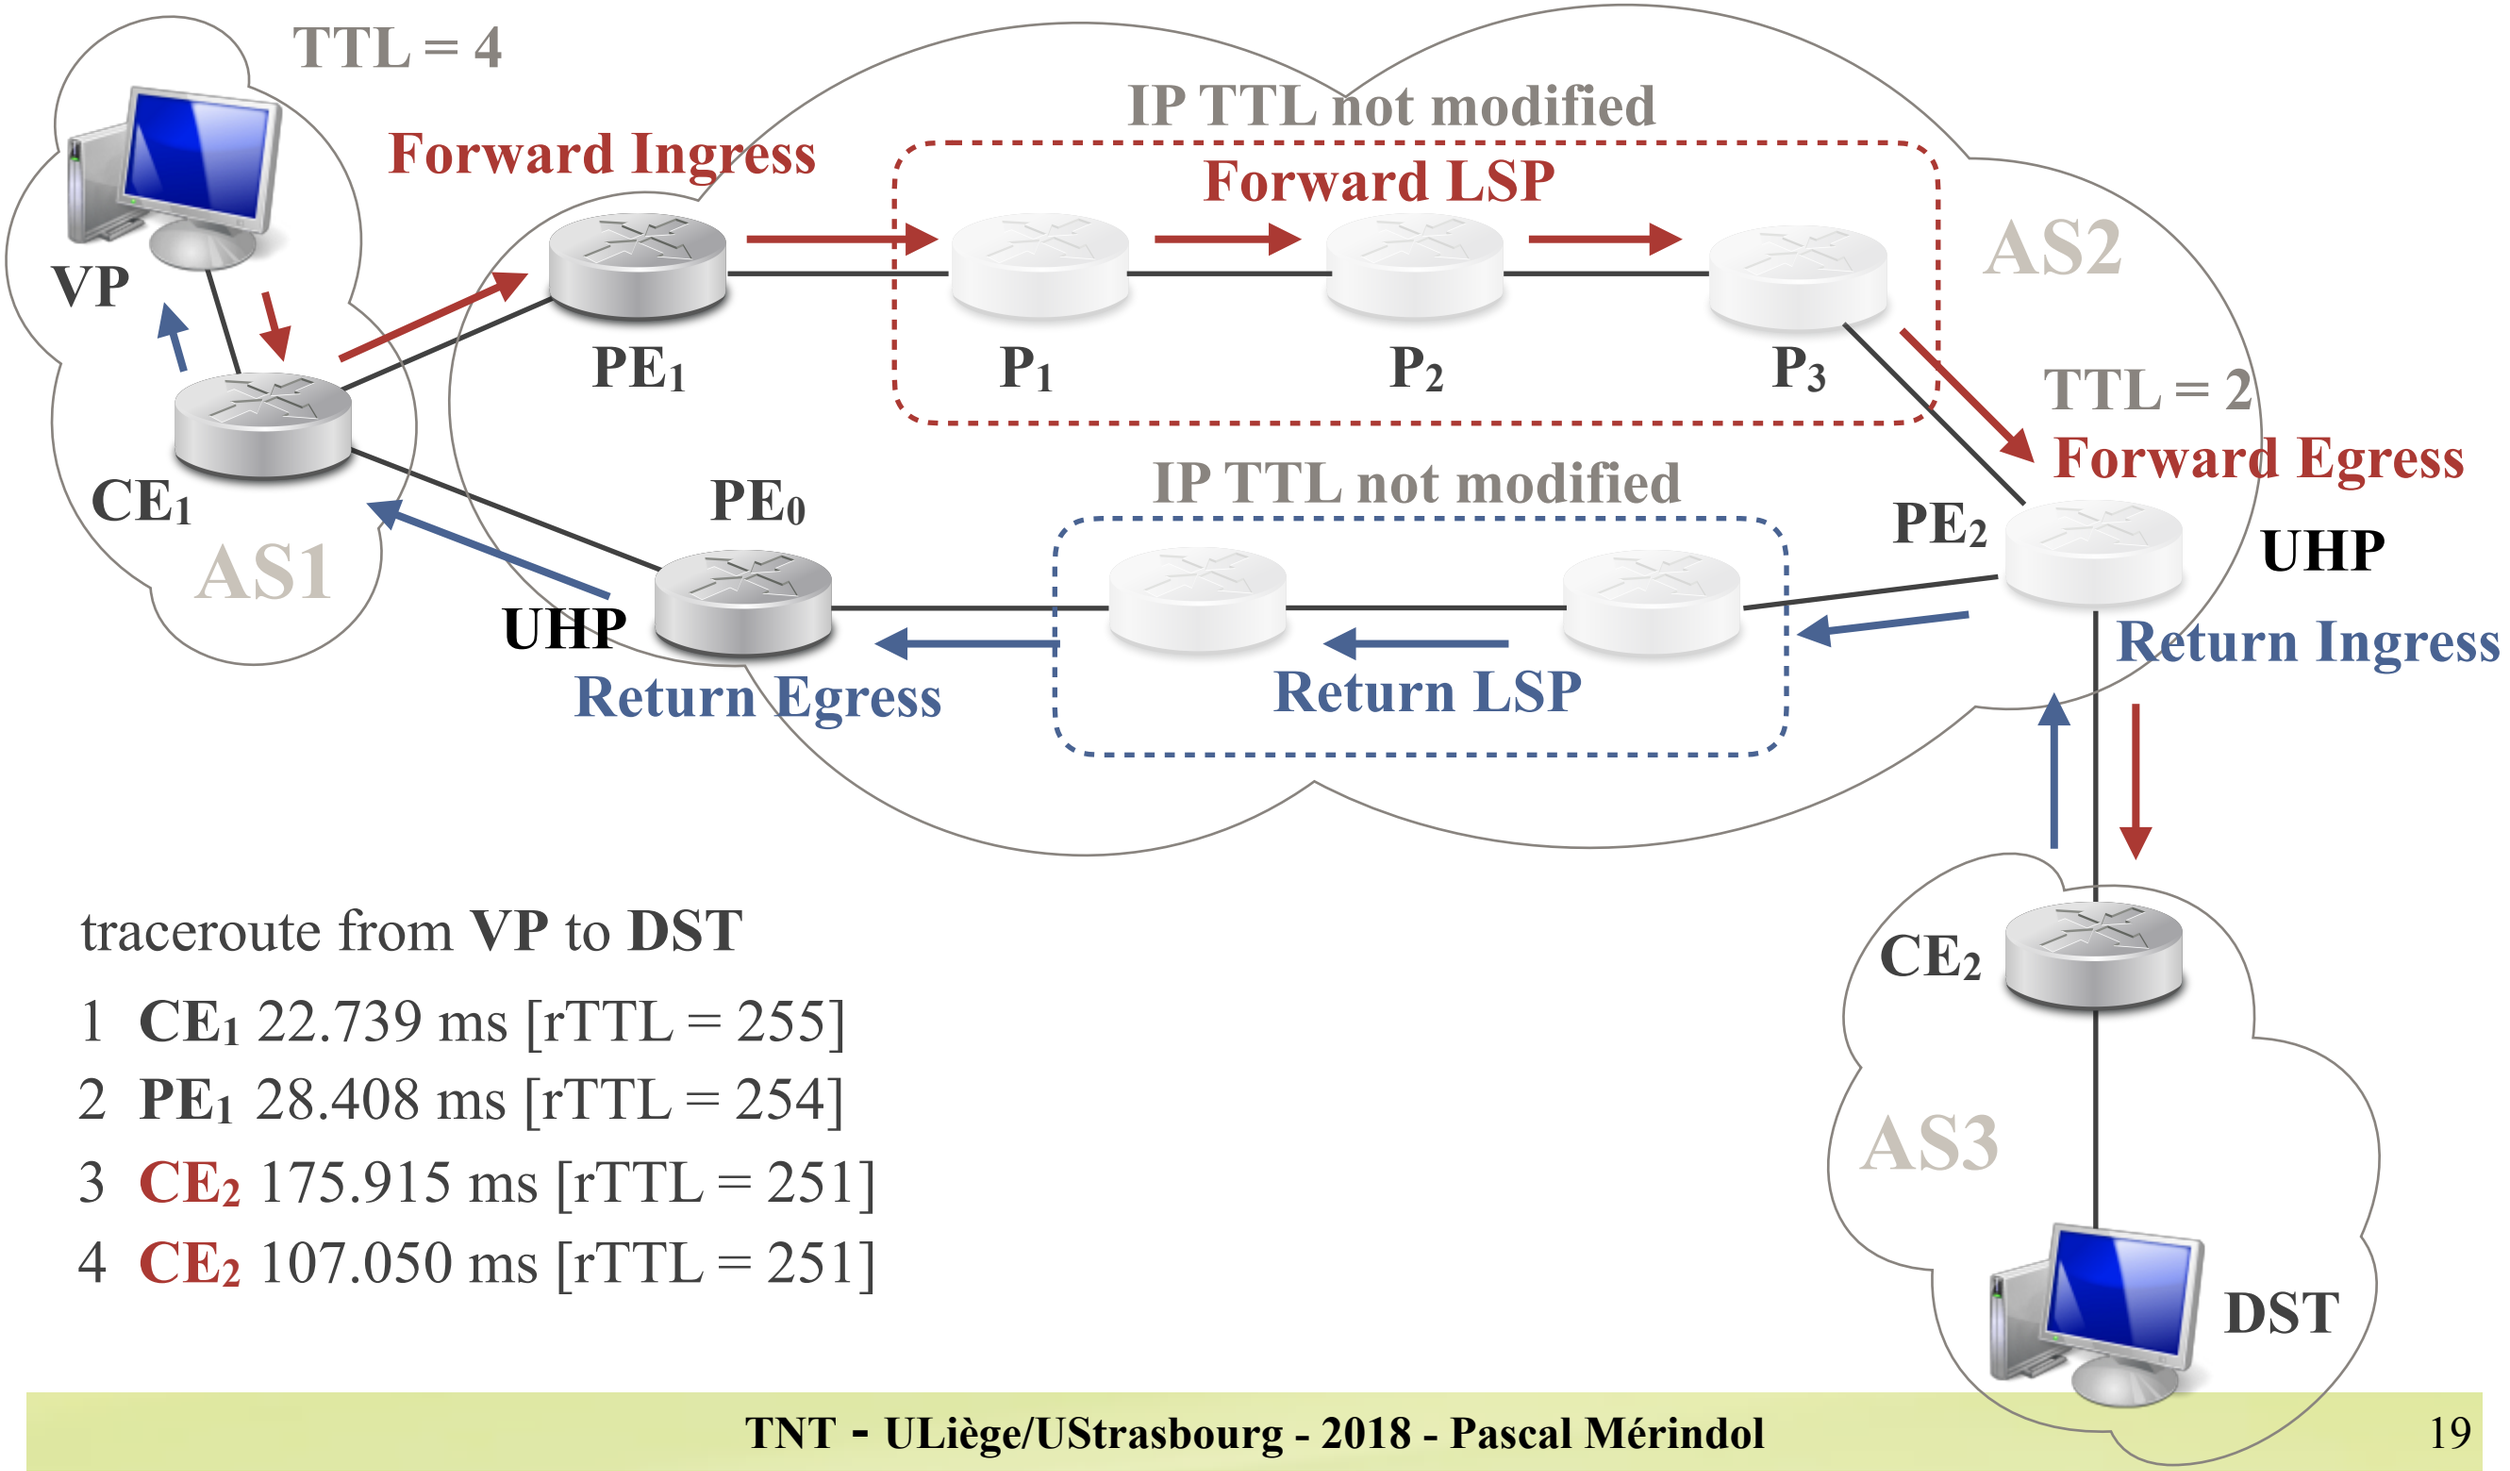

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]
2 **PE1** 28.408 ms [rTTL = 254]
3 **CE2** 175.915 ms [rTTL = 251]
4 **CE2** 107.050 ms [rTTL = 251]

Invisible Tunnels (2)

- Trigger 1: Duplicate IP address

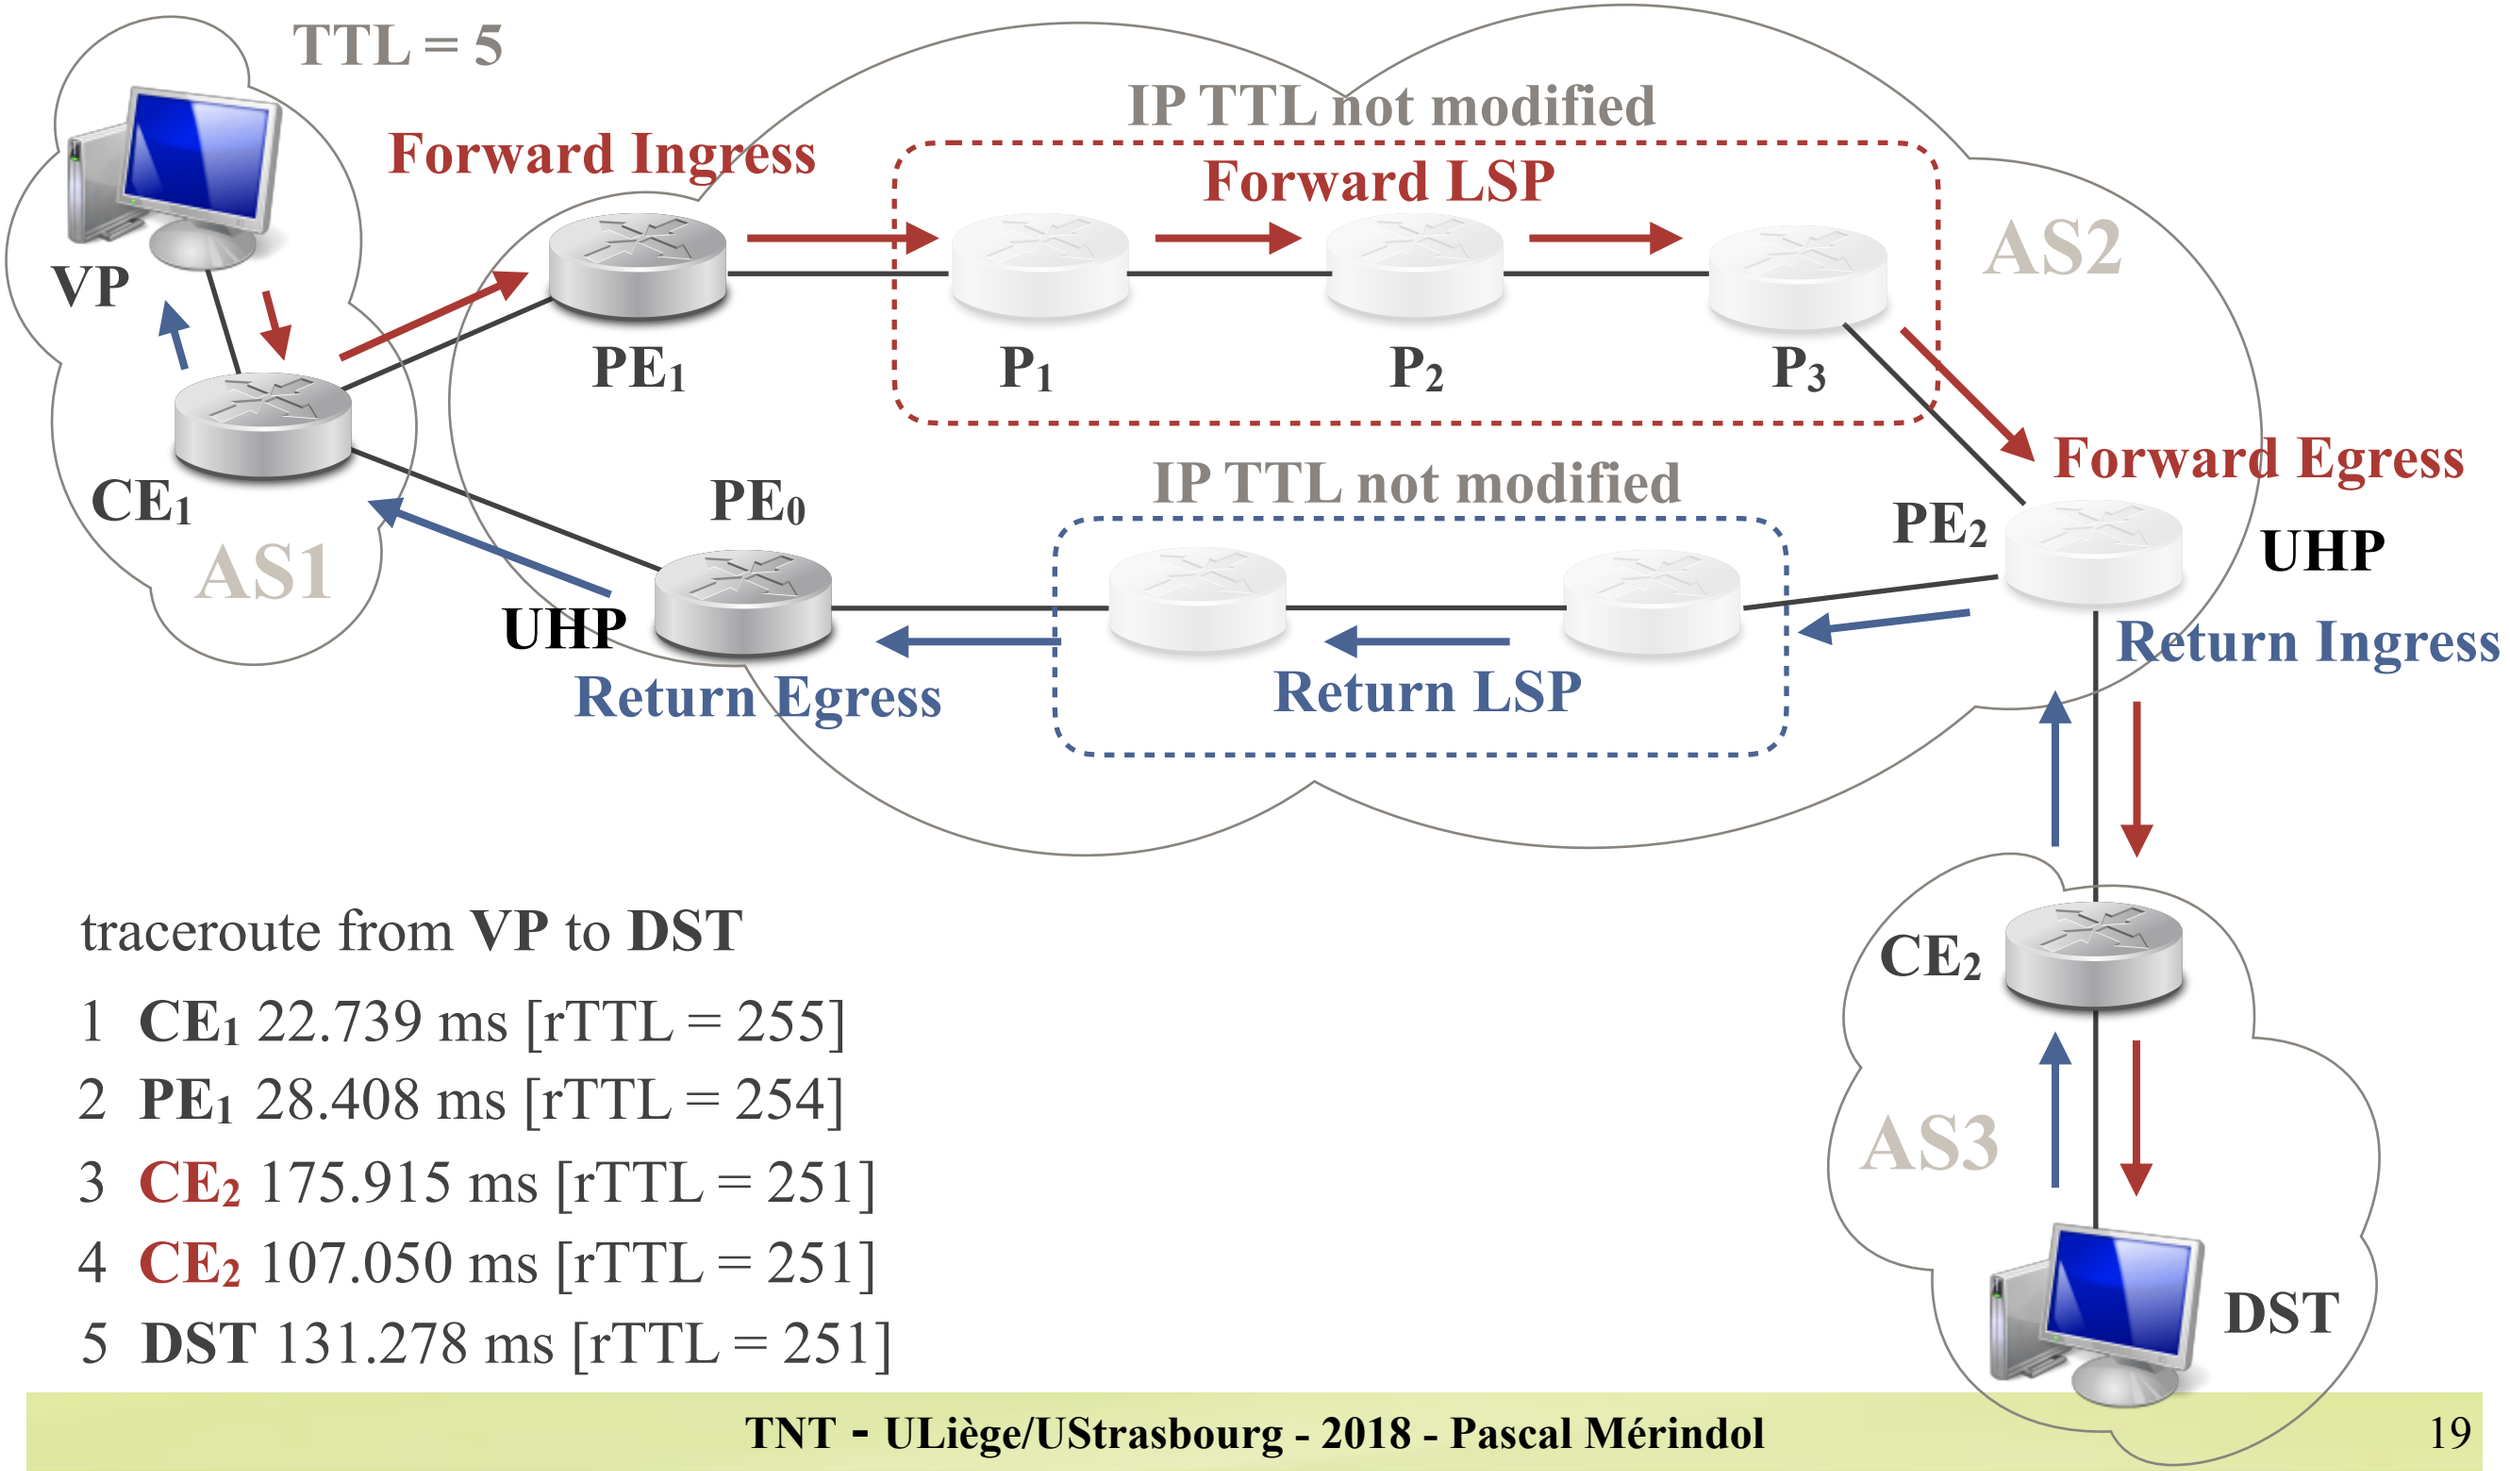

traceroute from **VP** to **DST**

1 **CE1** 22.739 ms [rTTL = 255]
2 **PE1** 28.408 ms [rTTL = 254]
3 **CE2** 175.915 ms [rTTL = 251]
4 **CE2** 107.050 ms [rTTL = 251]
5 **DST** 131.278 ms [rTTL = 251]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

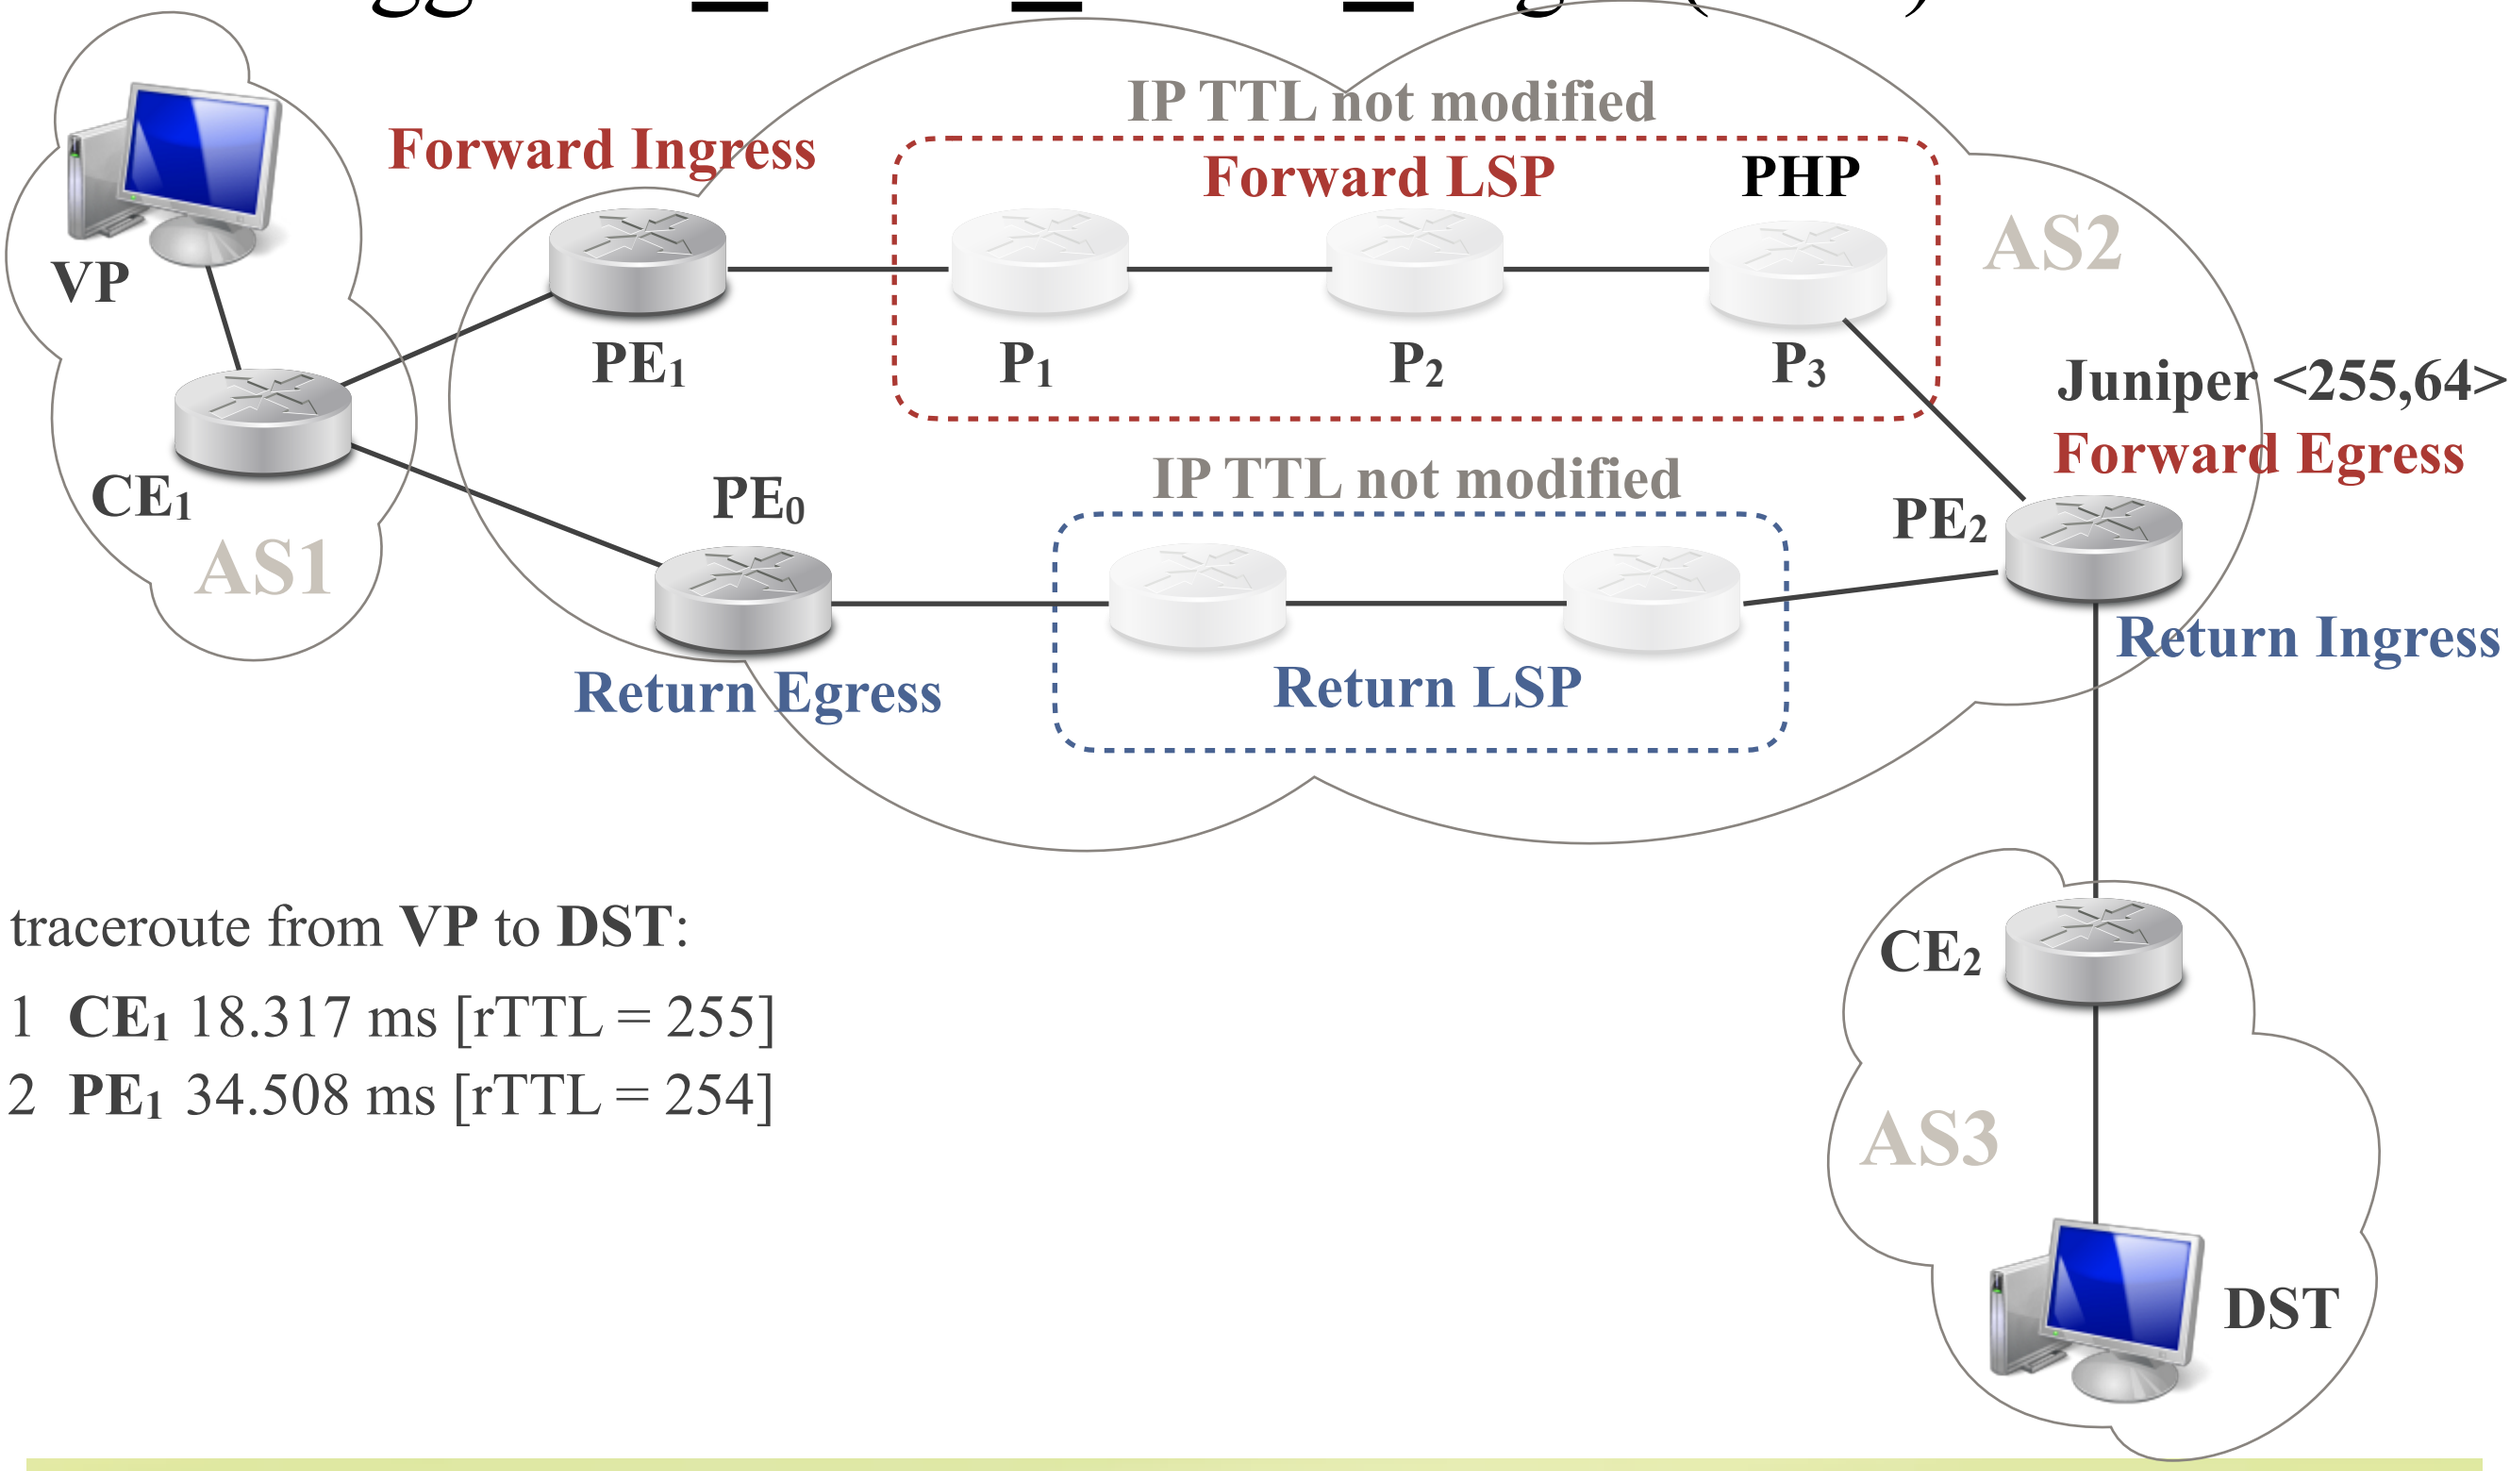

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

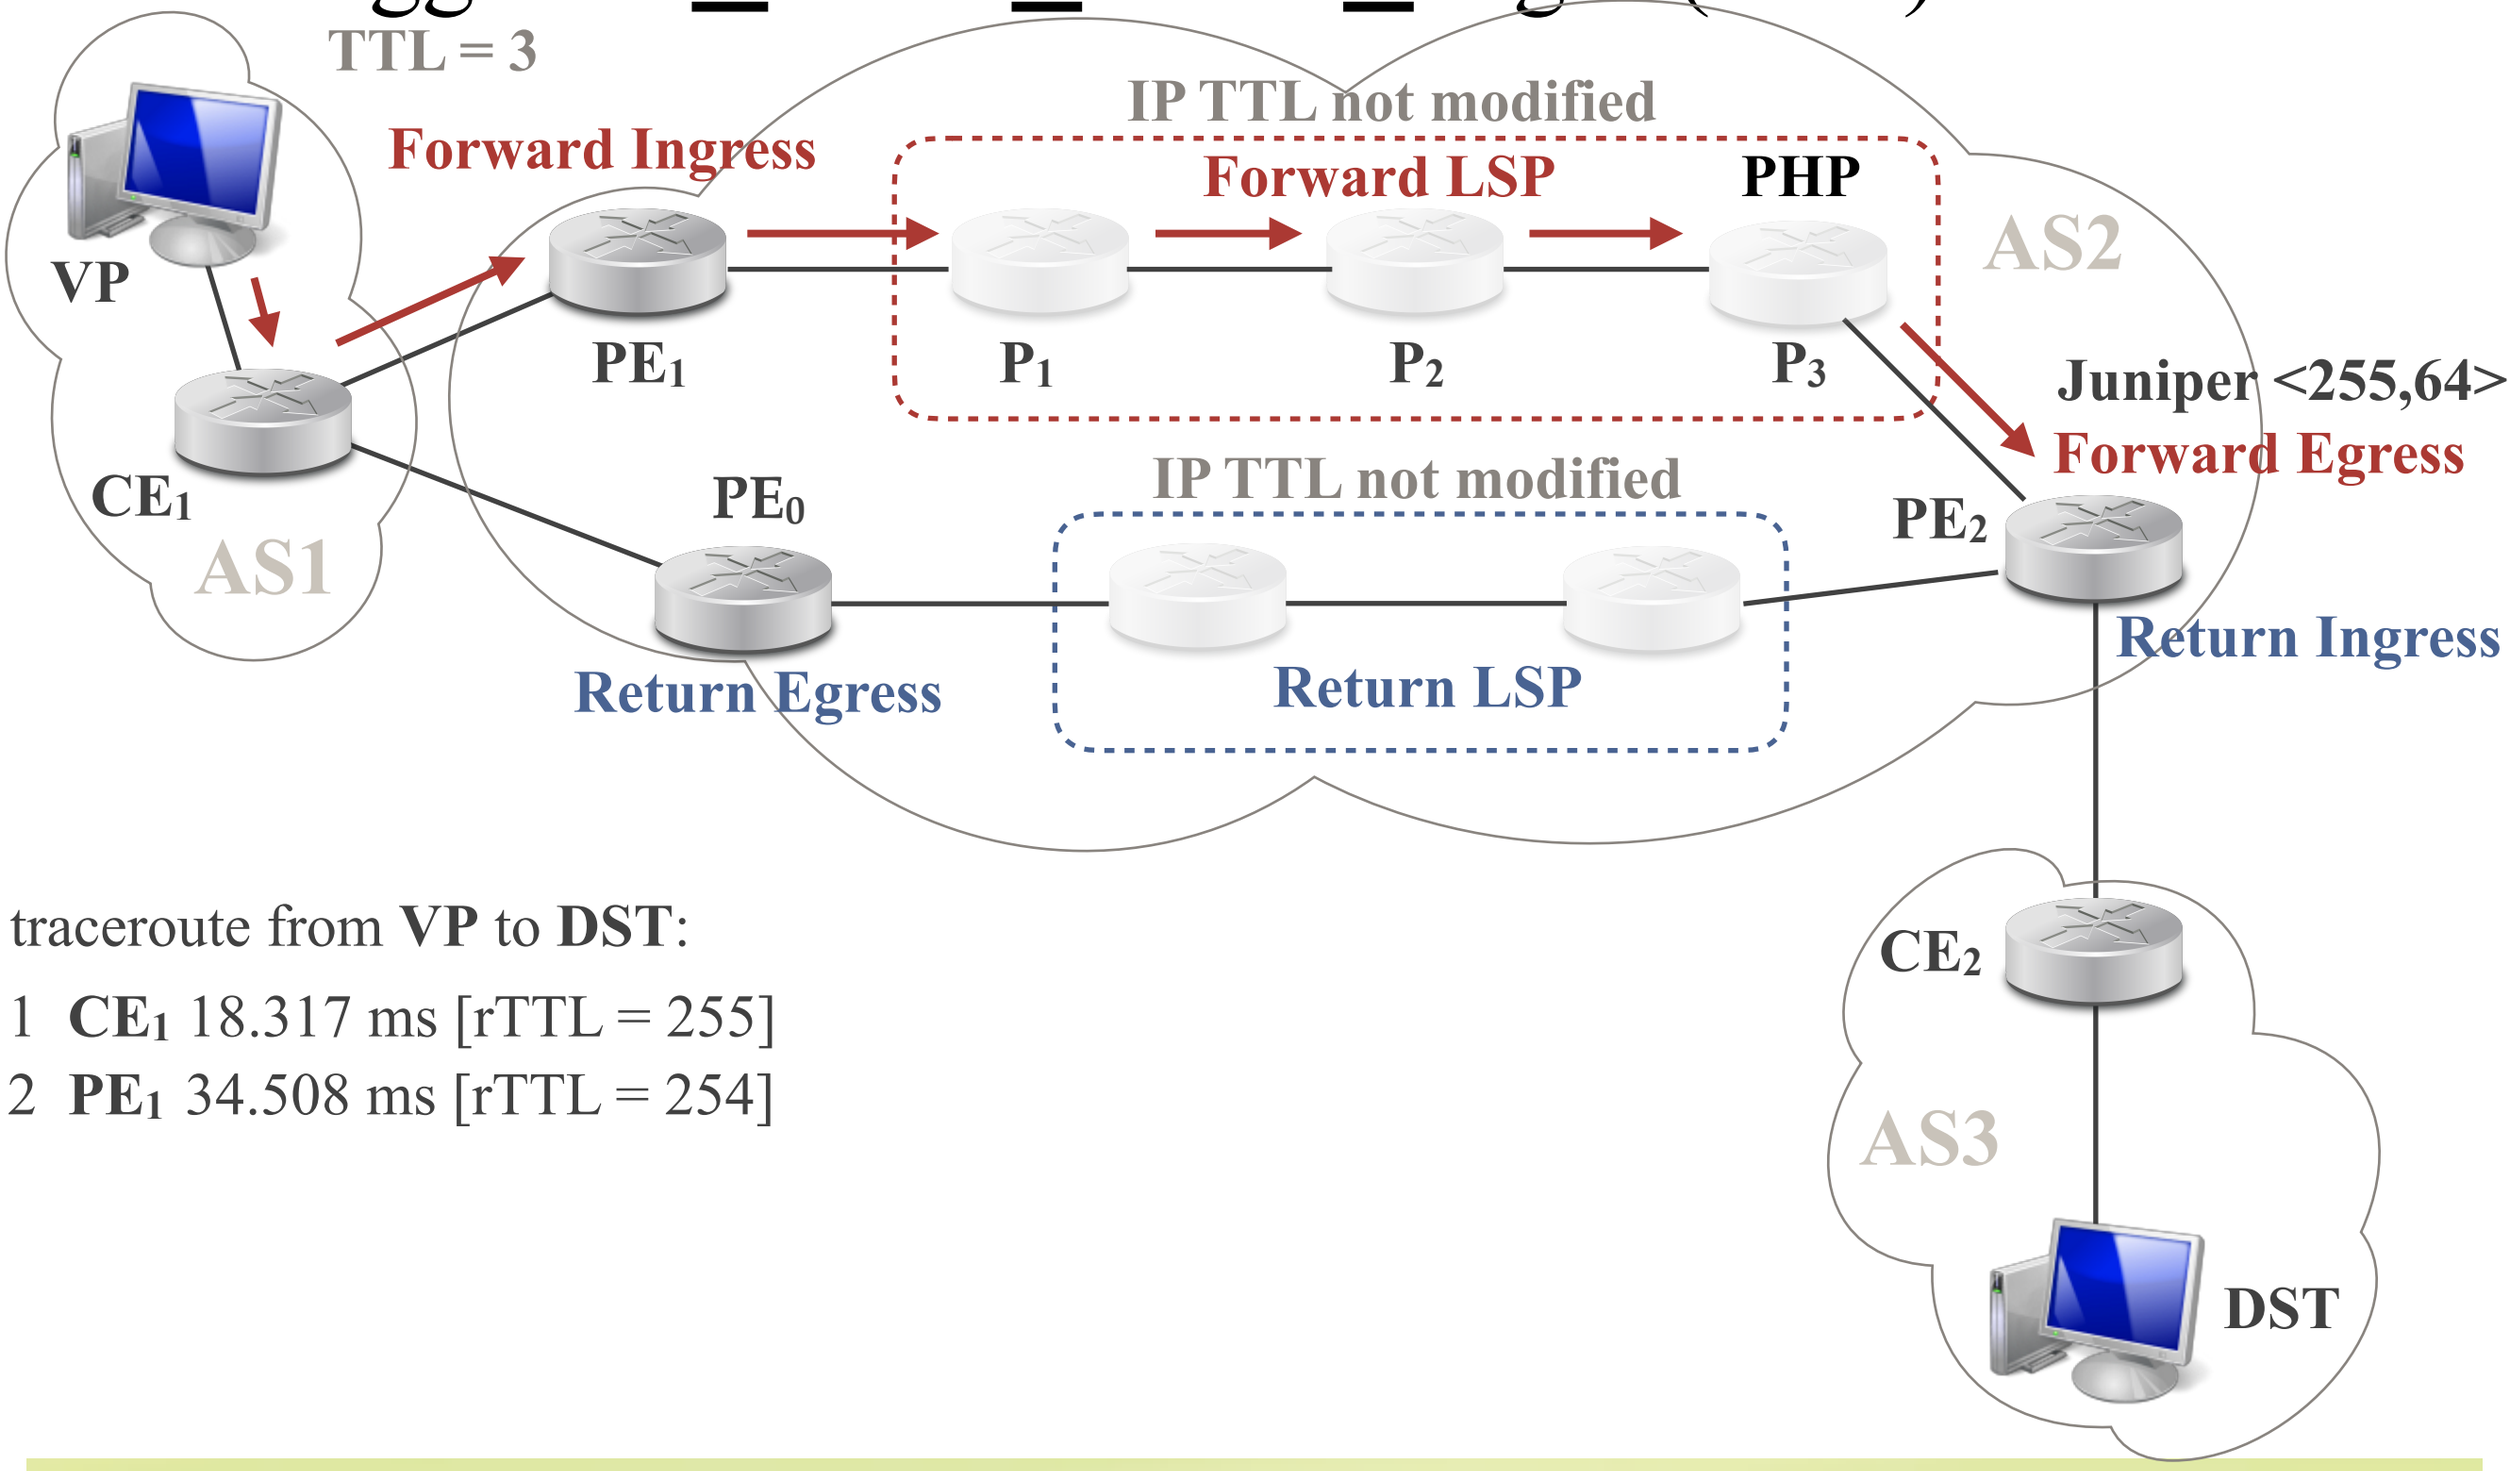

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

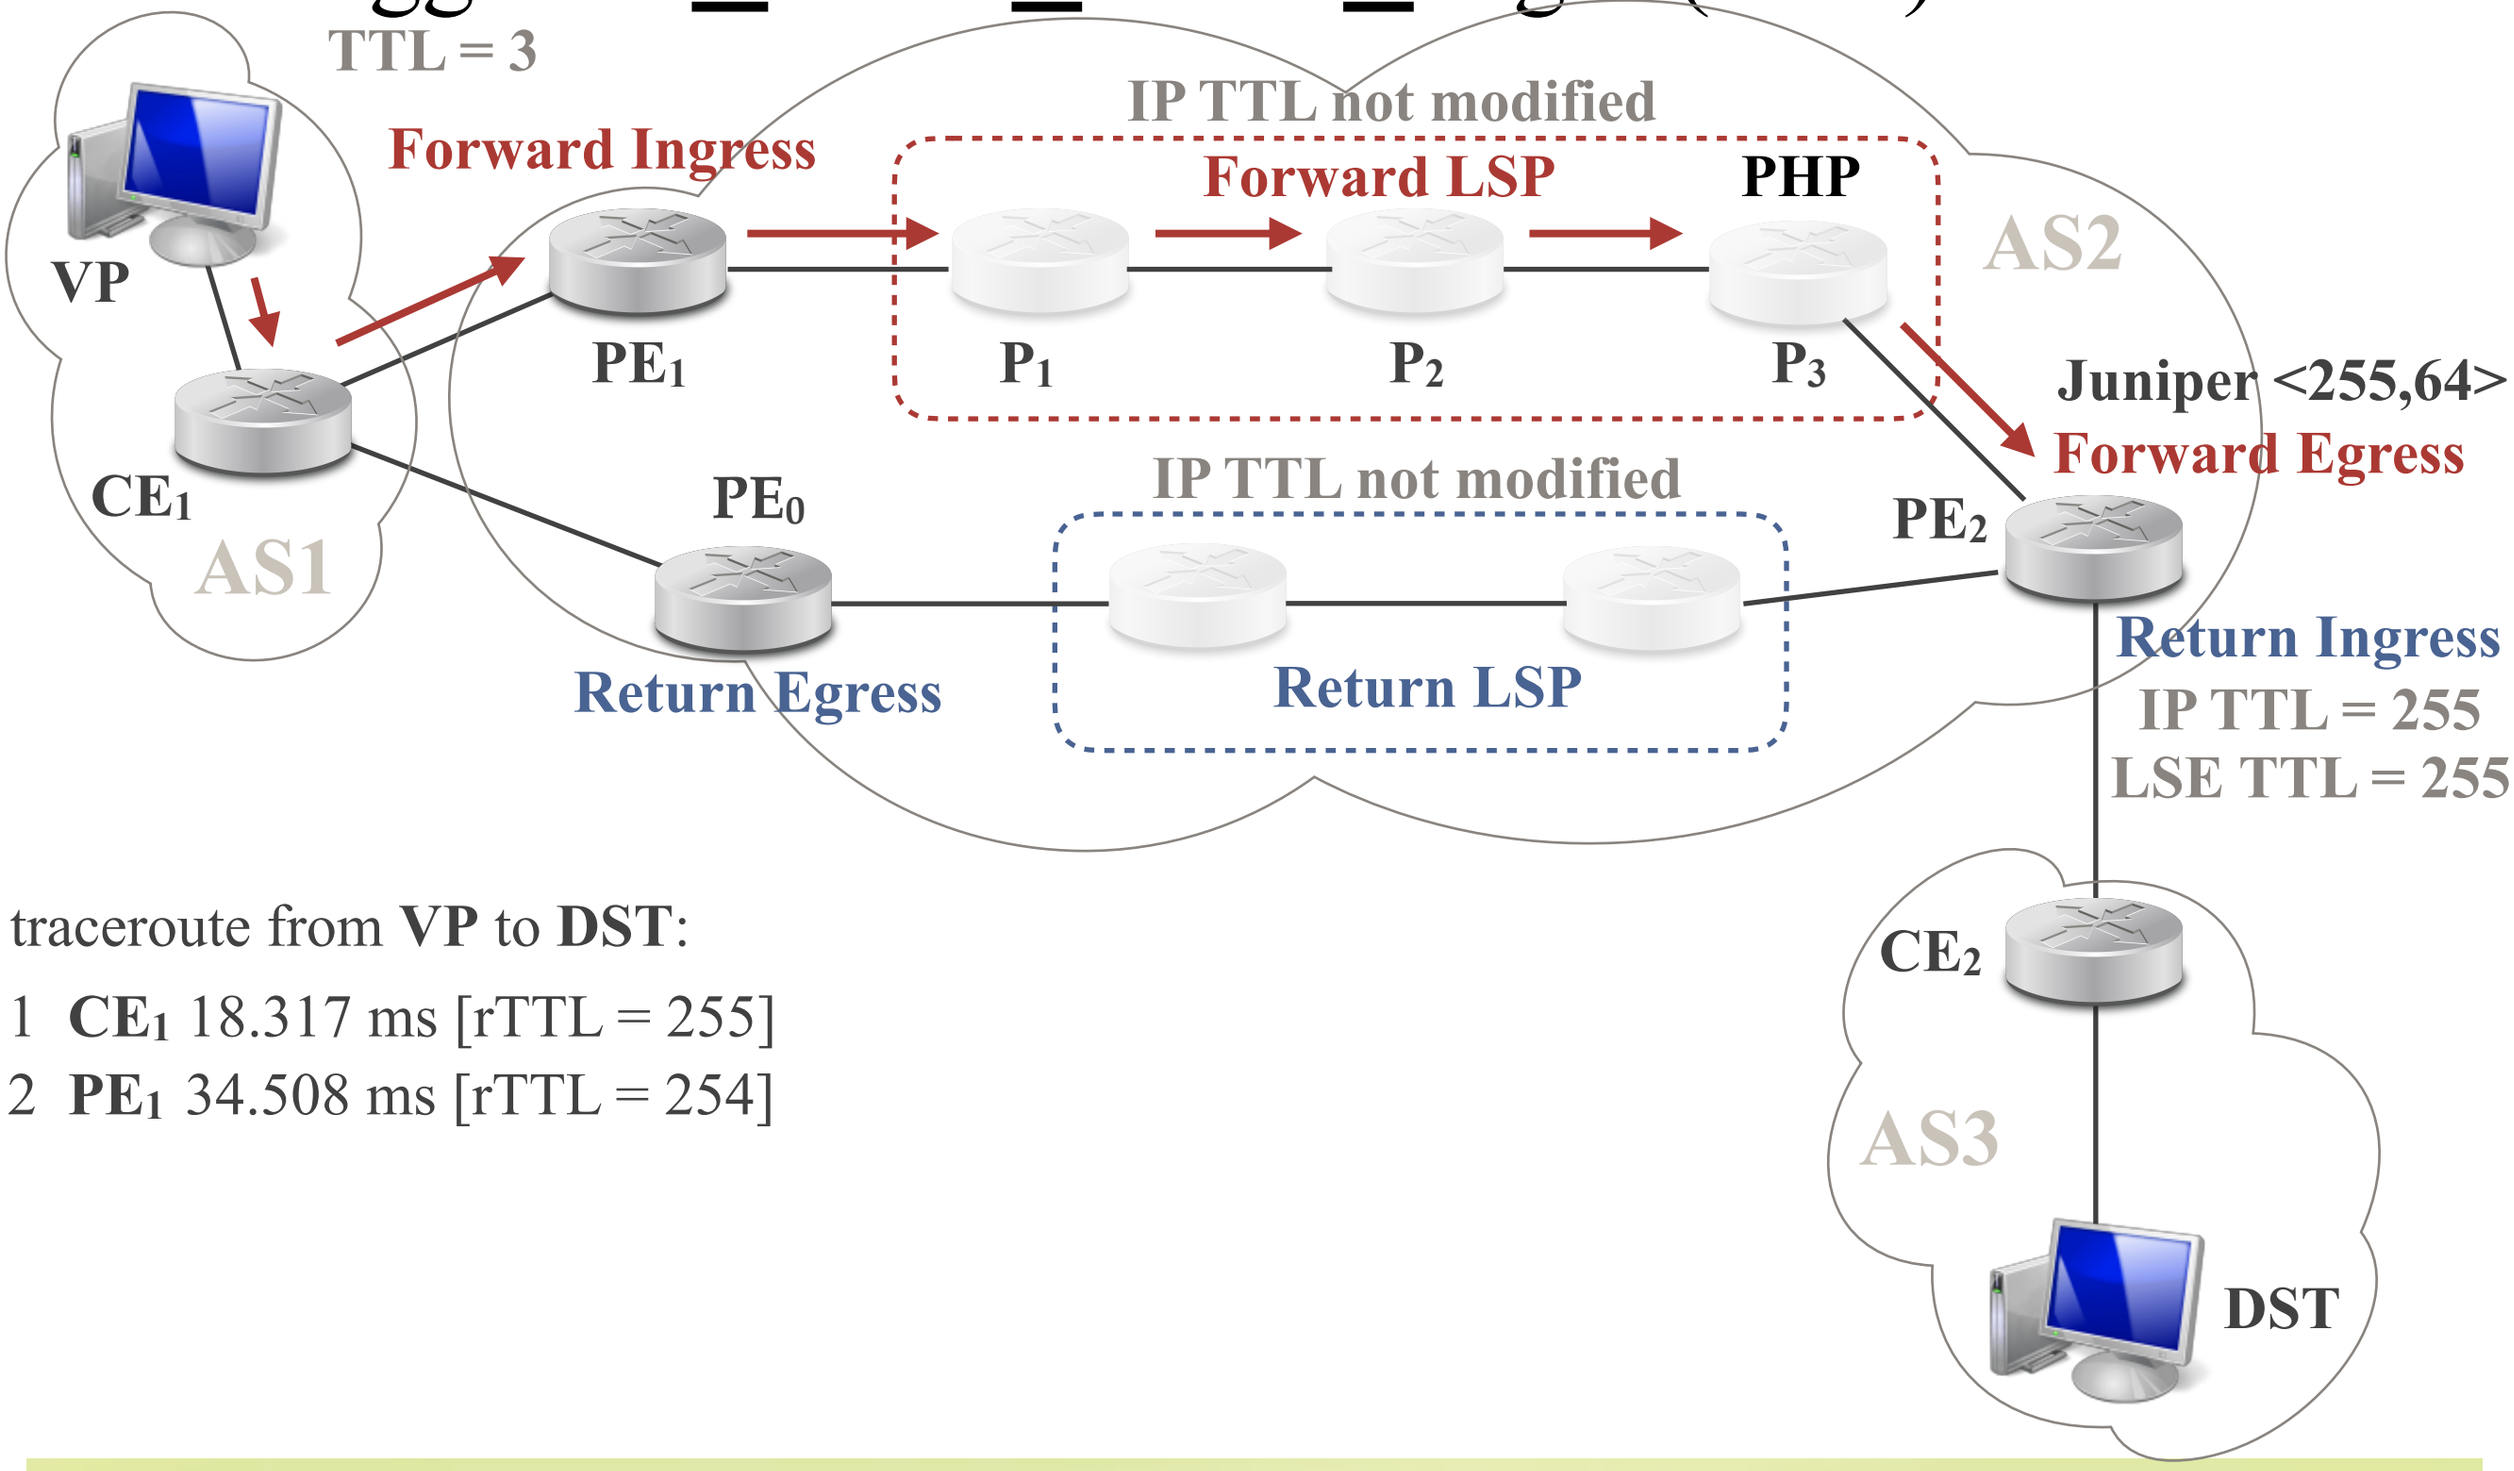

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

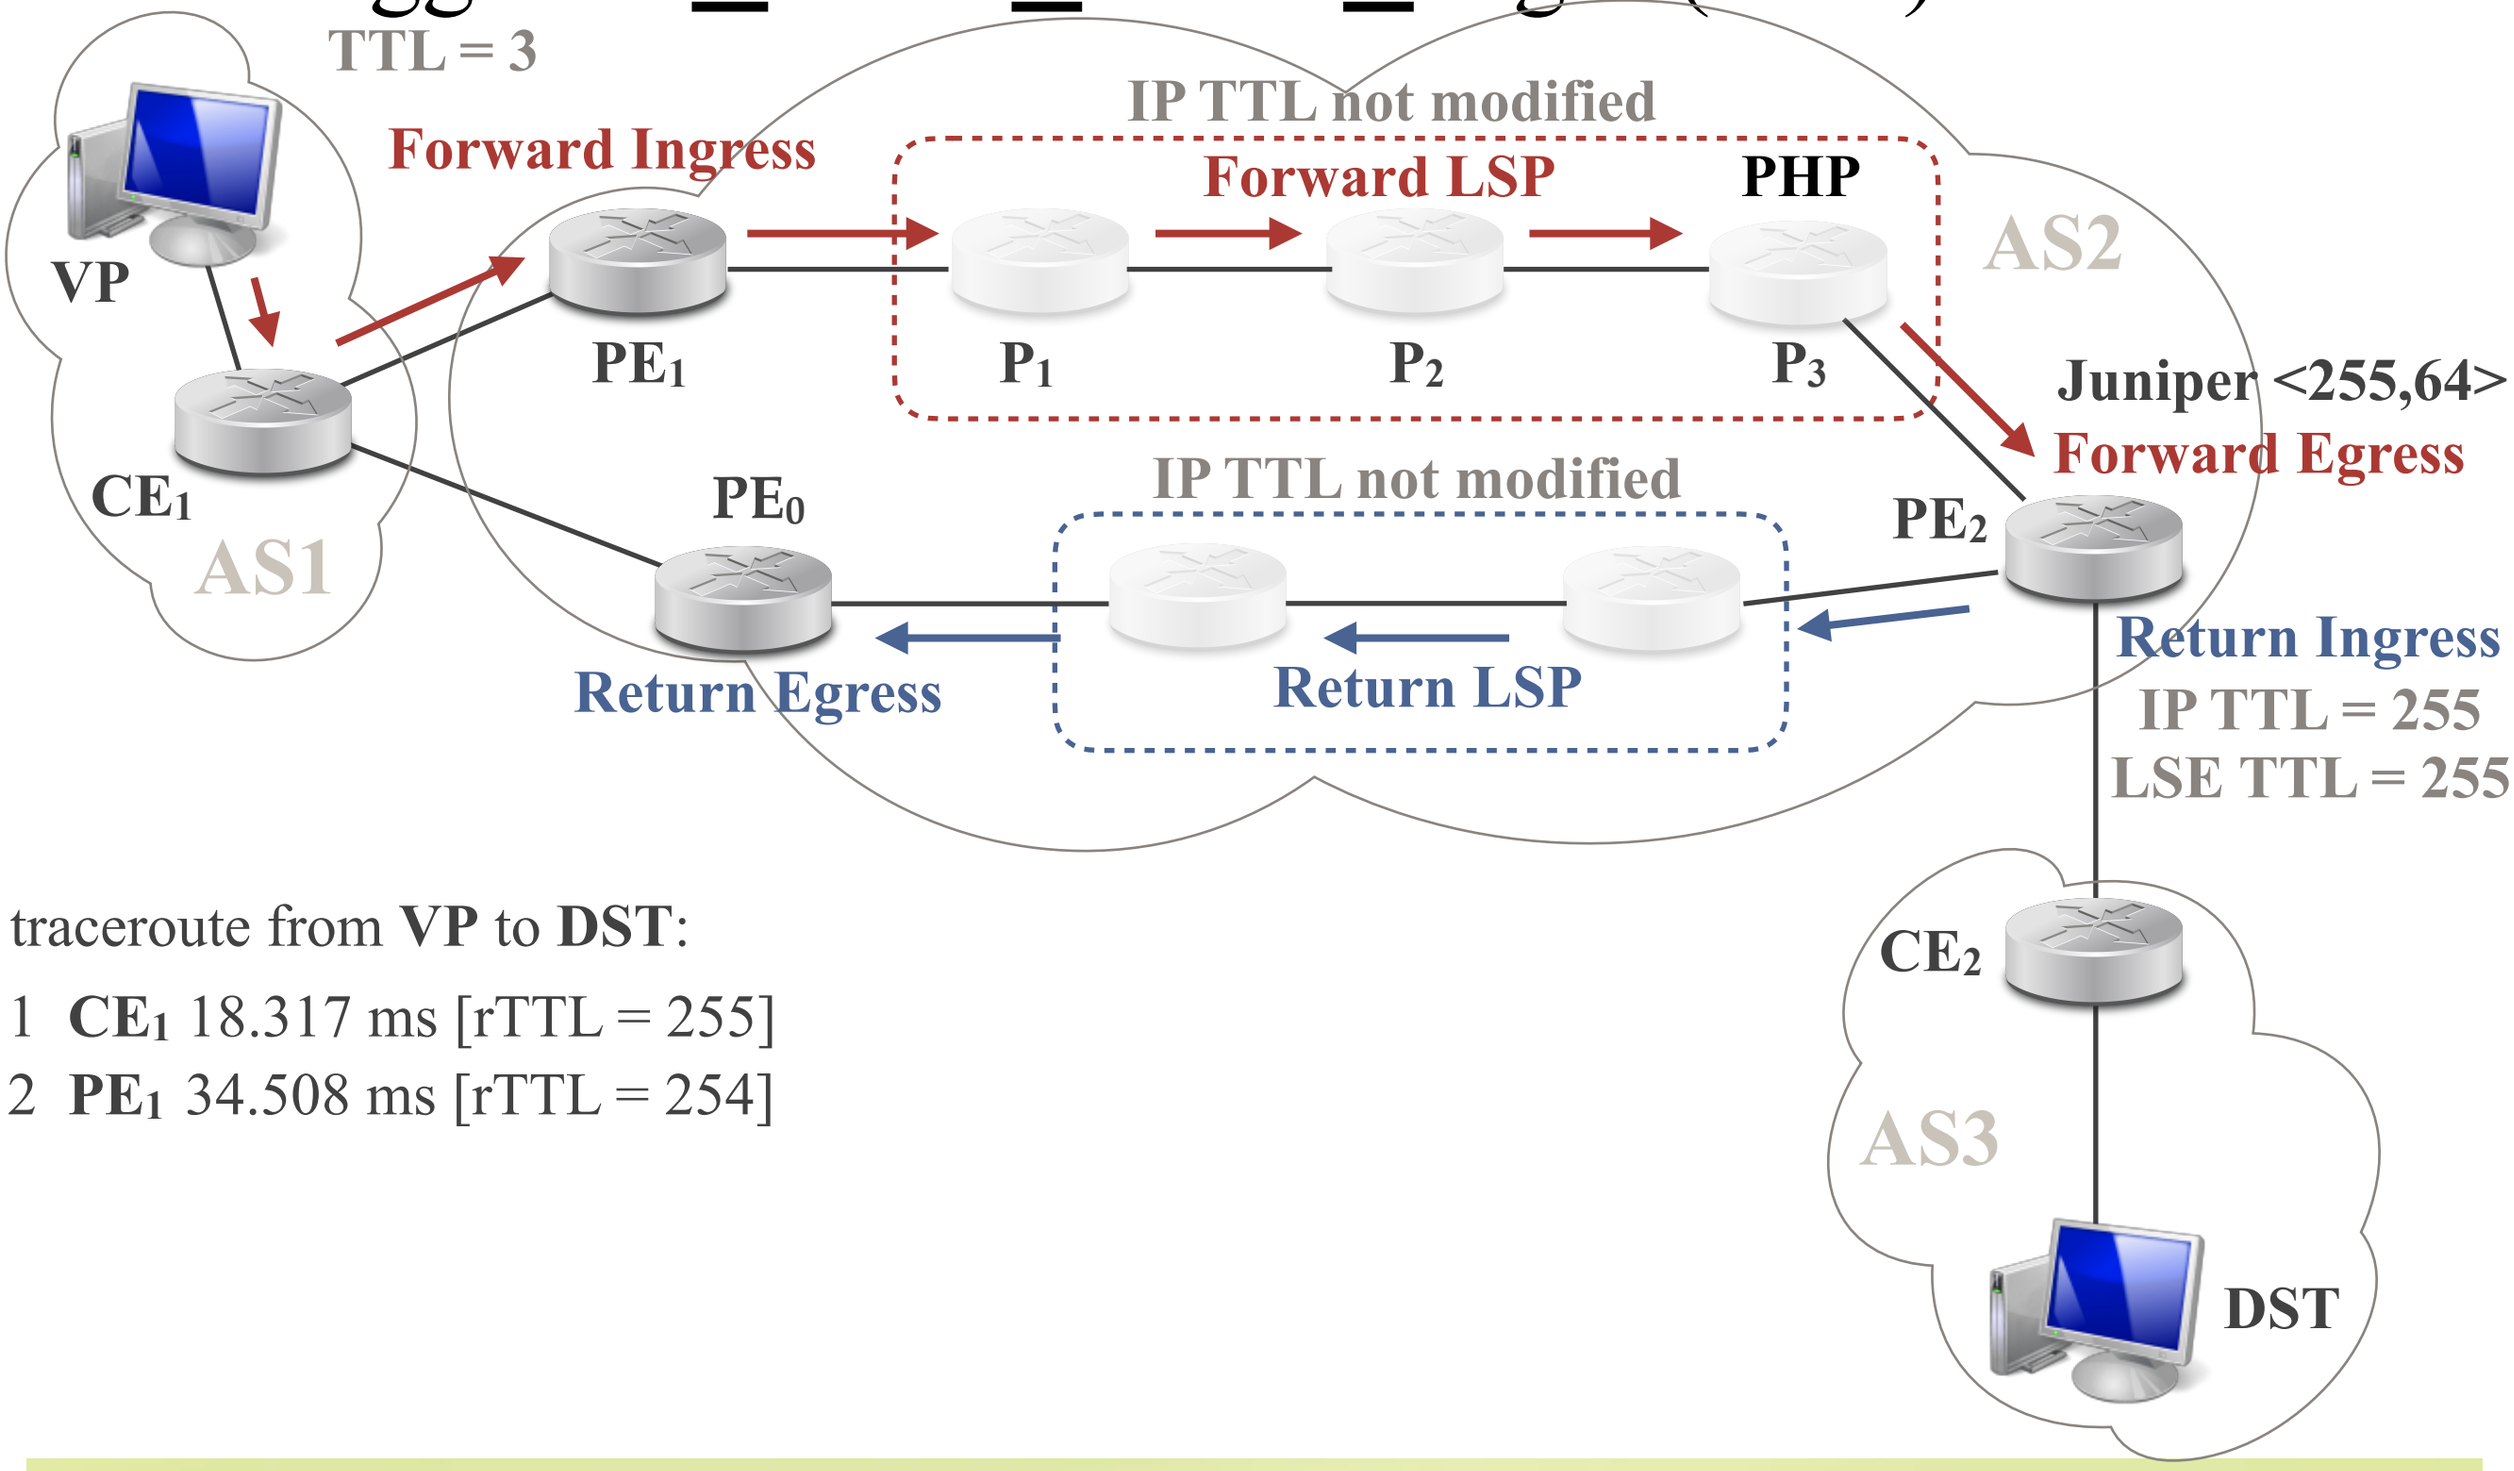

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

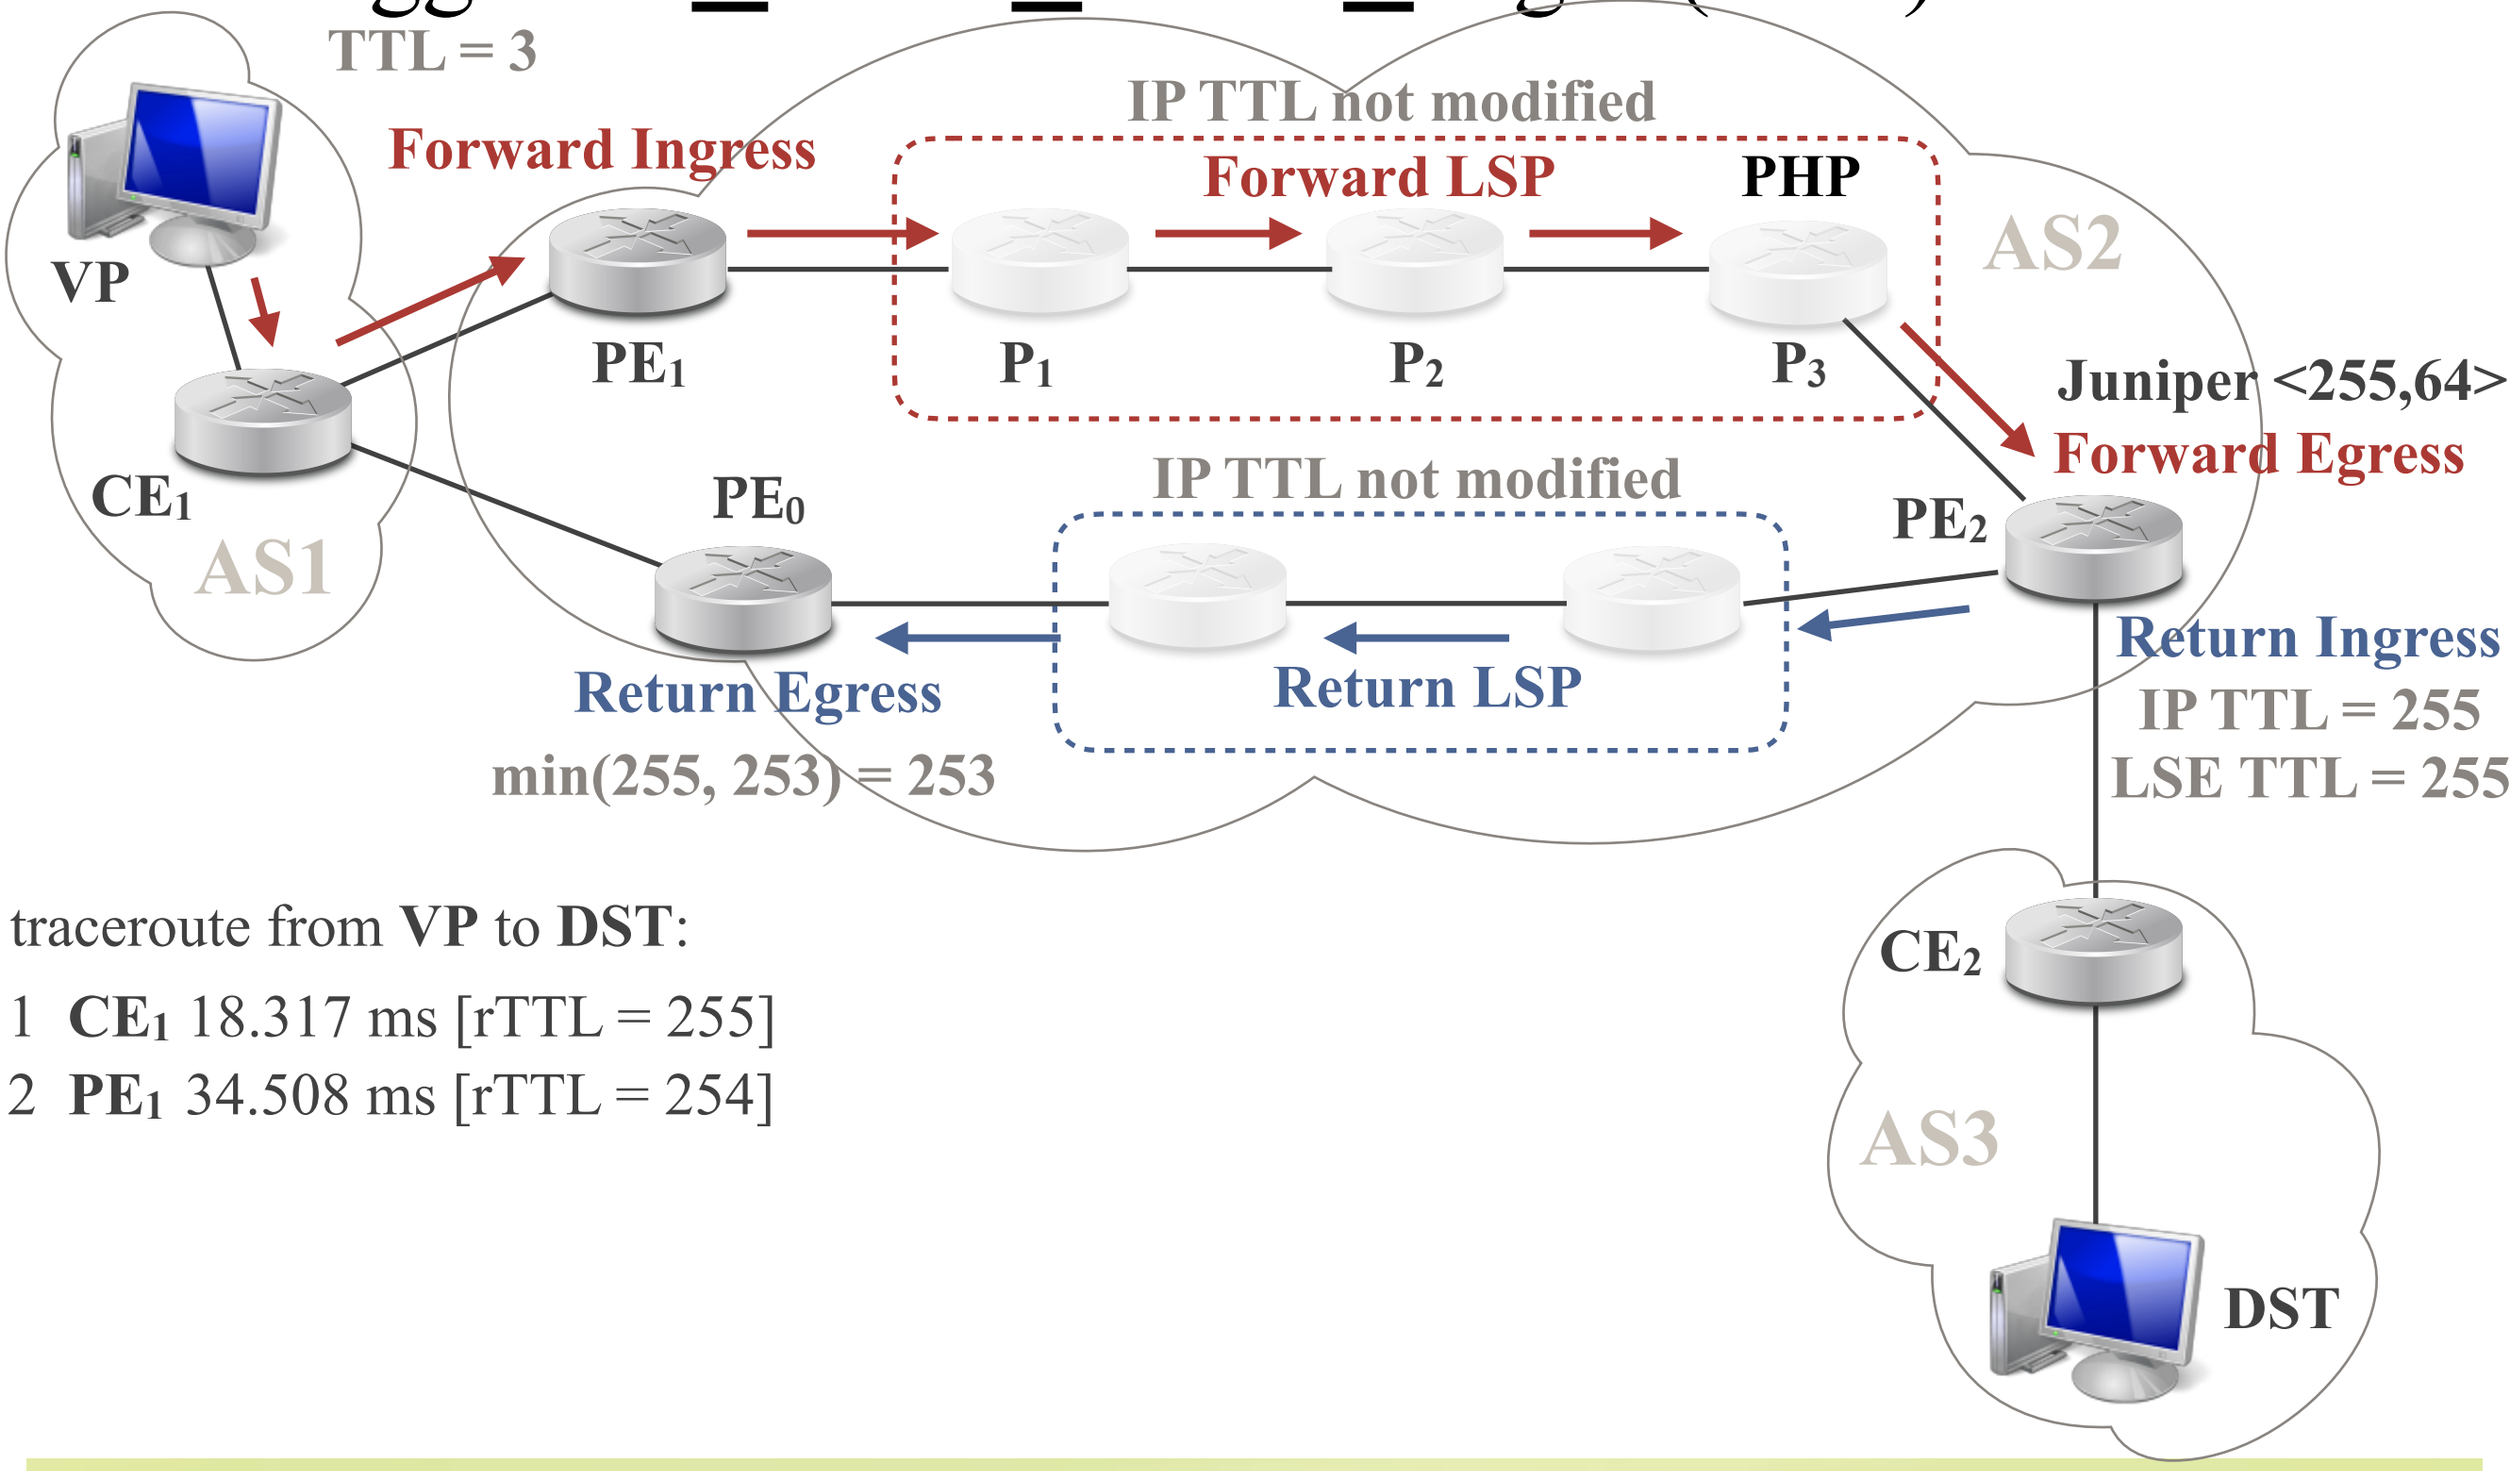

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

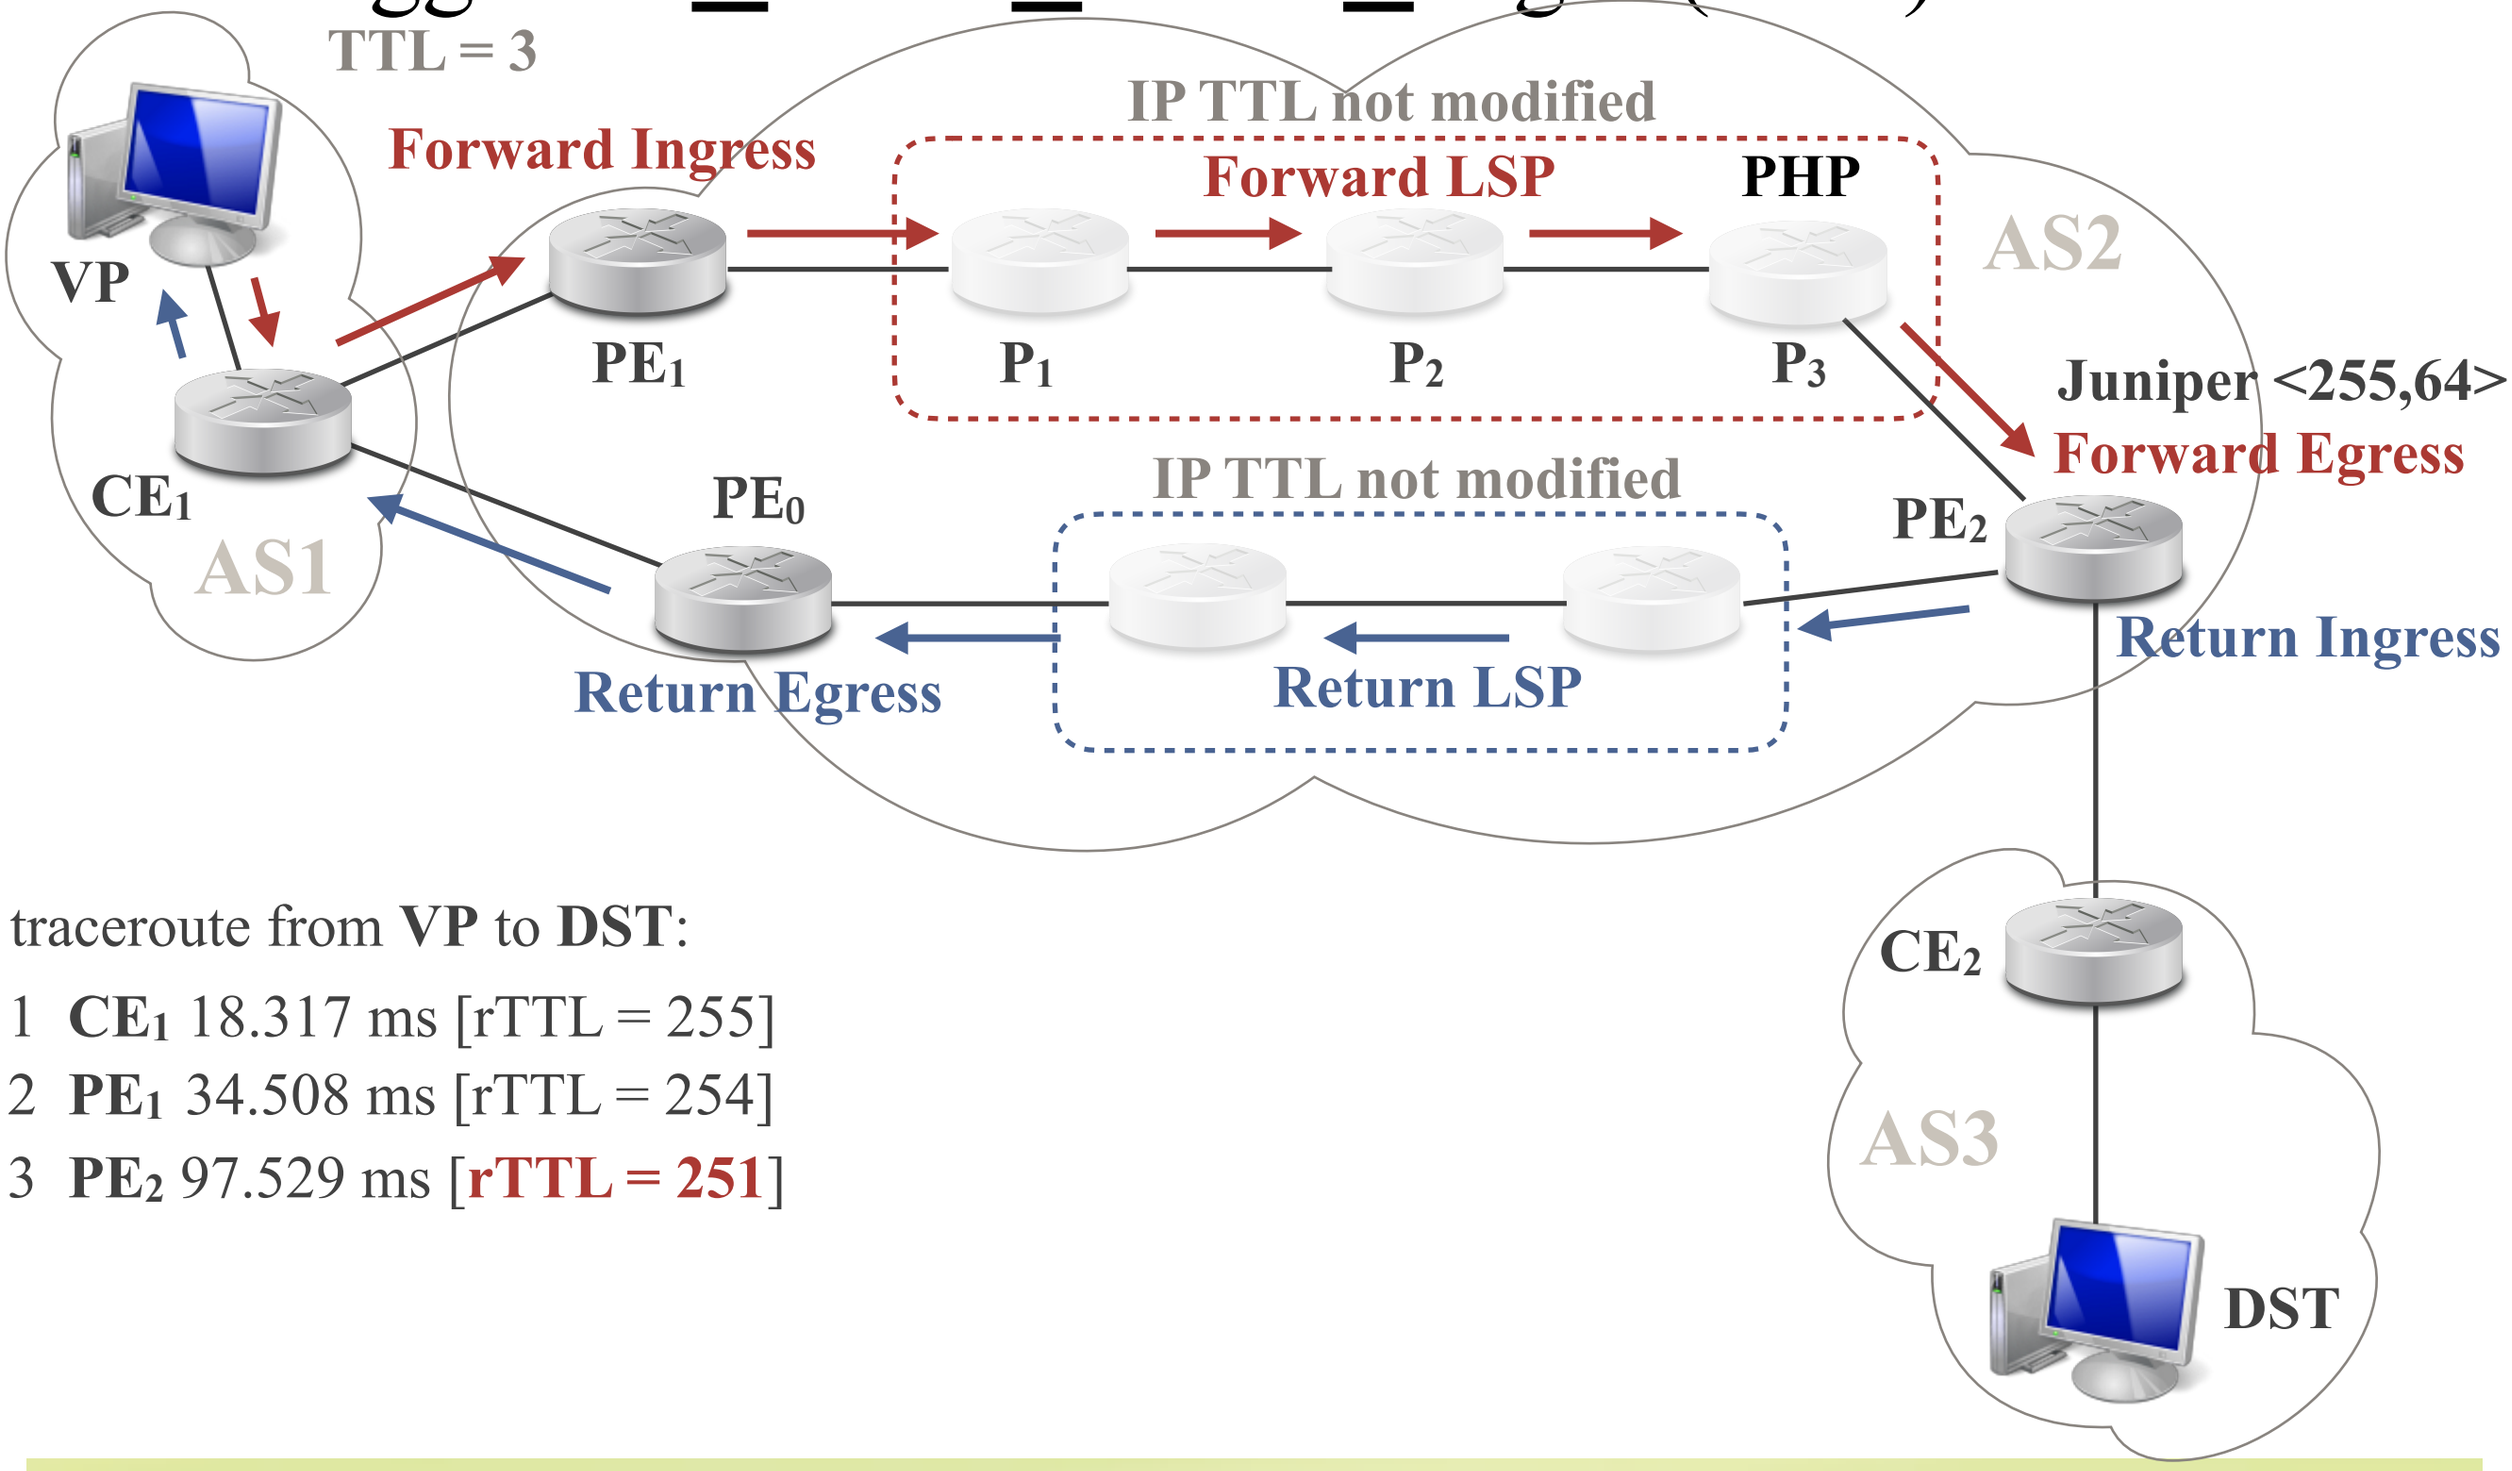

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

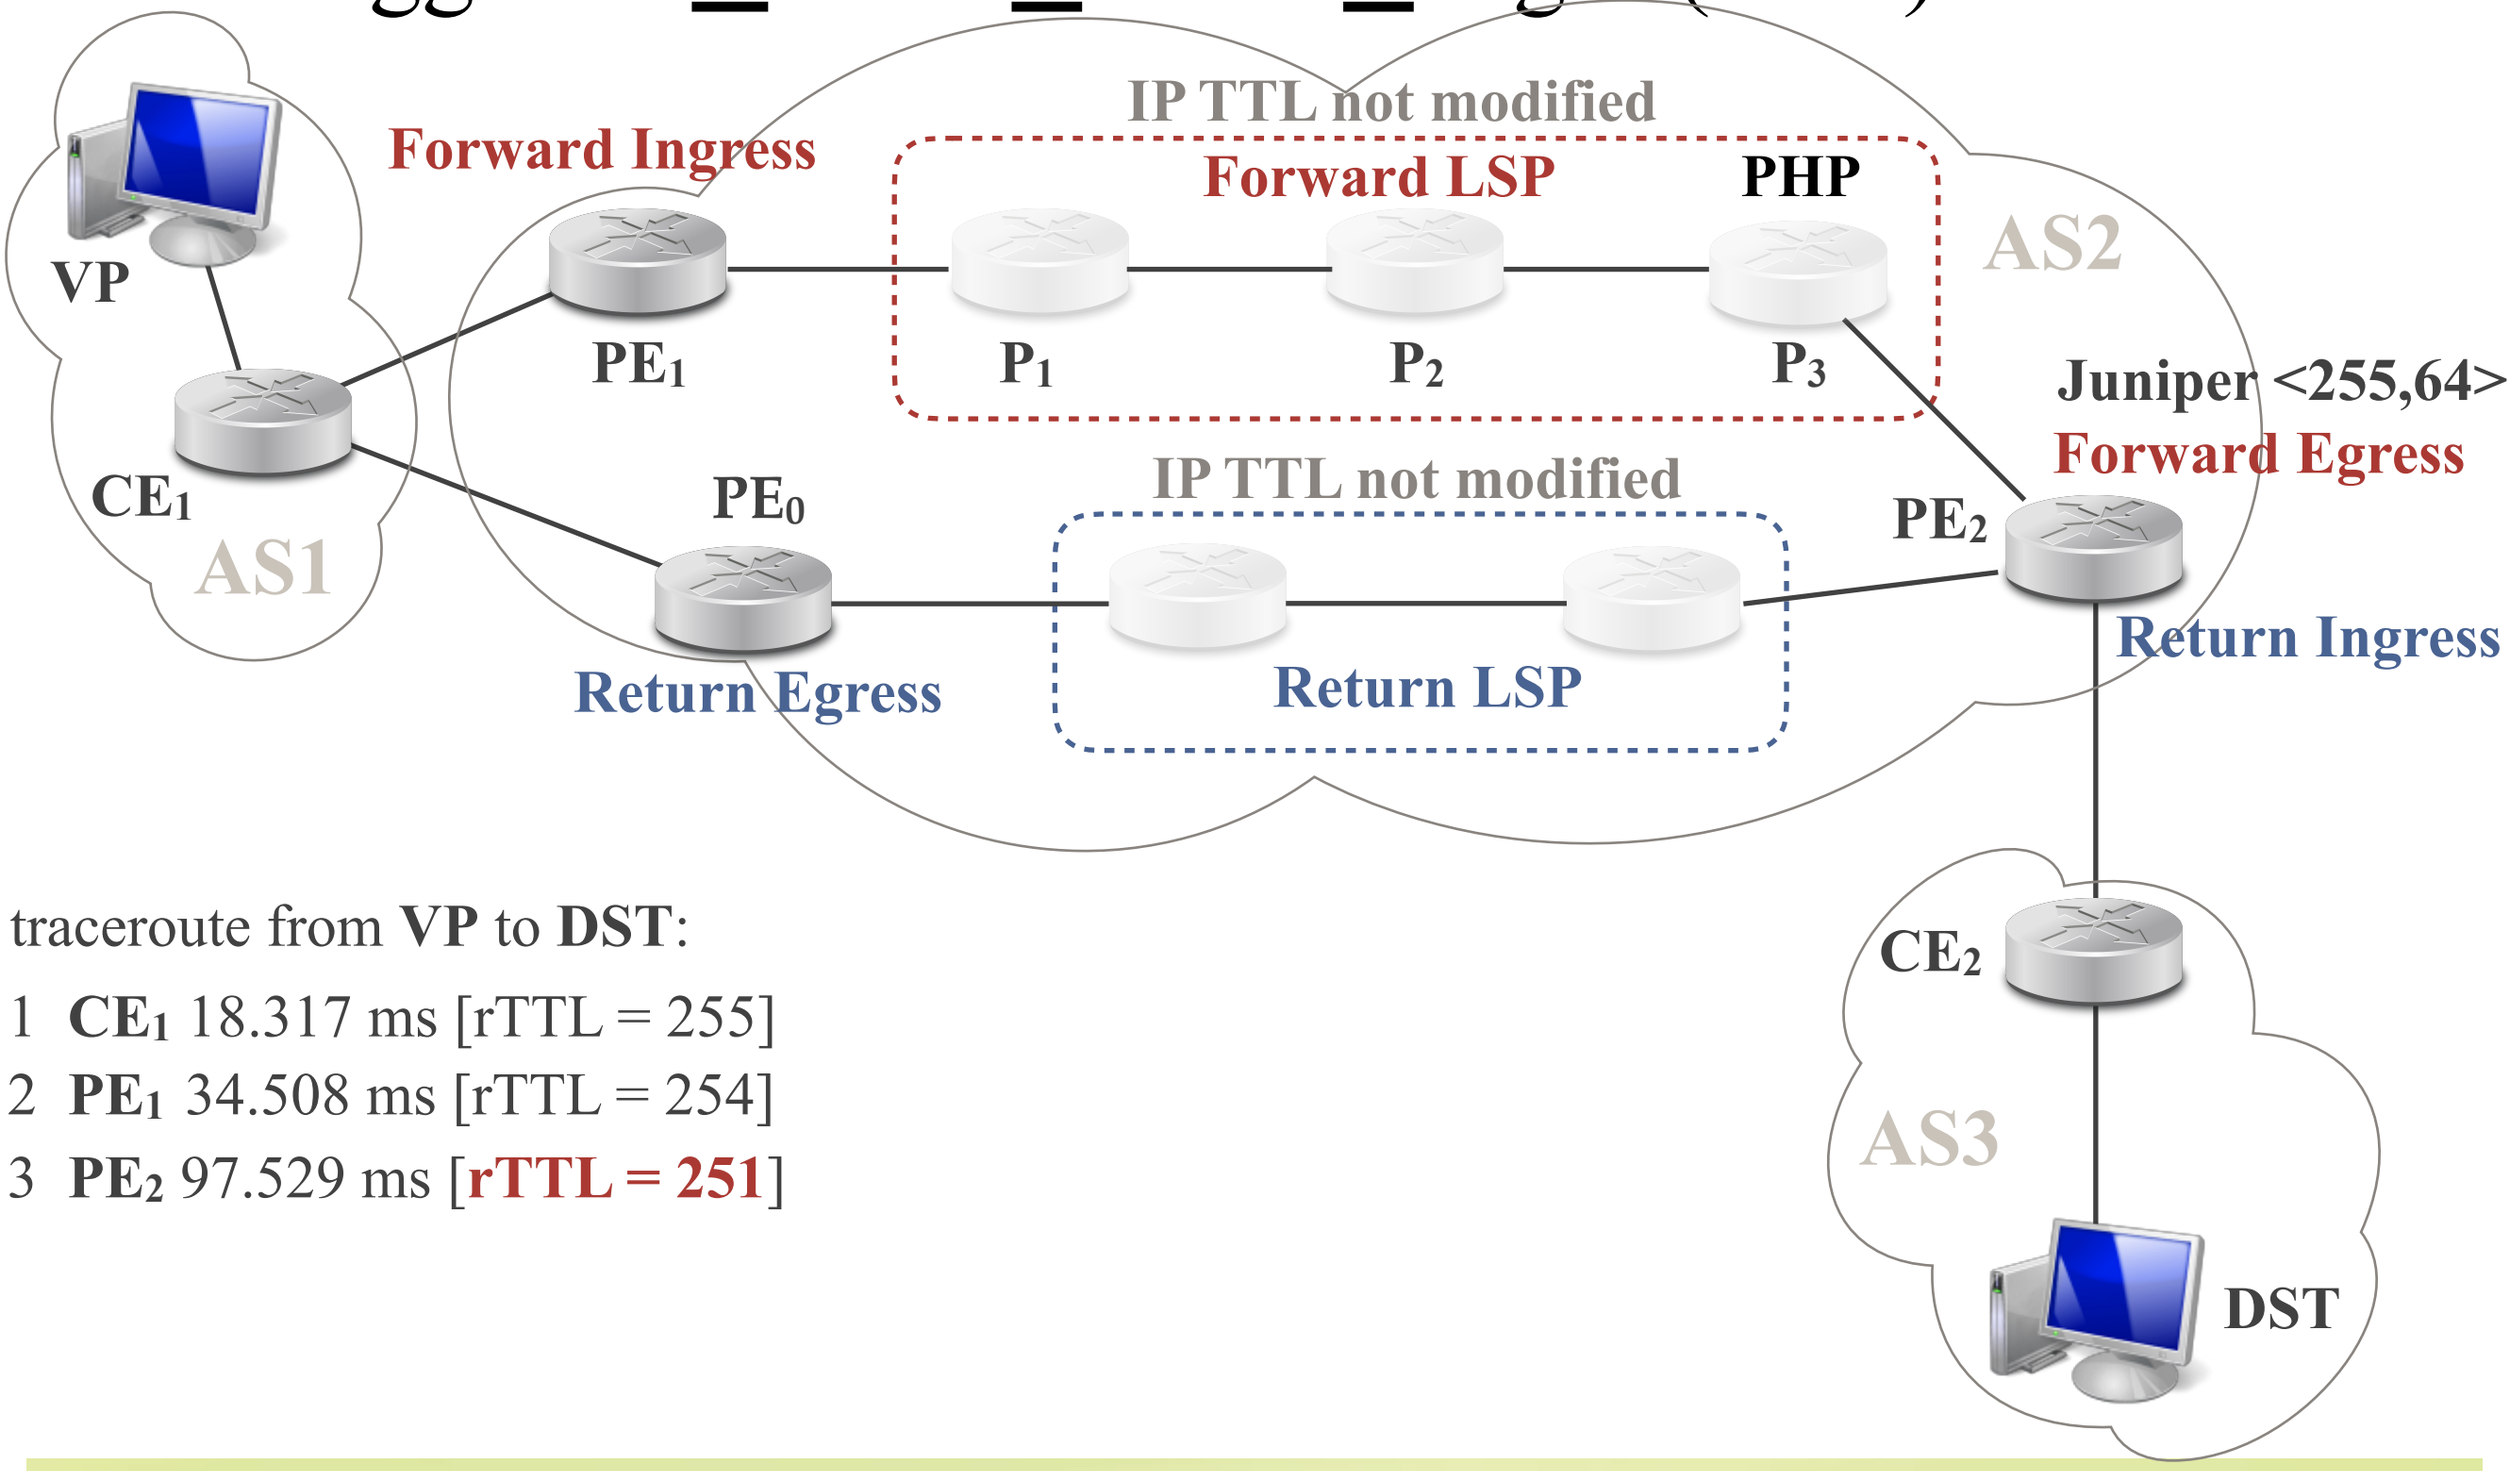

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

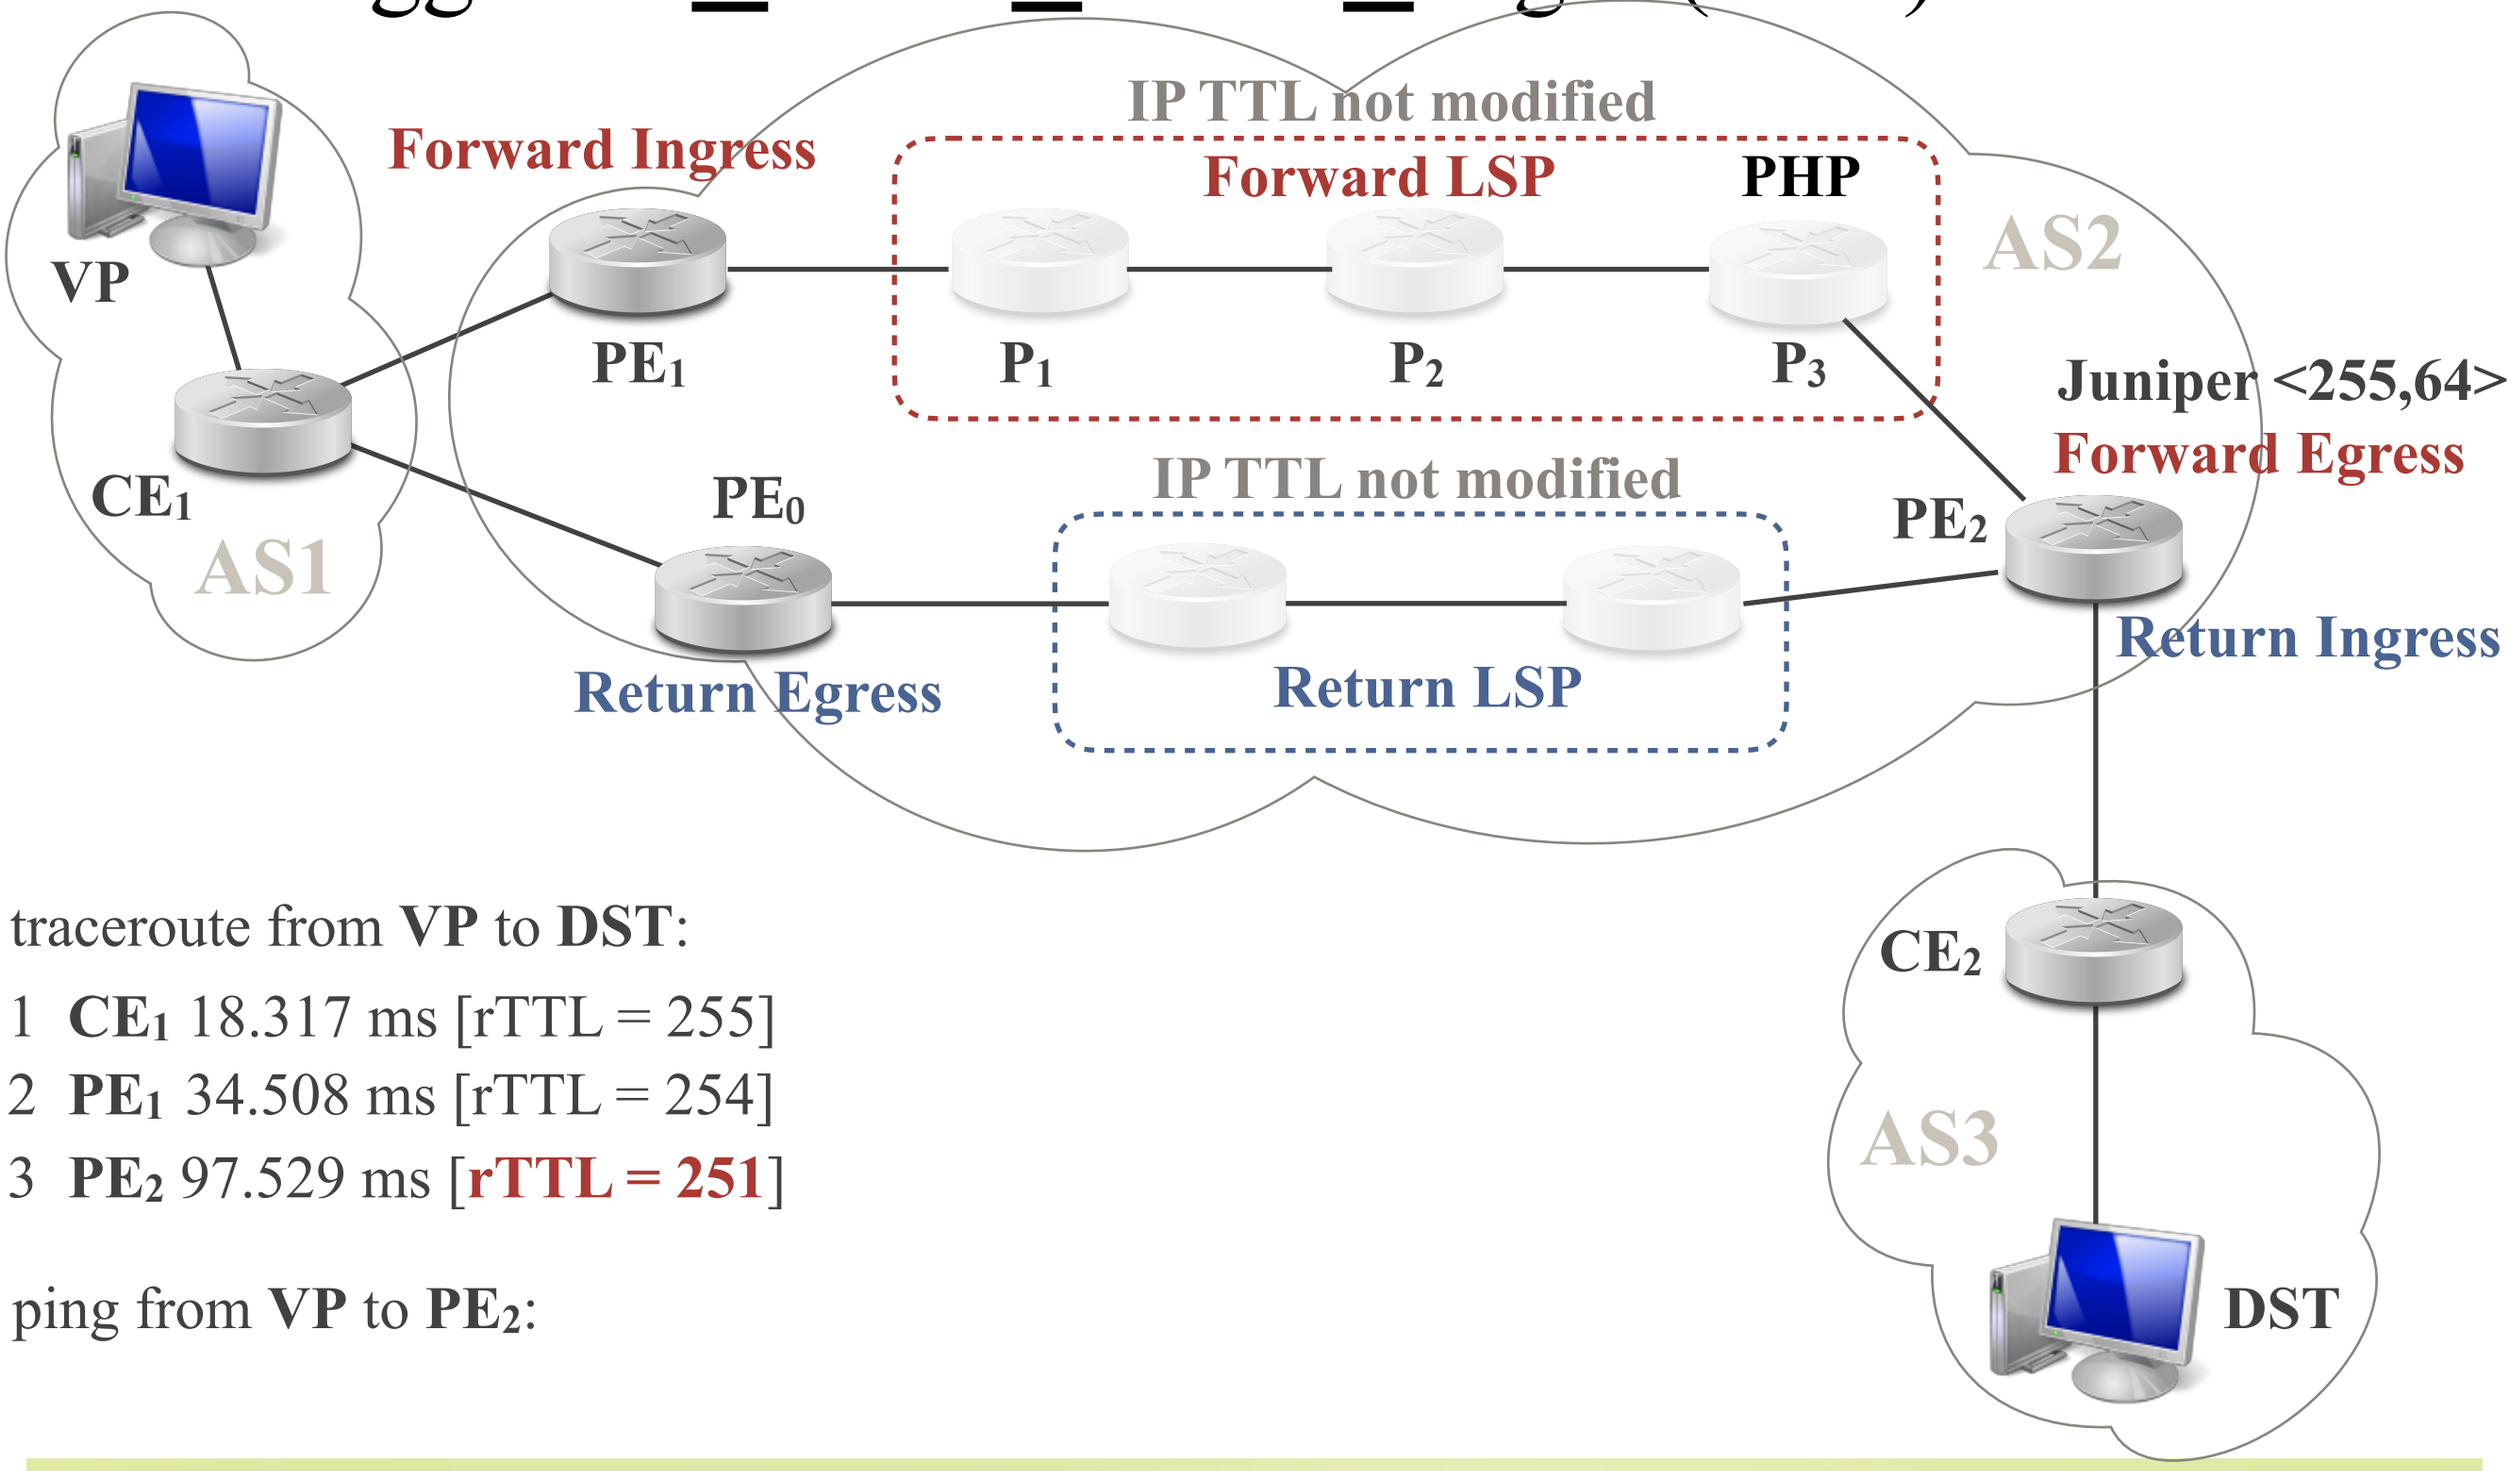

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

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

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

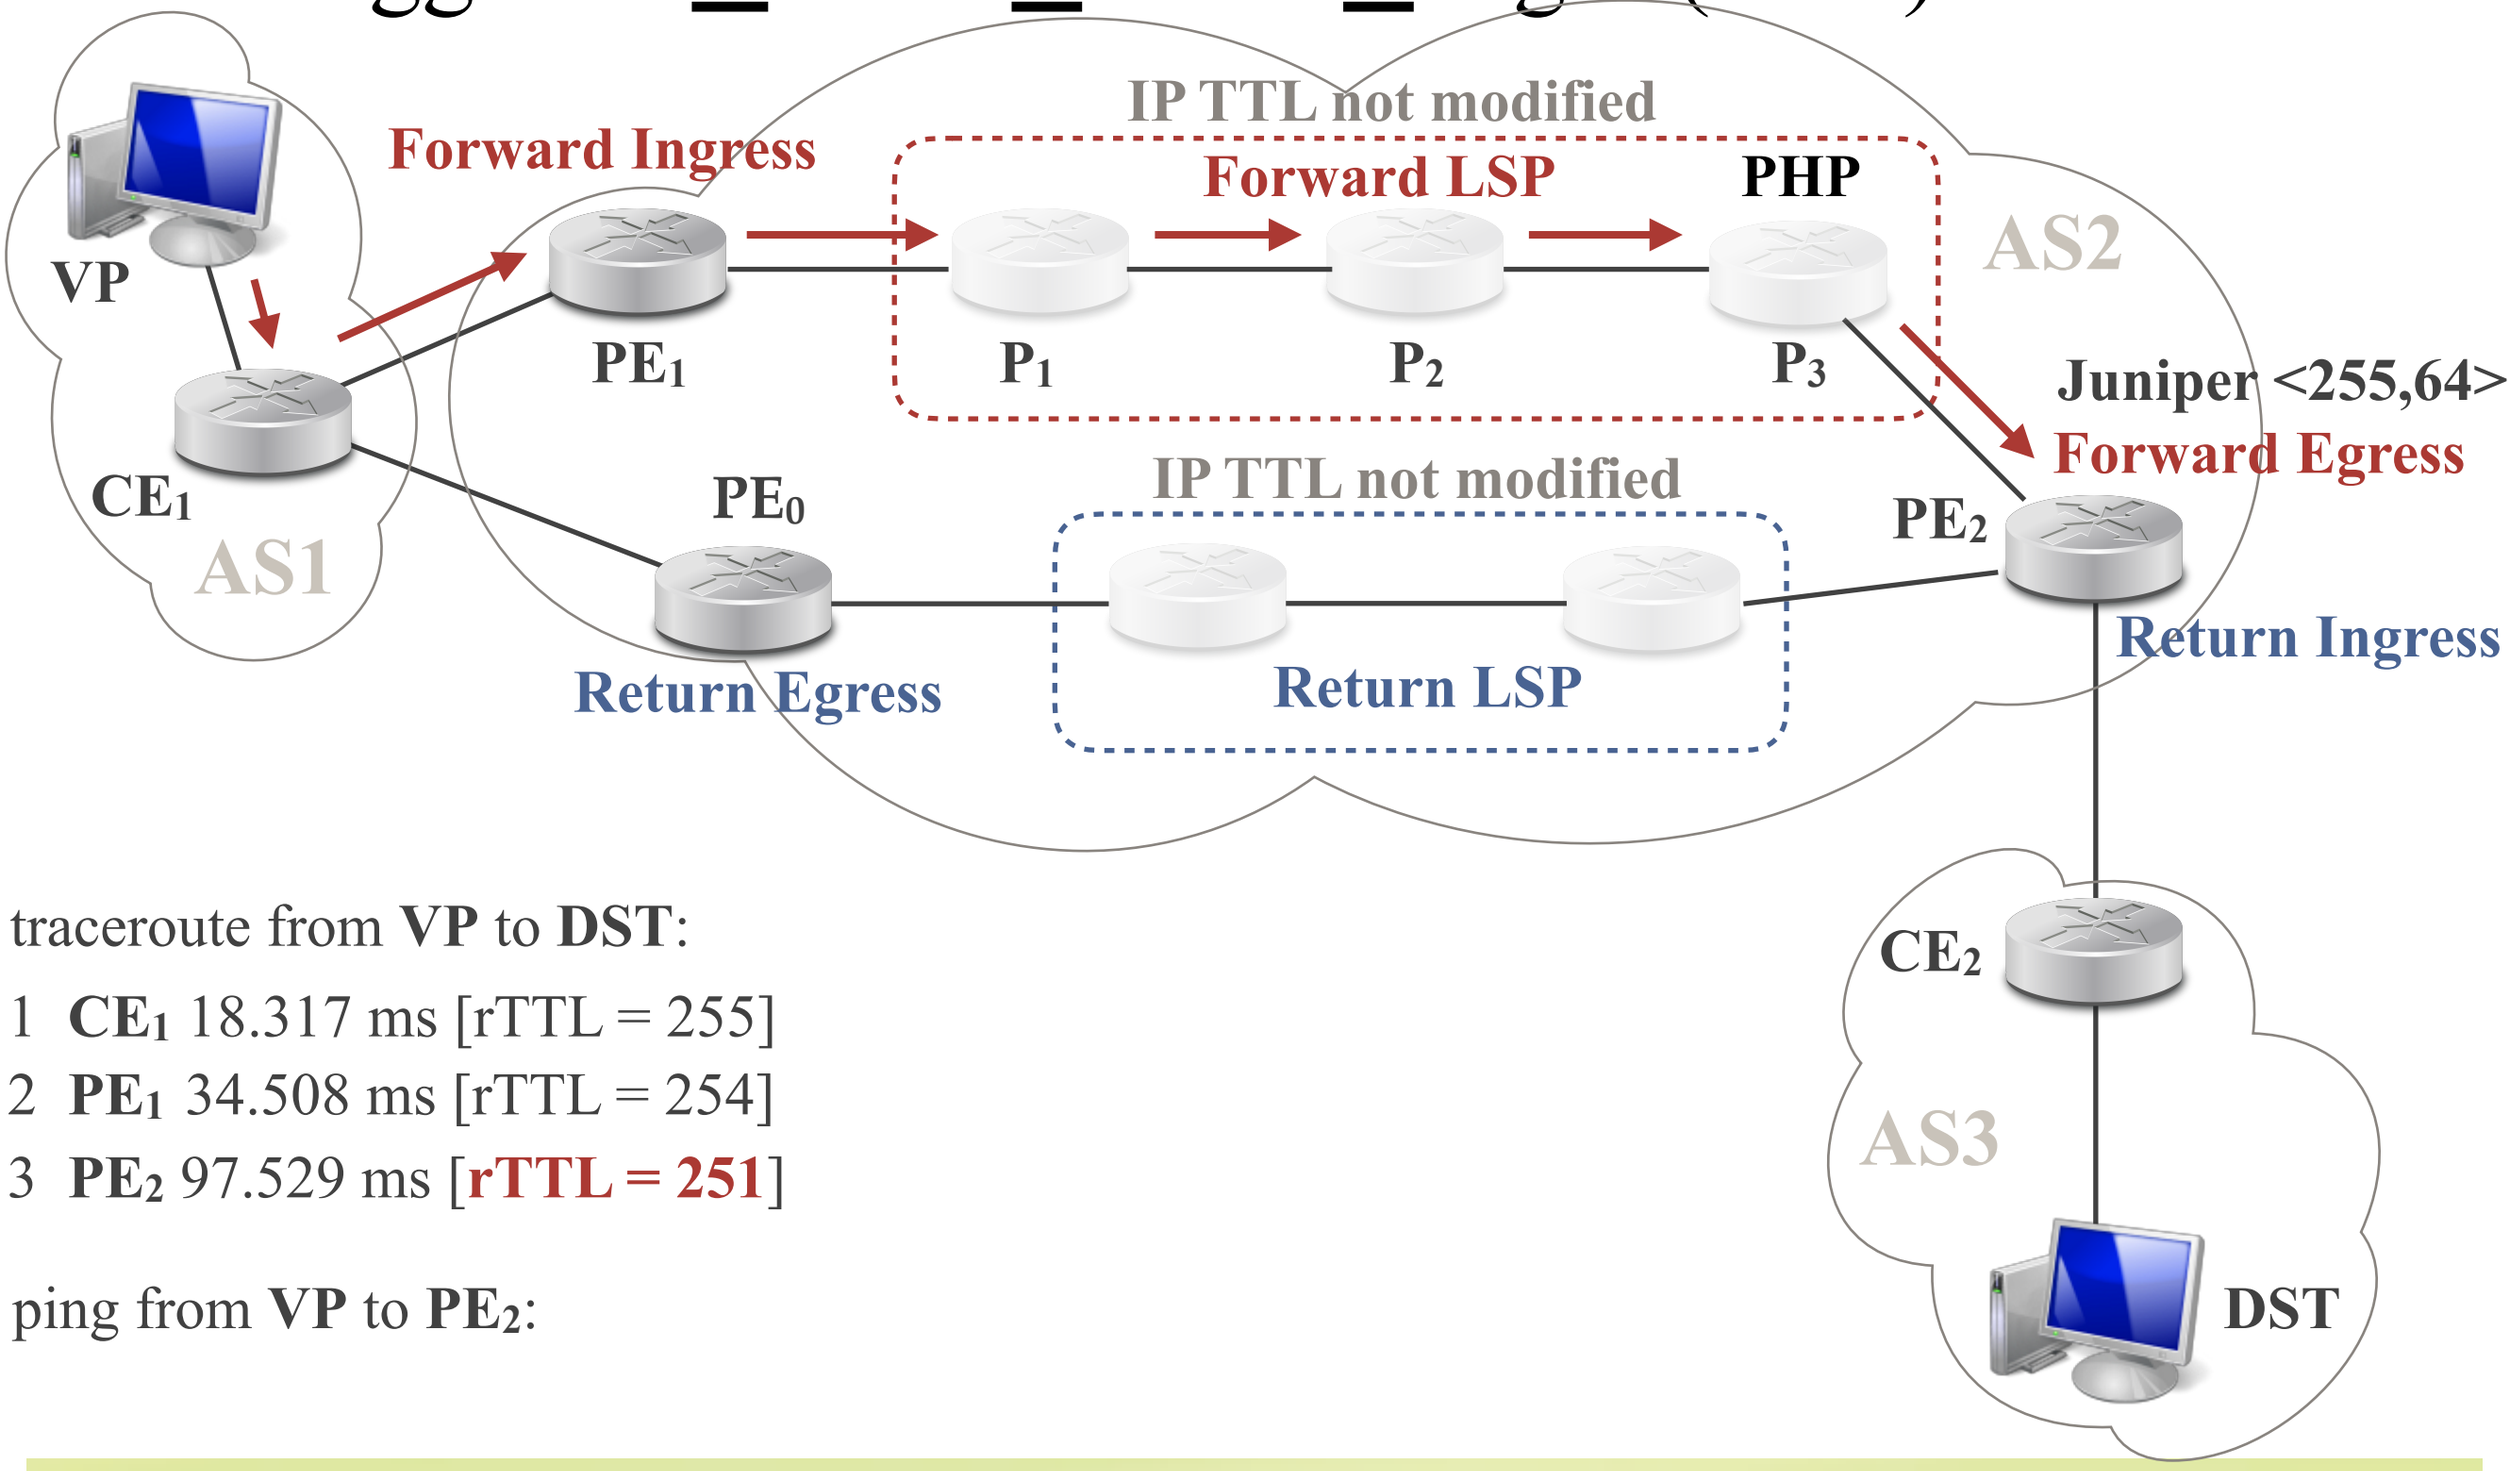

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

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

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

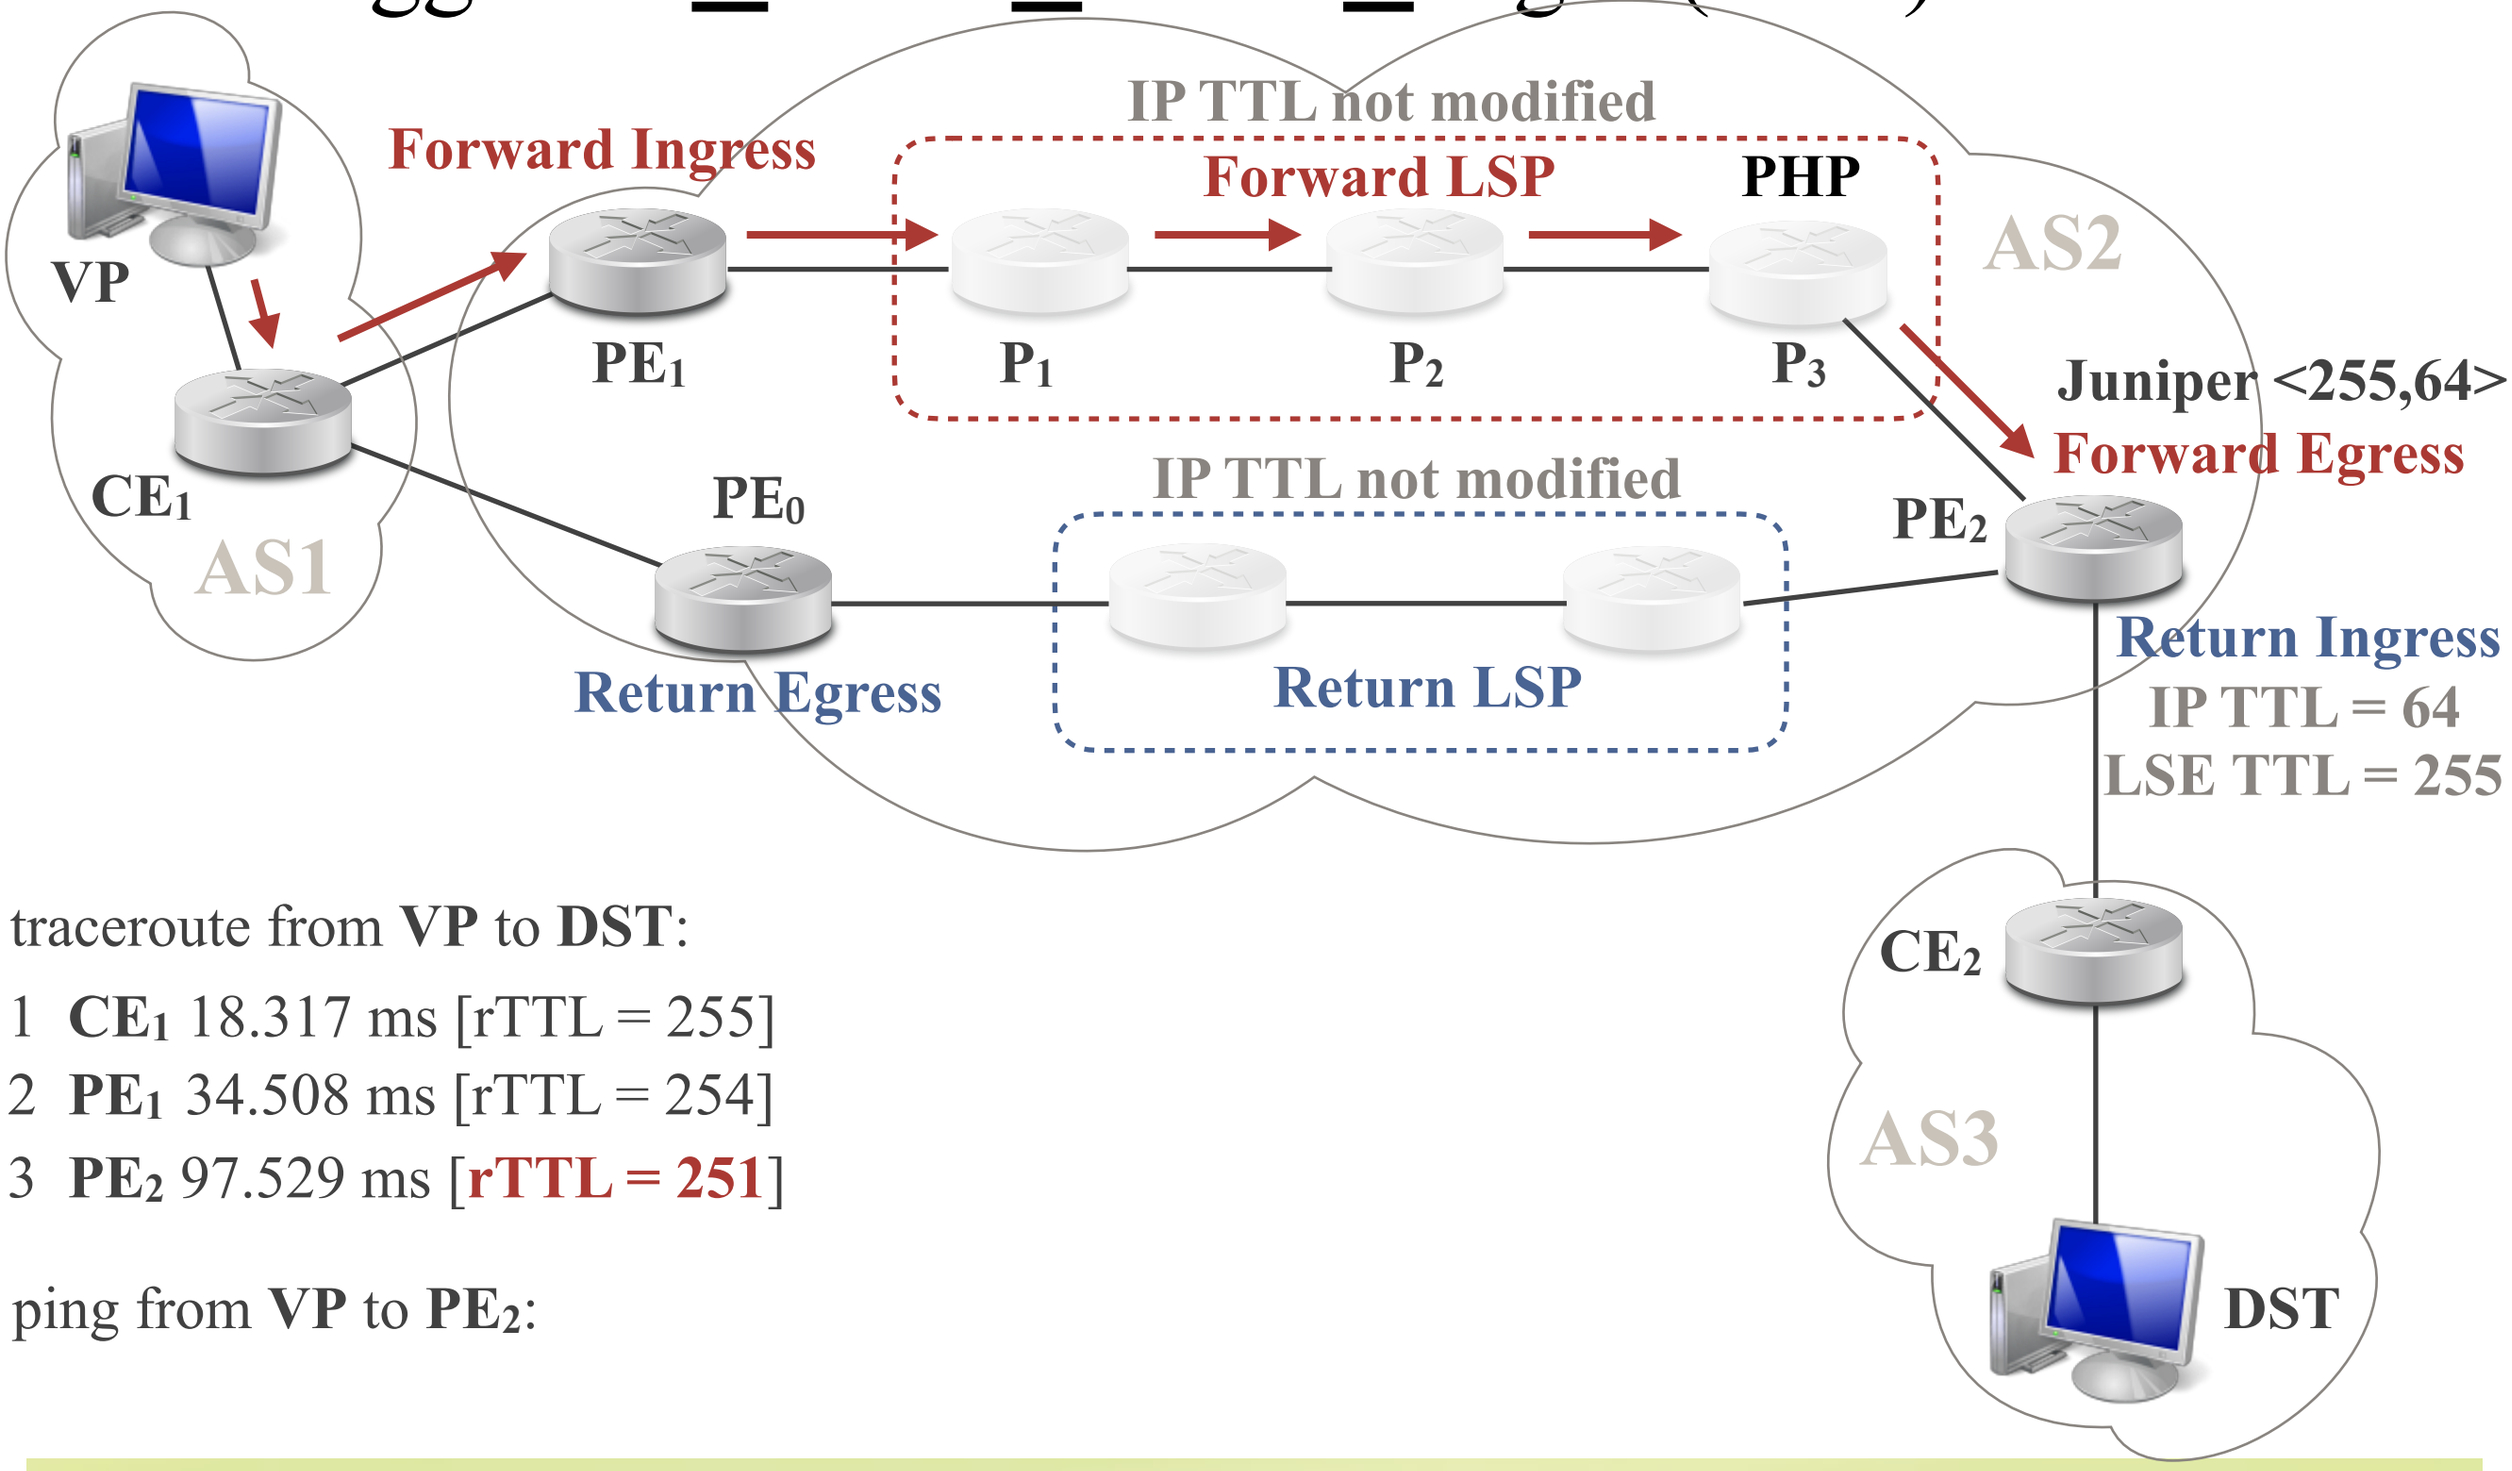

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

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

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

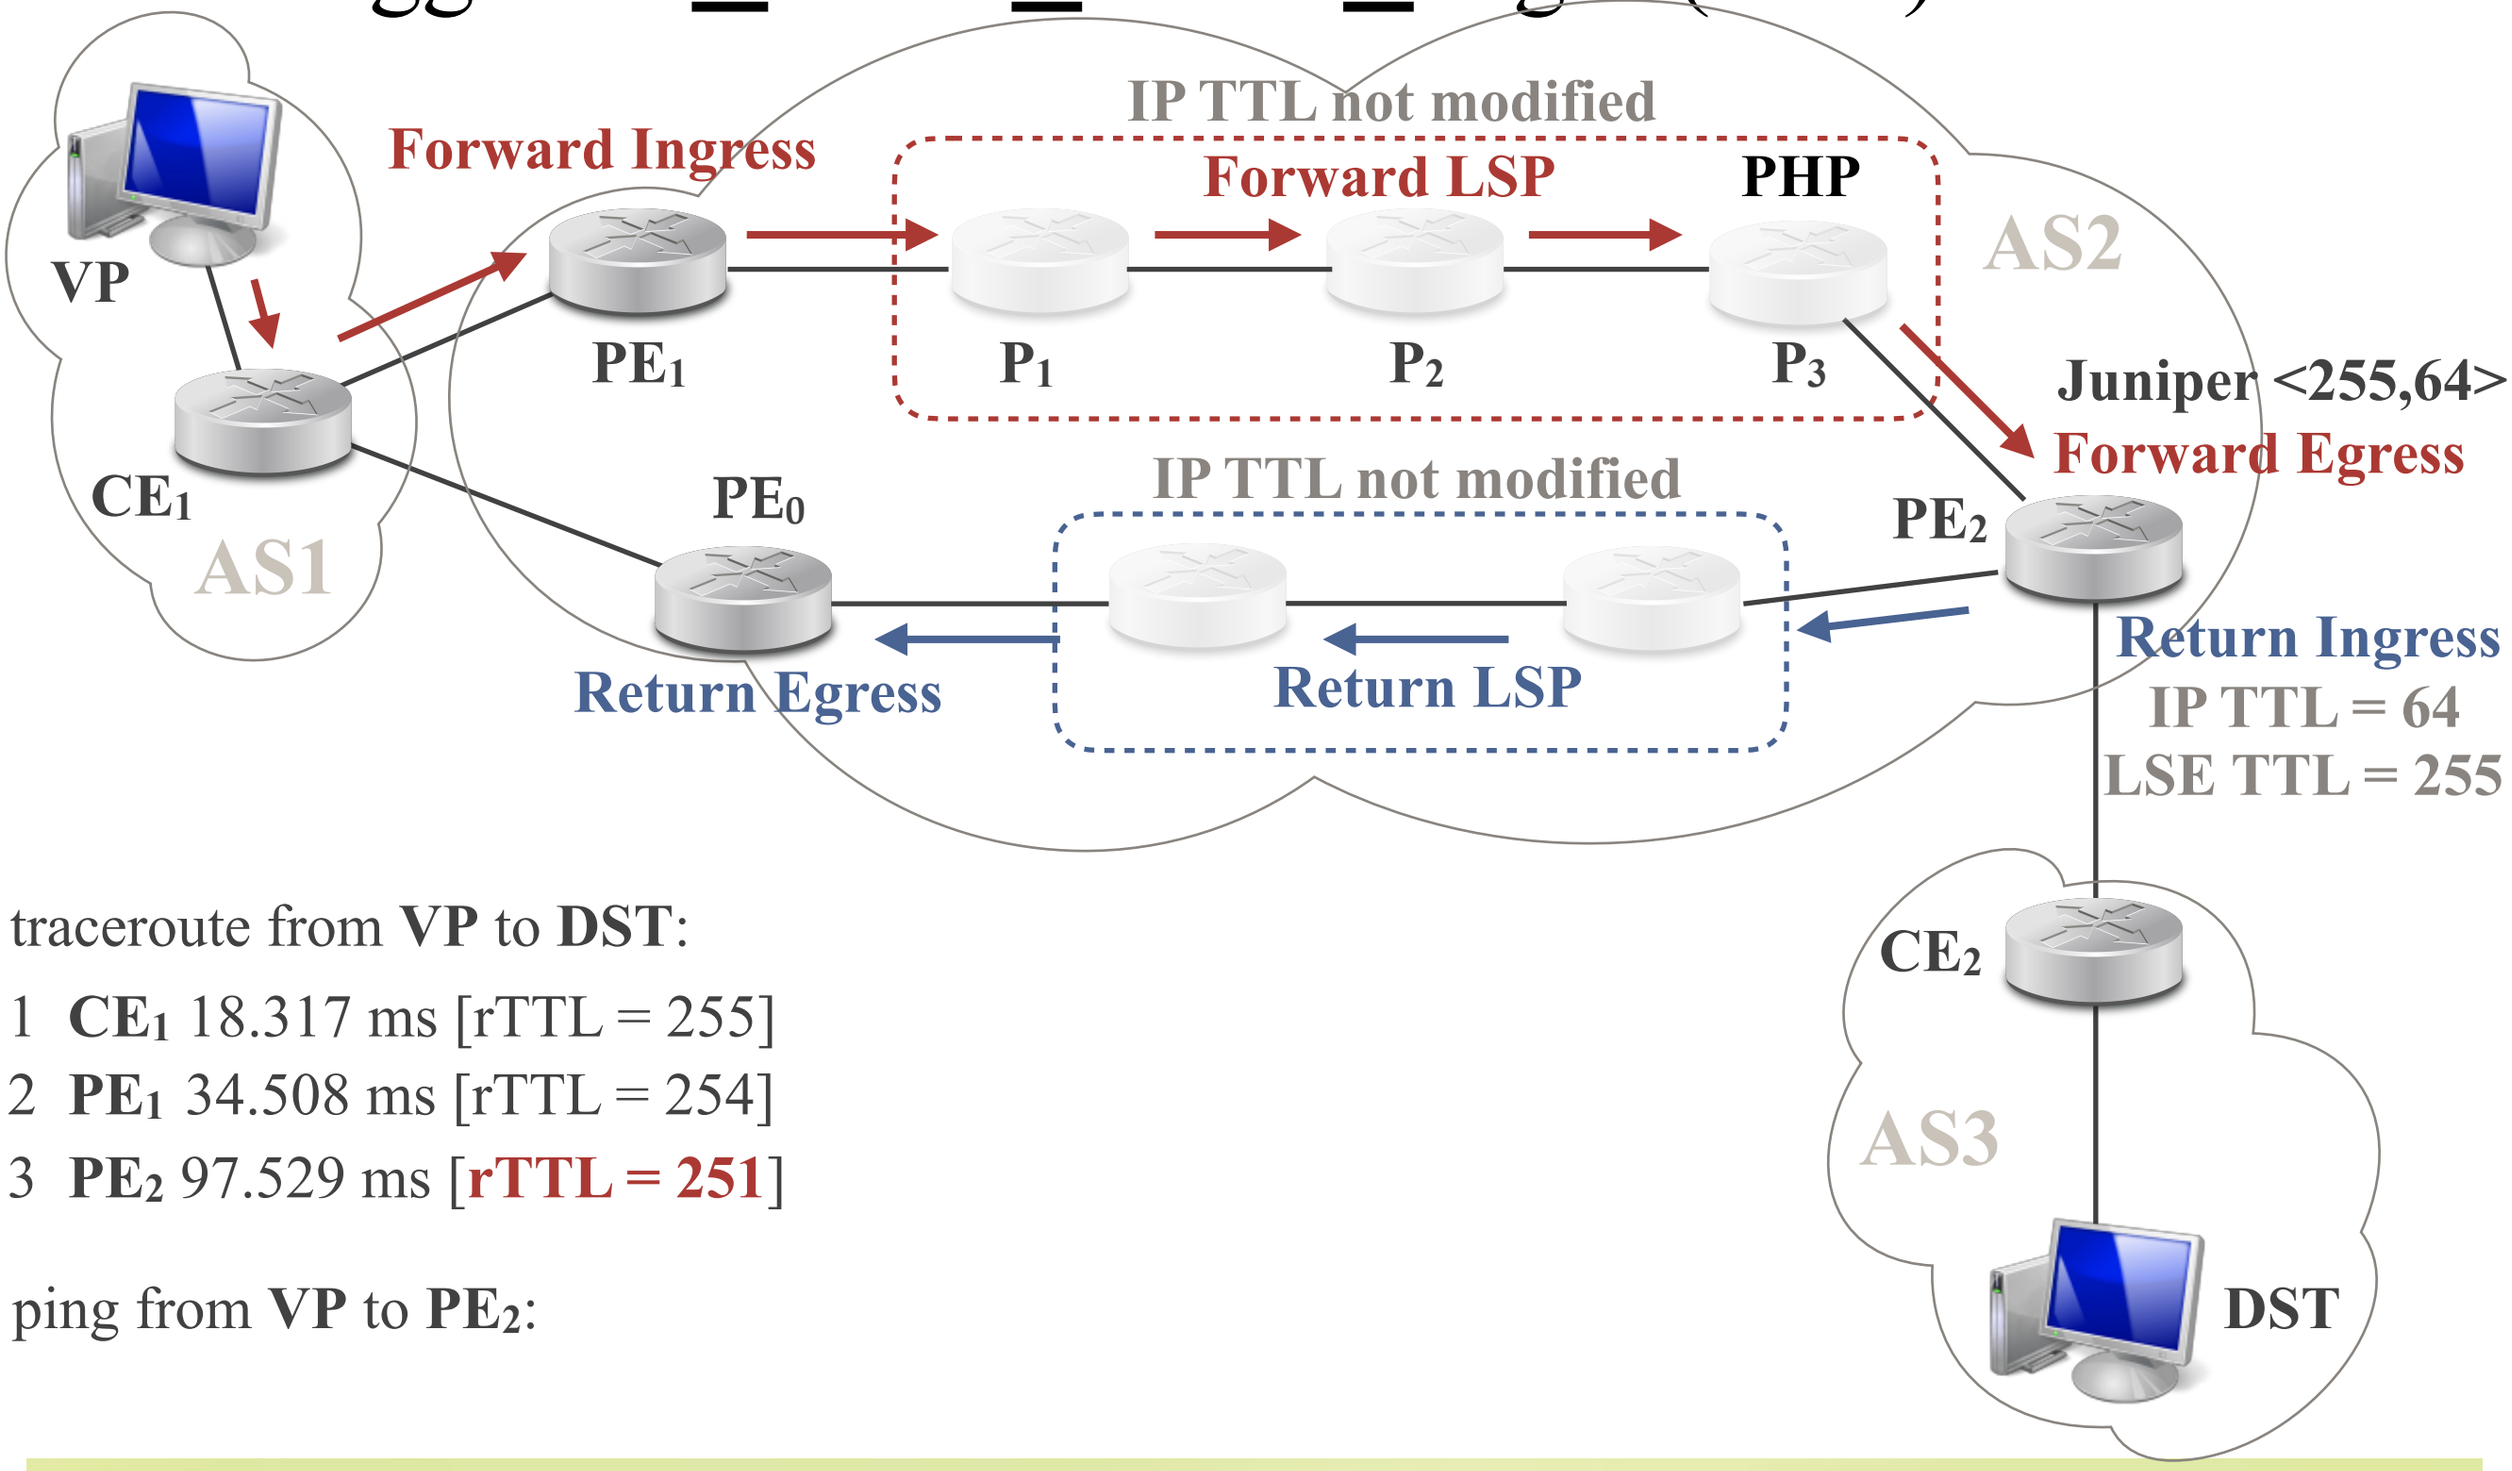

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

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

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

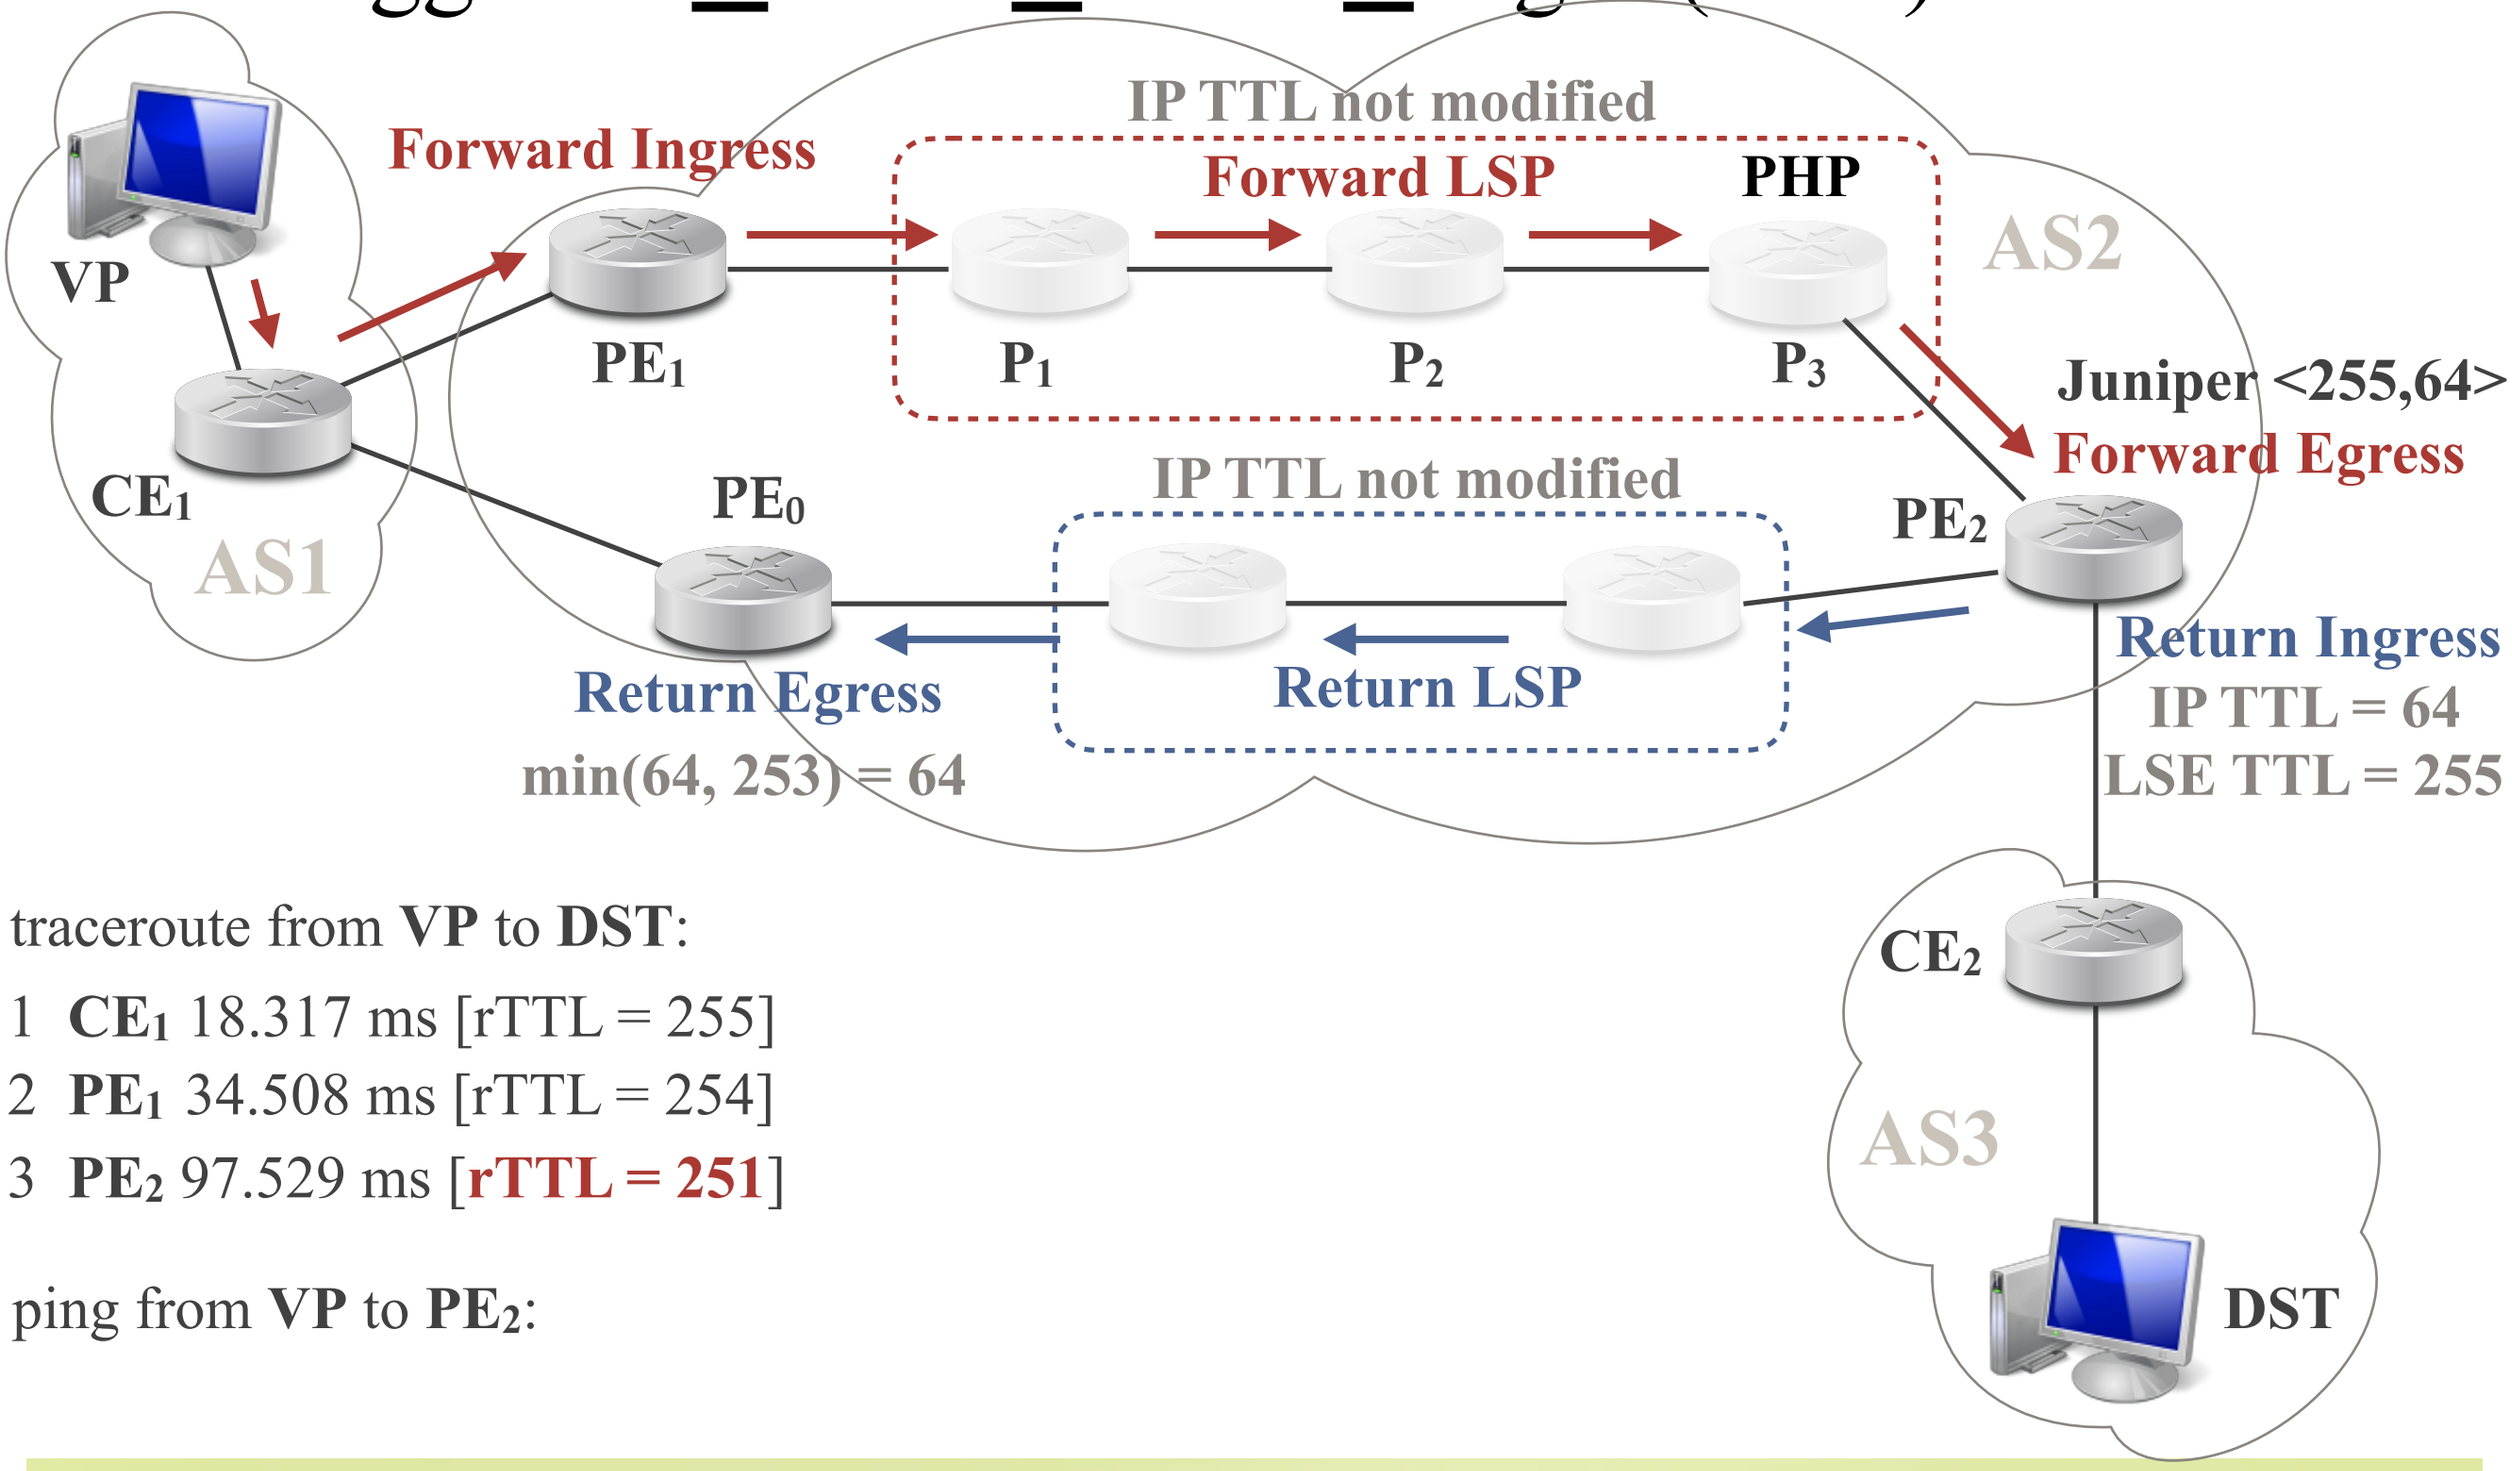

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

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

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

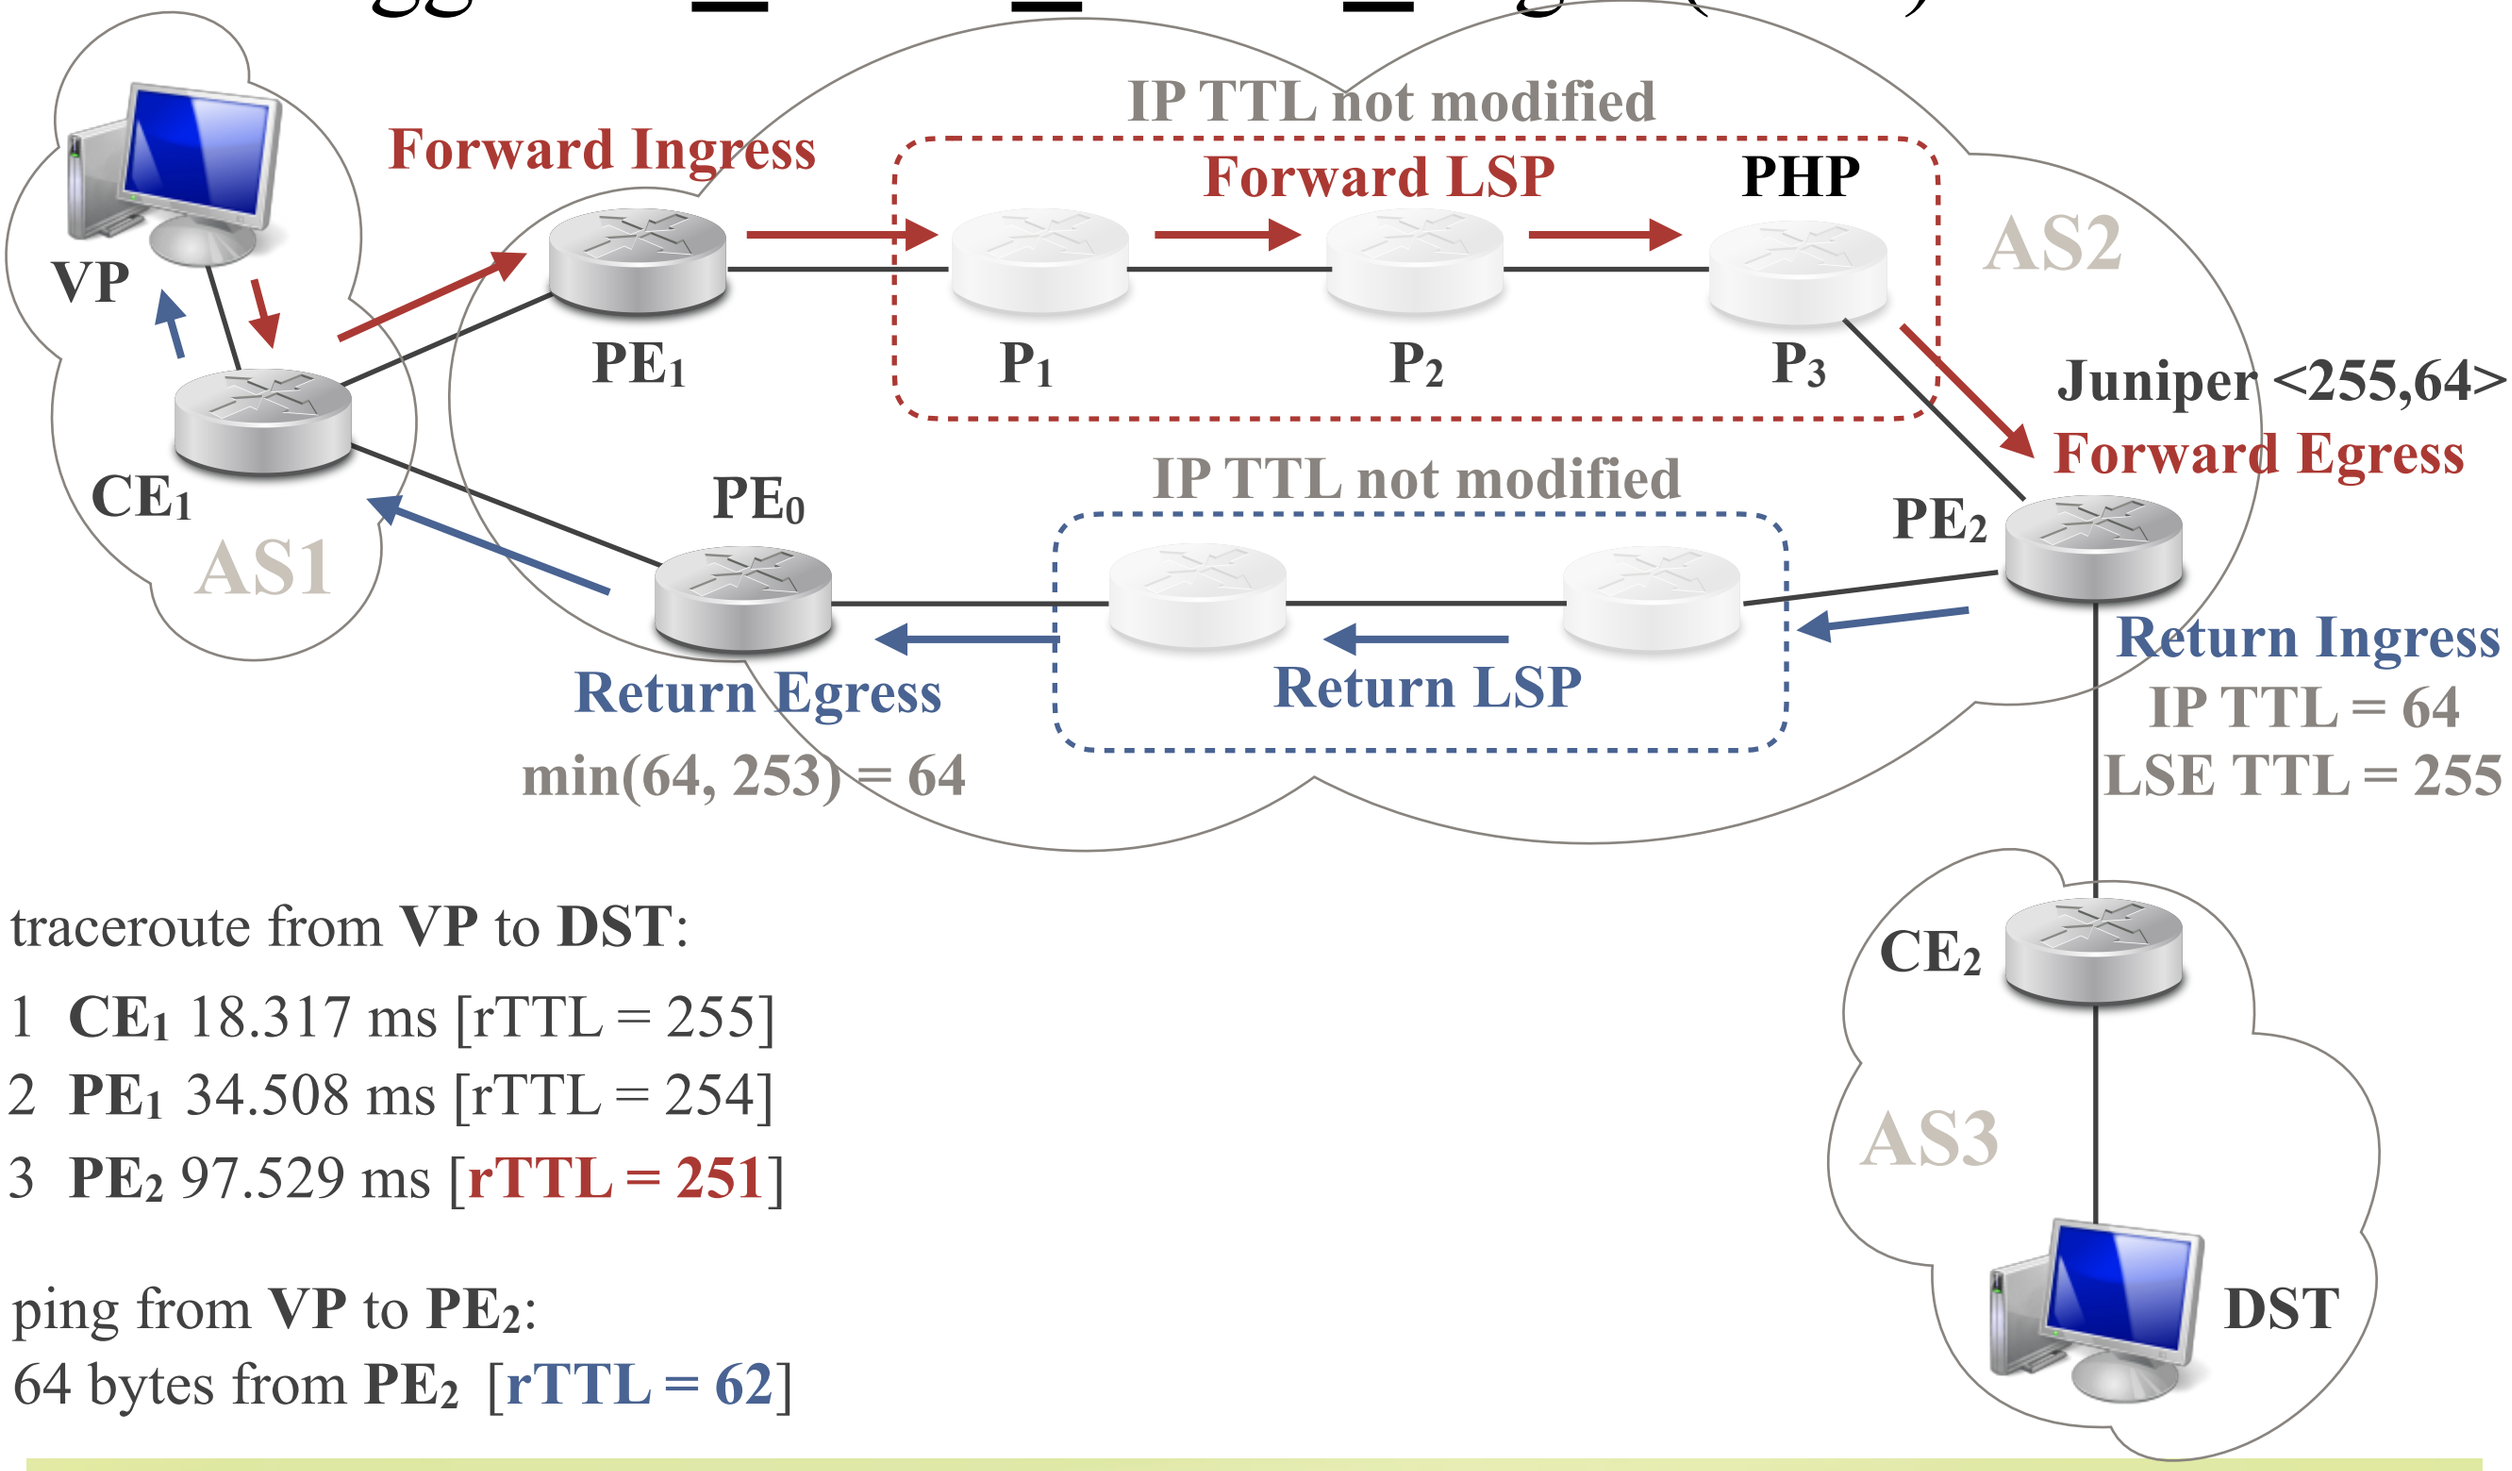

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 **PE2** 97.529 ms [**rTTL = 251**]

ping from **VP** to **PE2**:
64 bytes from **PE2** [**rTTL = 62**]

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

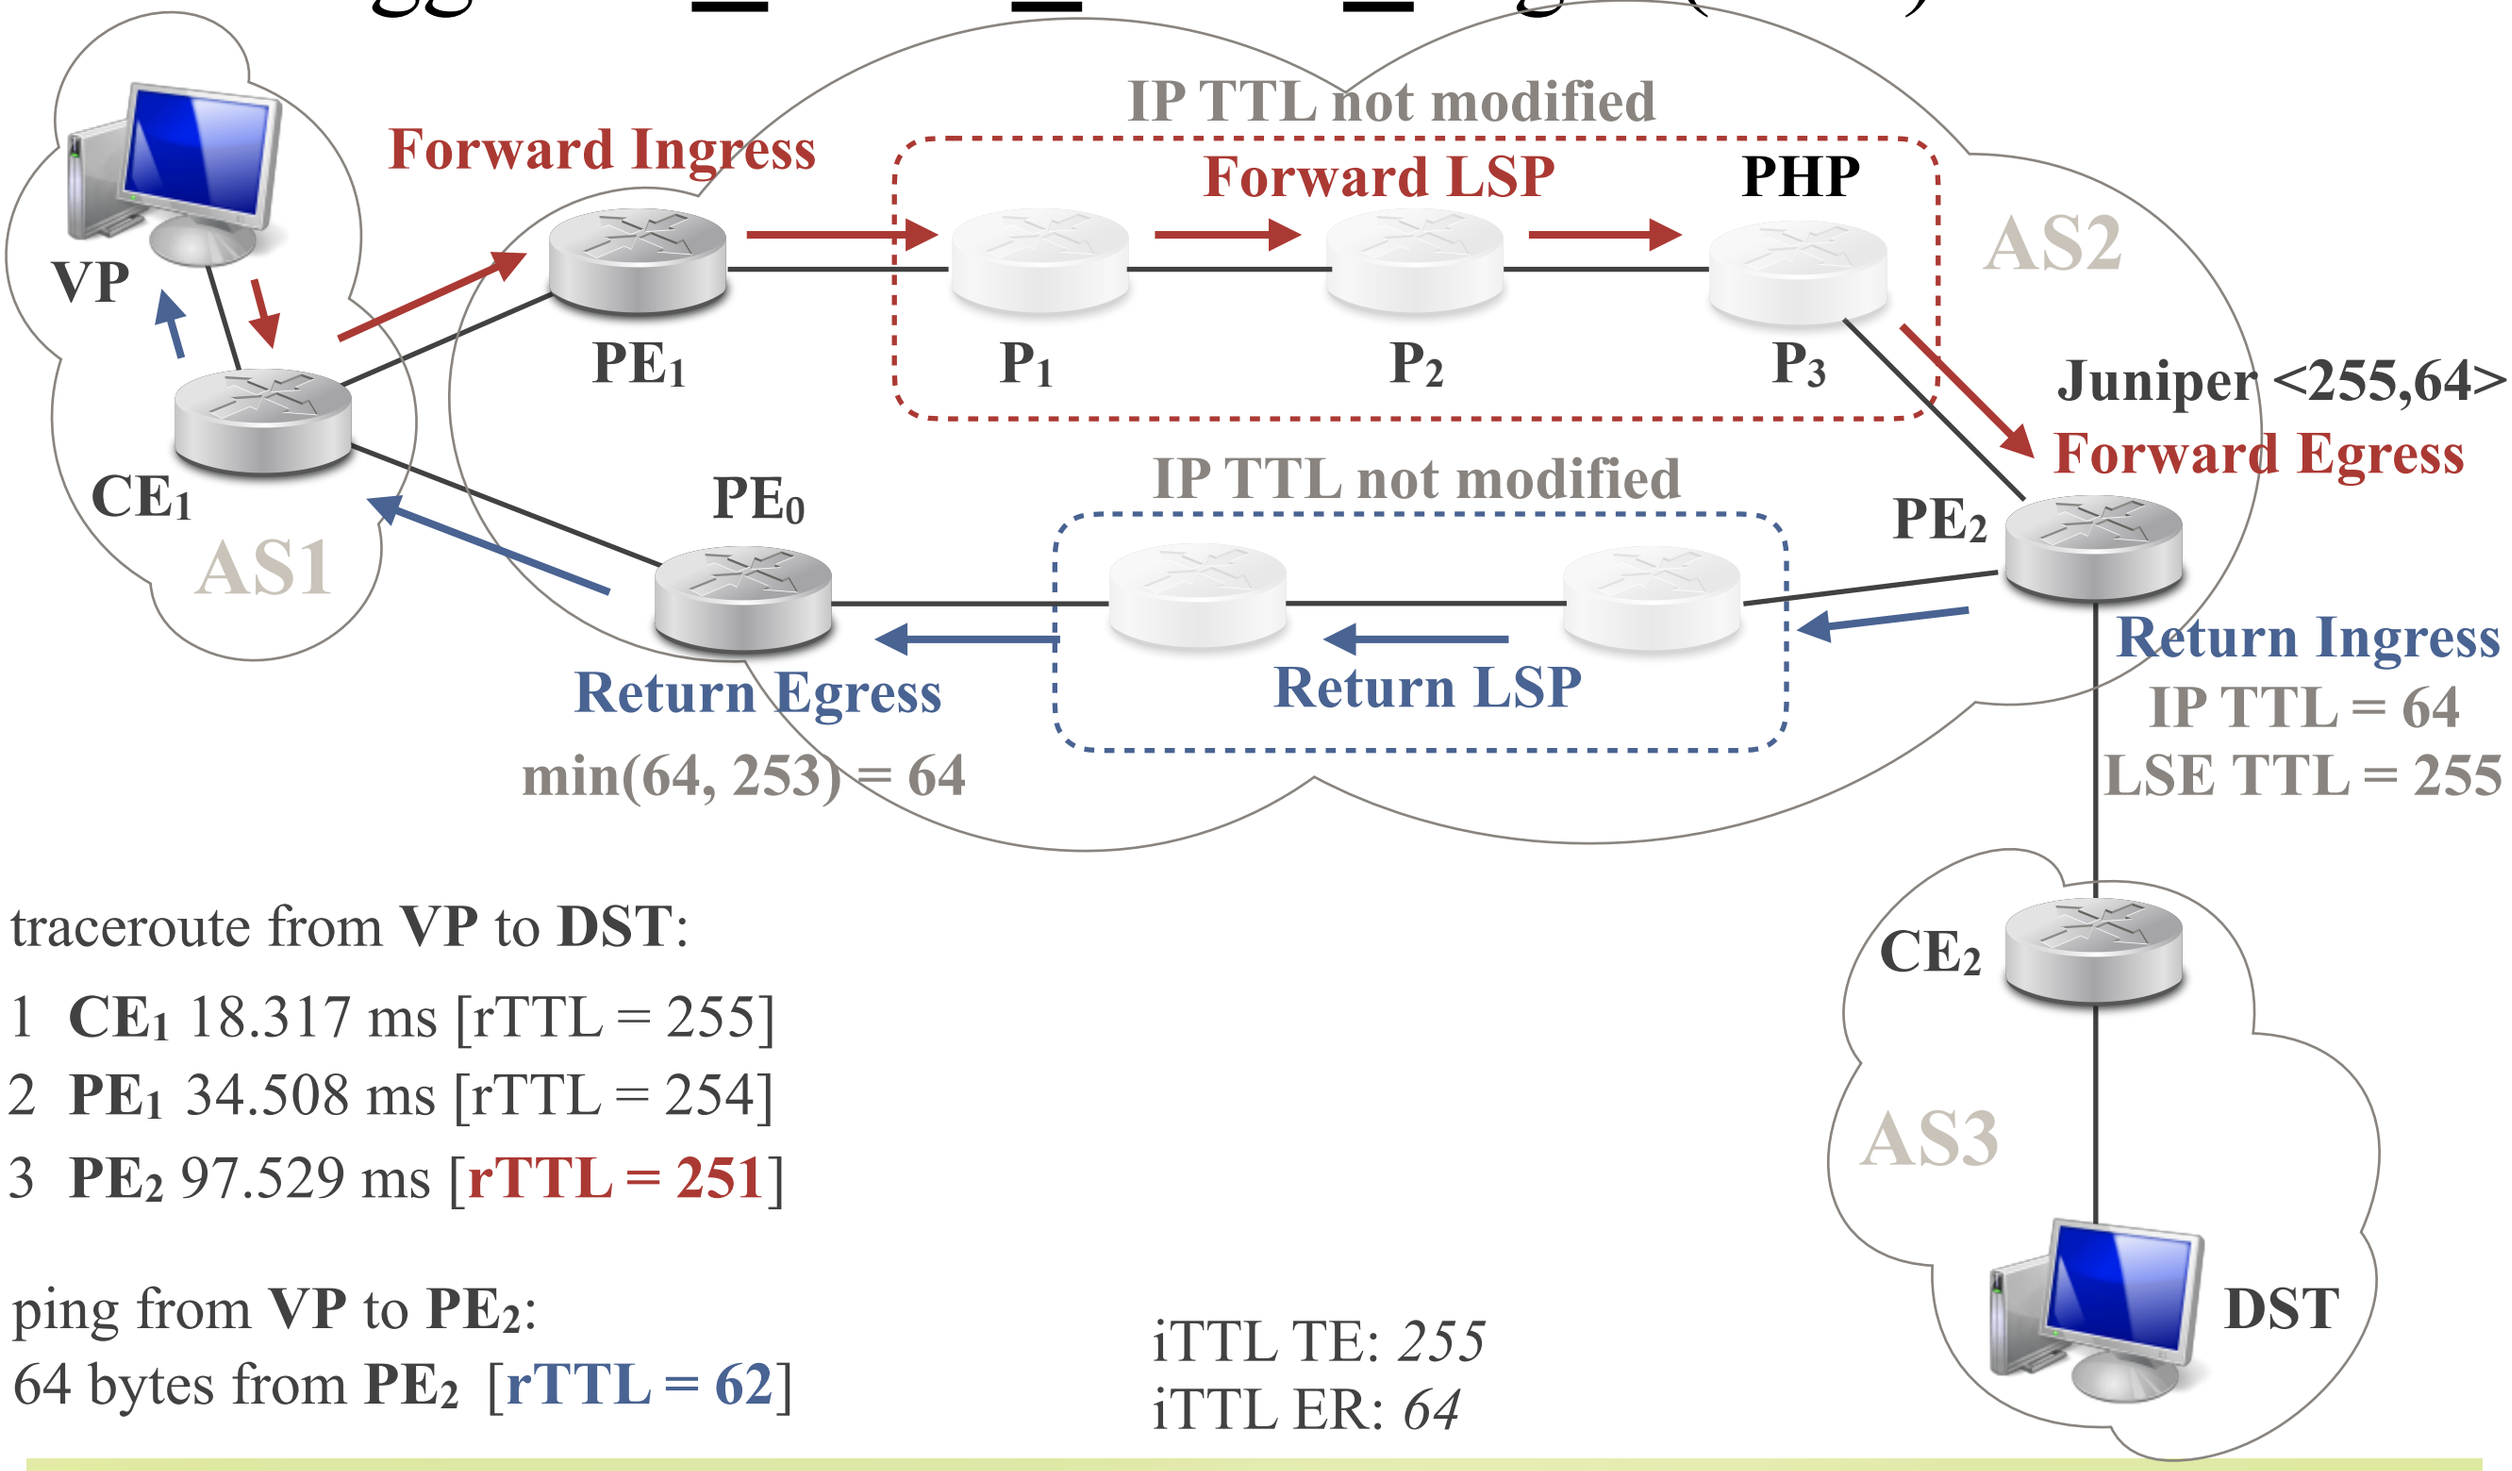

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 **PE2** 97.529 ms [**rTTL = 251**]

ping from **VP** to **PE2**:
64 bytes from **PE2** [**rTTL = 62**]

iTTL TE: *255*
iTTL ER: *64*

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

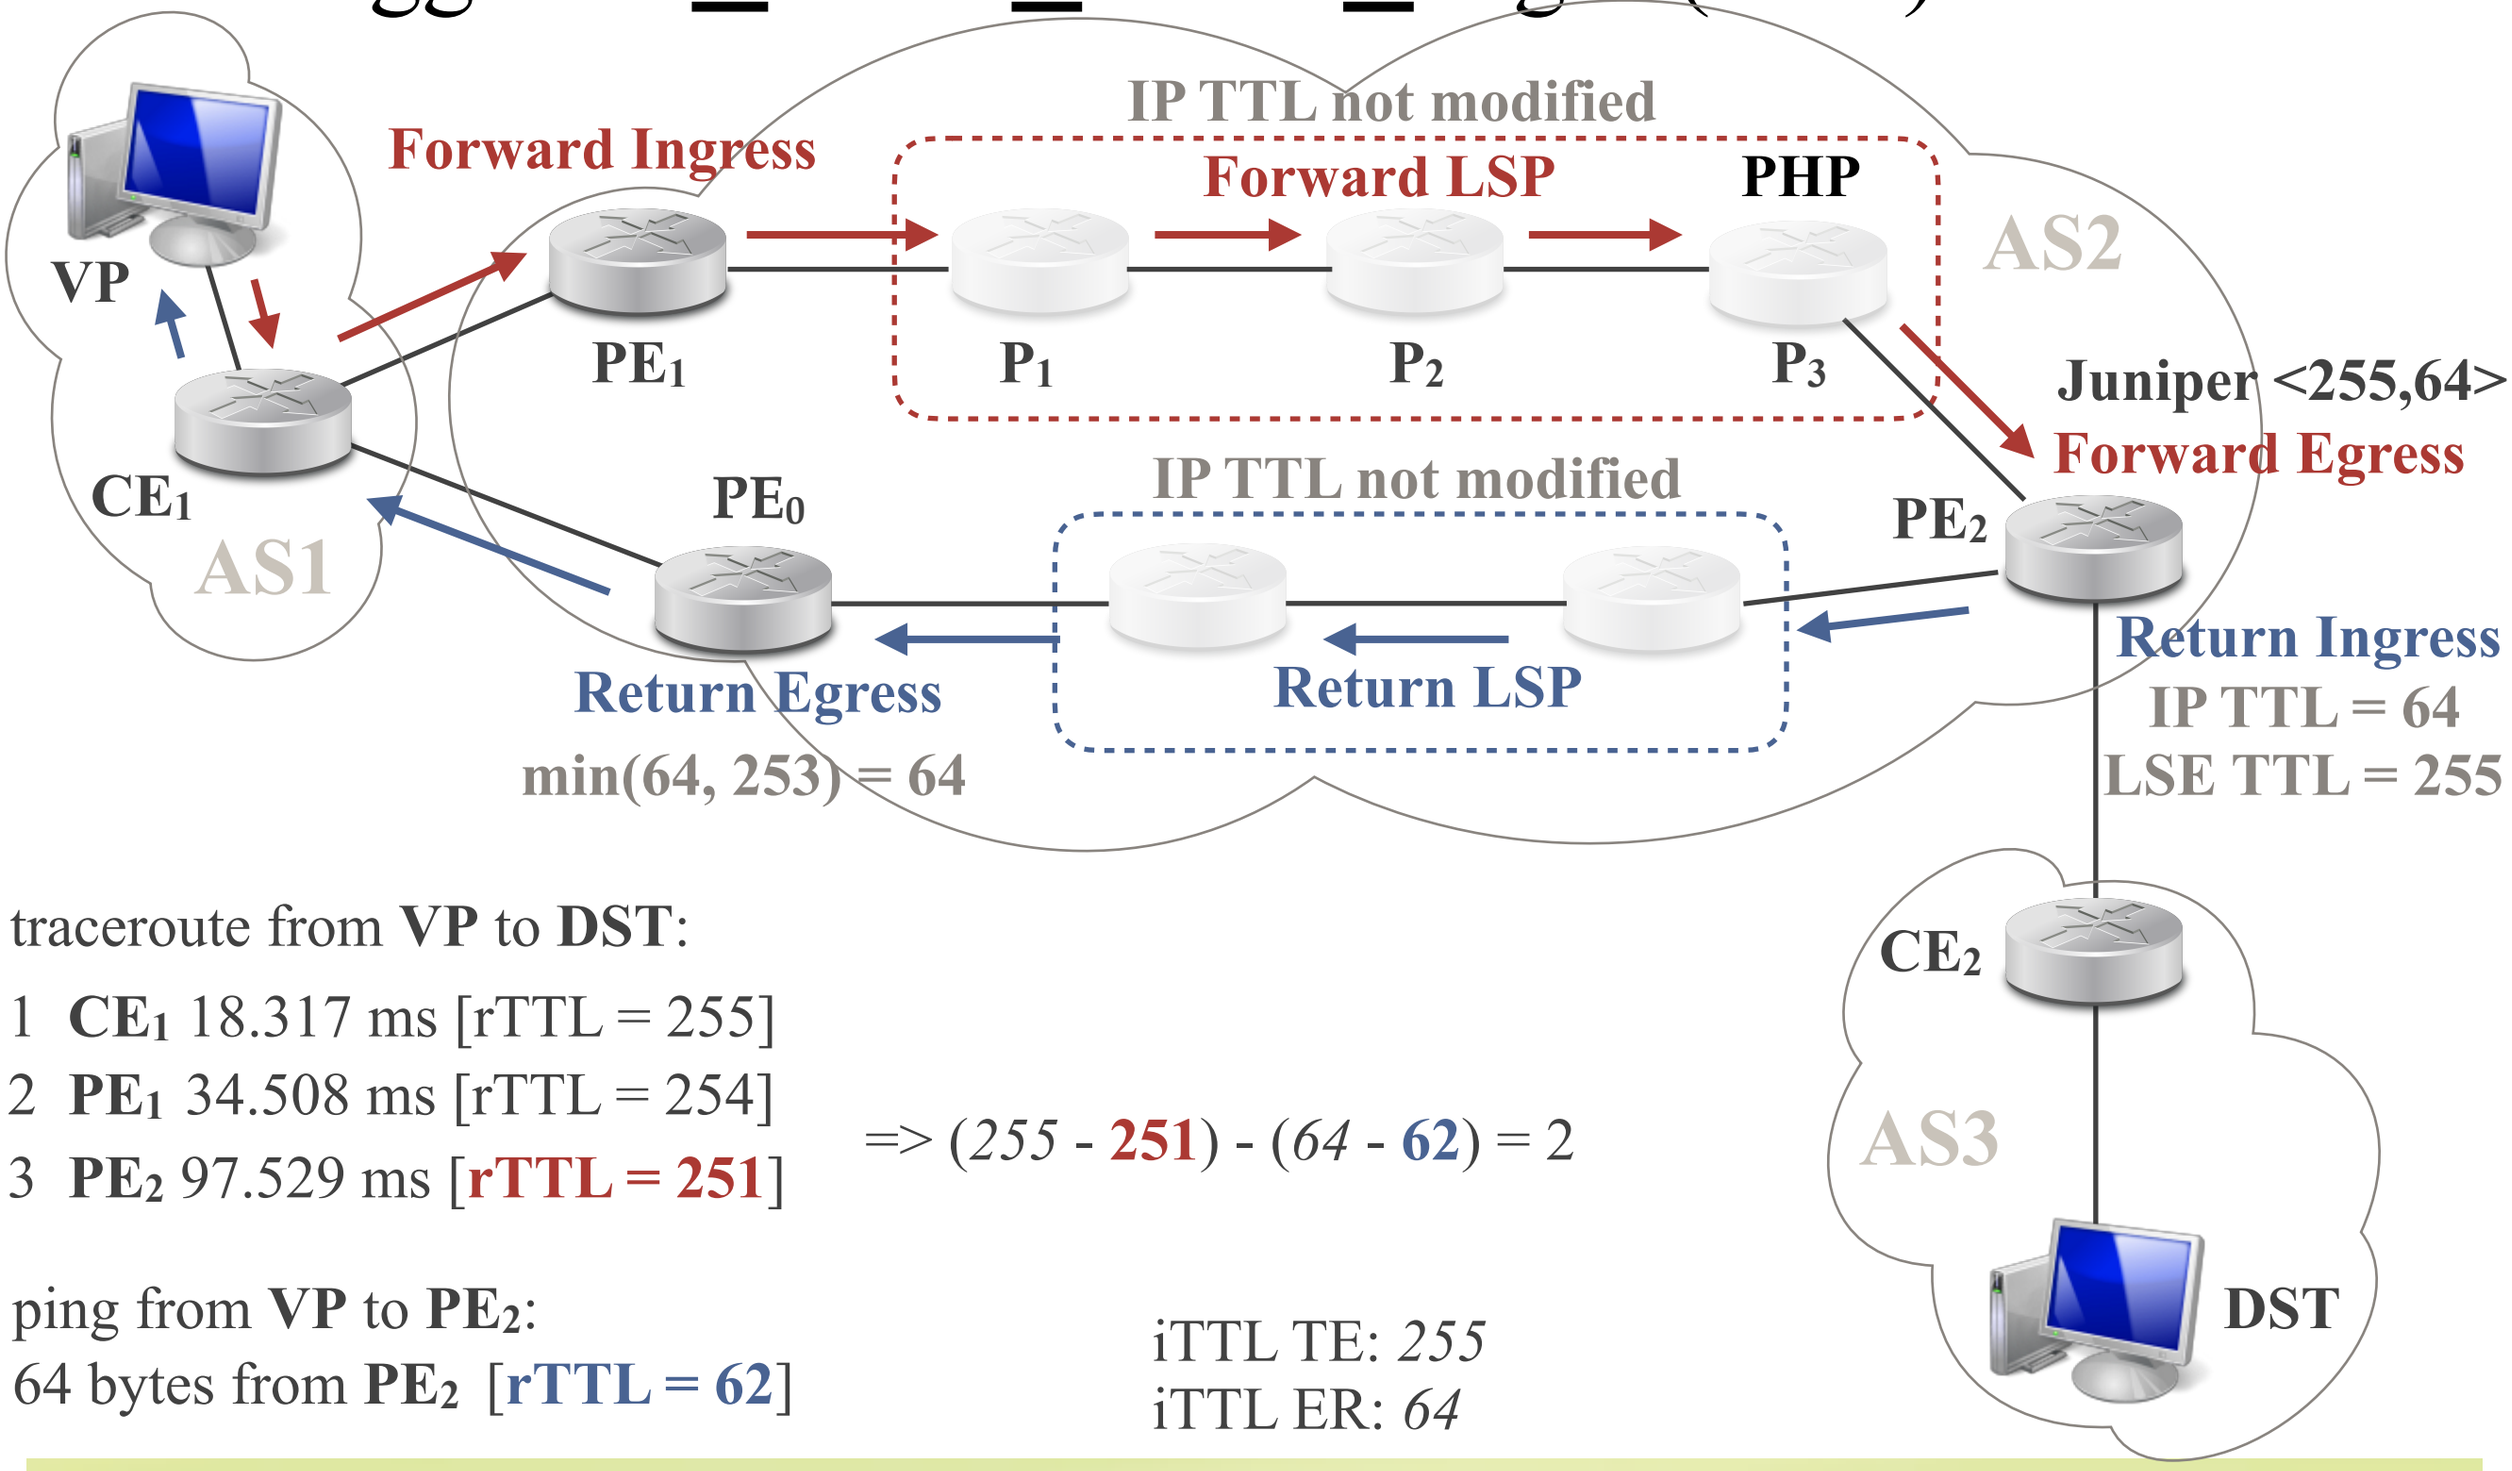

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 **PE2** 97.529 ms [**rTTL = 251**]

=> (*255* - **251**) - (*64* - **62**) = 2

ping from **VP** to **PE2**:
64 bytes from **PE2** [**rTTL = 62**]

iTTL TE: *255*
iTTL ER: *64*

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

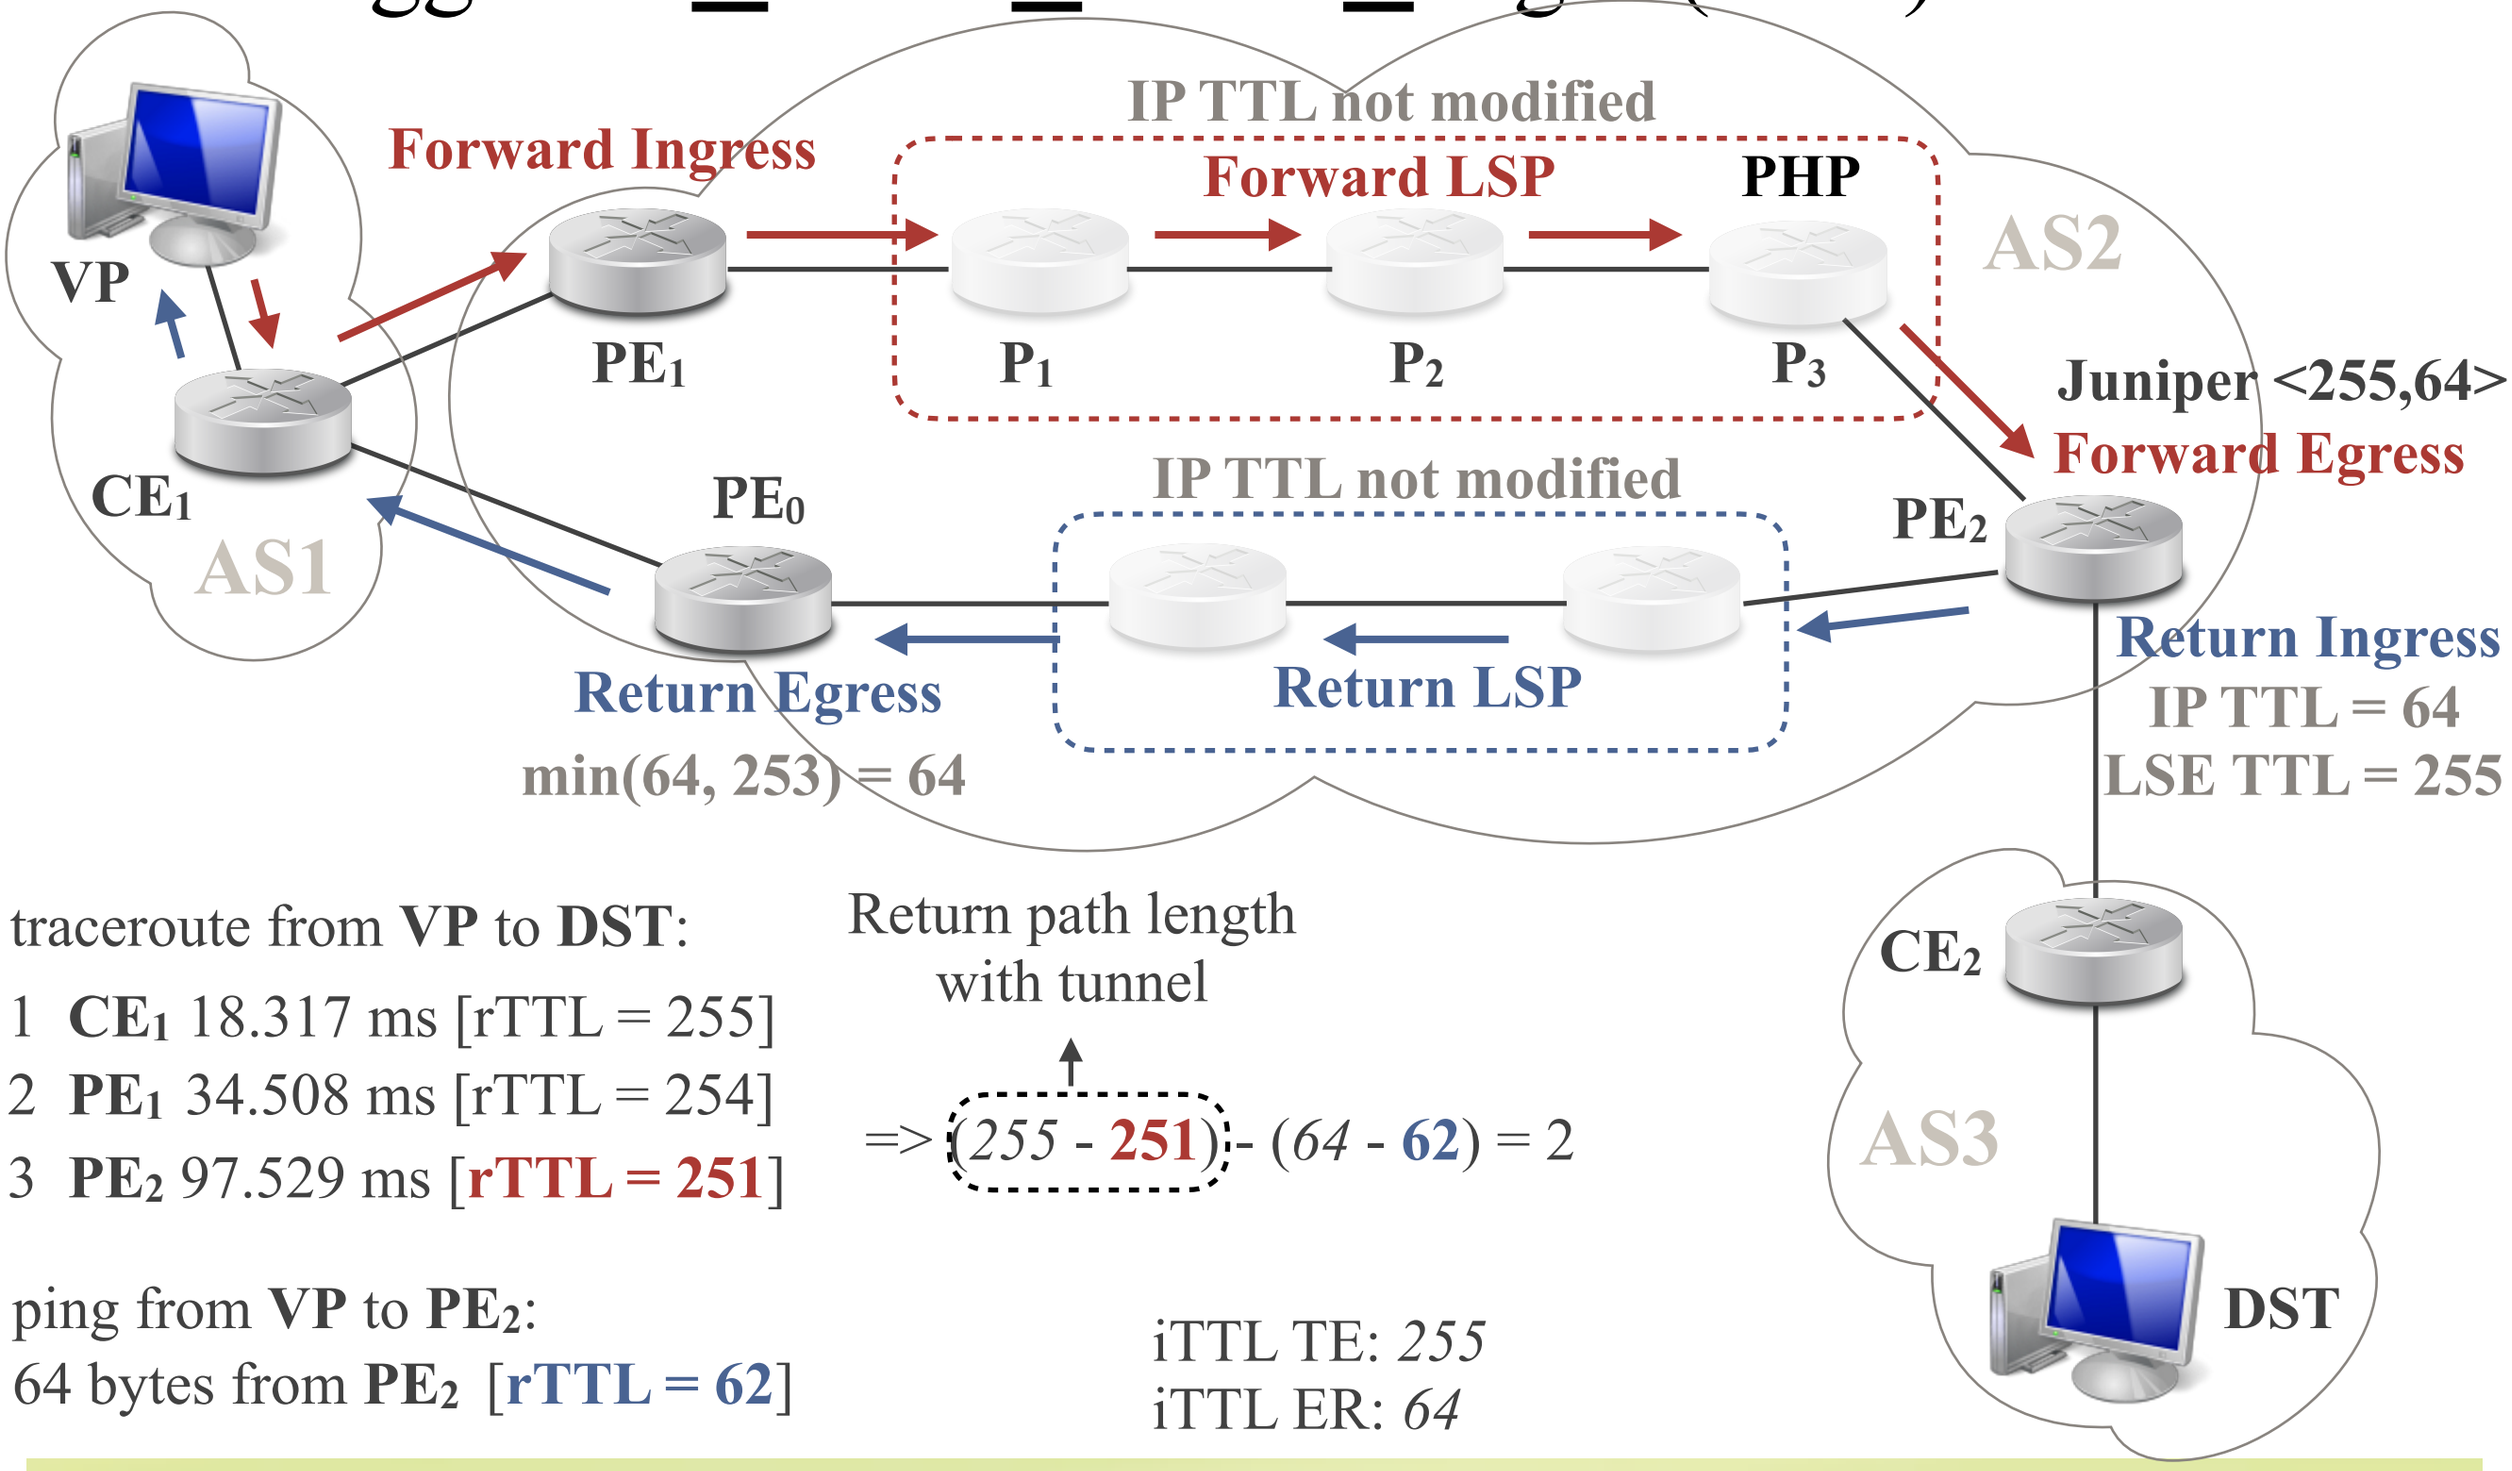

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

ping from **VP** to **PE$_2$**:
64 bytes from **PE$_2$** [**rTTL = 62**]

Return path length with tunnel

=> (*255* - **251**) - (*64* - **62**) = 2

iTTL TE: *255*
iTTL ER: *64*

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

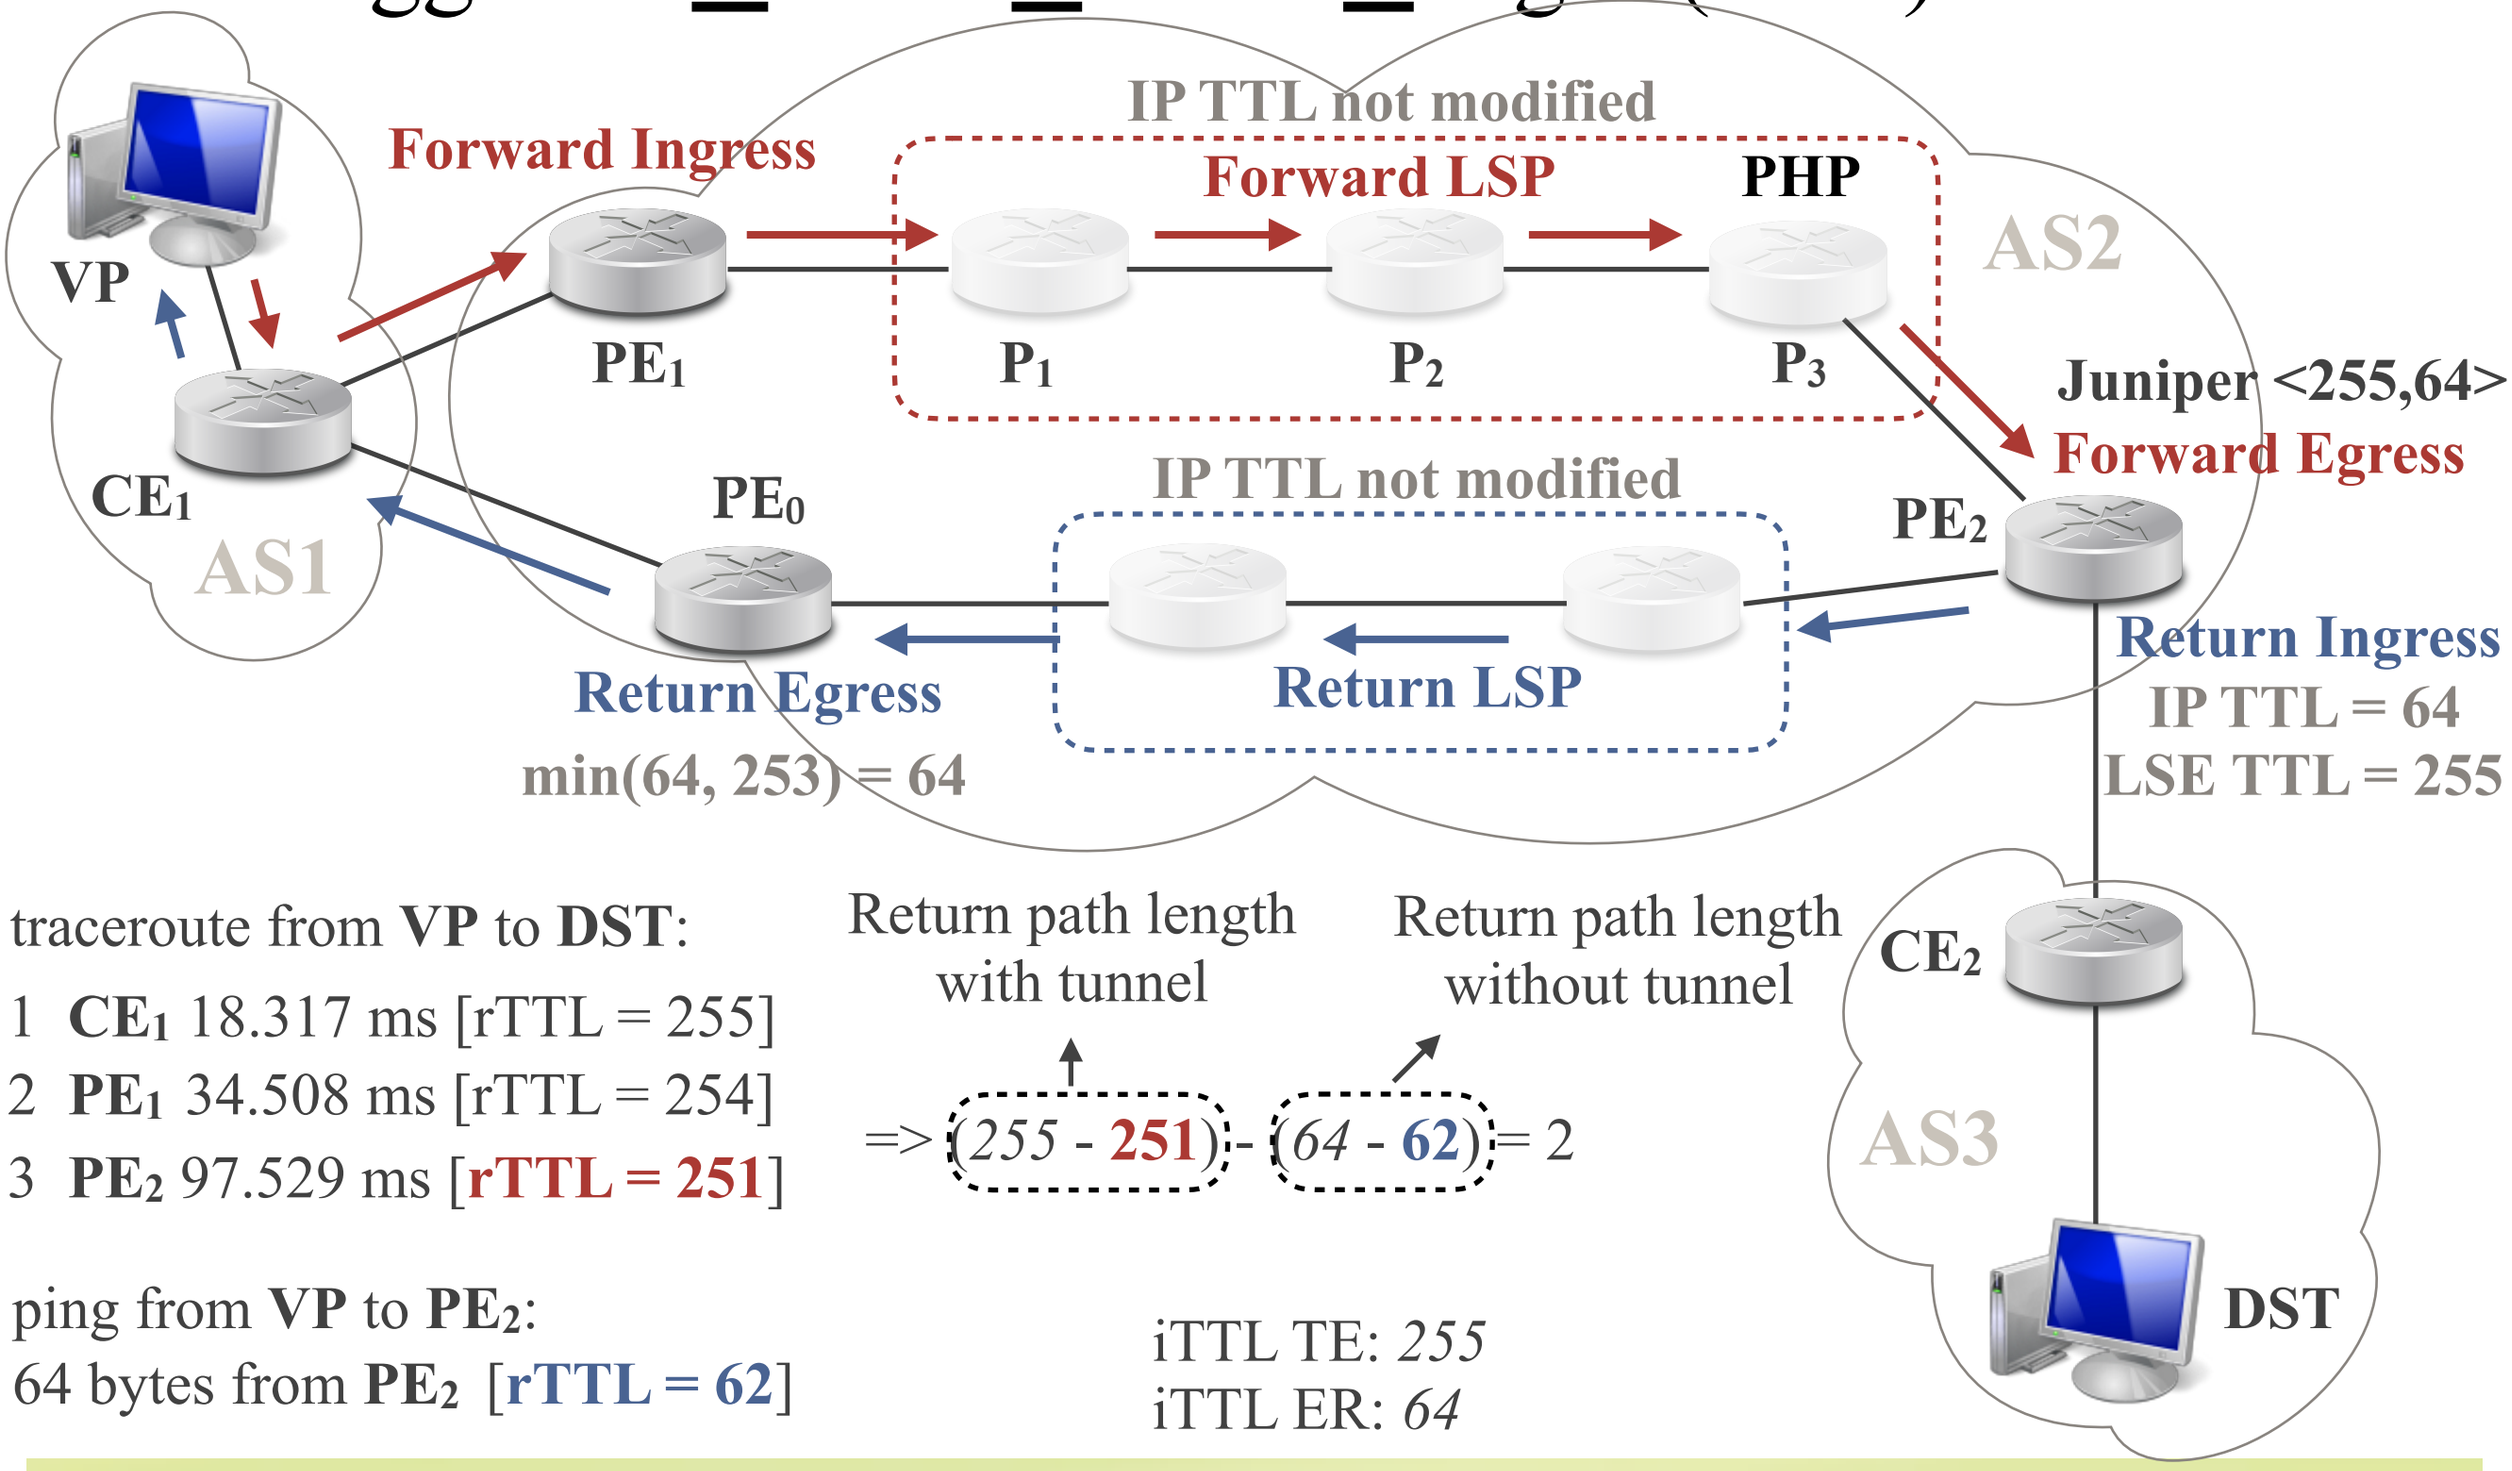

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 **PE2** 97.529 ms [**rTTL = 251**]

ping from **VP** to **PE2**:
64 bytes from **PE2** [**rTTL = 62**]

Return path length with tunnel

Return path length without tunnel

=> (*255* - **251**) - (*64* - **62**) = 2

iTTL TE: *255*
iTTL ER: *64*

Invisible Tunnels (3)

- Trigger 2: **R**eturn **T**unnel **L**ength (RTL)

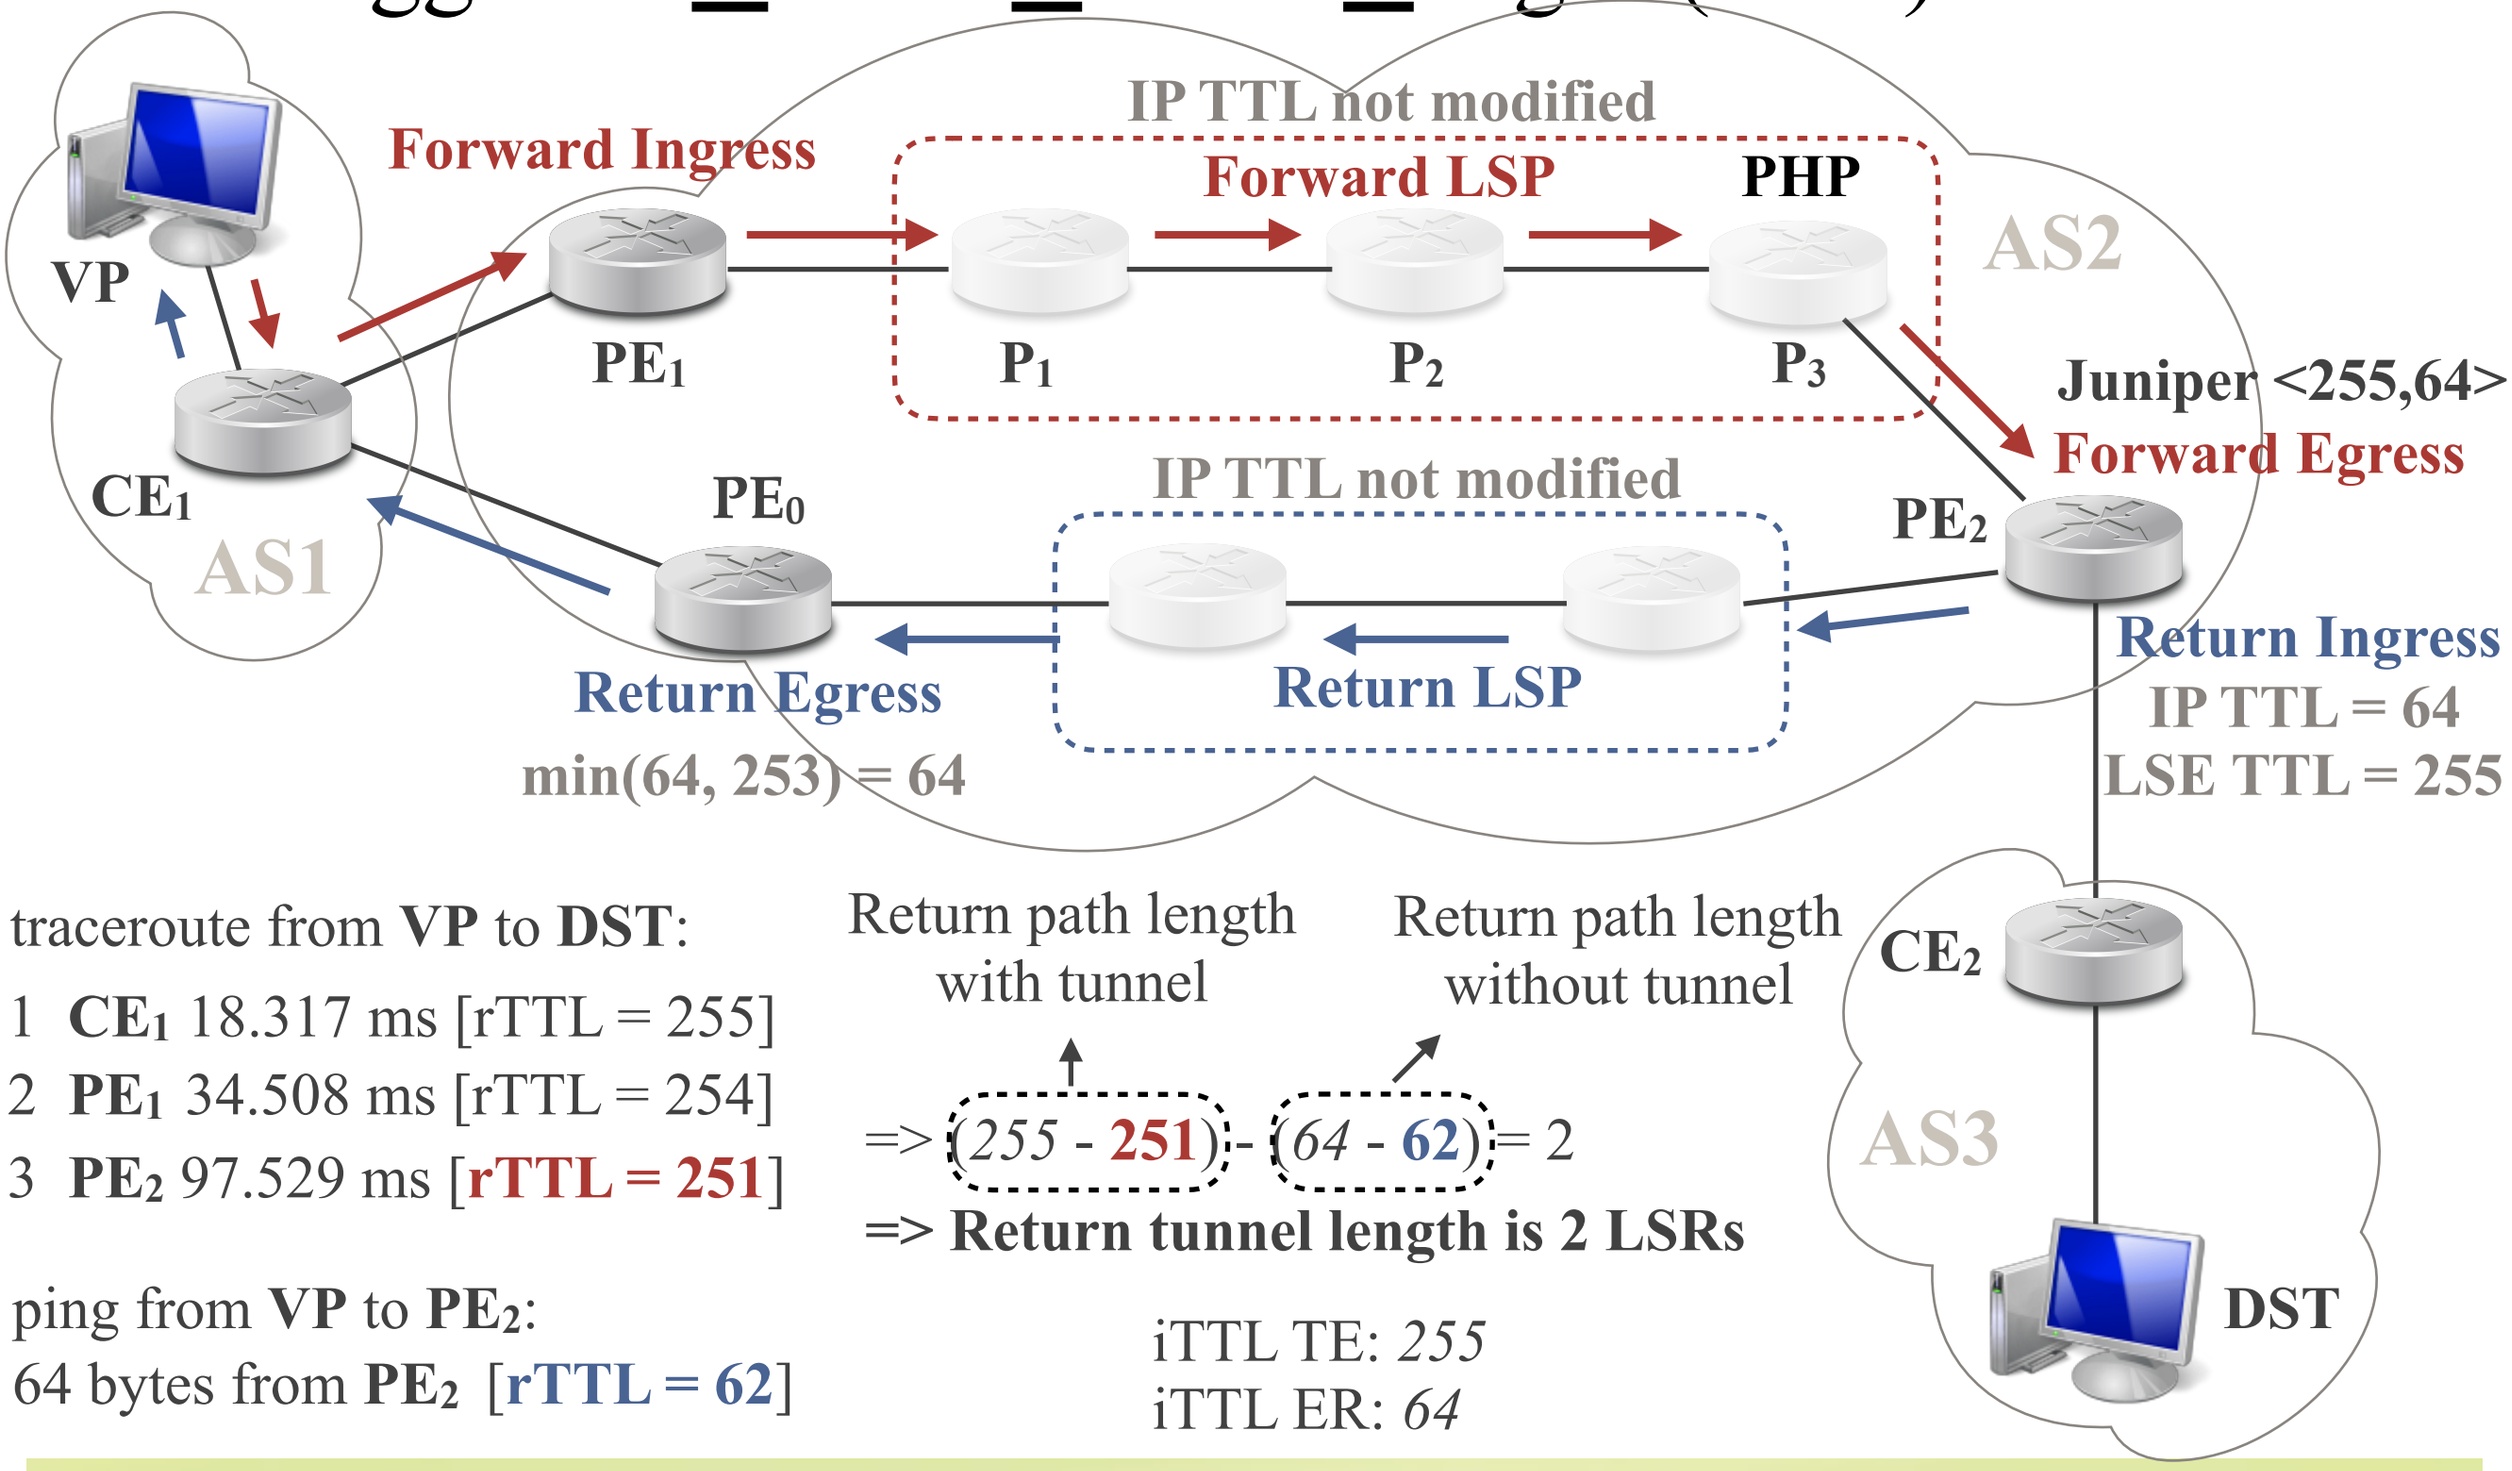

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]

ping from **VP** to **PE$_2$**:
64 bytes from **PE$_2$** [**rTTL = 62**]

Return path length with tunnel

Return path length without tunnel

=> (*255* - **251**) - (*64* - **62**) = 2

=> Return tunnel length is 2 LSRs

iTTL TE: *255*
iTTL ER: *64*

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

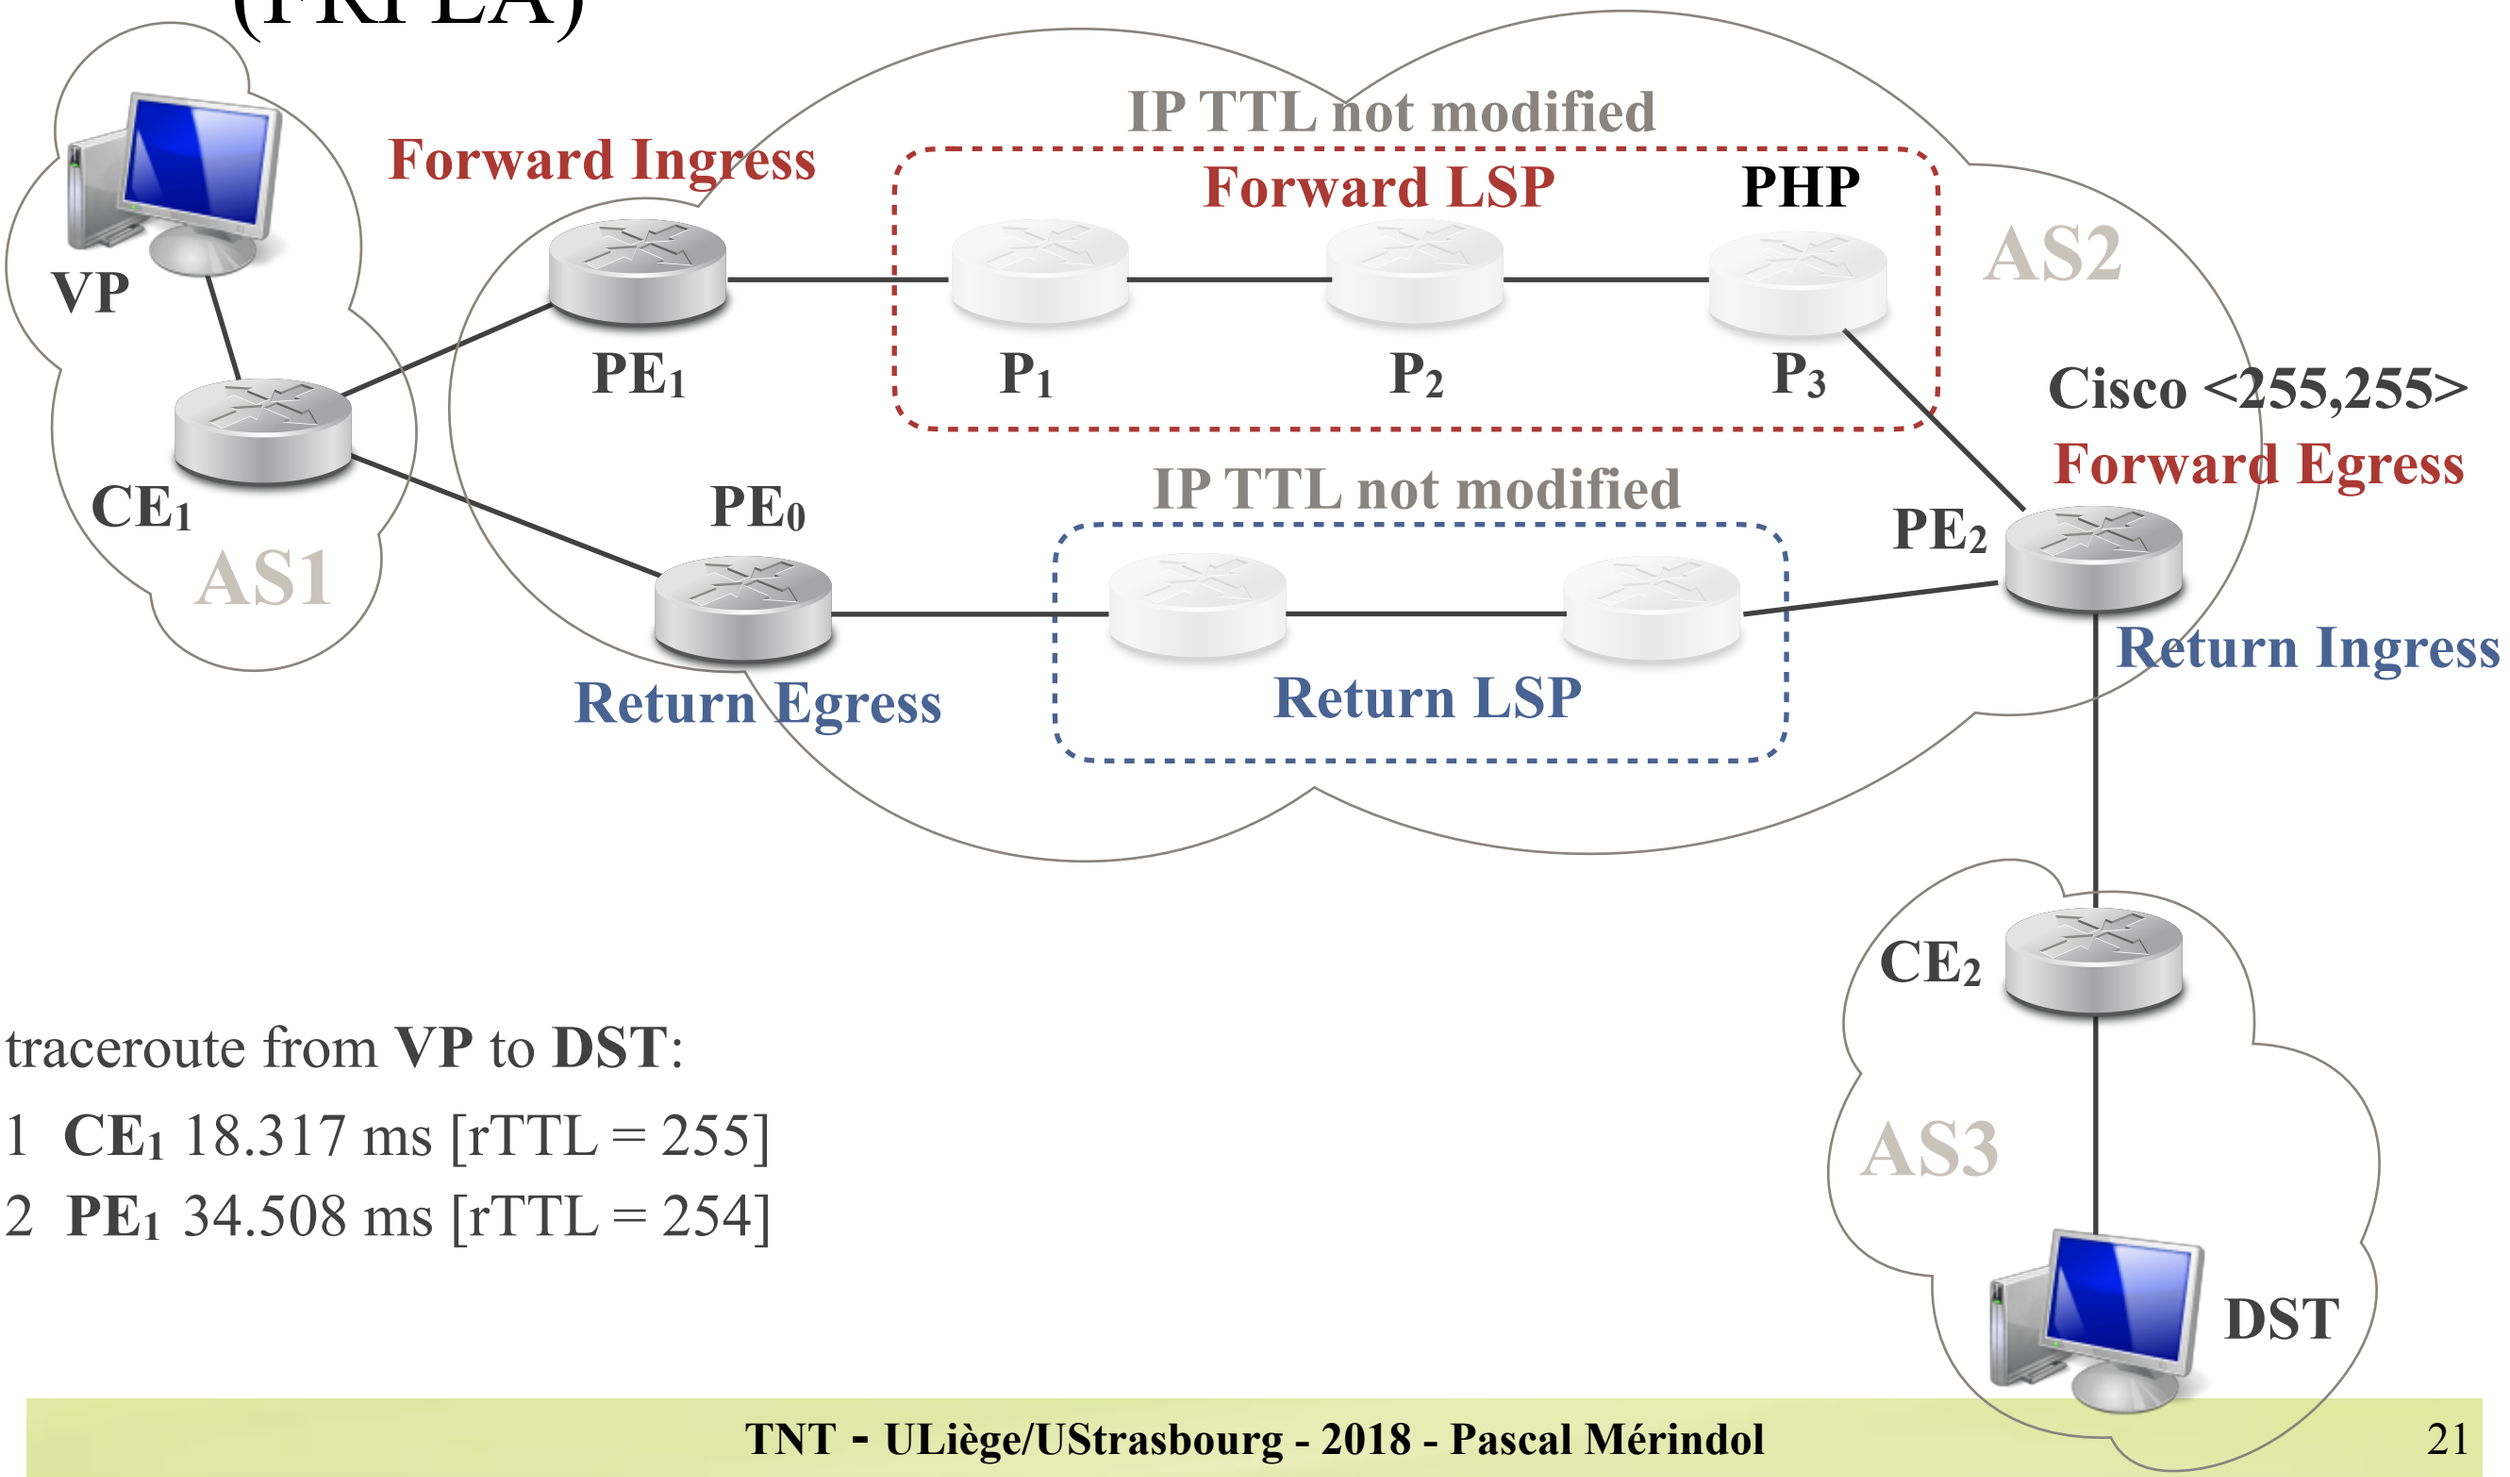

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

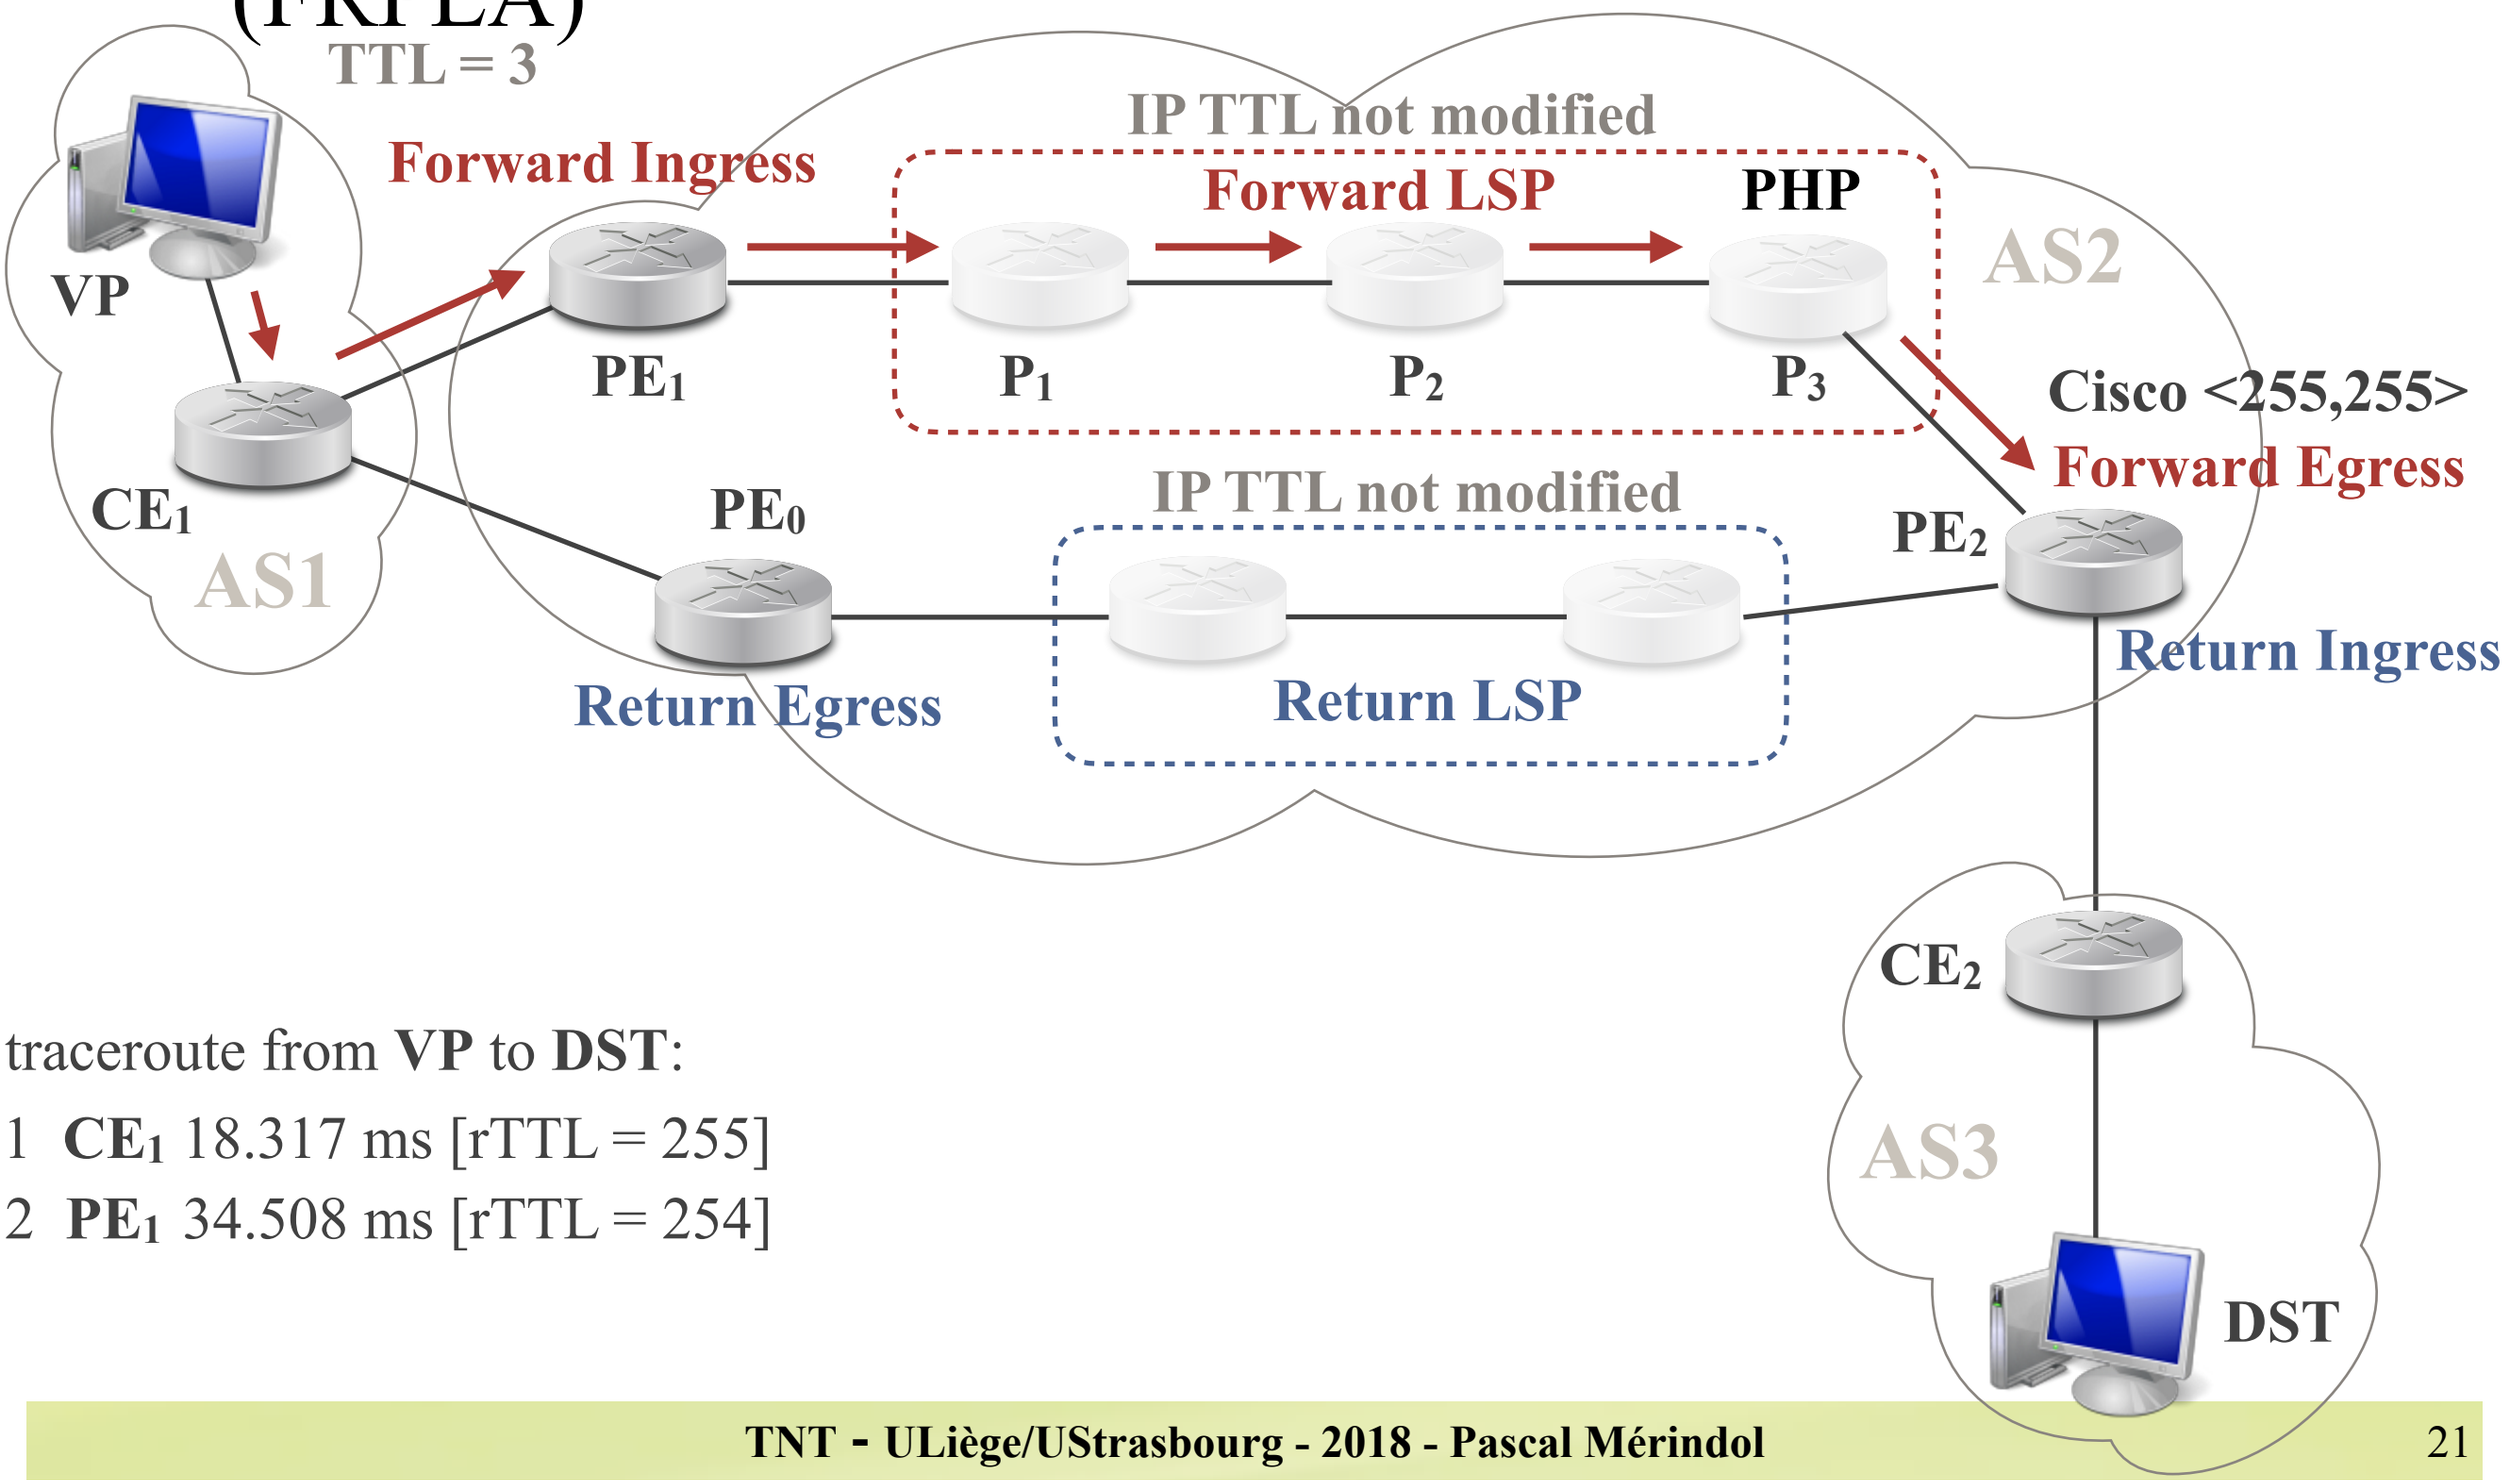

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

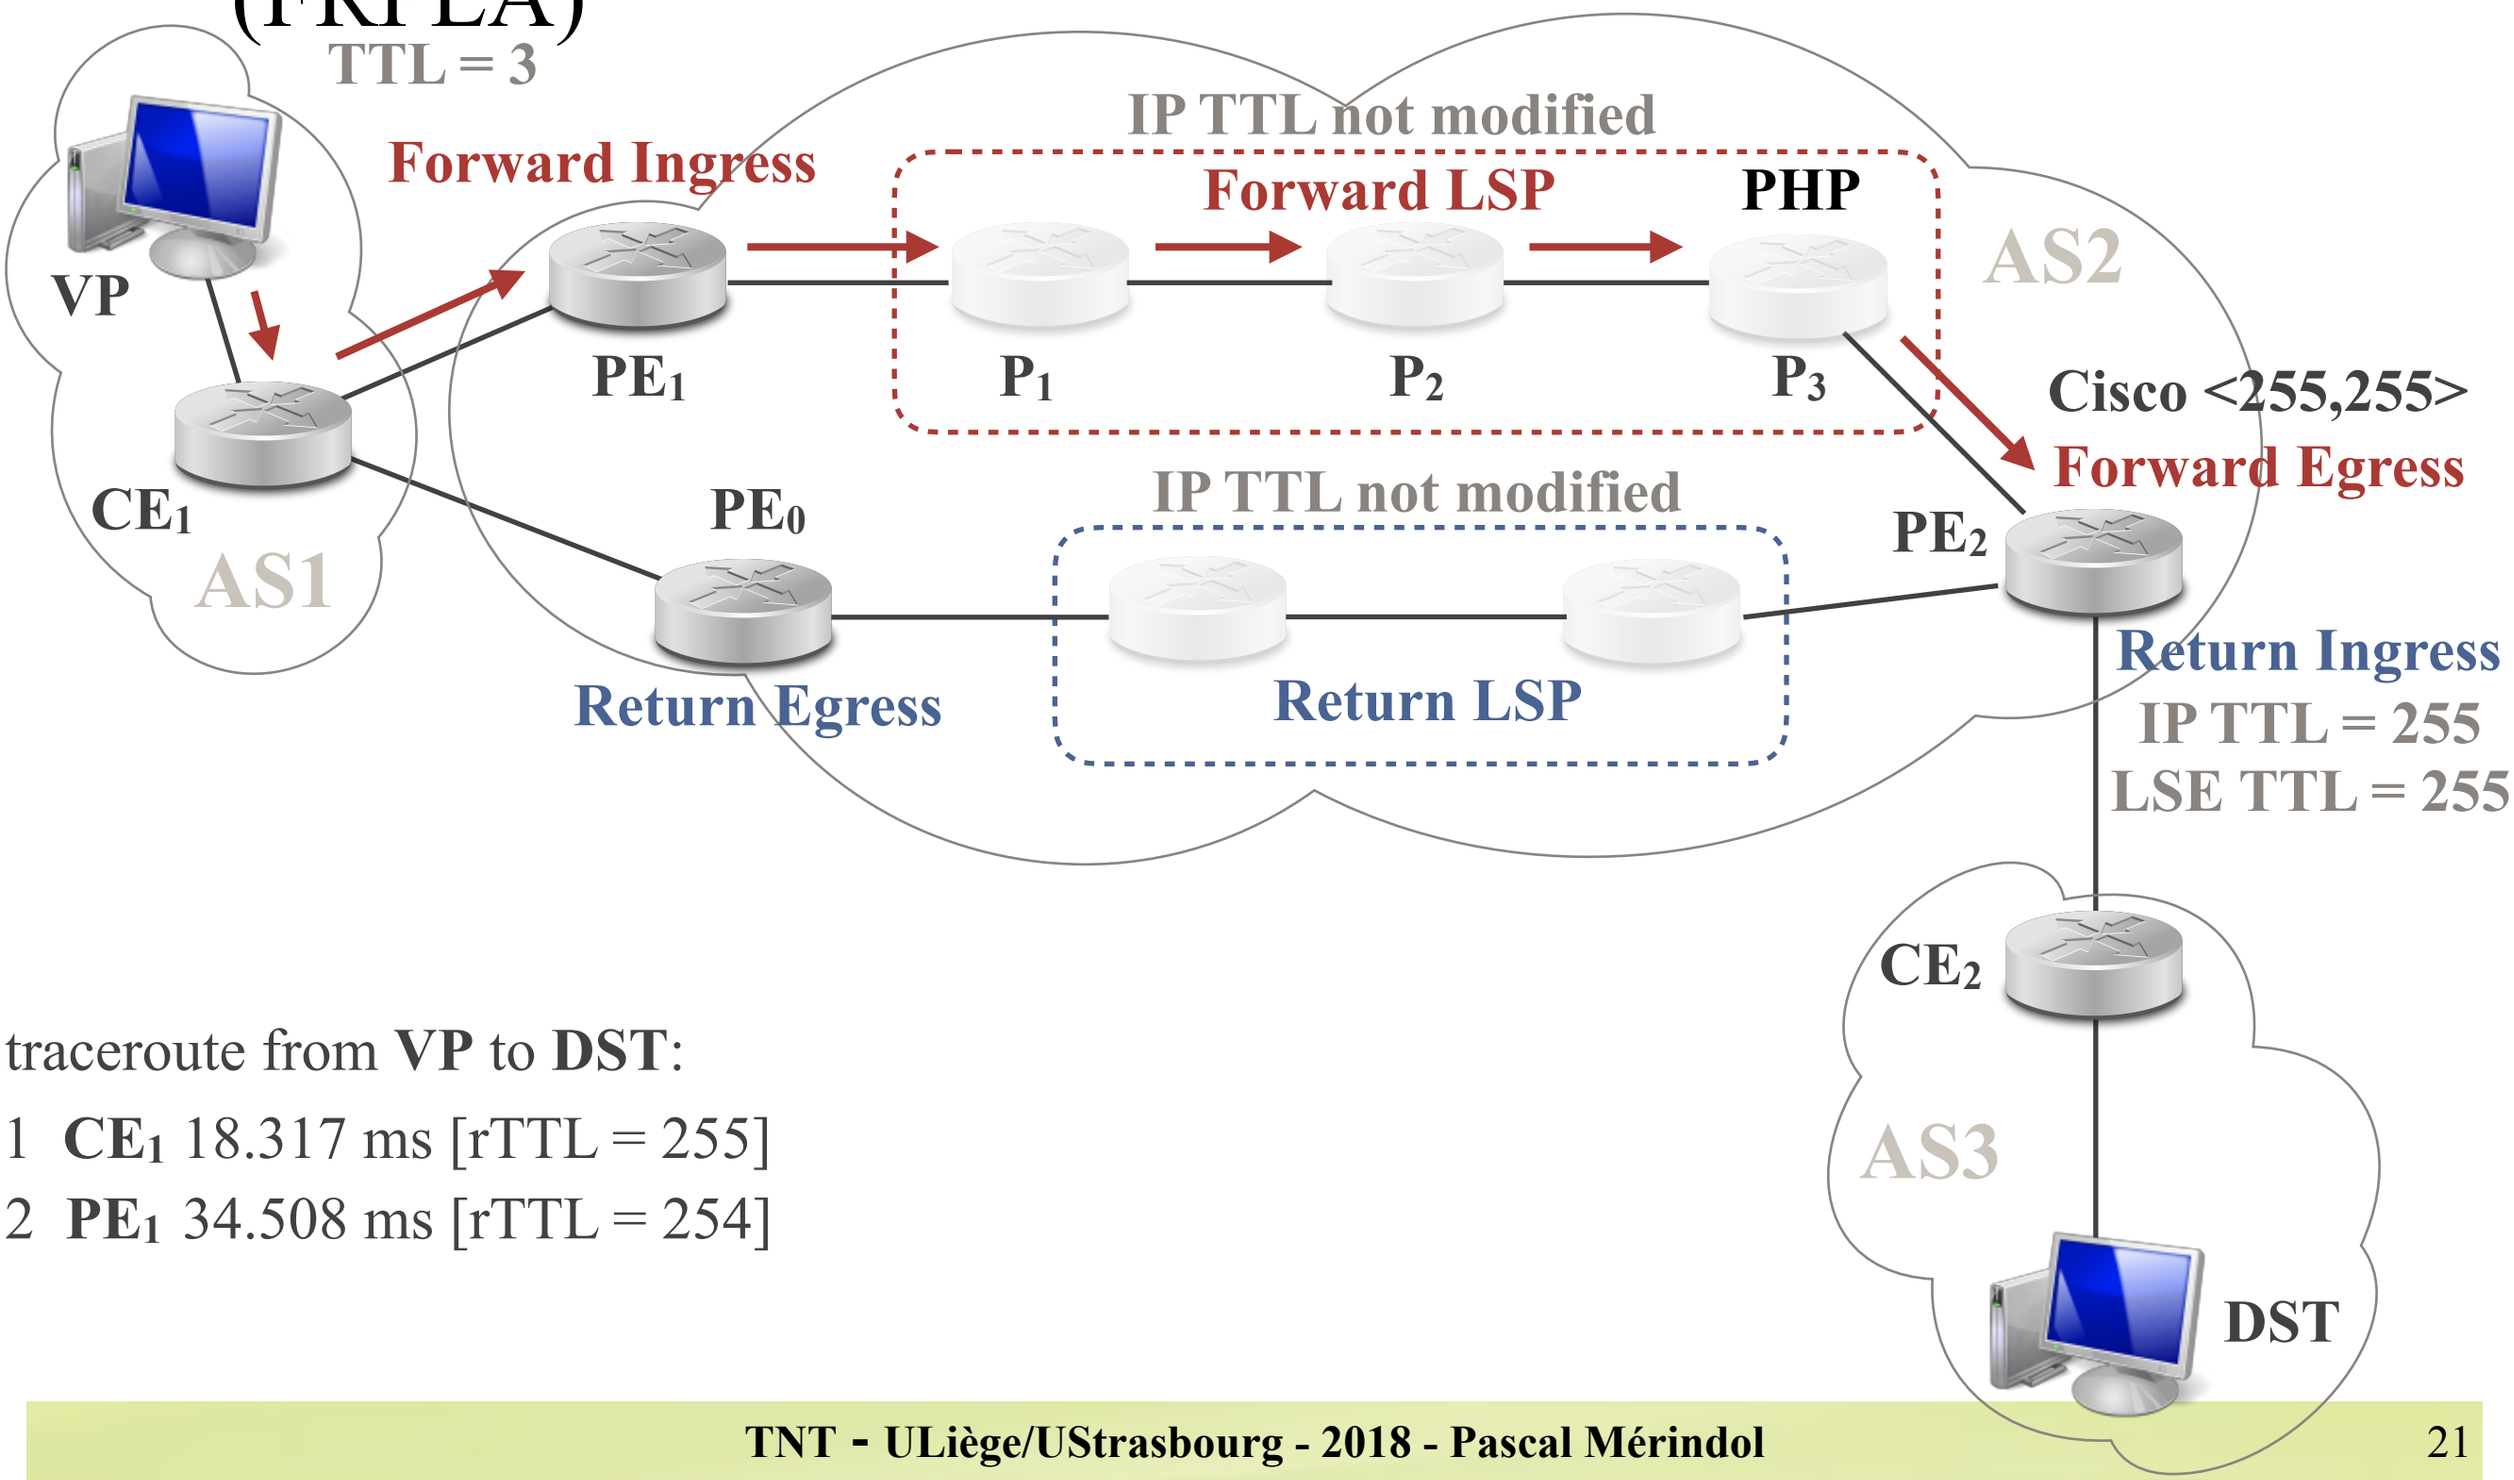

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

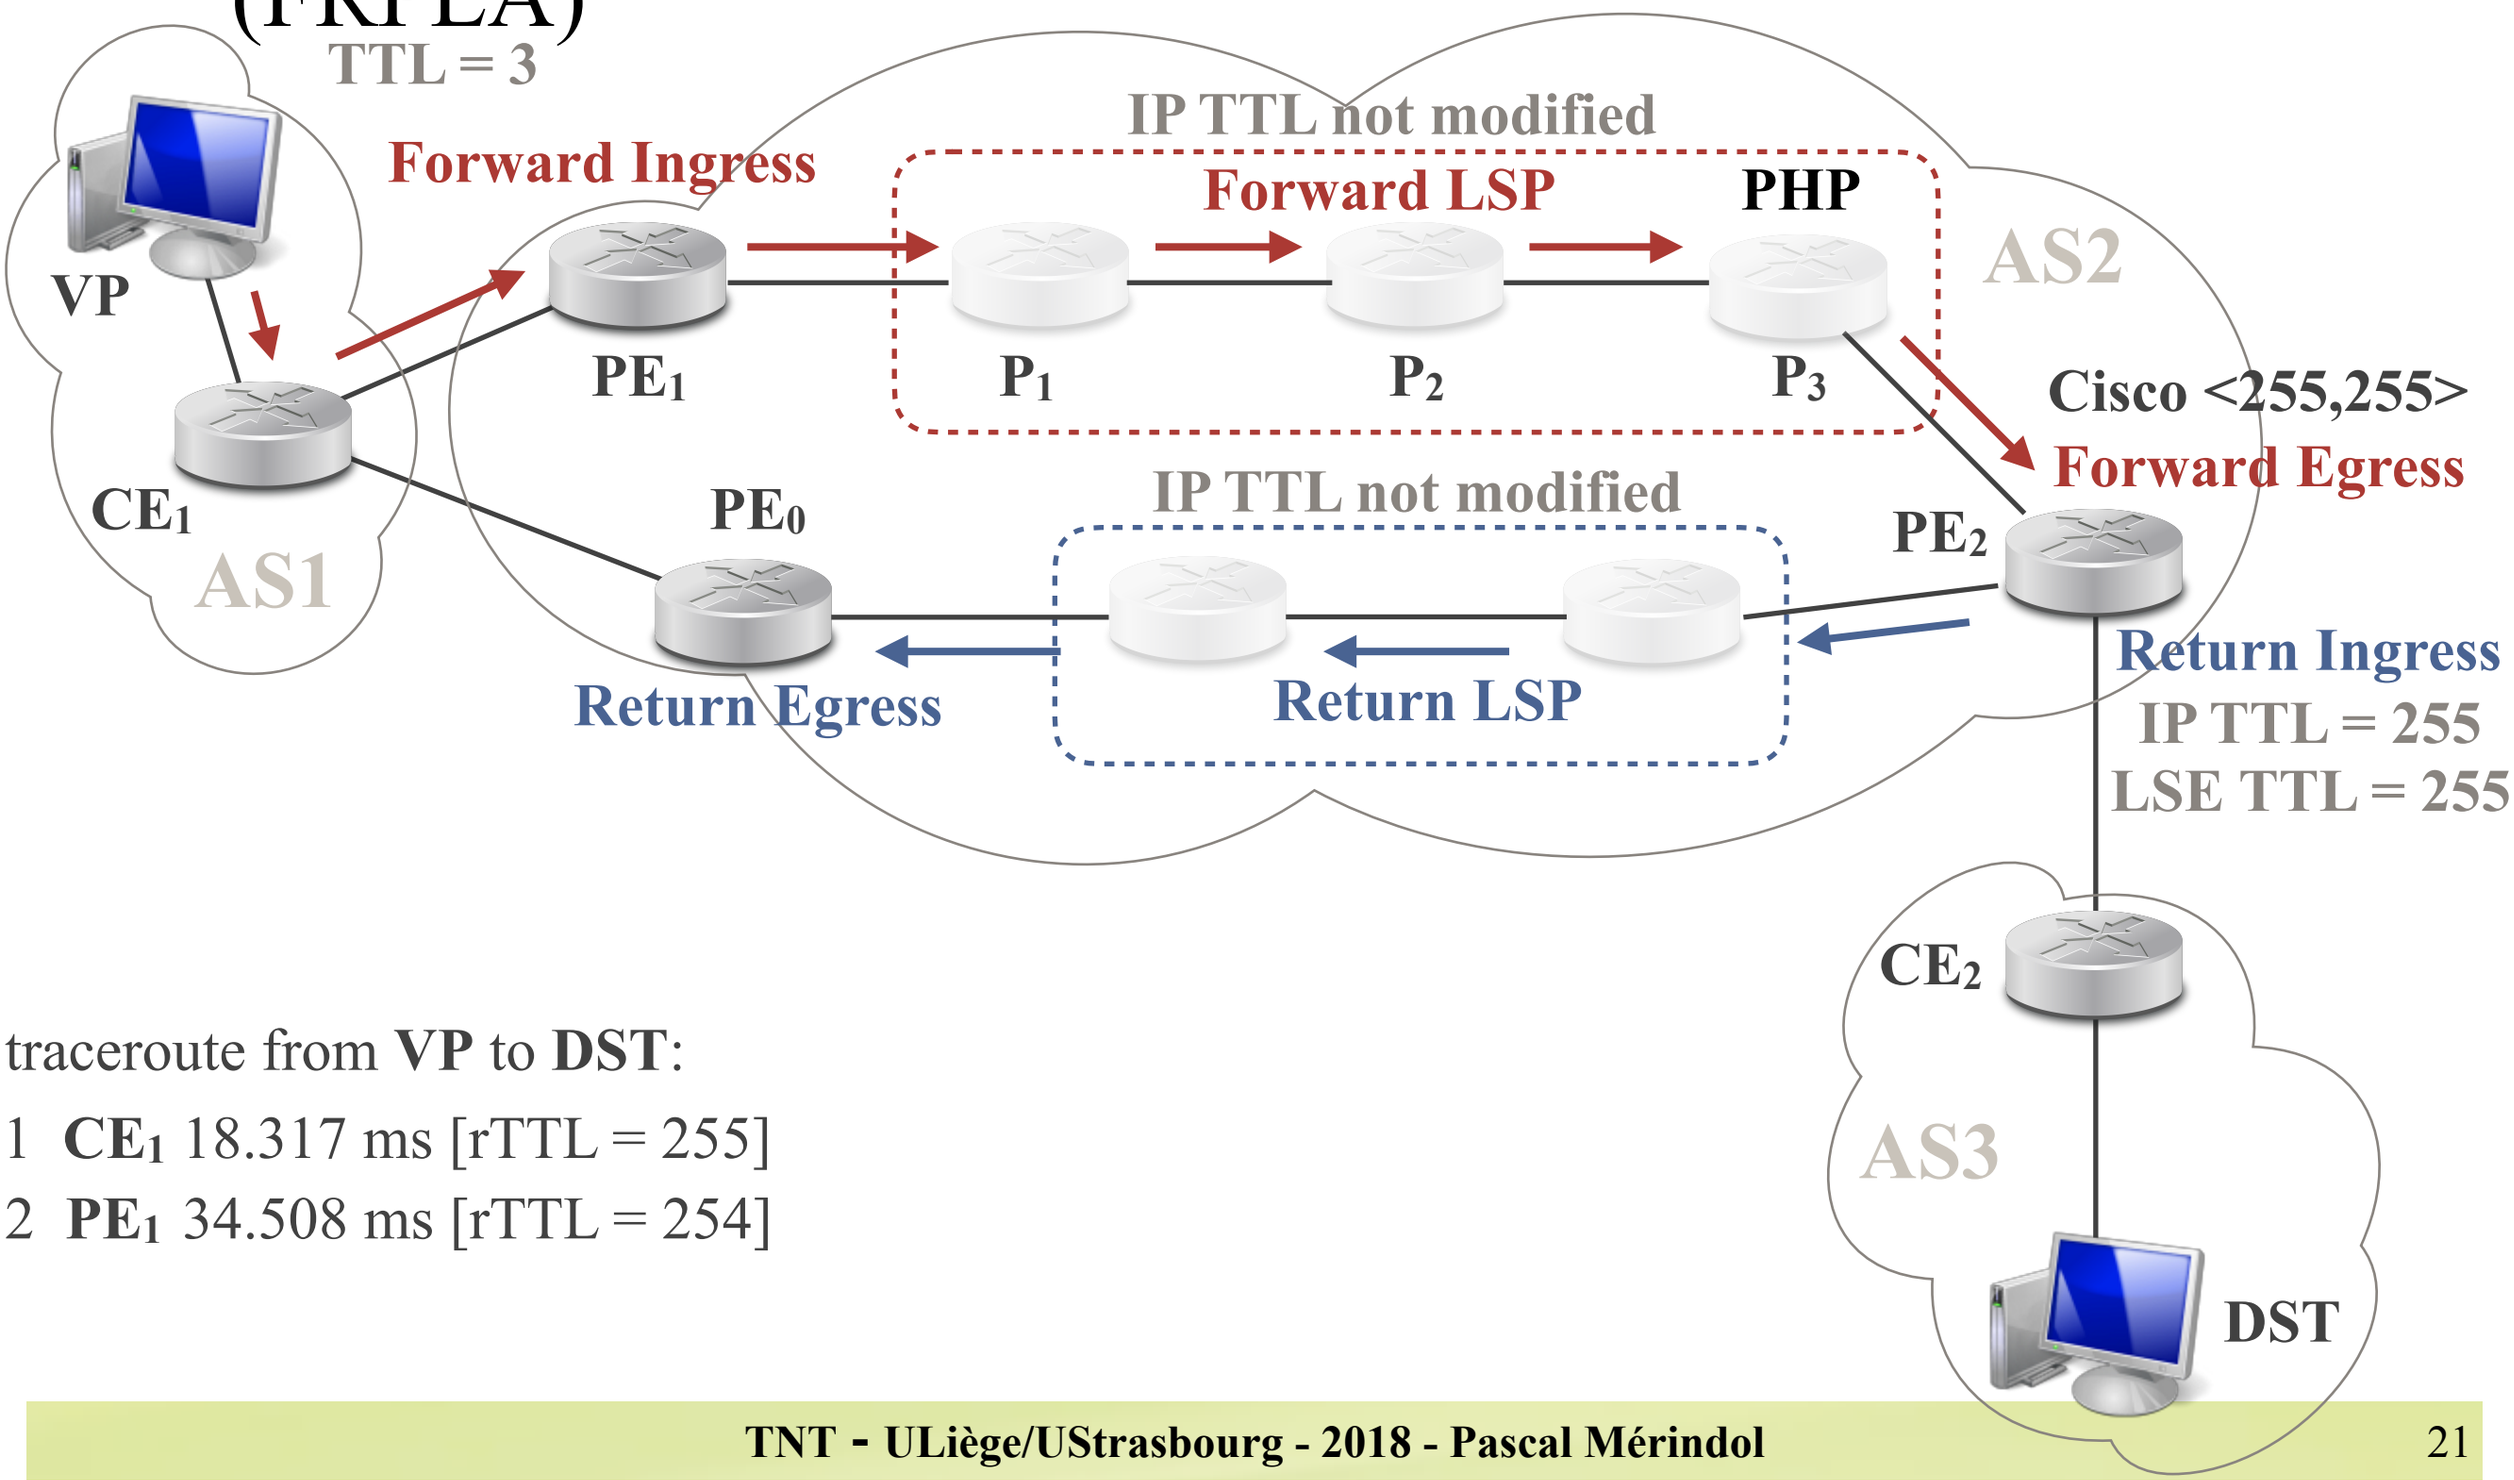

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

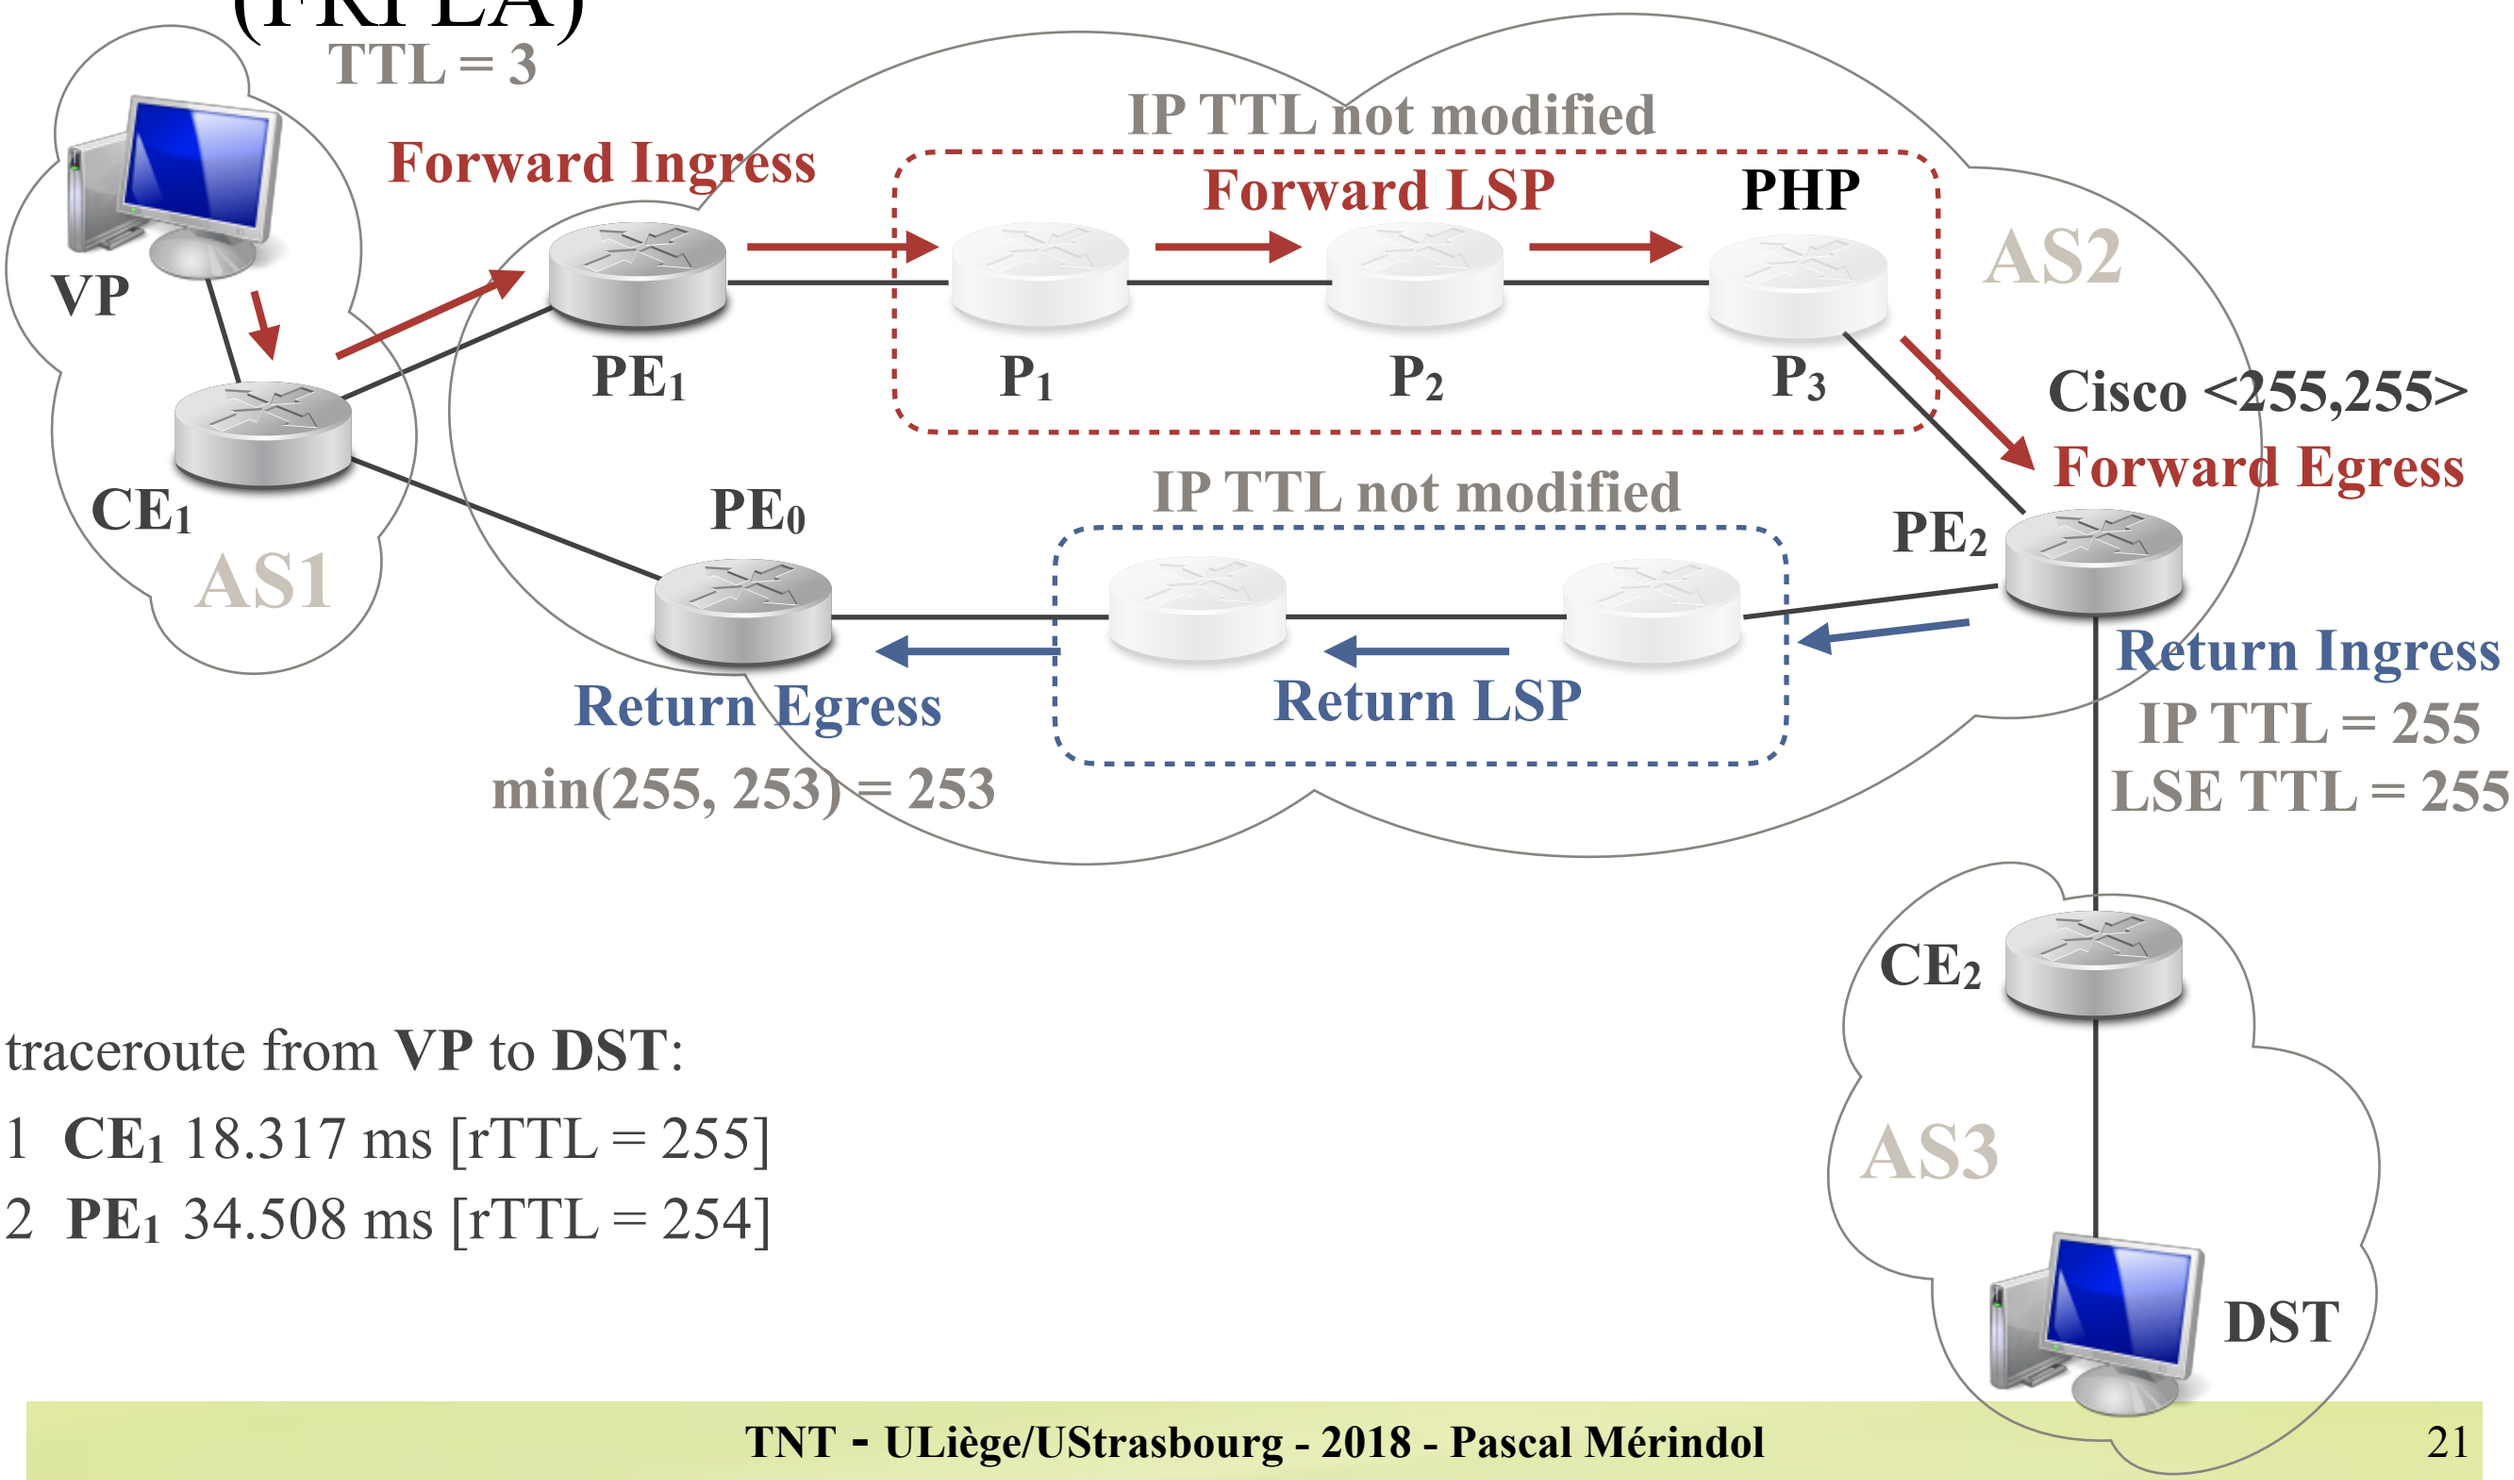

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

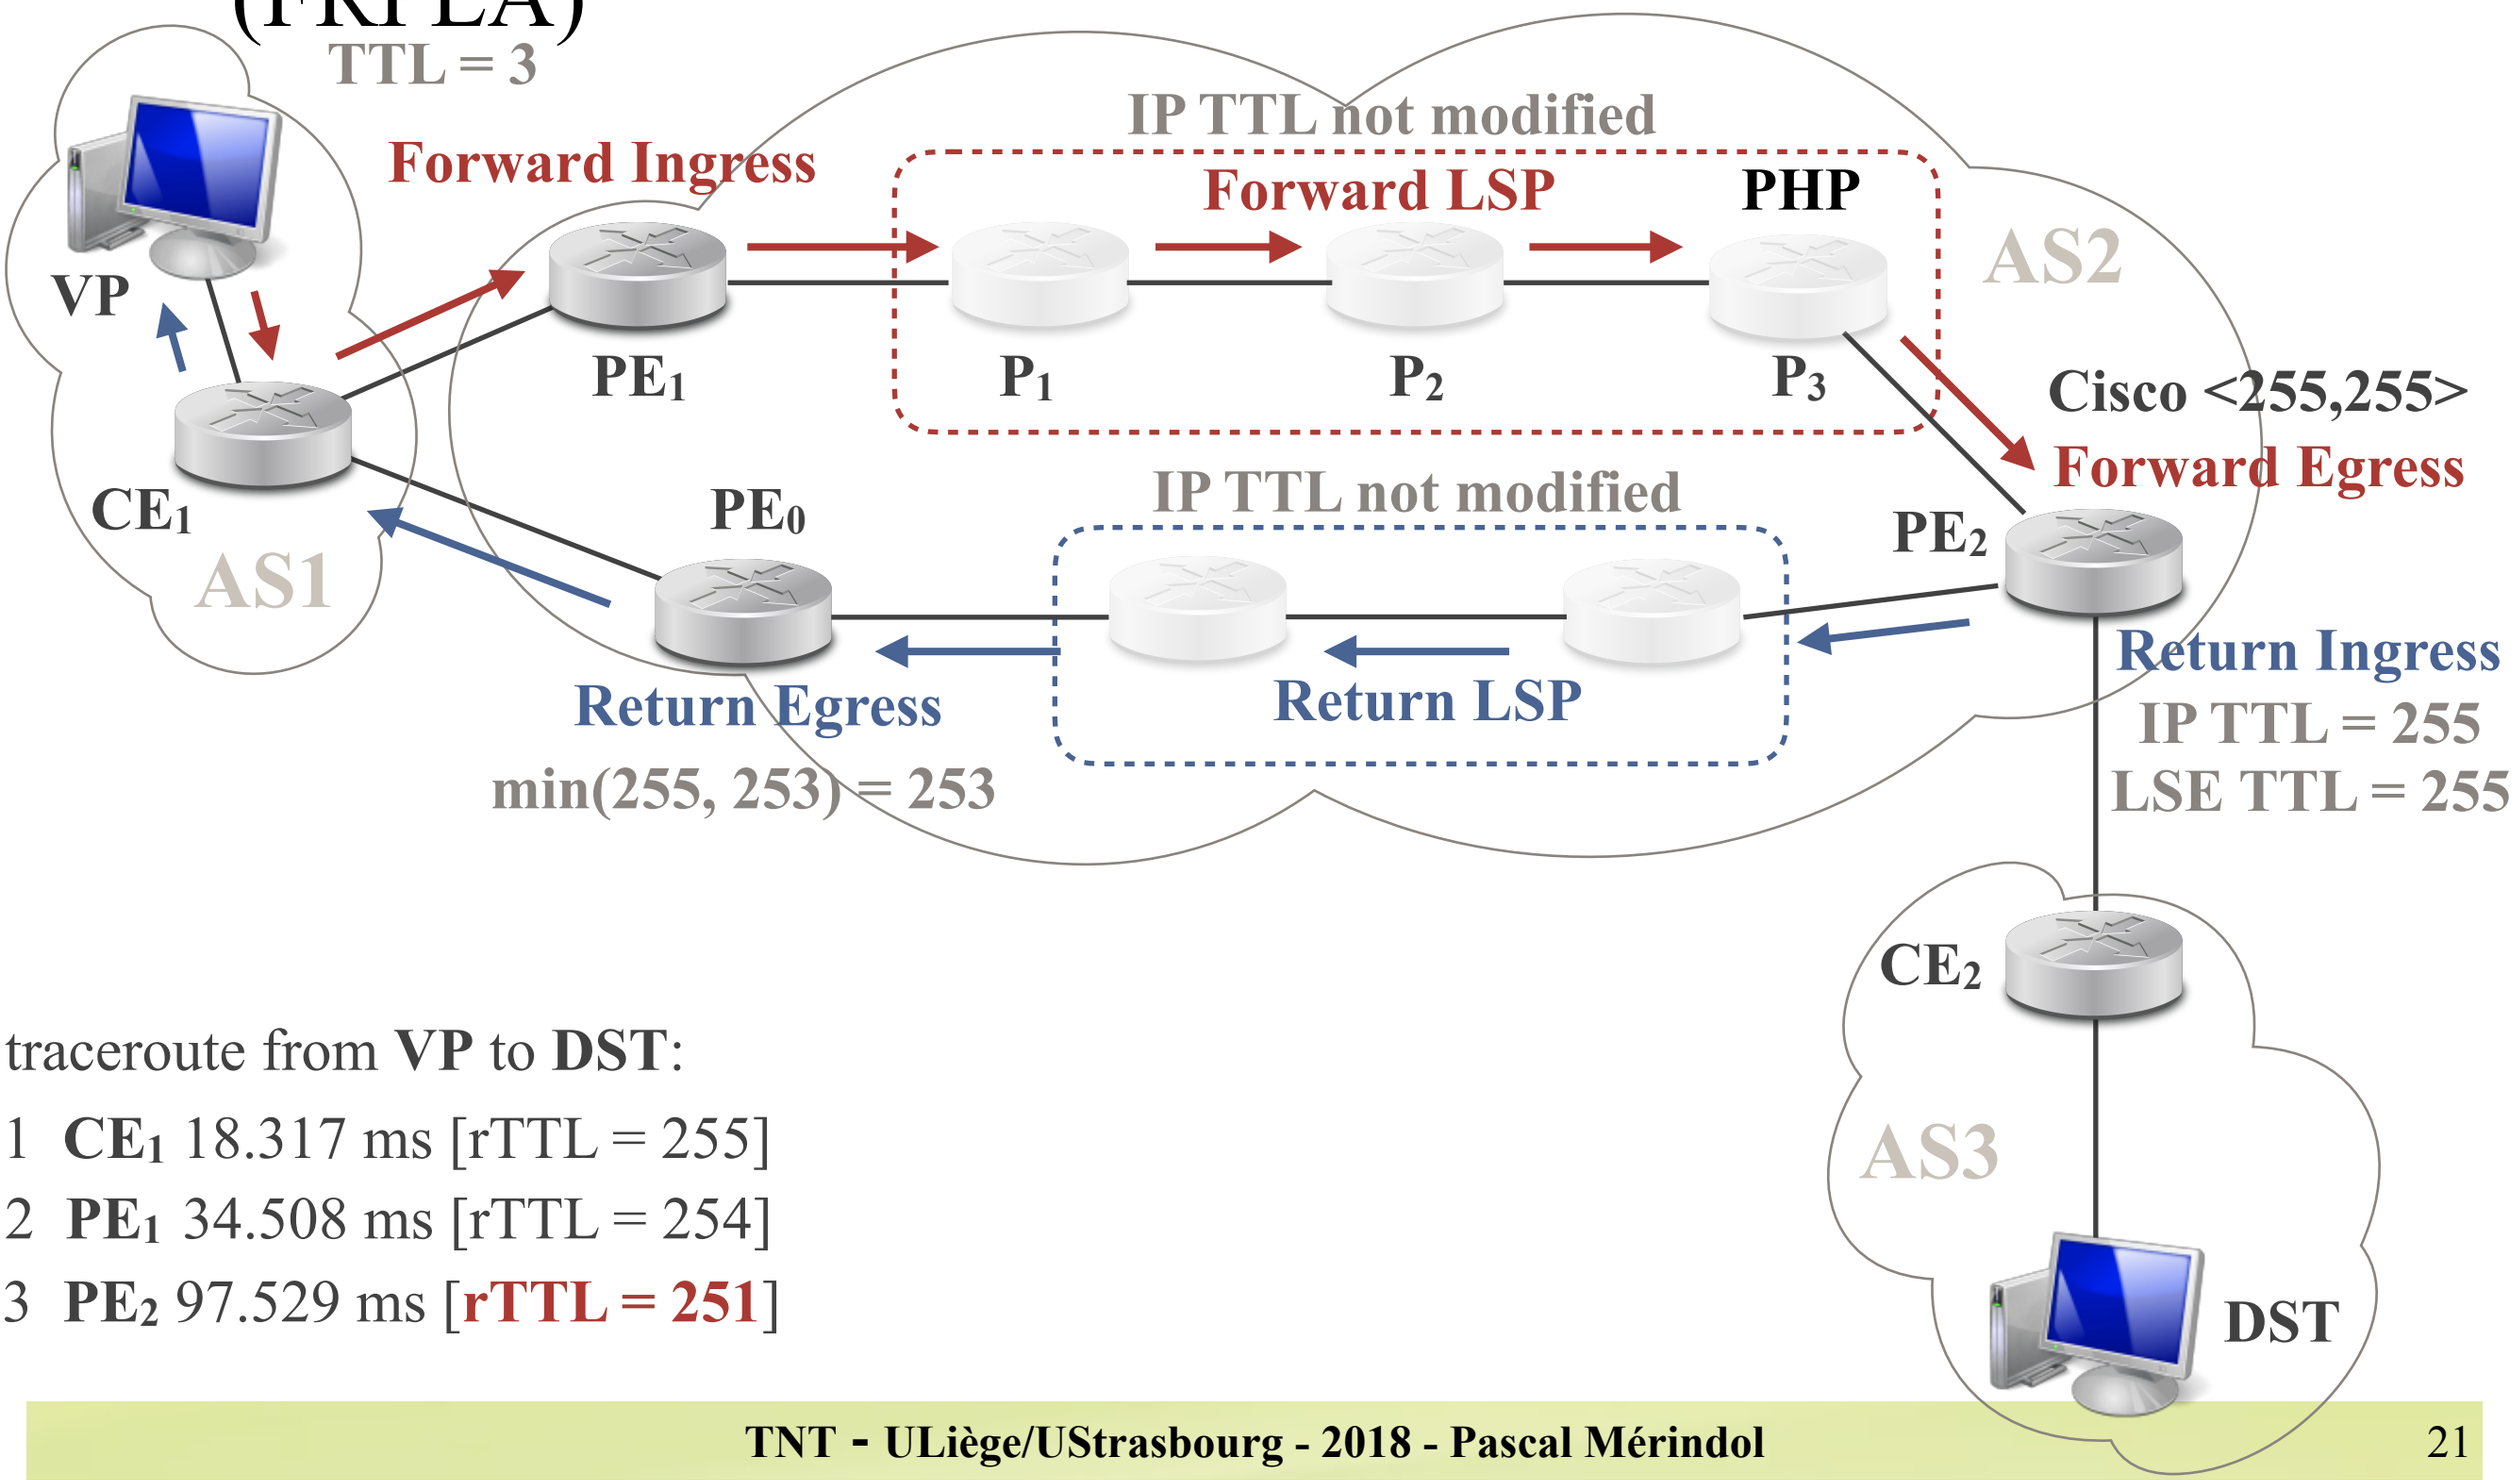

traceroute from **VP** to **DST**:

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

3 **PE2** 97.529 ms [**rTTL = 251**]

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

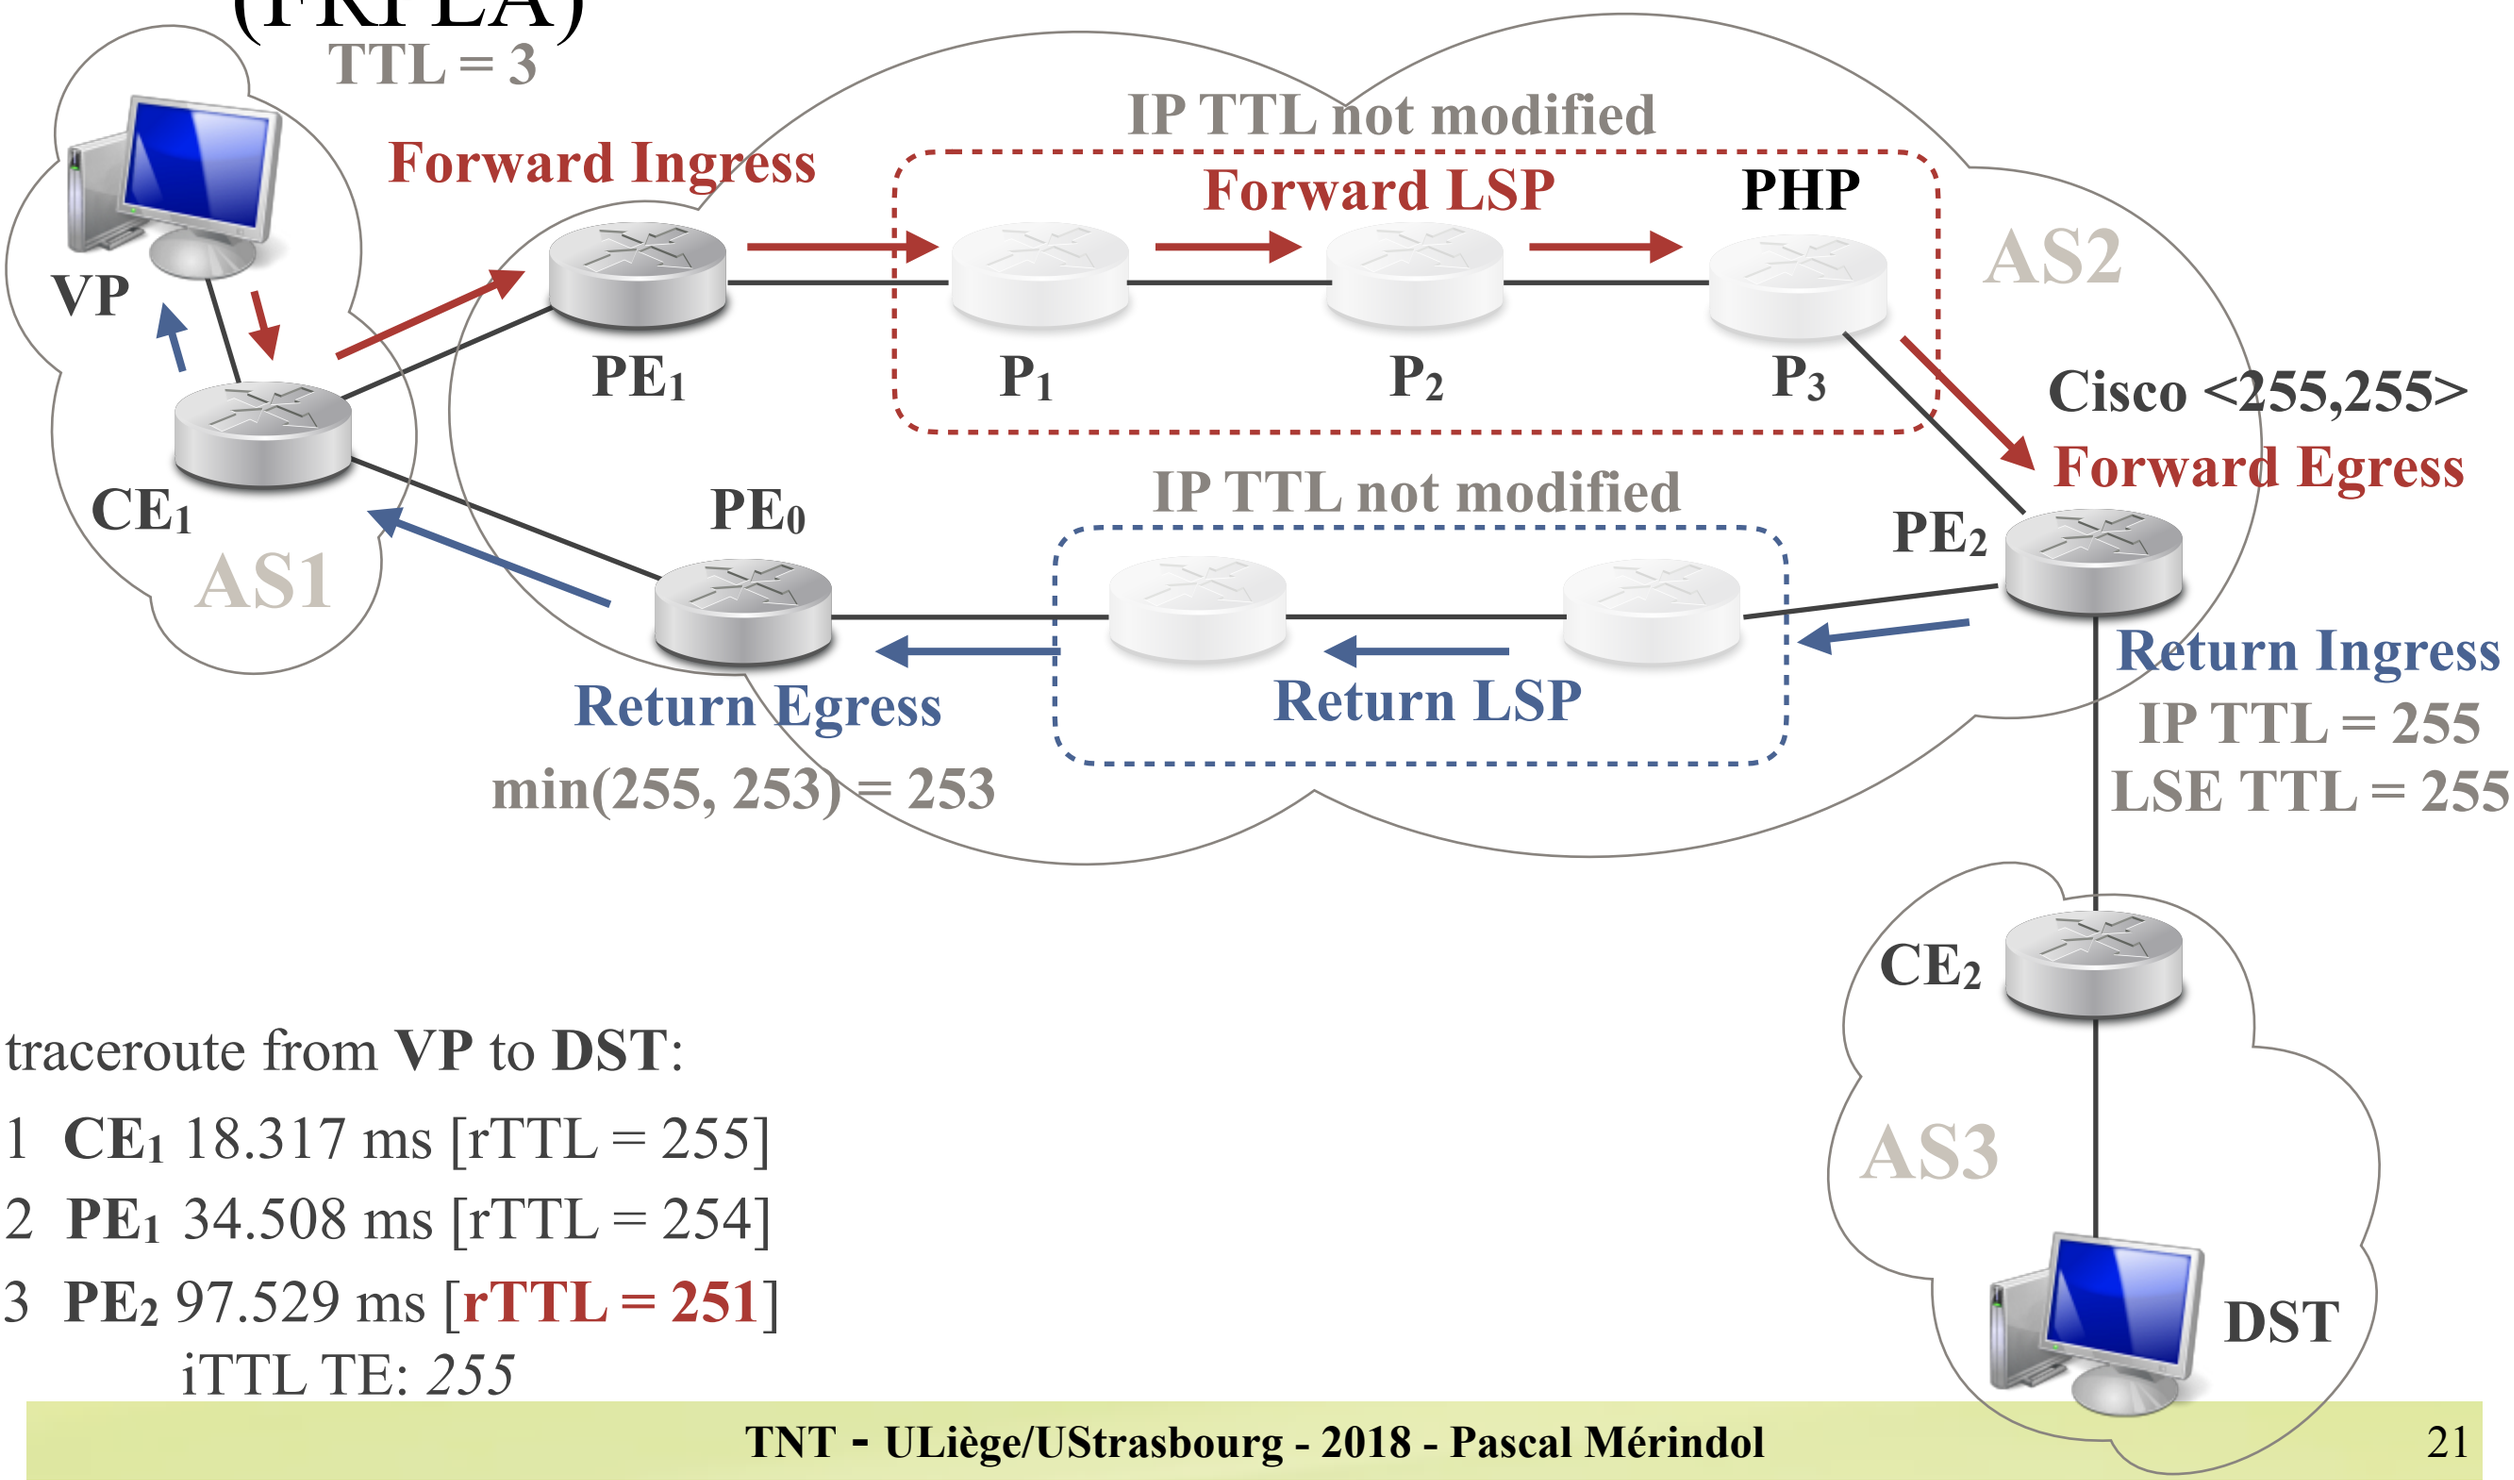

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

3 **PE$_2$** 97.529 ms [**rTTL = 251**]

iTTL TE: *255*

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

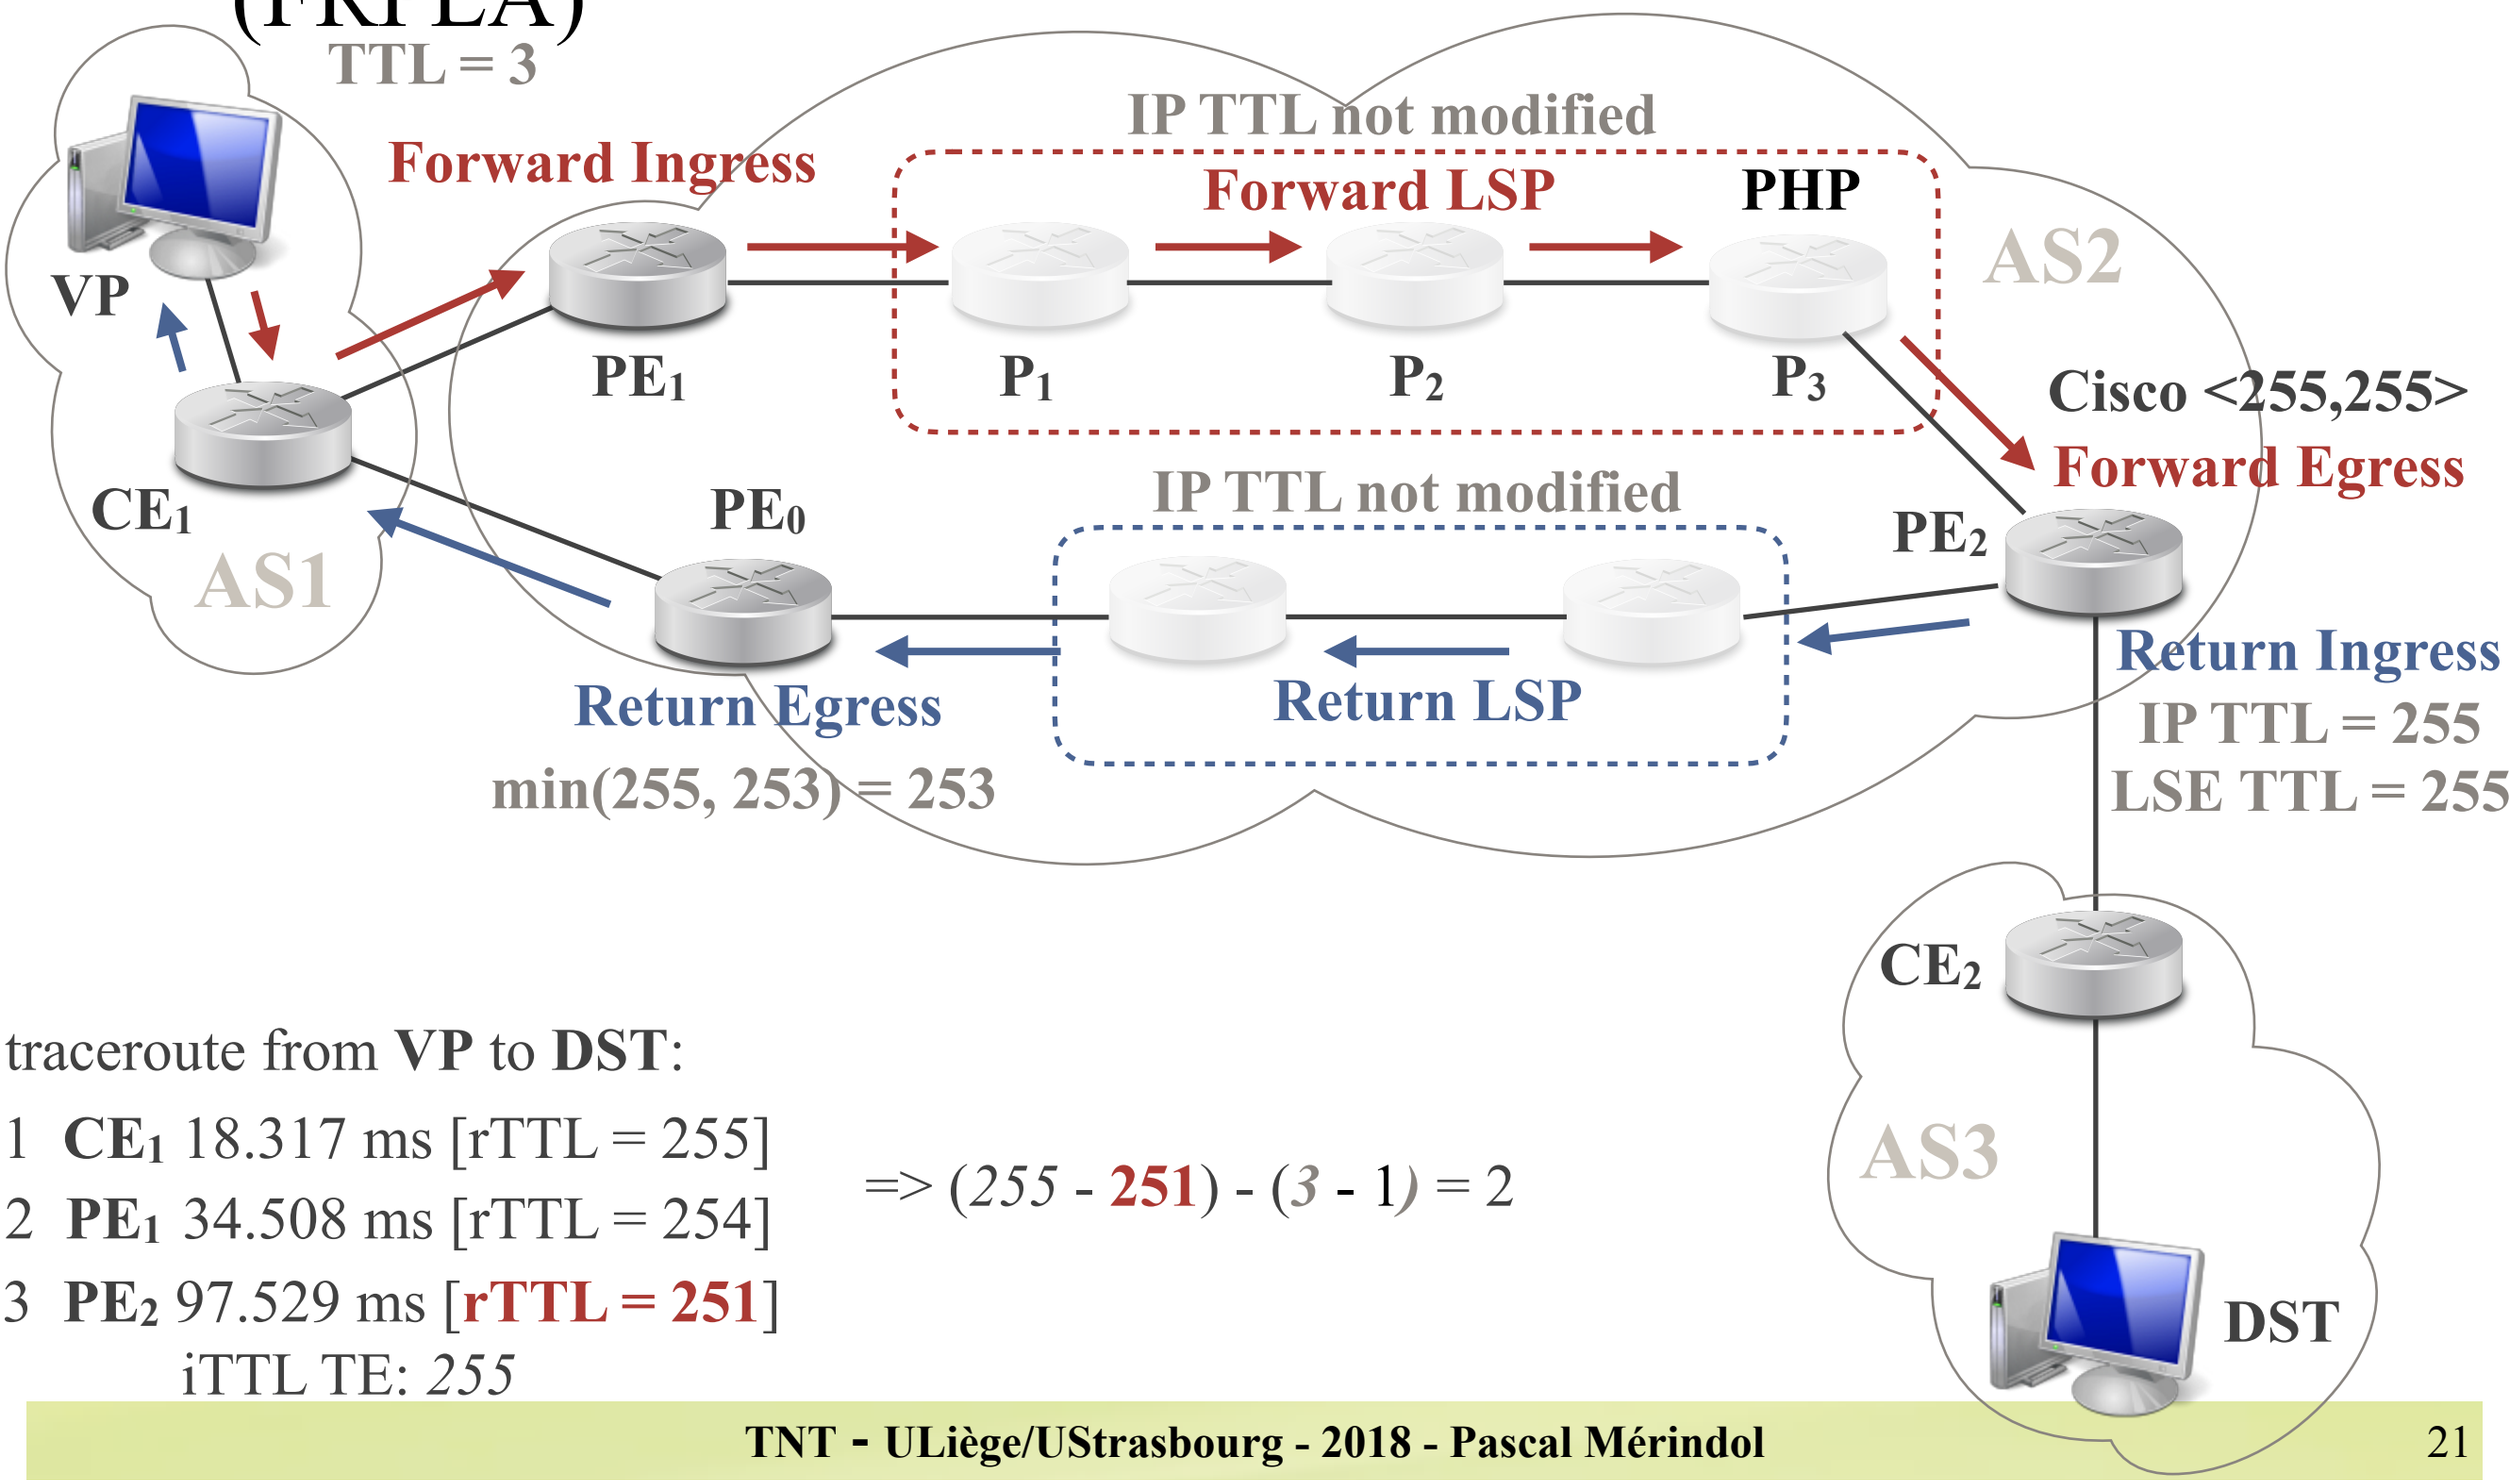

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 251**]
iTTL TE: *255*

=> (*255* - **251**) - (*3* - 1) = 2

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

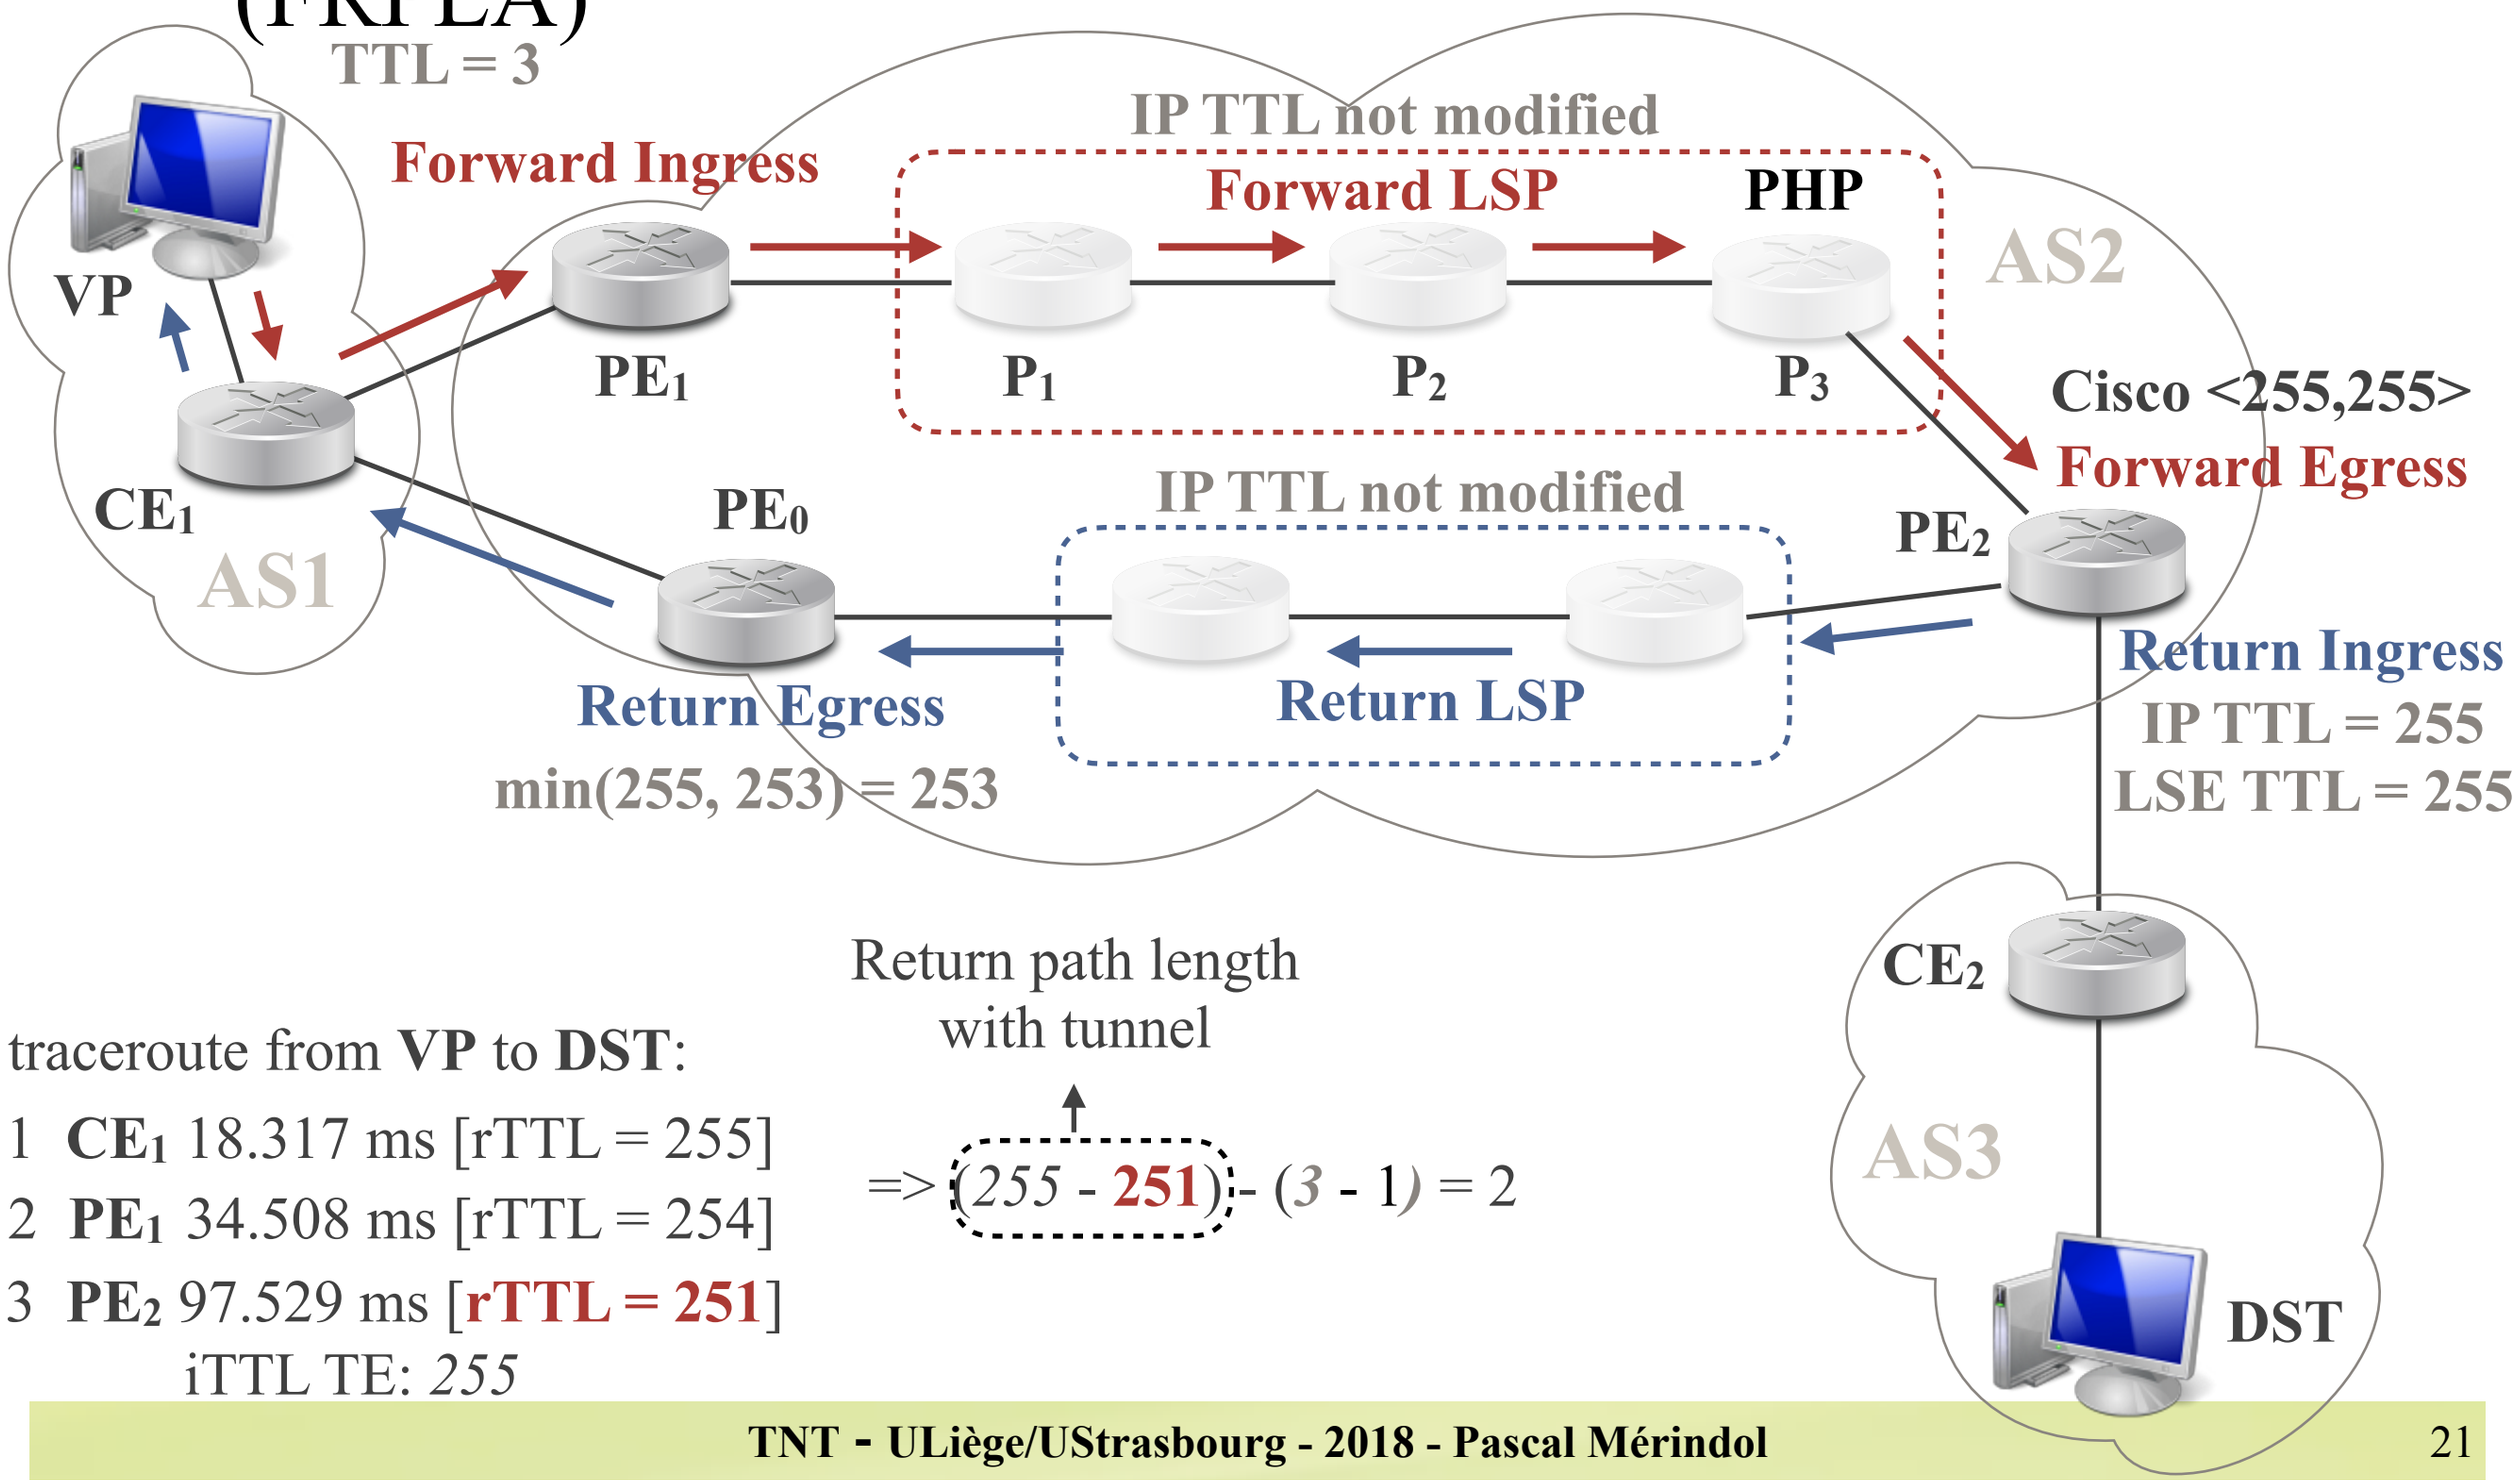

Return path length with tunnel

=> (*255* - **251**) - (*3* - 1*)* = 2

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

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

3 **PE2** 97.529 ms [**rTTL = 251**]
iTTL TE: *255*

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

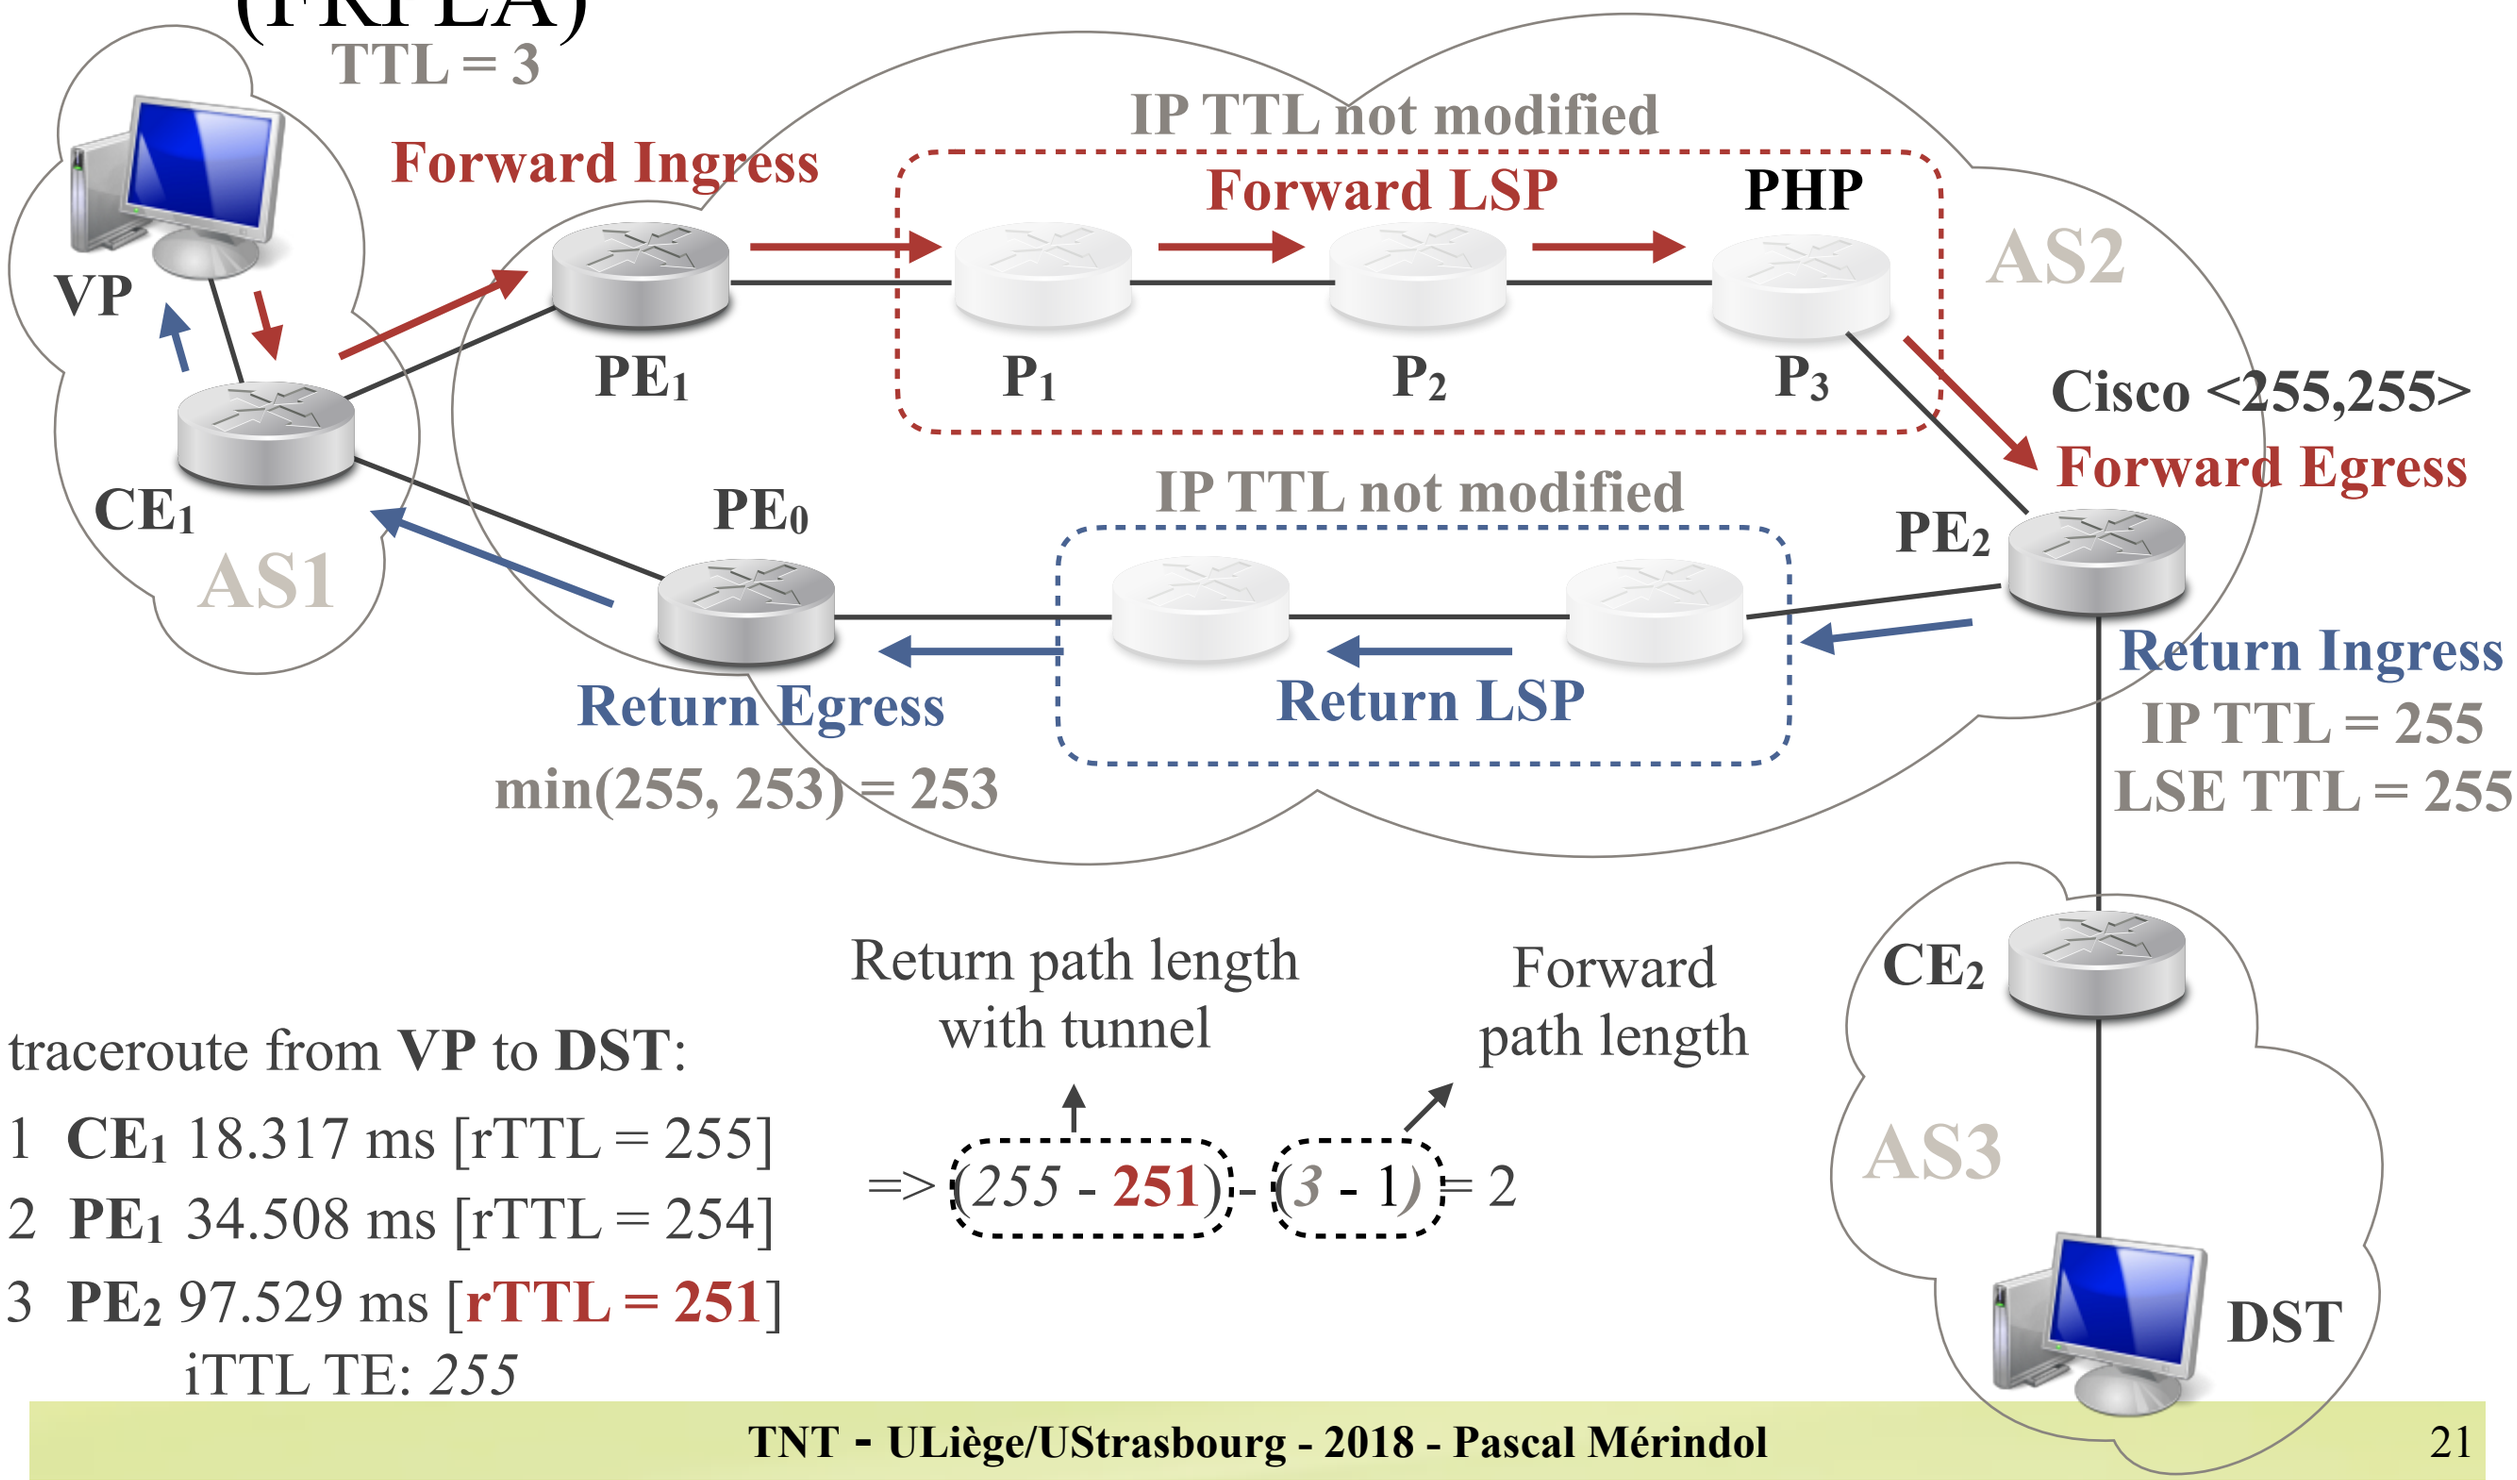

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

3 **PE$_2$** 97.529 ms [**rTTL = 251**]

iTTL TE: *255*

Return path length with tunnel

Forward path length

=> (*255* - **251**) - (*3* - 1) = 2

Invisible Tunnels (4)

- Trigger 3: **F**orward and **R**eturn **P**ath **L**ength **A**symmetry (FRPLA)

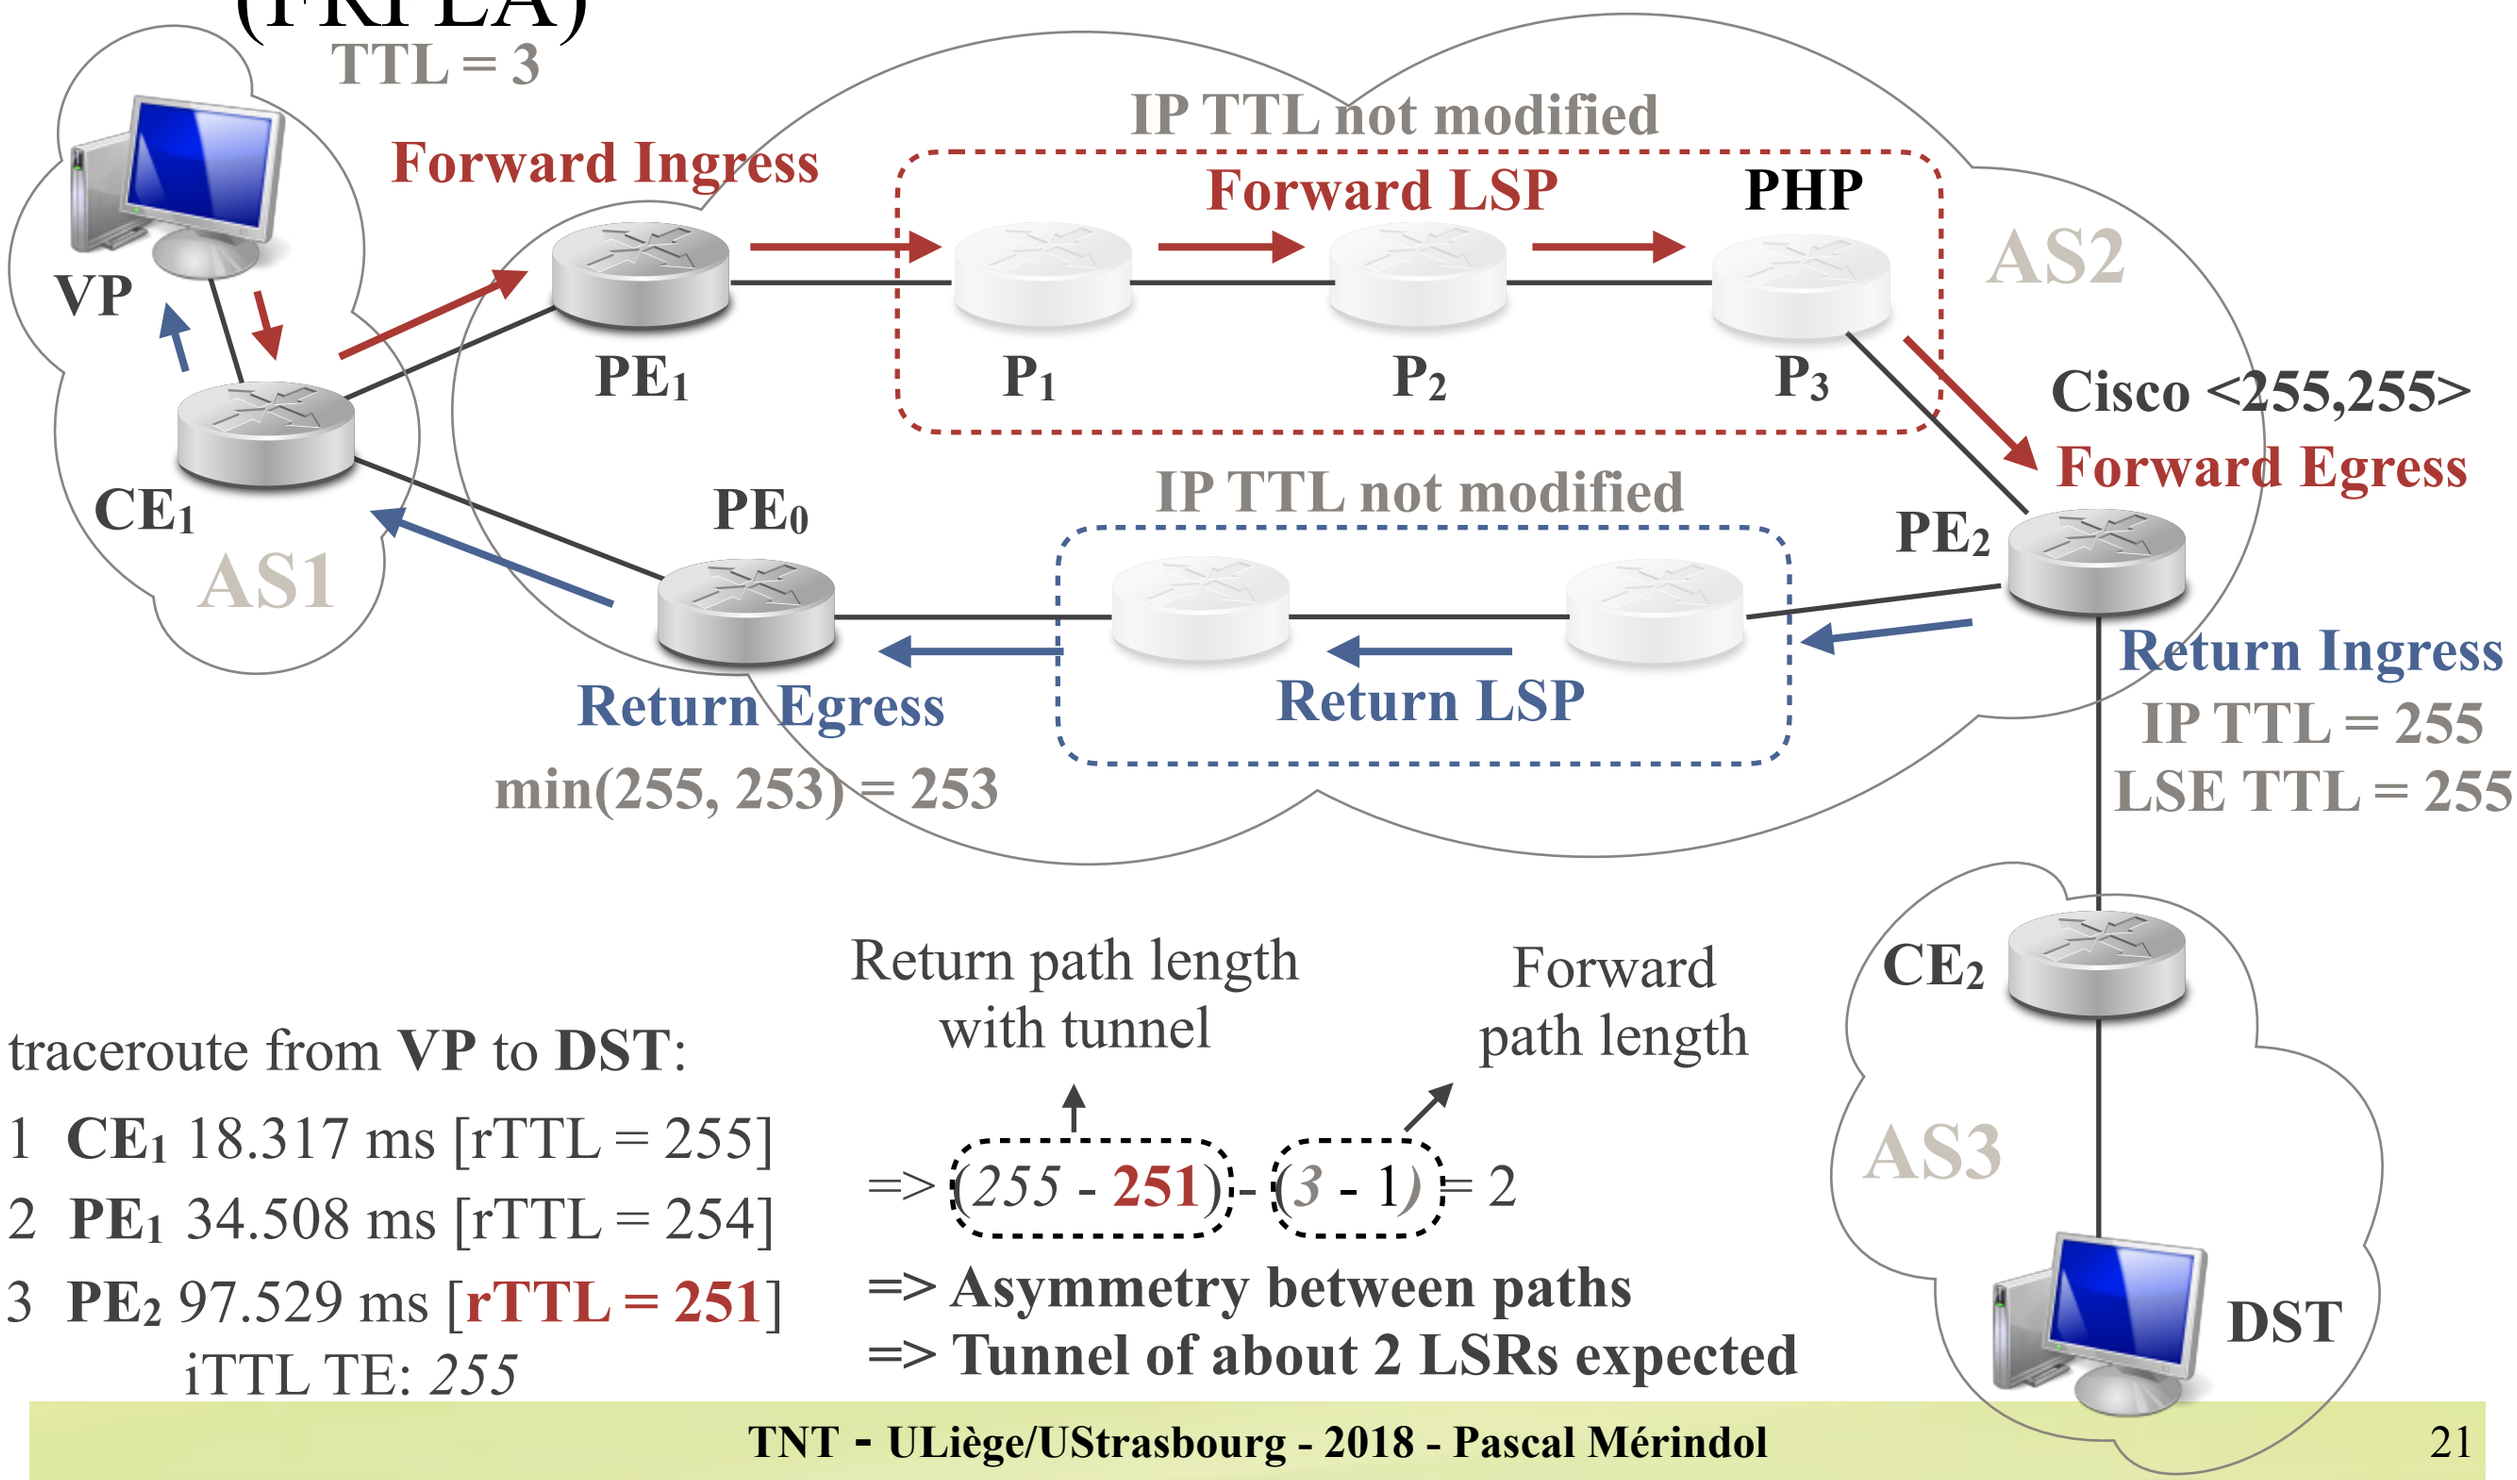

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]

2 **PE$_1$** 34.508 ms [rTTL = 254]

3 **PE$_2$** 97.529 ms [**rTTL = 251**]

iTTL TE: *255*

Return path length with tunnel

Forward path length

=> (*255* - **251**) - (*3* - 1) = 2

=> Asymmetry between paths

=> Tunnel of about 2 LSRs expected

Invisible Tunnels (5)

- How can we reveal the content of invisible tunnels?
  - **<u>Direct Path Revelation</u>** (DPR)
    - ✓ for networks using MPLS for transit, not for internal traffic
    - ✓ mostly Juniper devices
  - **<u>Backward Recursive Path Revelation</u>** (BRPR)
    - ✓ for networks announcing all prefixes (external and internal) with LDP
    - ✓ mostly CISCO devices

Invisible Tunnels (6)

- Direct Path Revelation

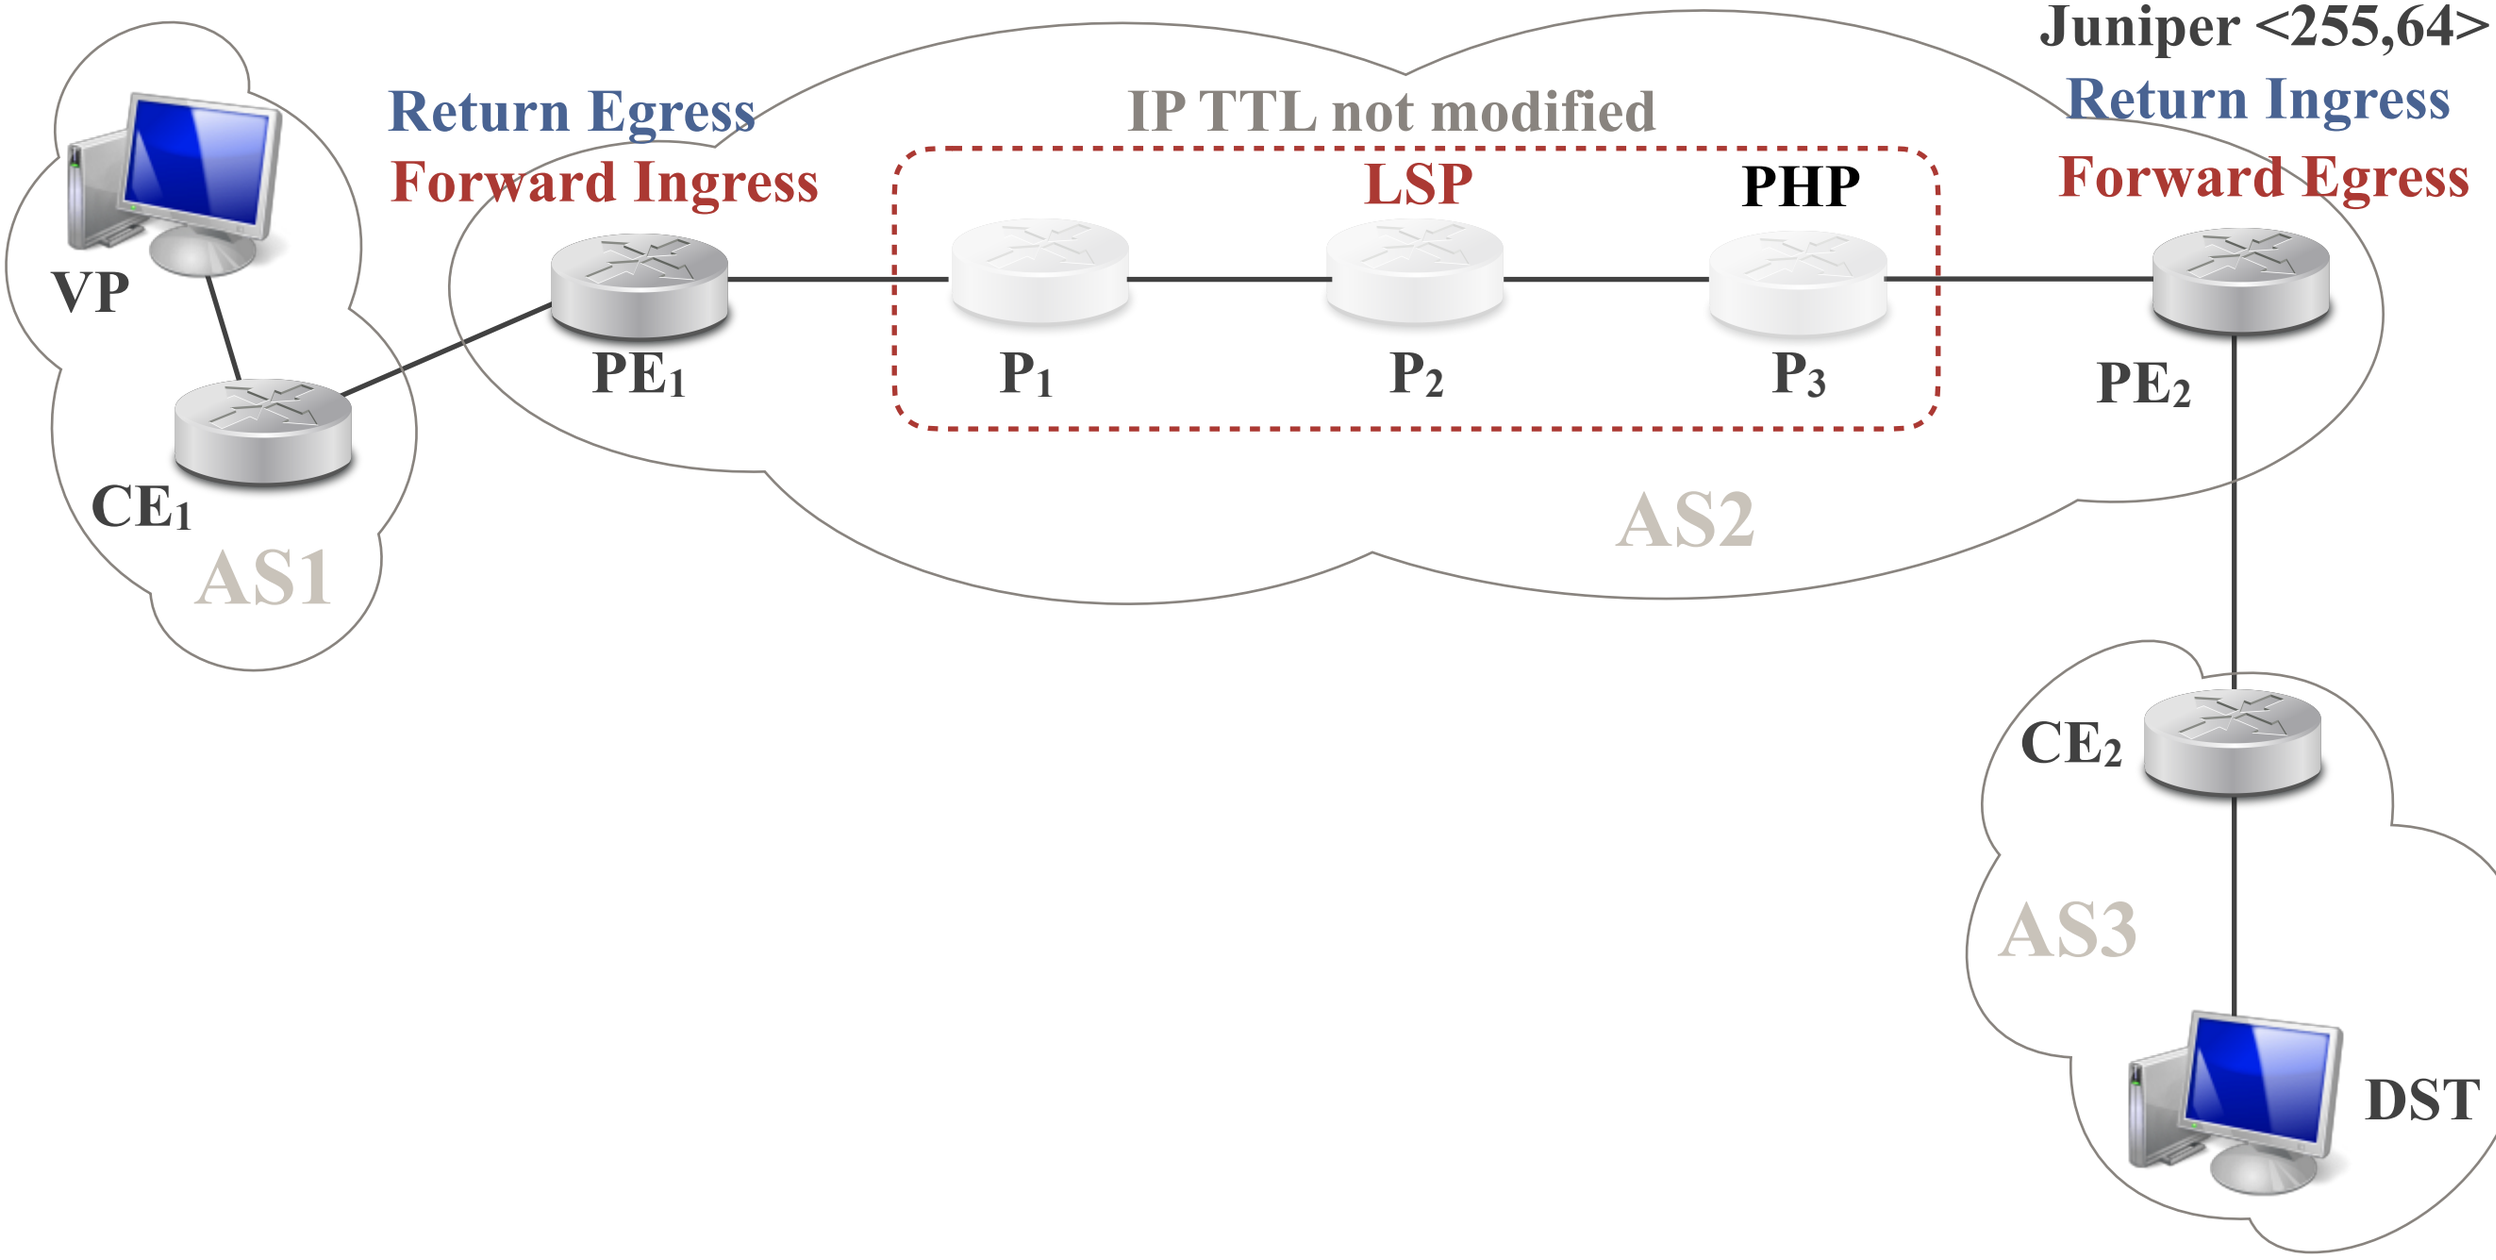

Invisible Tunnels (6)

- Direct Path Revelation

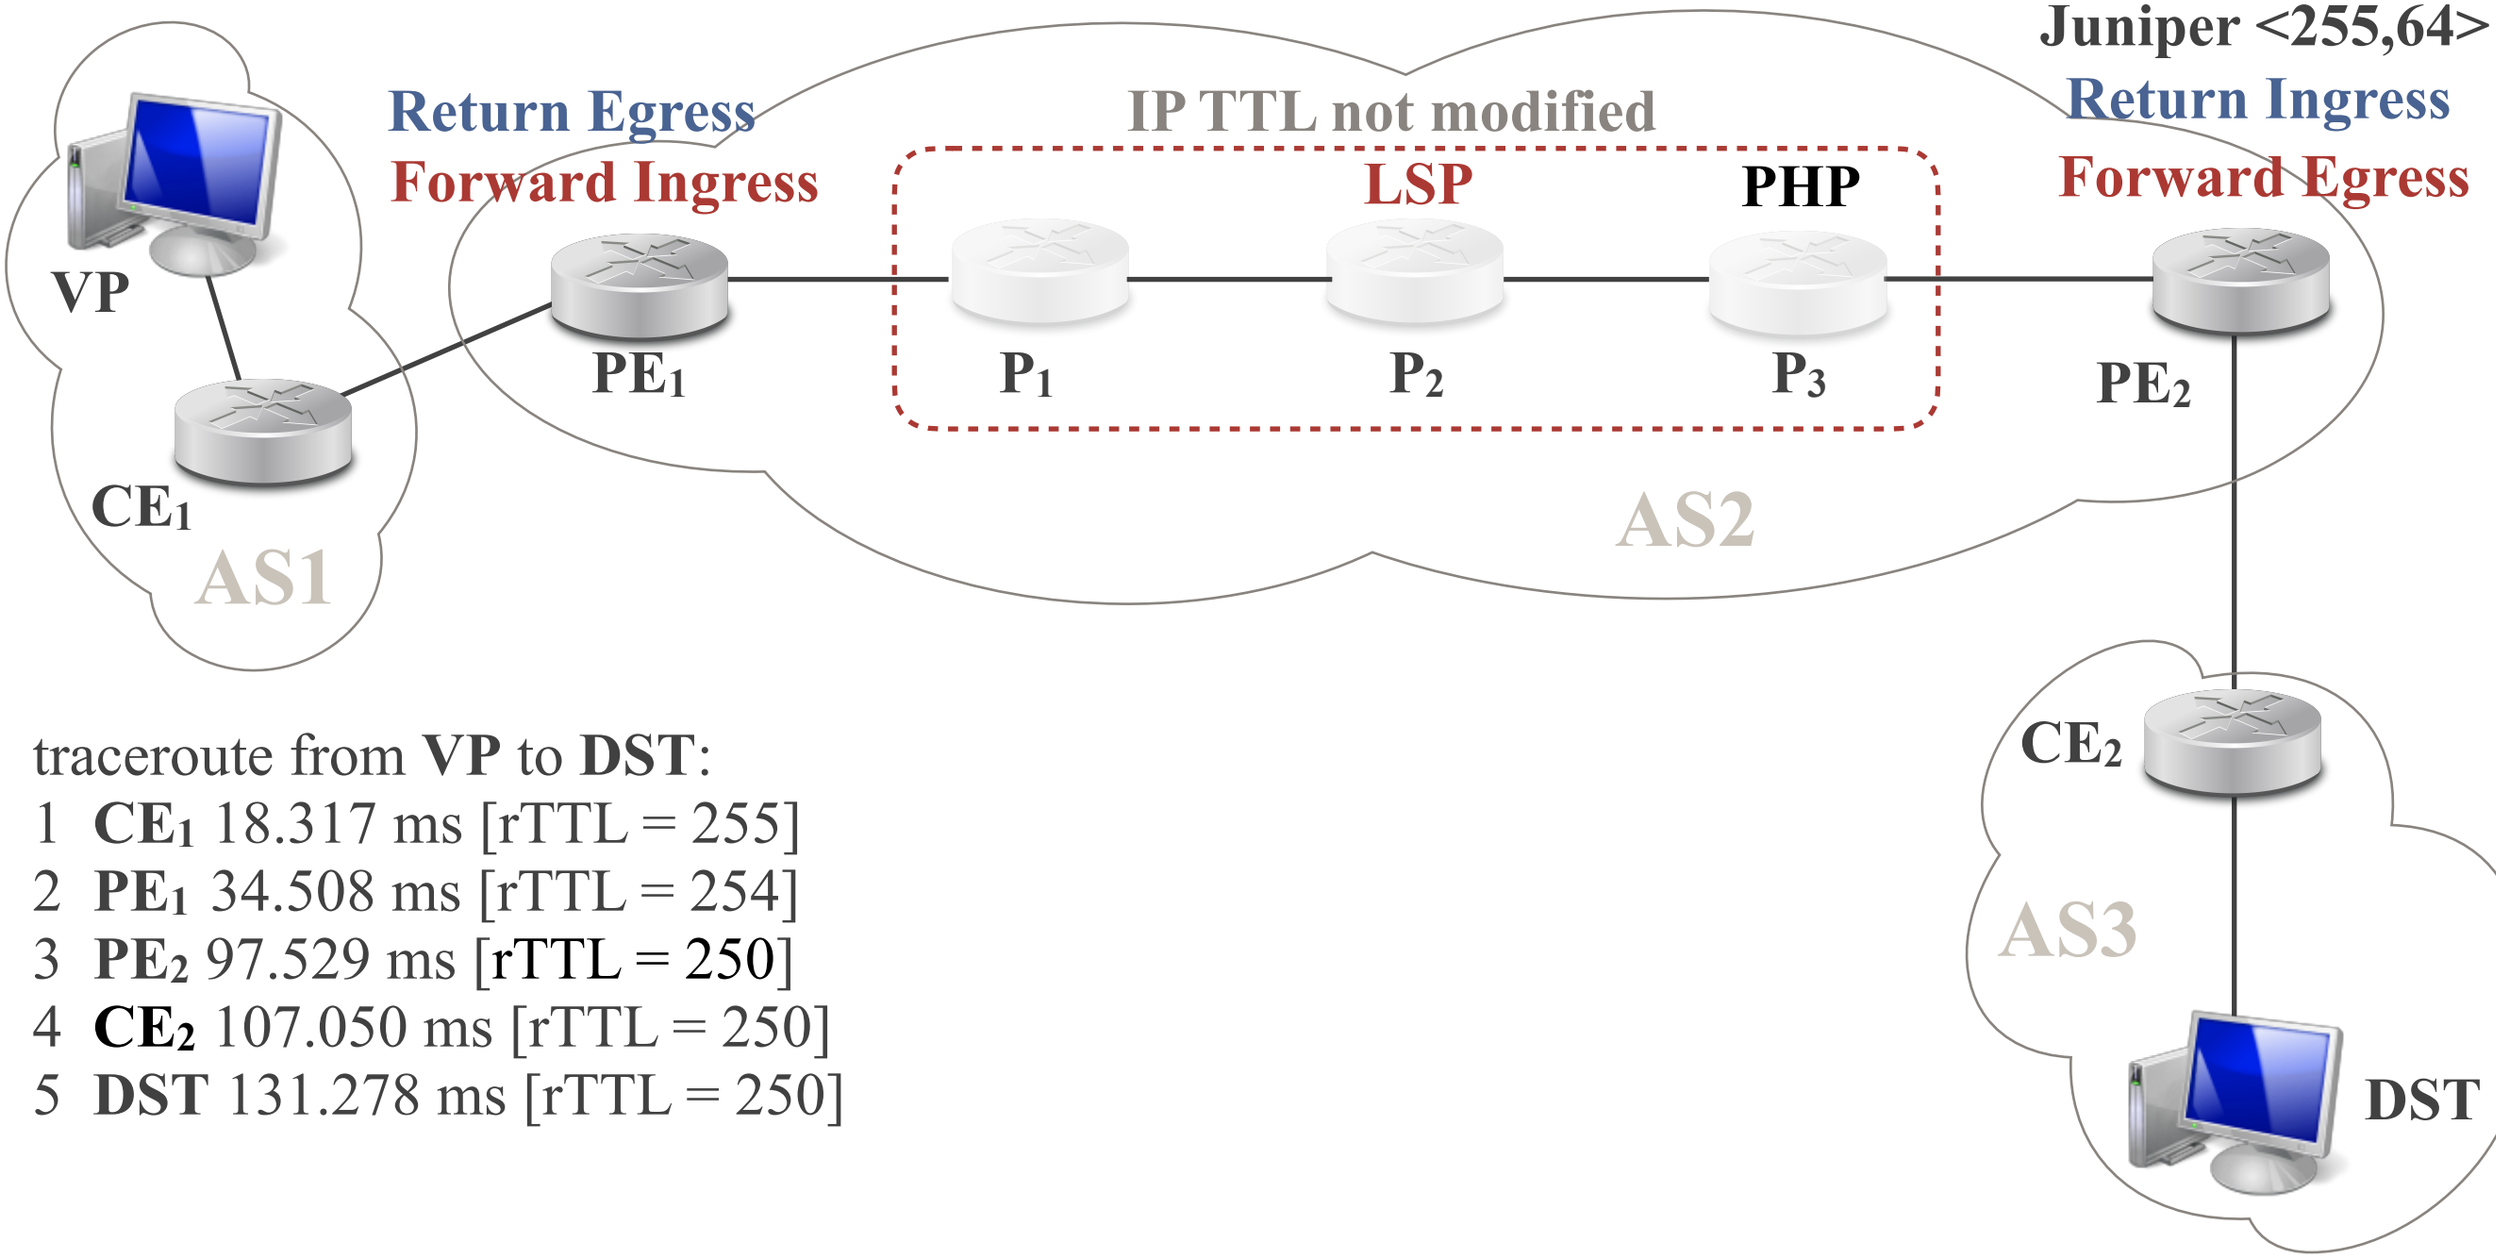

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 250**]
4 **CE$_2$** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

Invisible Tunnels (6)

- Direct Path Revelation

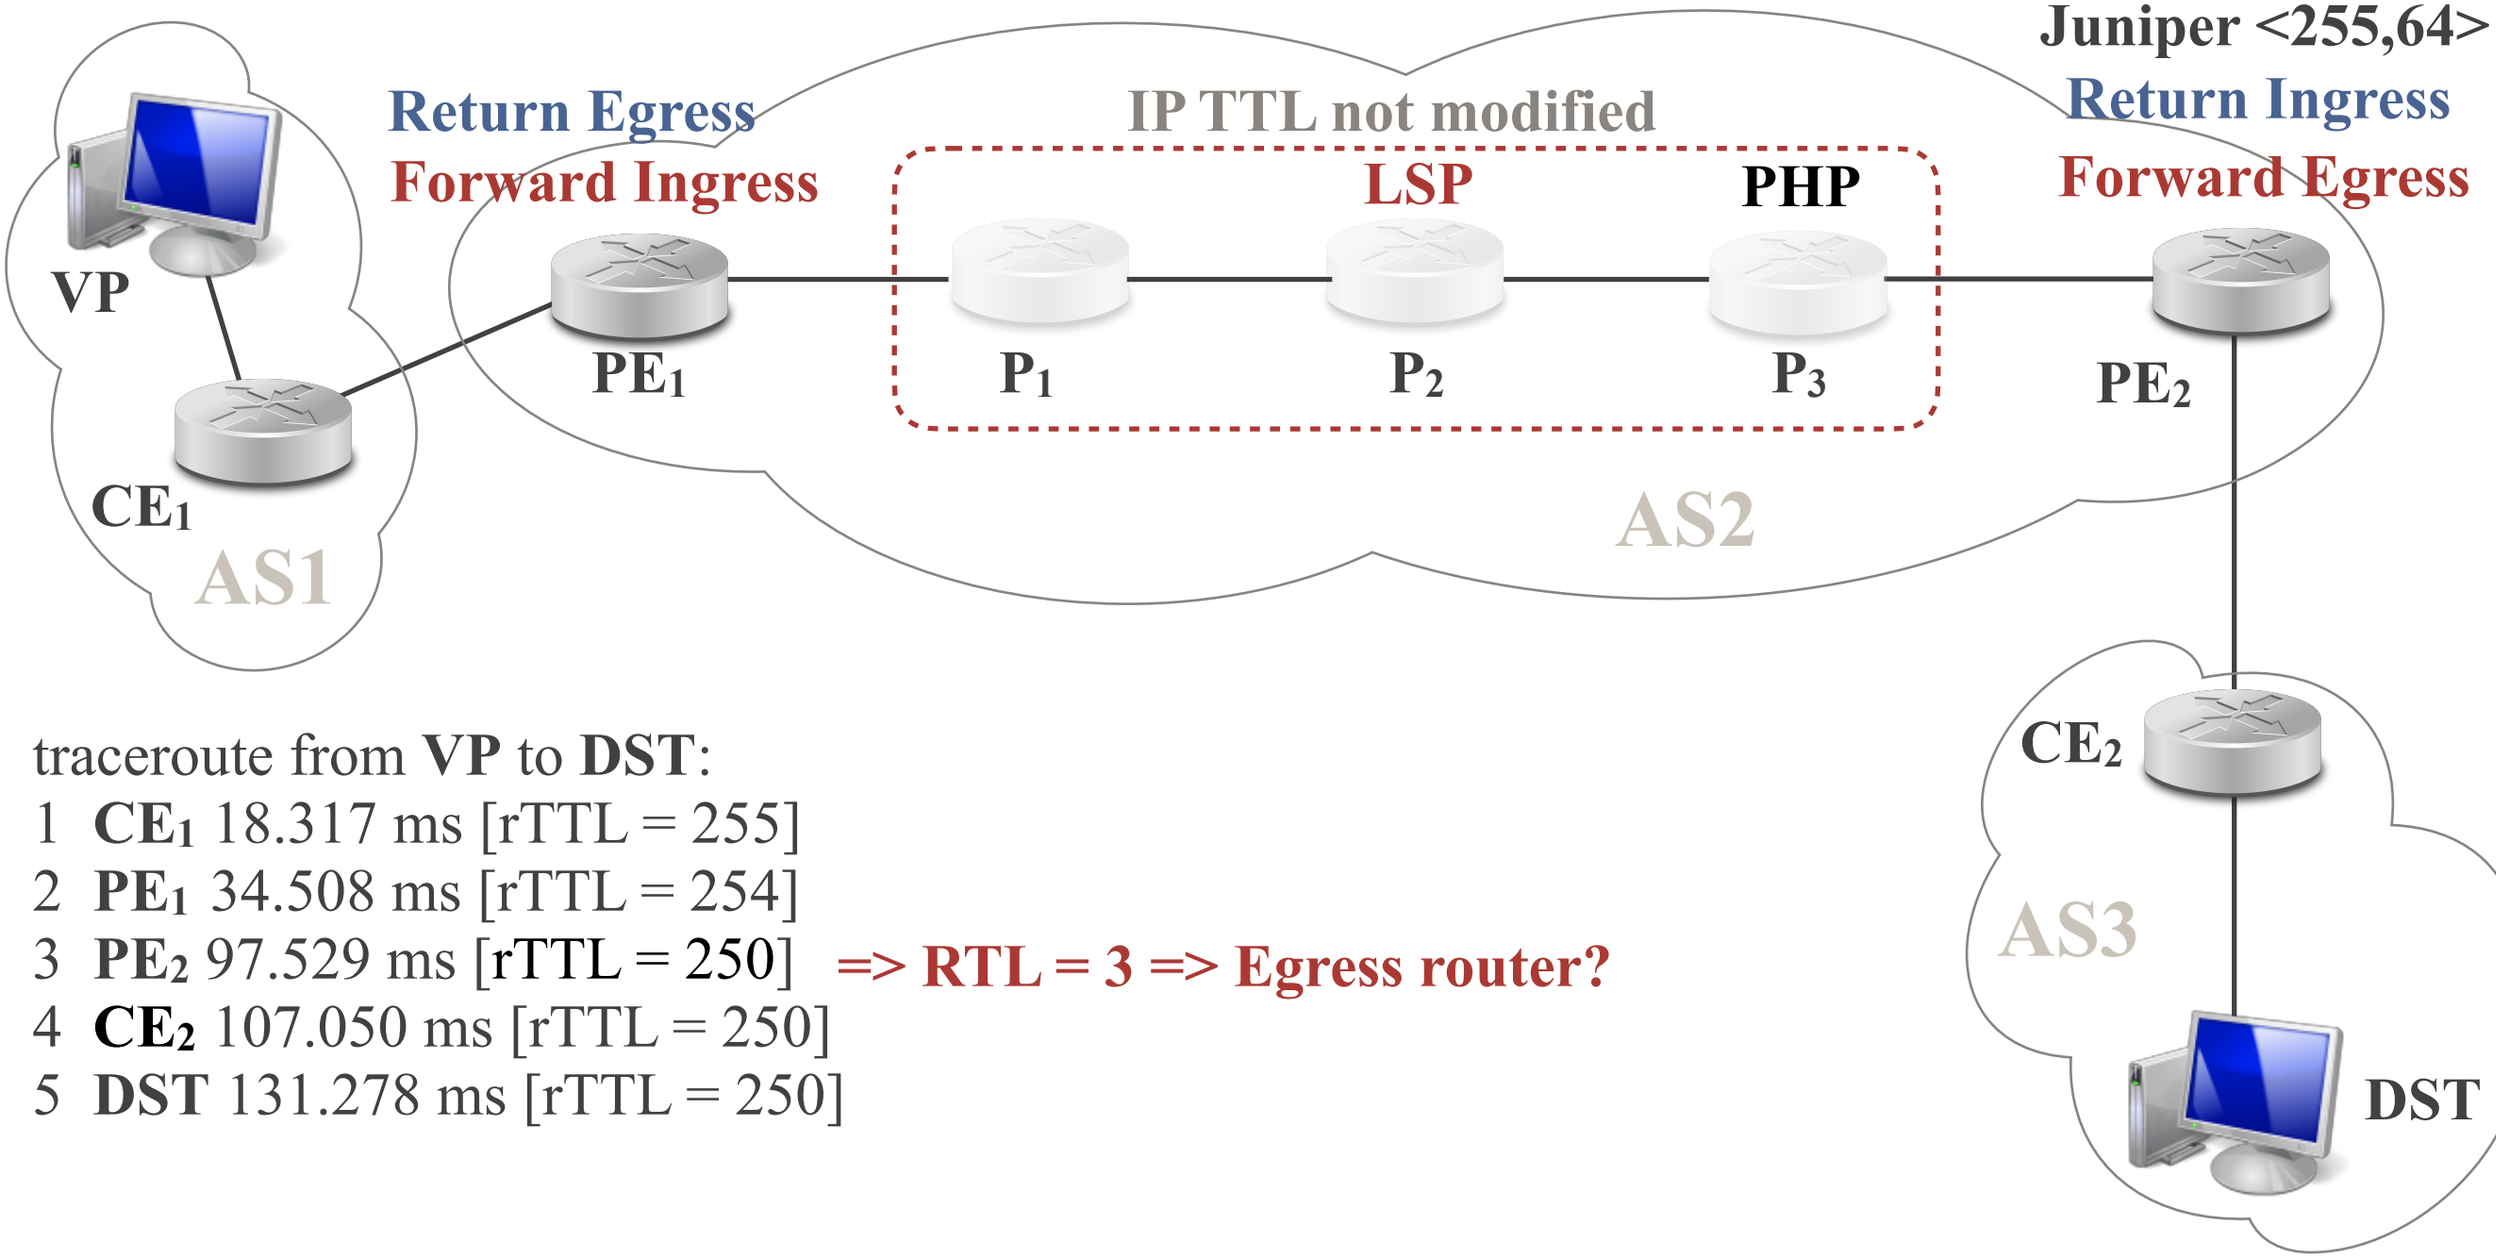

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [rTTL = 250] **=> RTL = 3 => Egress router?**
4 **CE$_2$** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

Invisible Tunnels (6)

- Direct Path Revelation

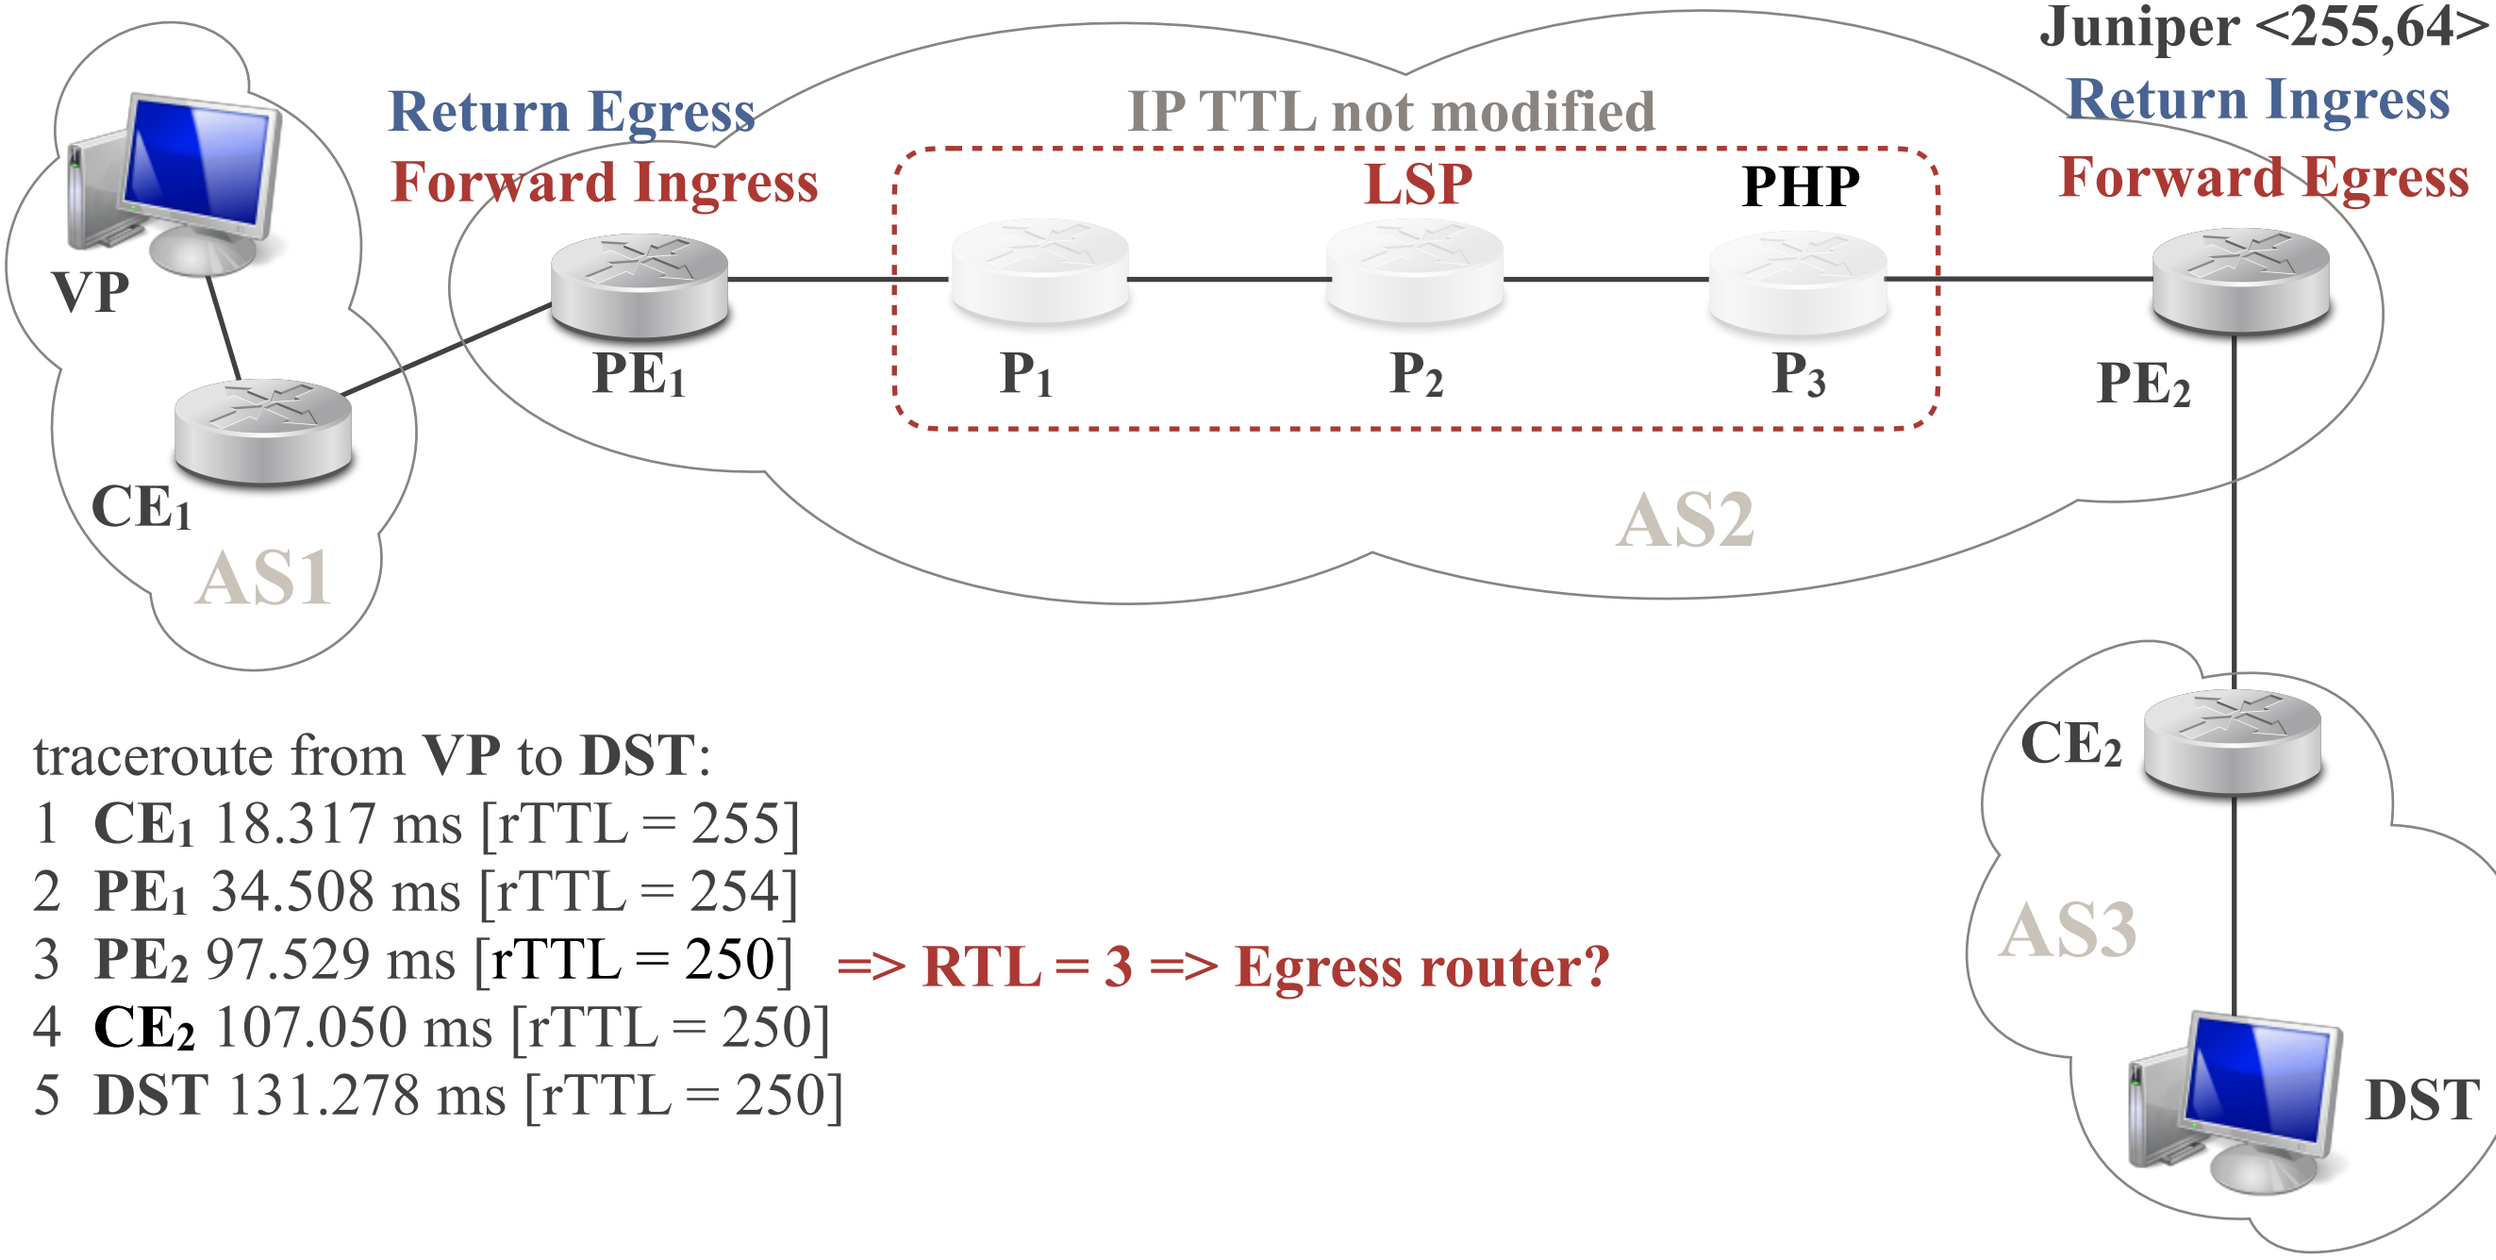

traceroute from **VP** to **DST**:
1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 **PE2** 97.529 ms [rTTL = 250] **=> RTL = 3 => Egress router?**
4 **CE2** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

Simple IP forwarding if MPLS not used for internal traffic

Invisible Tunnels (6)

- Direct Path Revelation

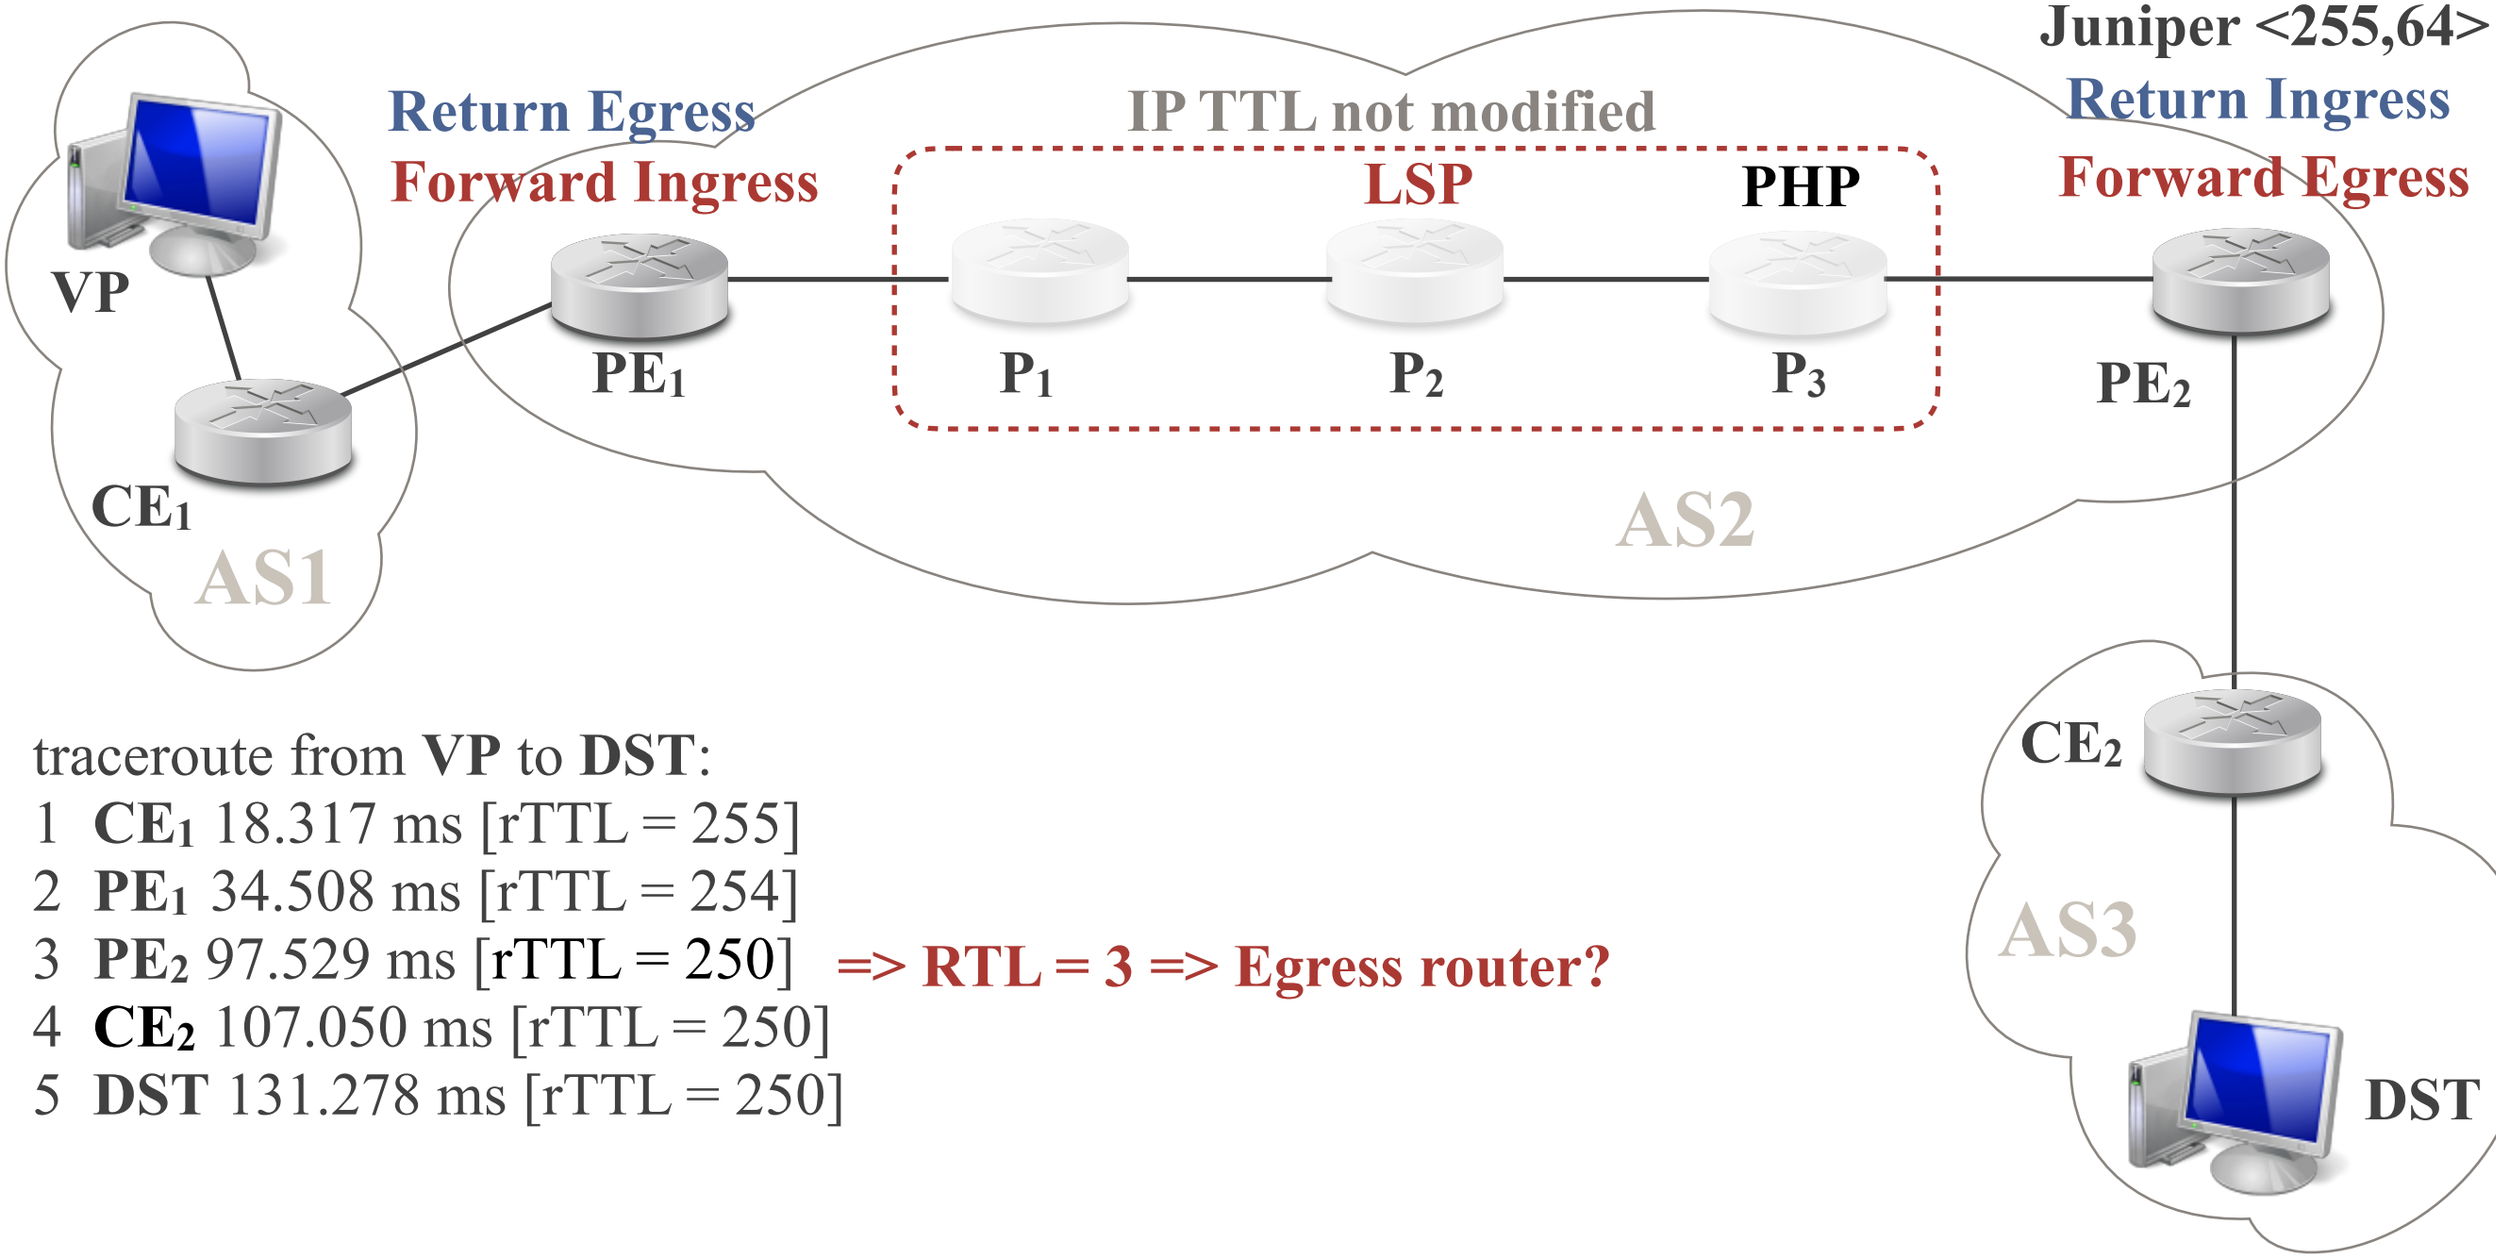

traceroute from **VP** to **DST**:
1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [rTTL = 250] **=> RTL = 3 => Egress router?**
4 **CE$_2$** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

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

Invisible Tunnels (6)

- Direct Path Revelation

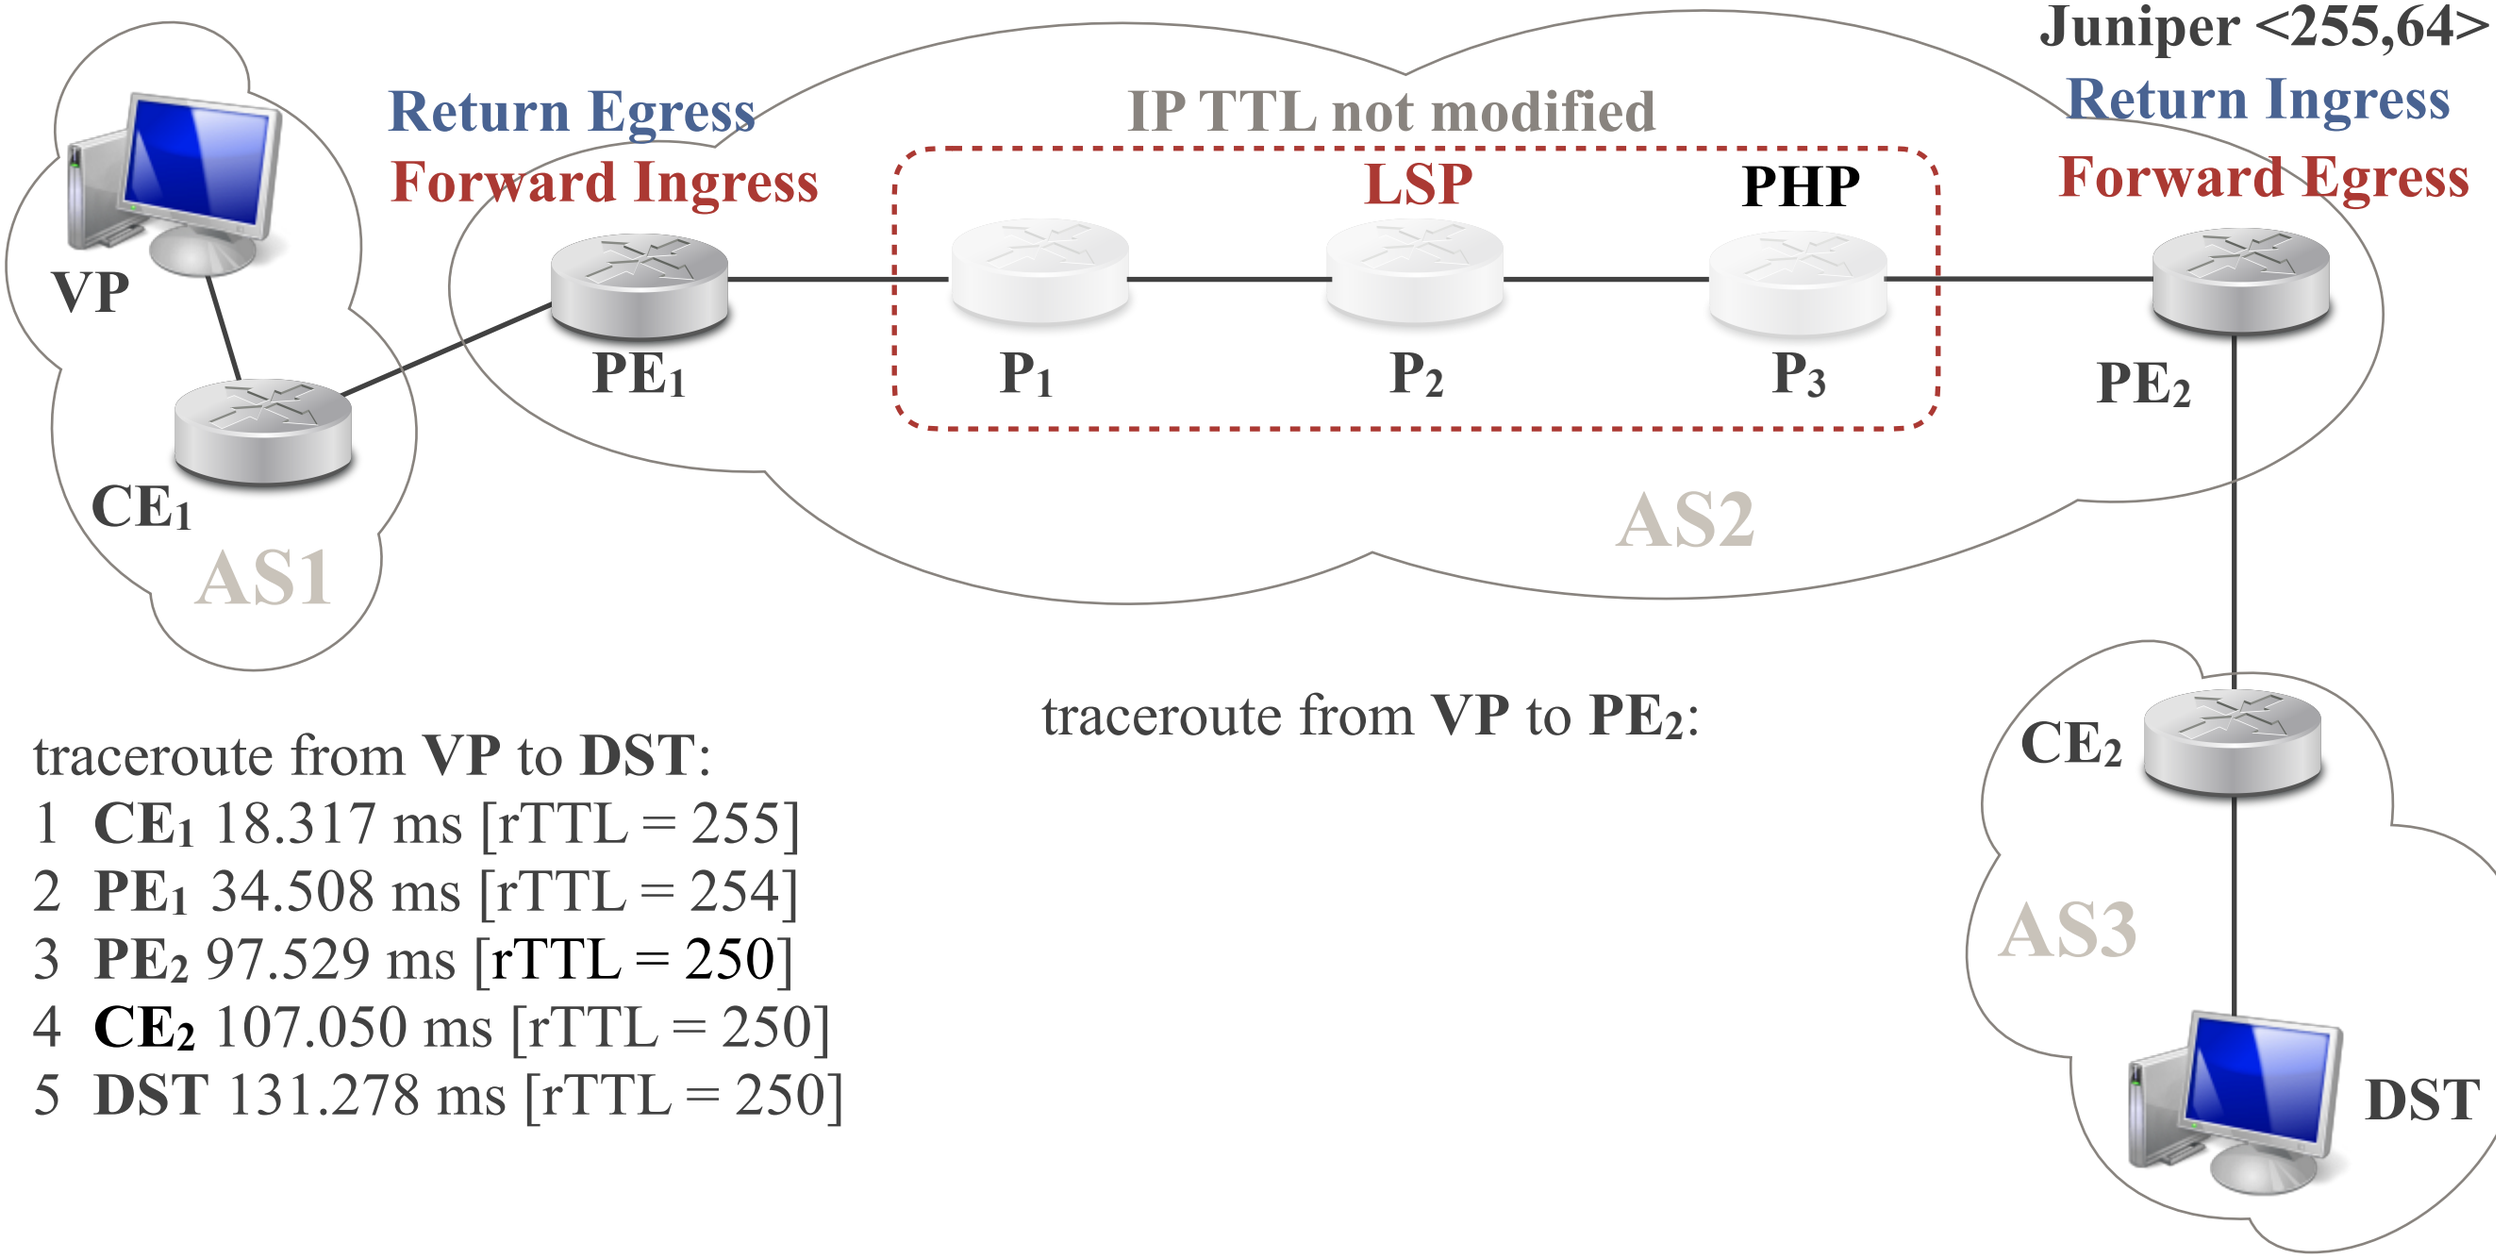

traceroute from **VP** to **DST**:

1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 250**]
4 **CE$_2$** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

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

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

Invisible Tunnels (6)

- Direct Path Revelation

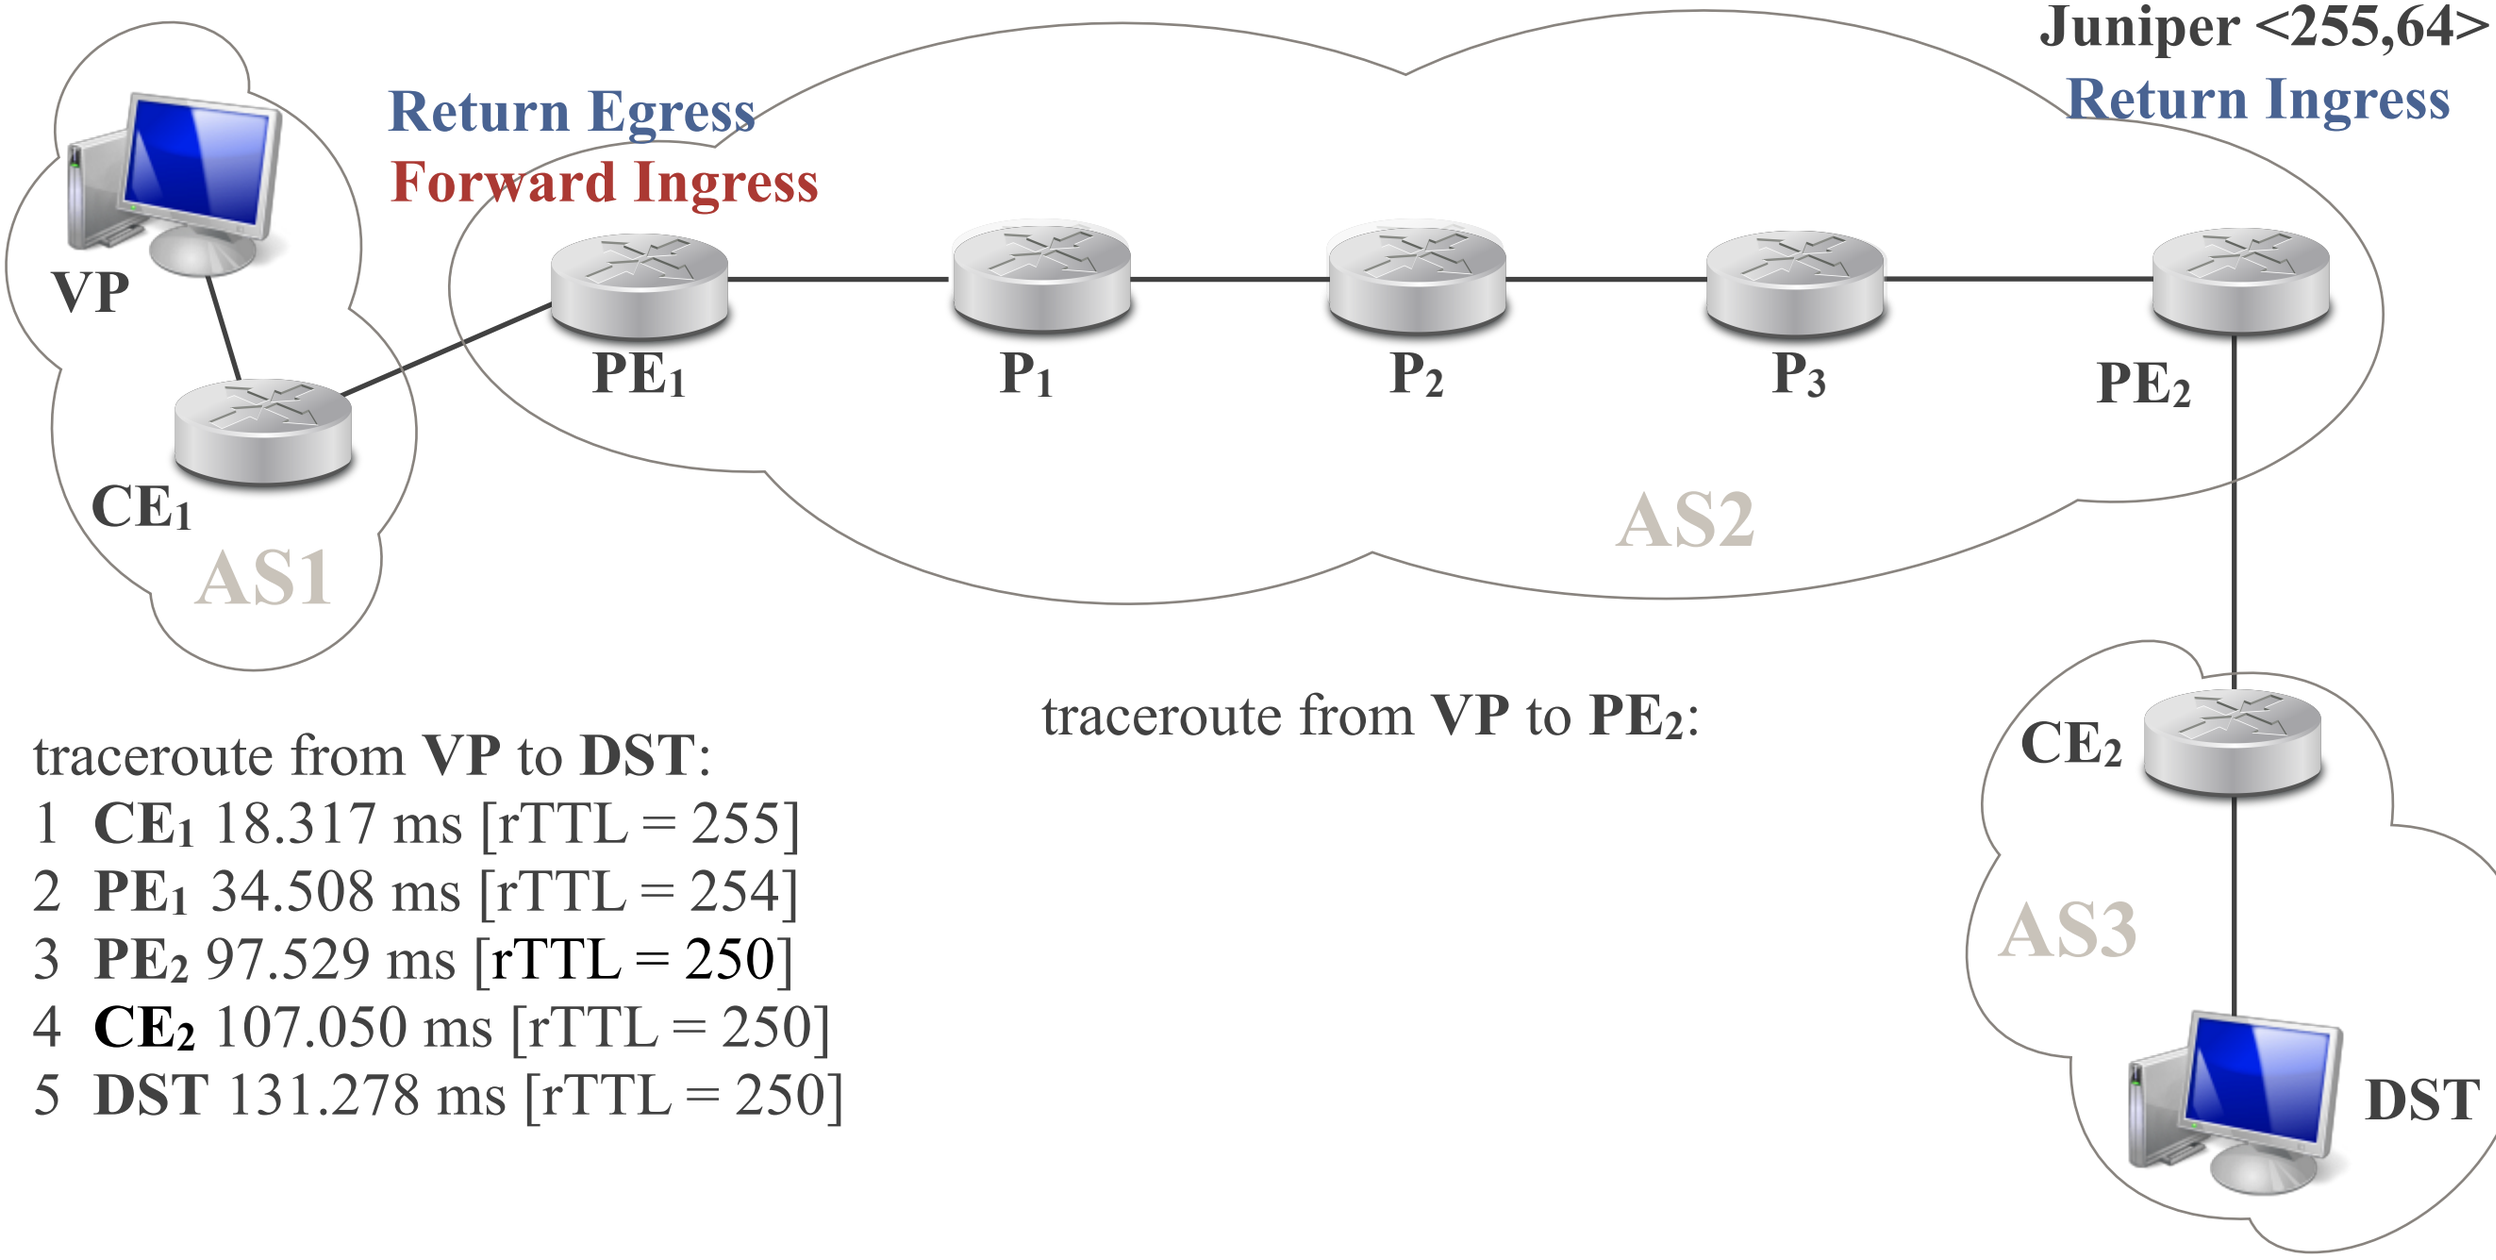

traceroute from **VP** to **DST**:
1 **CE$_1$** 18.317 ms [rTTL = 255]
2 **PE$_1$** 34.508 ms [rTTL = 254]
3 **PE$_2$** 97.529 ms [**rTTL = 250**]
4 **CE$_2$** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

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

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

Invisible Tunnels (6)

- Direct Path Revelation

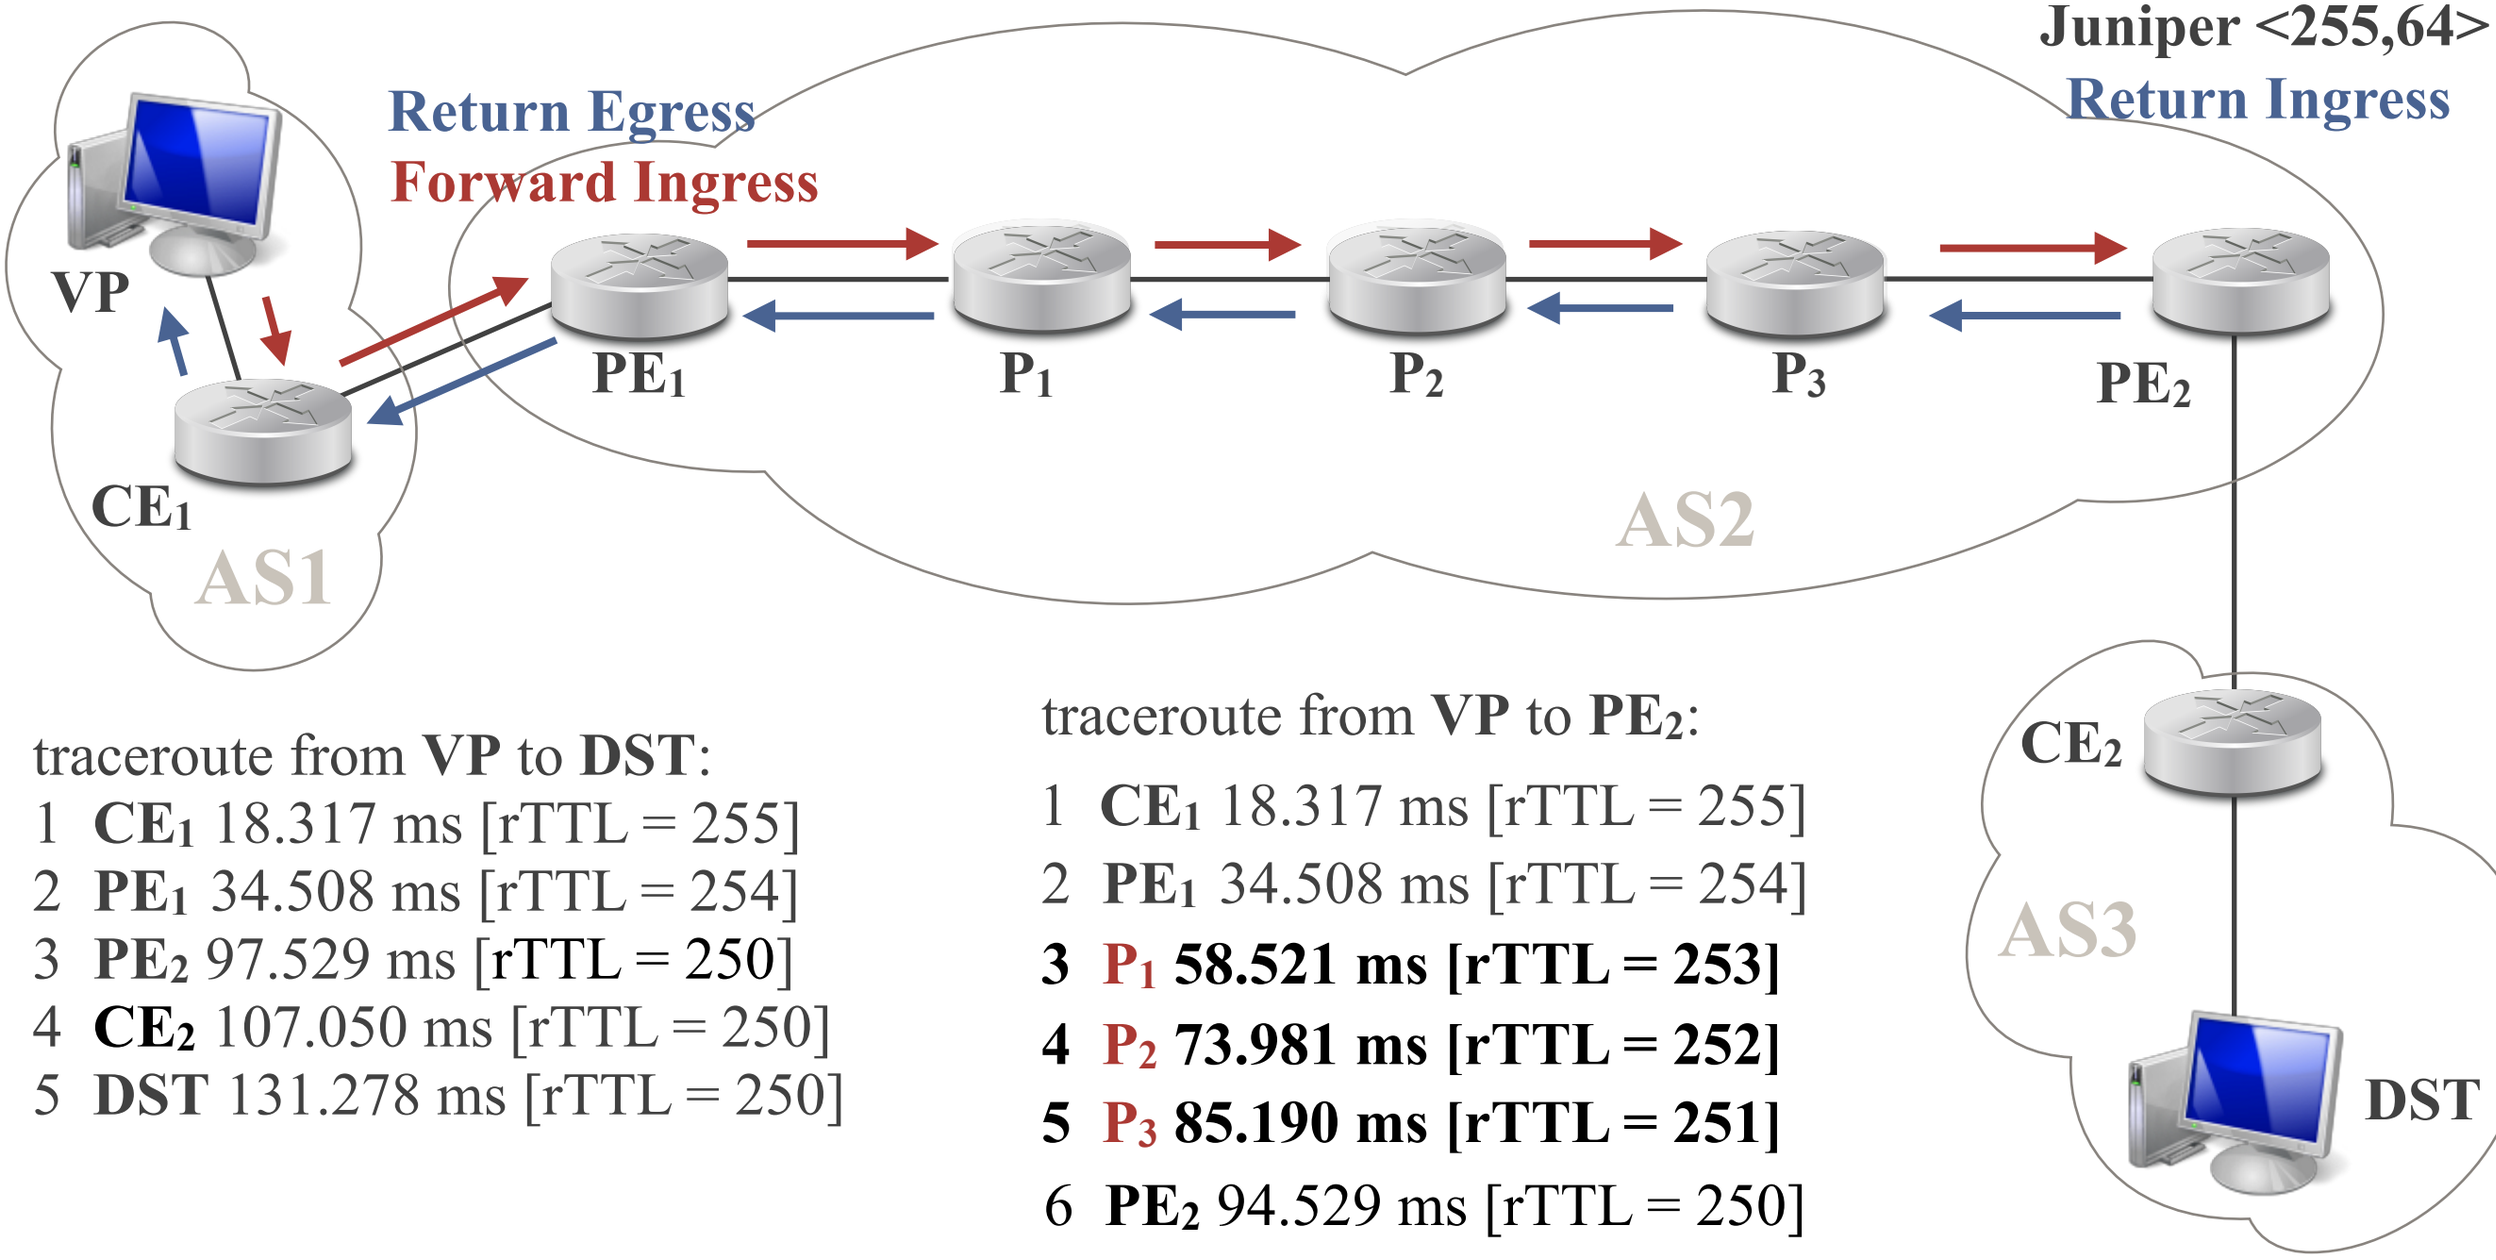

traceroute from **VP** to **DST**:
1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 **PE2** 97.529 ms [**rTTL = 250**]
4 **CE2** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

traceroute from **VP** to **PE2**:
1 **CE1** 18.317 ms [rTTL = 255]
2 **PE1** 34.508 ms [rTTL = 254]
3 P1 58.521 ms [rTTL = 253]
4 P2 73.981 ms [rTTL = 252]
5 P3 85.190 ms [rTTL = 251]
6 **PE2** 94.529 ms [rTTL = 250]

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

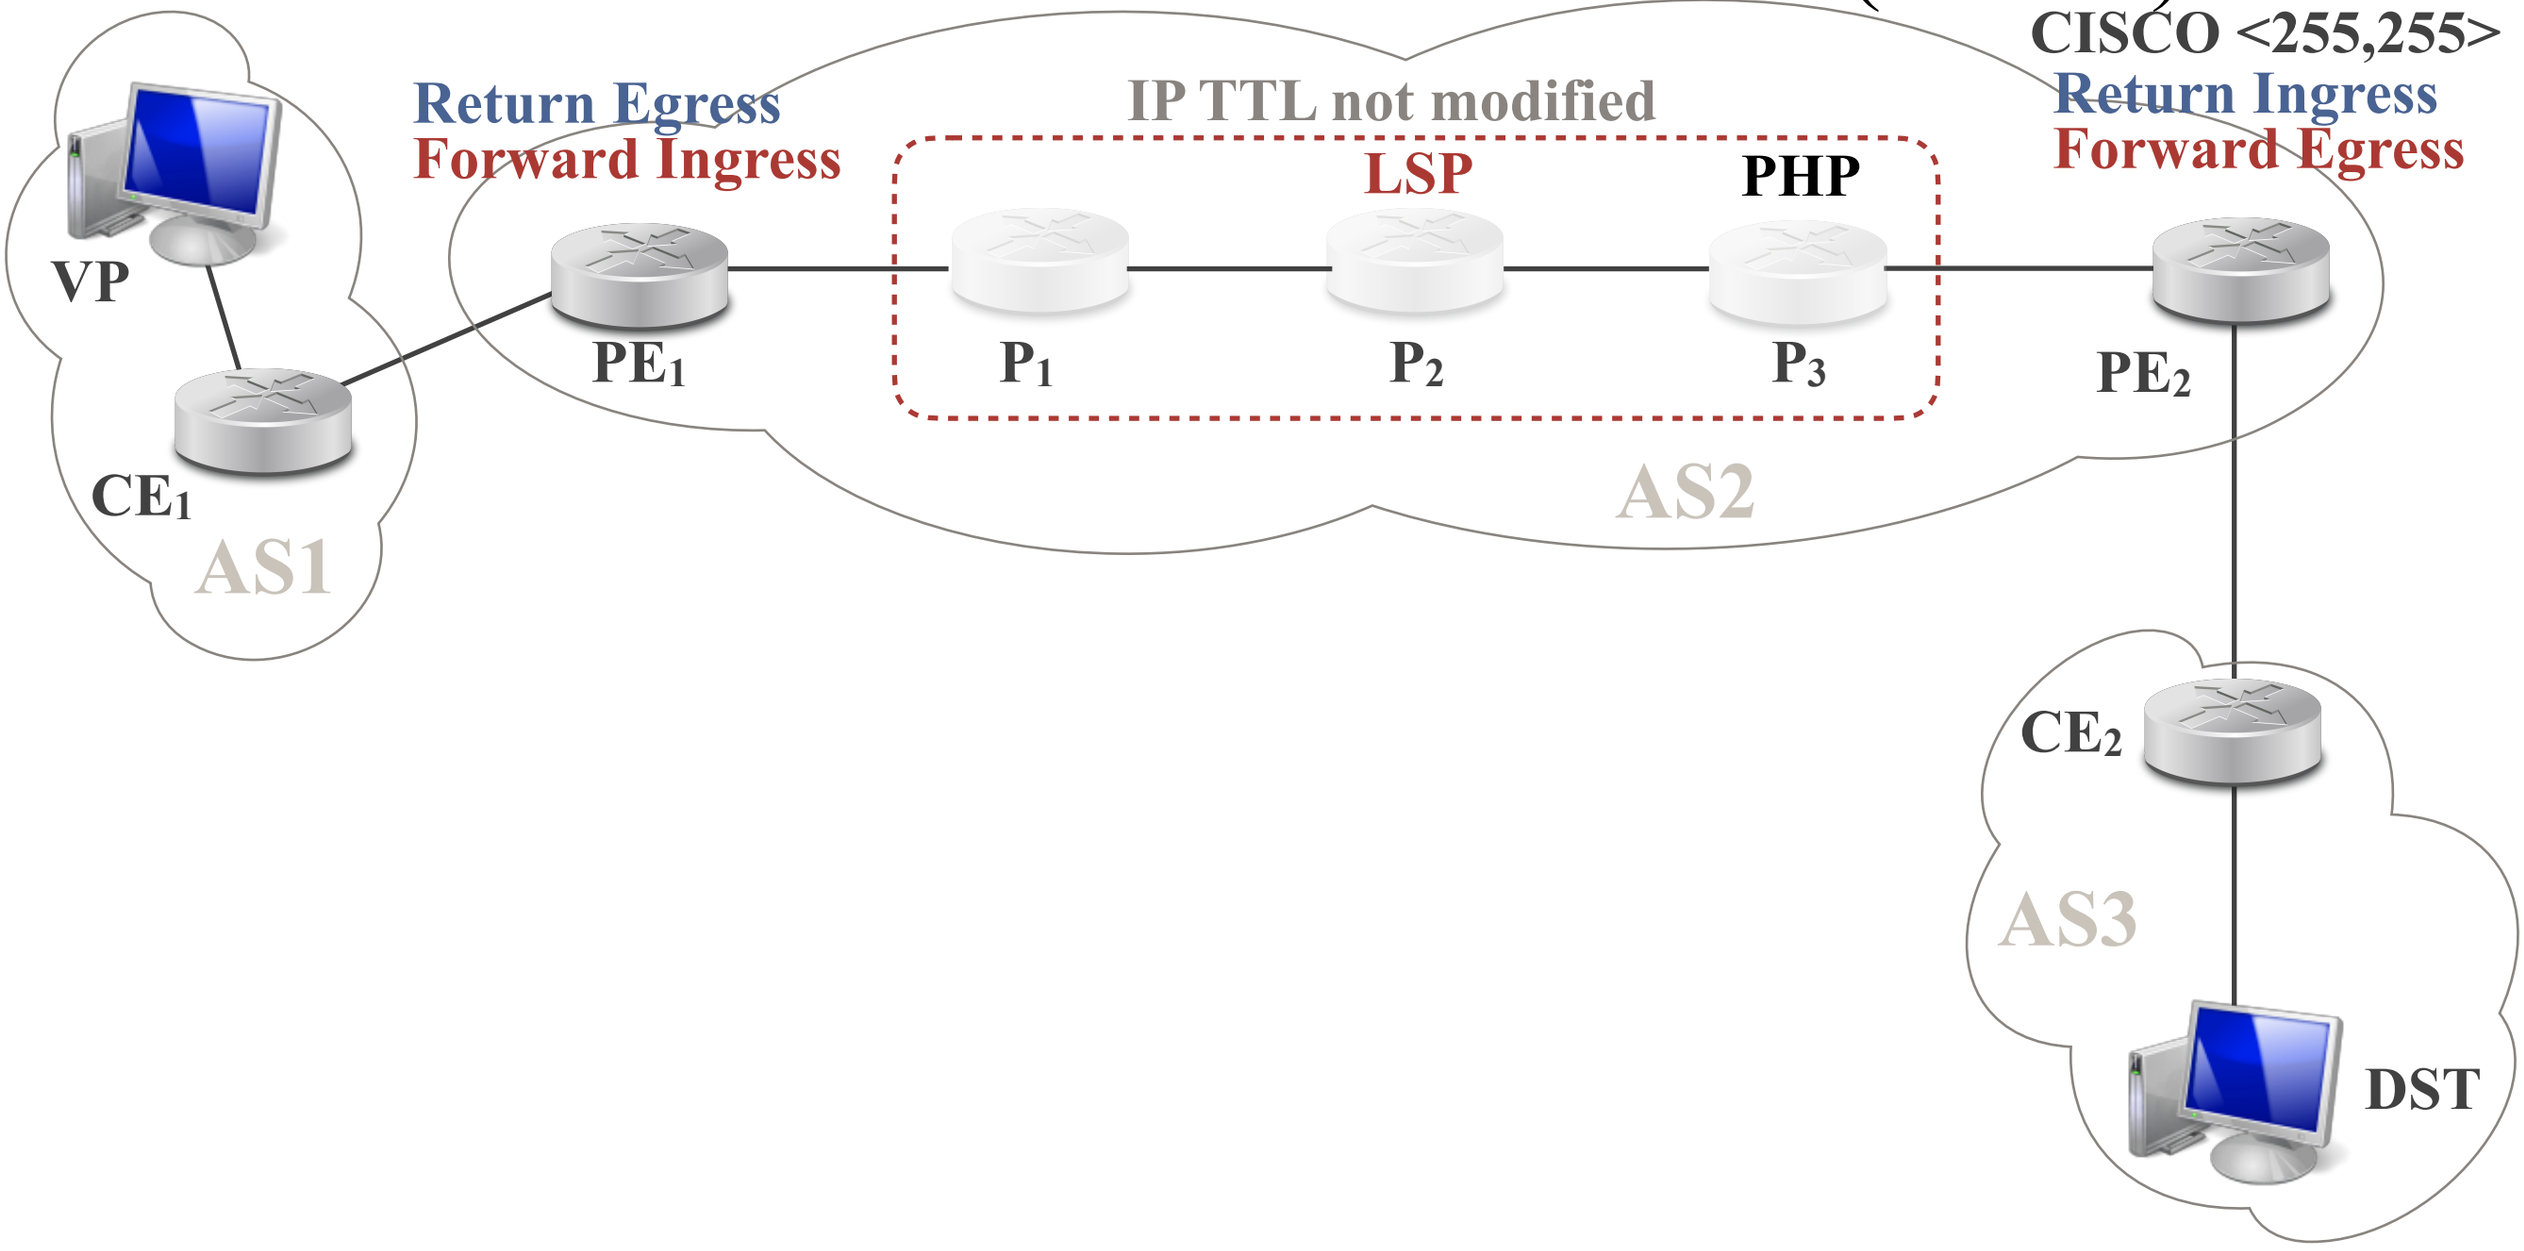

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

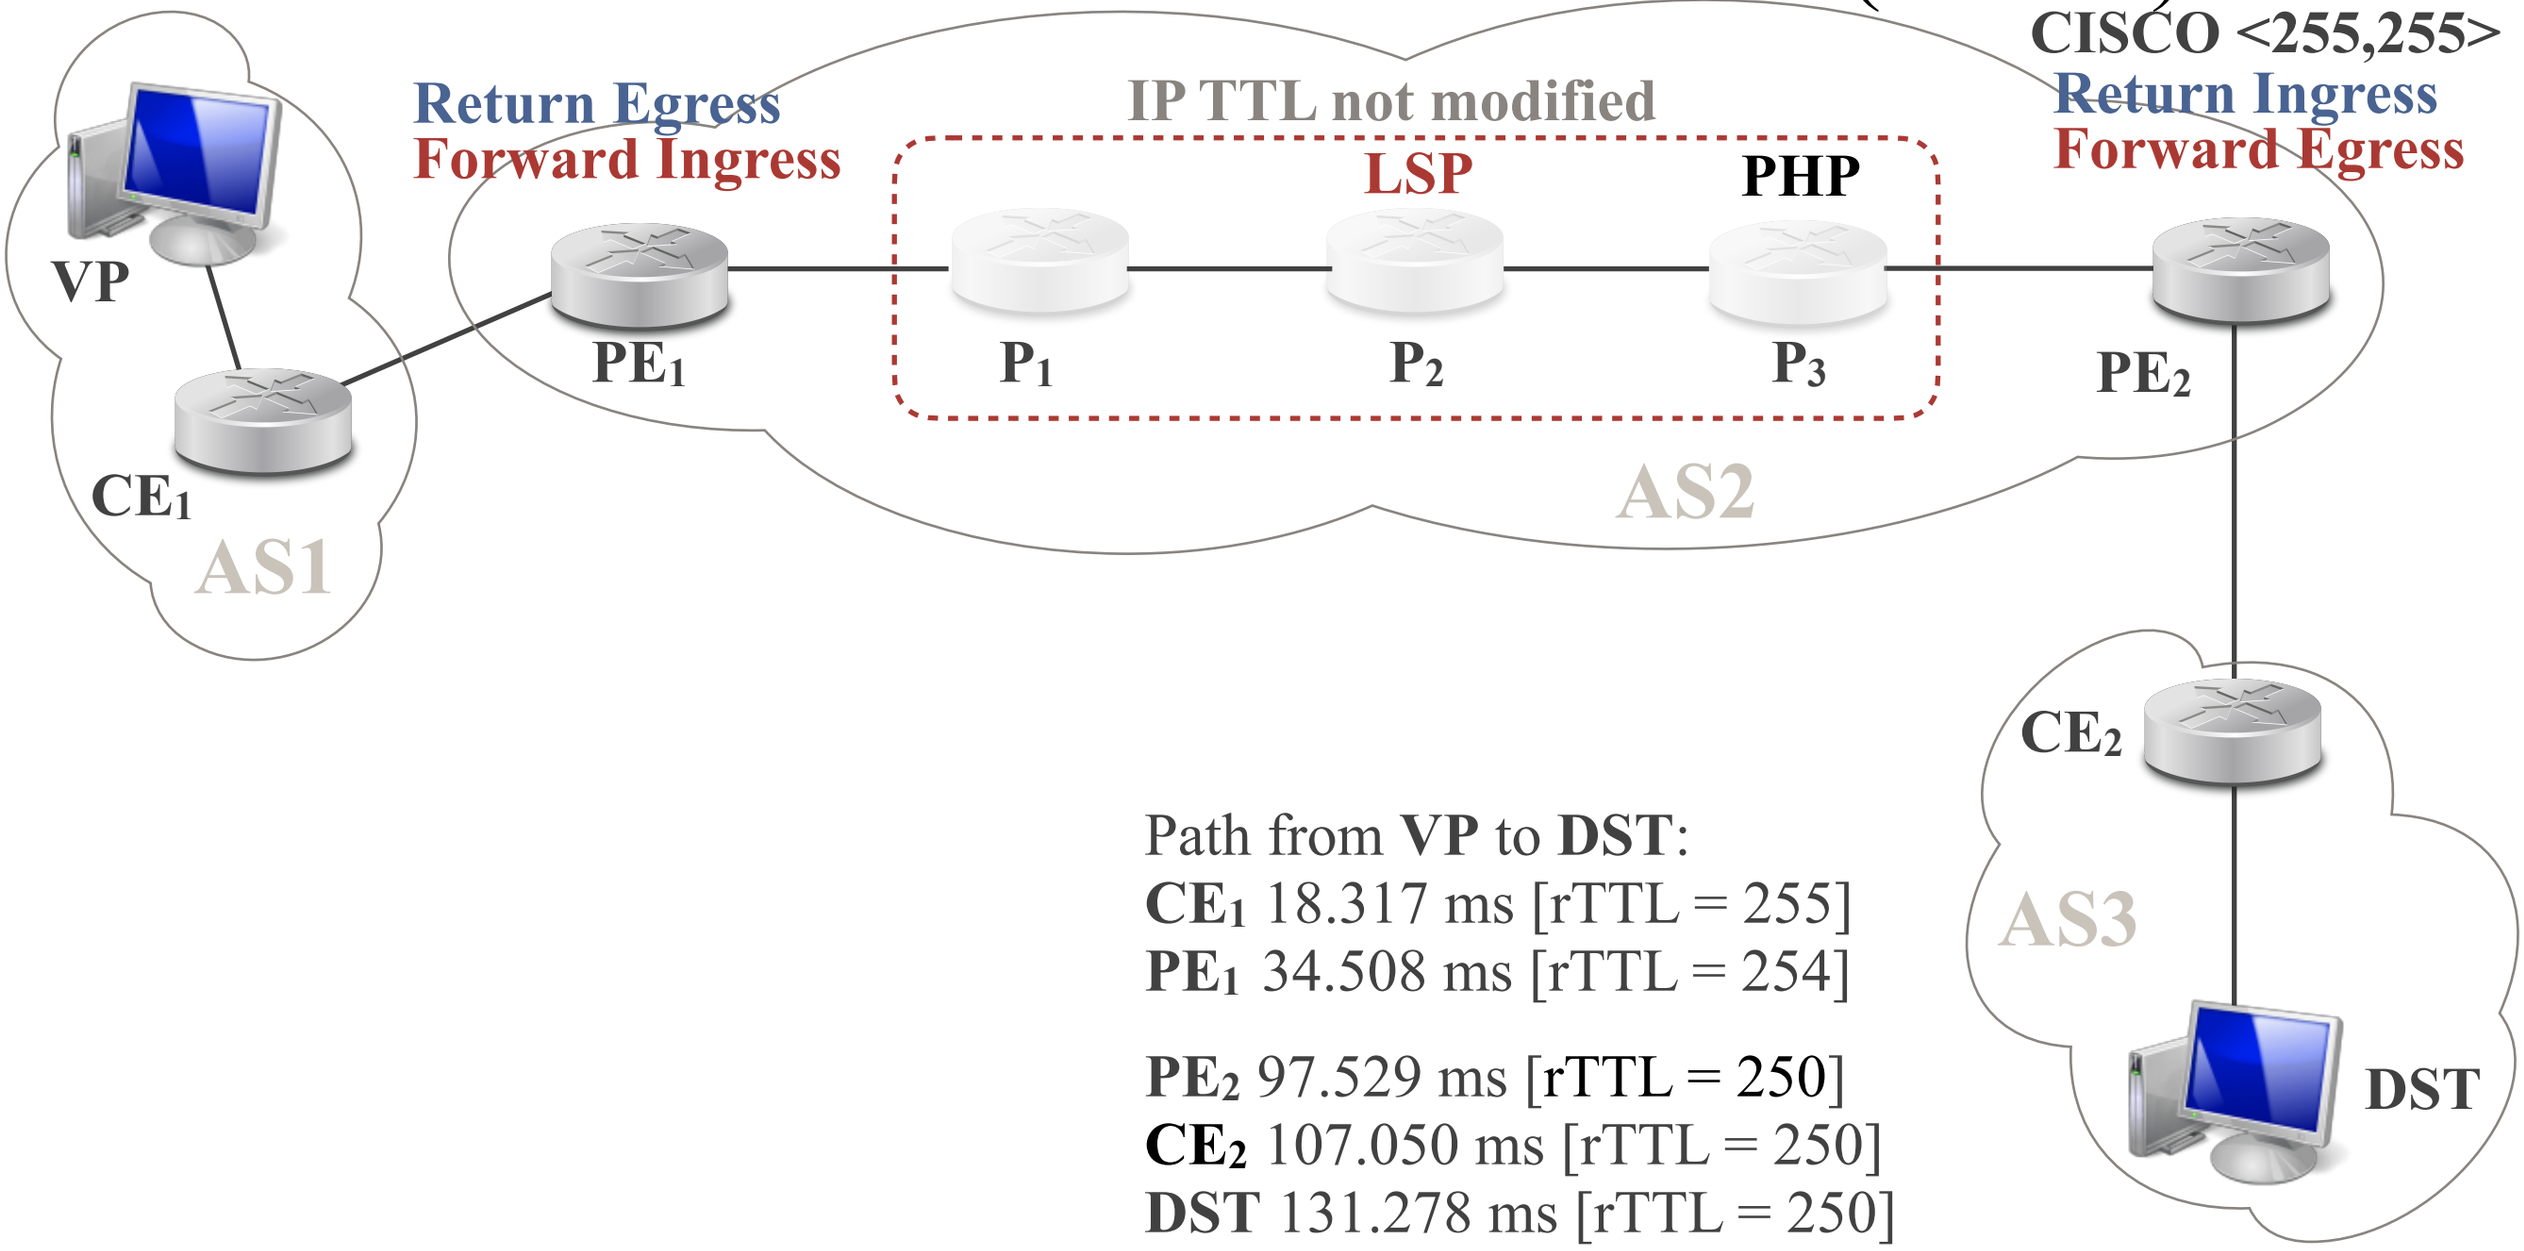

Path from **VP** to **DST**:
CE1 18.317 ms [rTTL = 255]
PE1 34.508 ms [rTTL = 254]

PE2 97.529 ms [**rTTL = 250**]
CE2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

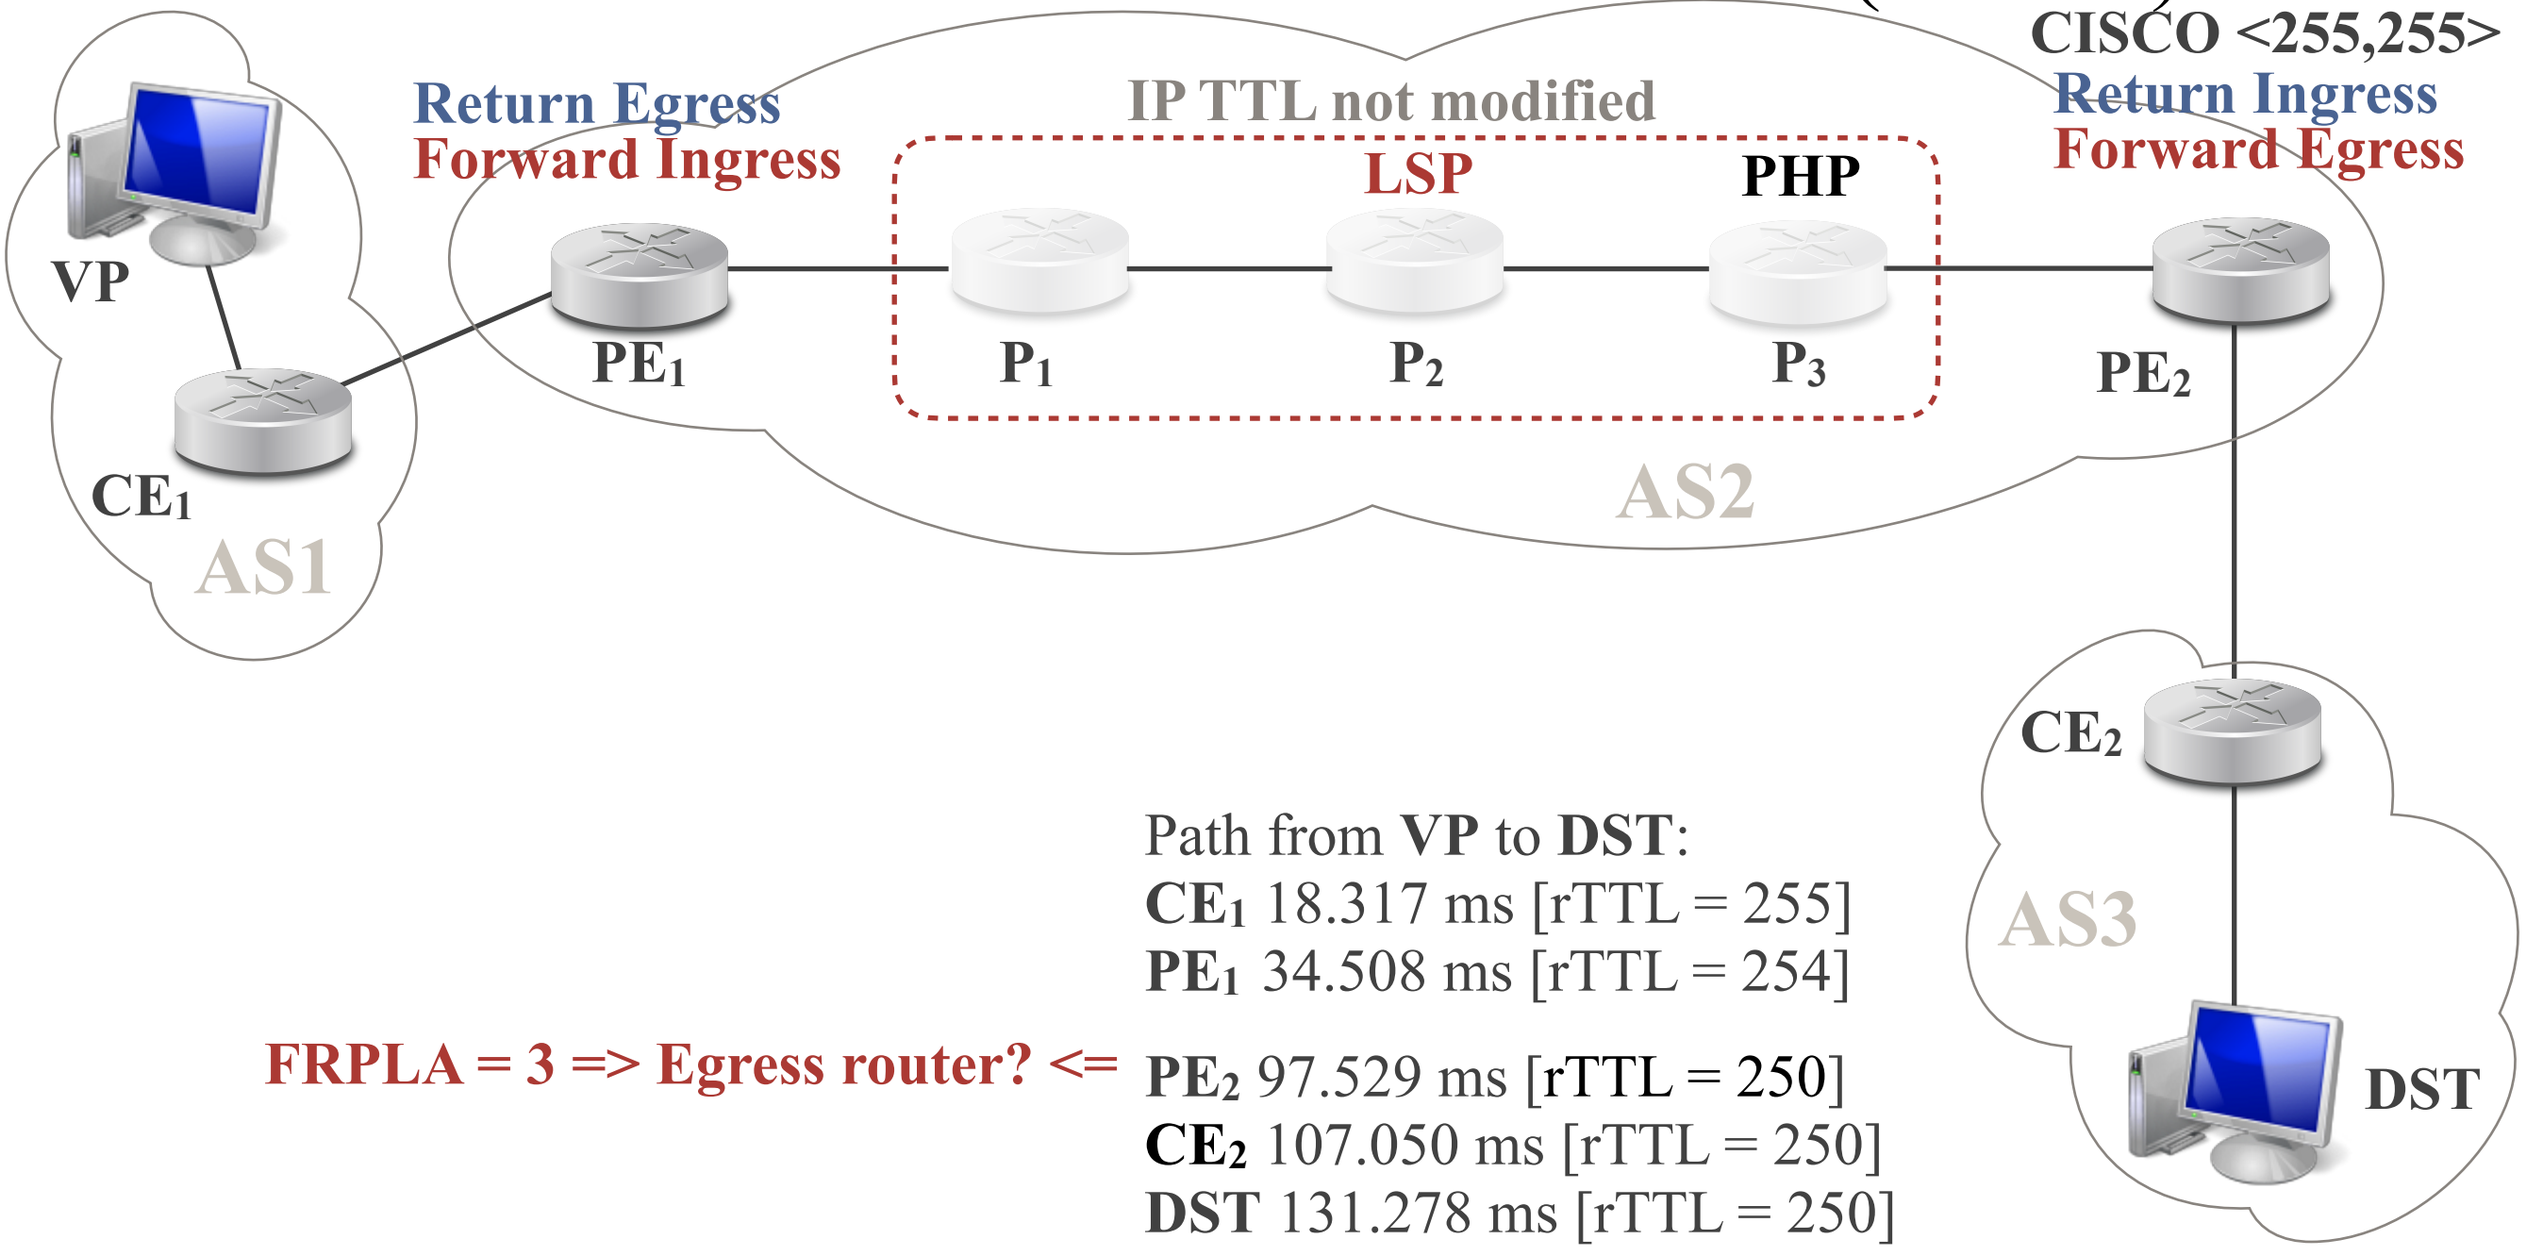

FRPLA = 3 => Egress router? <=

Path from **VP** to **DST**:
CE1 18.317 ms [rTTL = 255]
PE1 34.508 ms [rTTL = 254]

PE2 97.529 ms [rTTL = 250]
CE2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

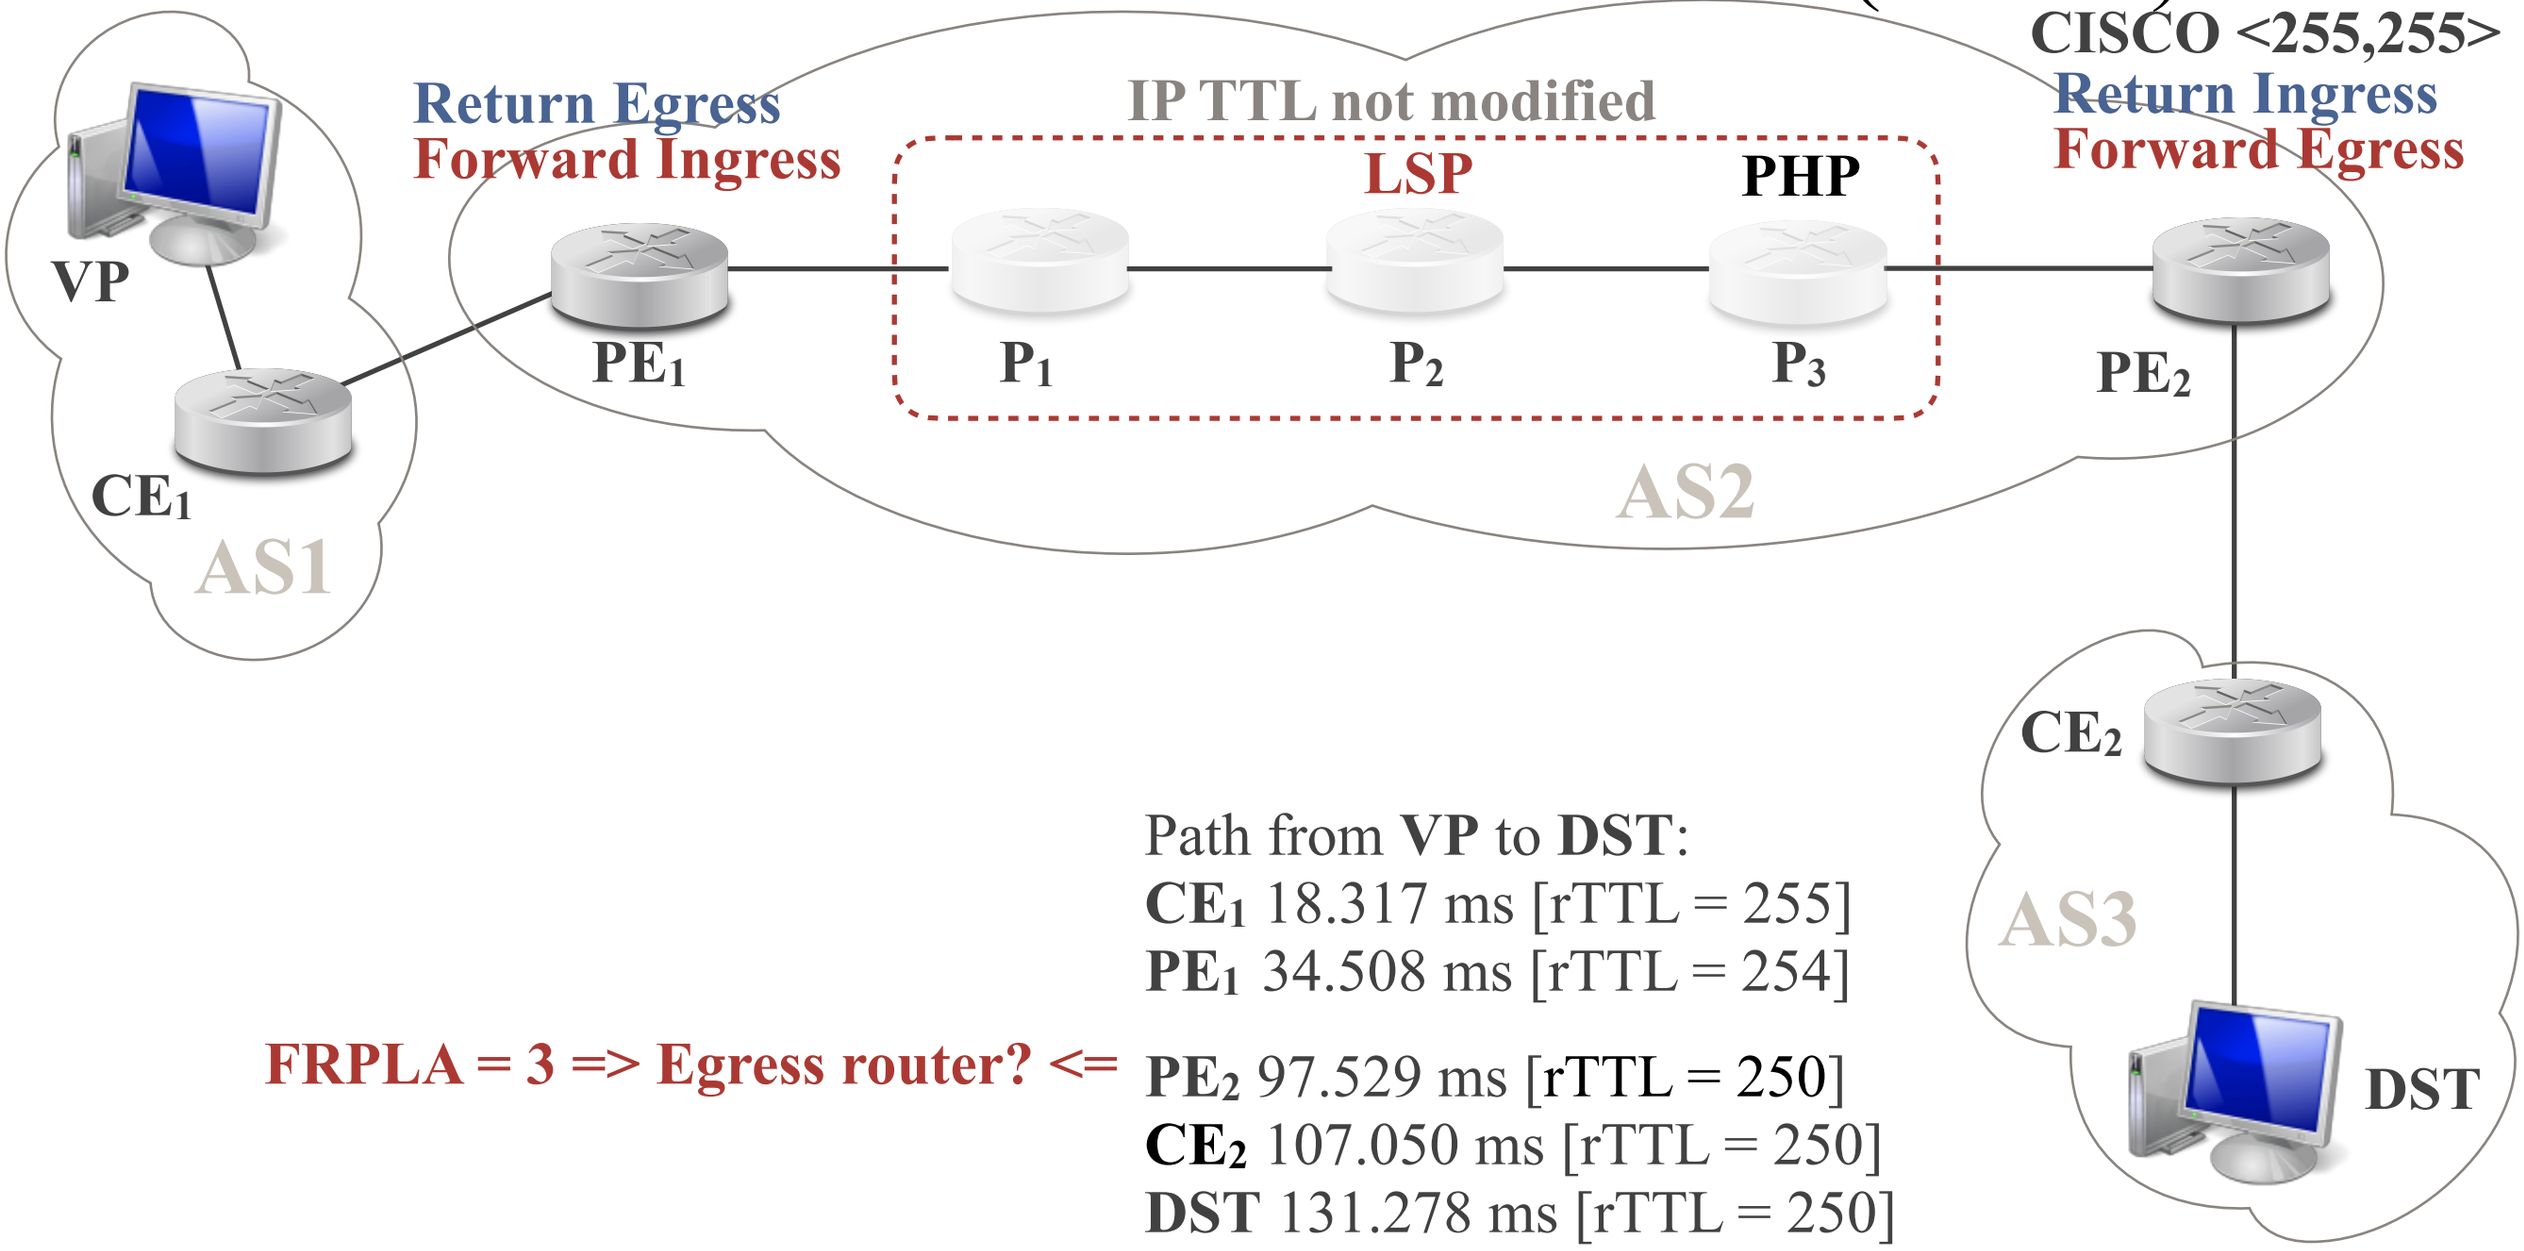

Path from **VP** to **DST**:
CE1 18.317 ms [rTTL = 255]
PE1 34.508 ms [rTTL = 254]

FRPLA = 3 => Egress router? <=

PE2 97.529 ms [rTTL = 250]
CE2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

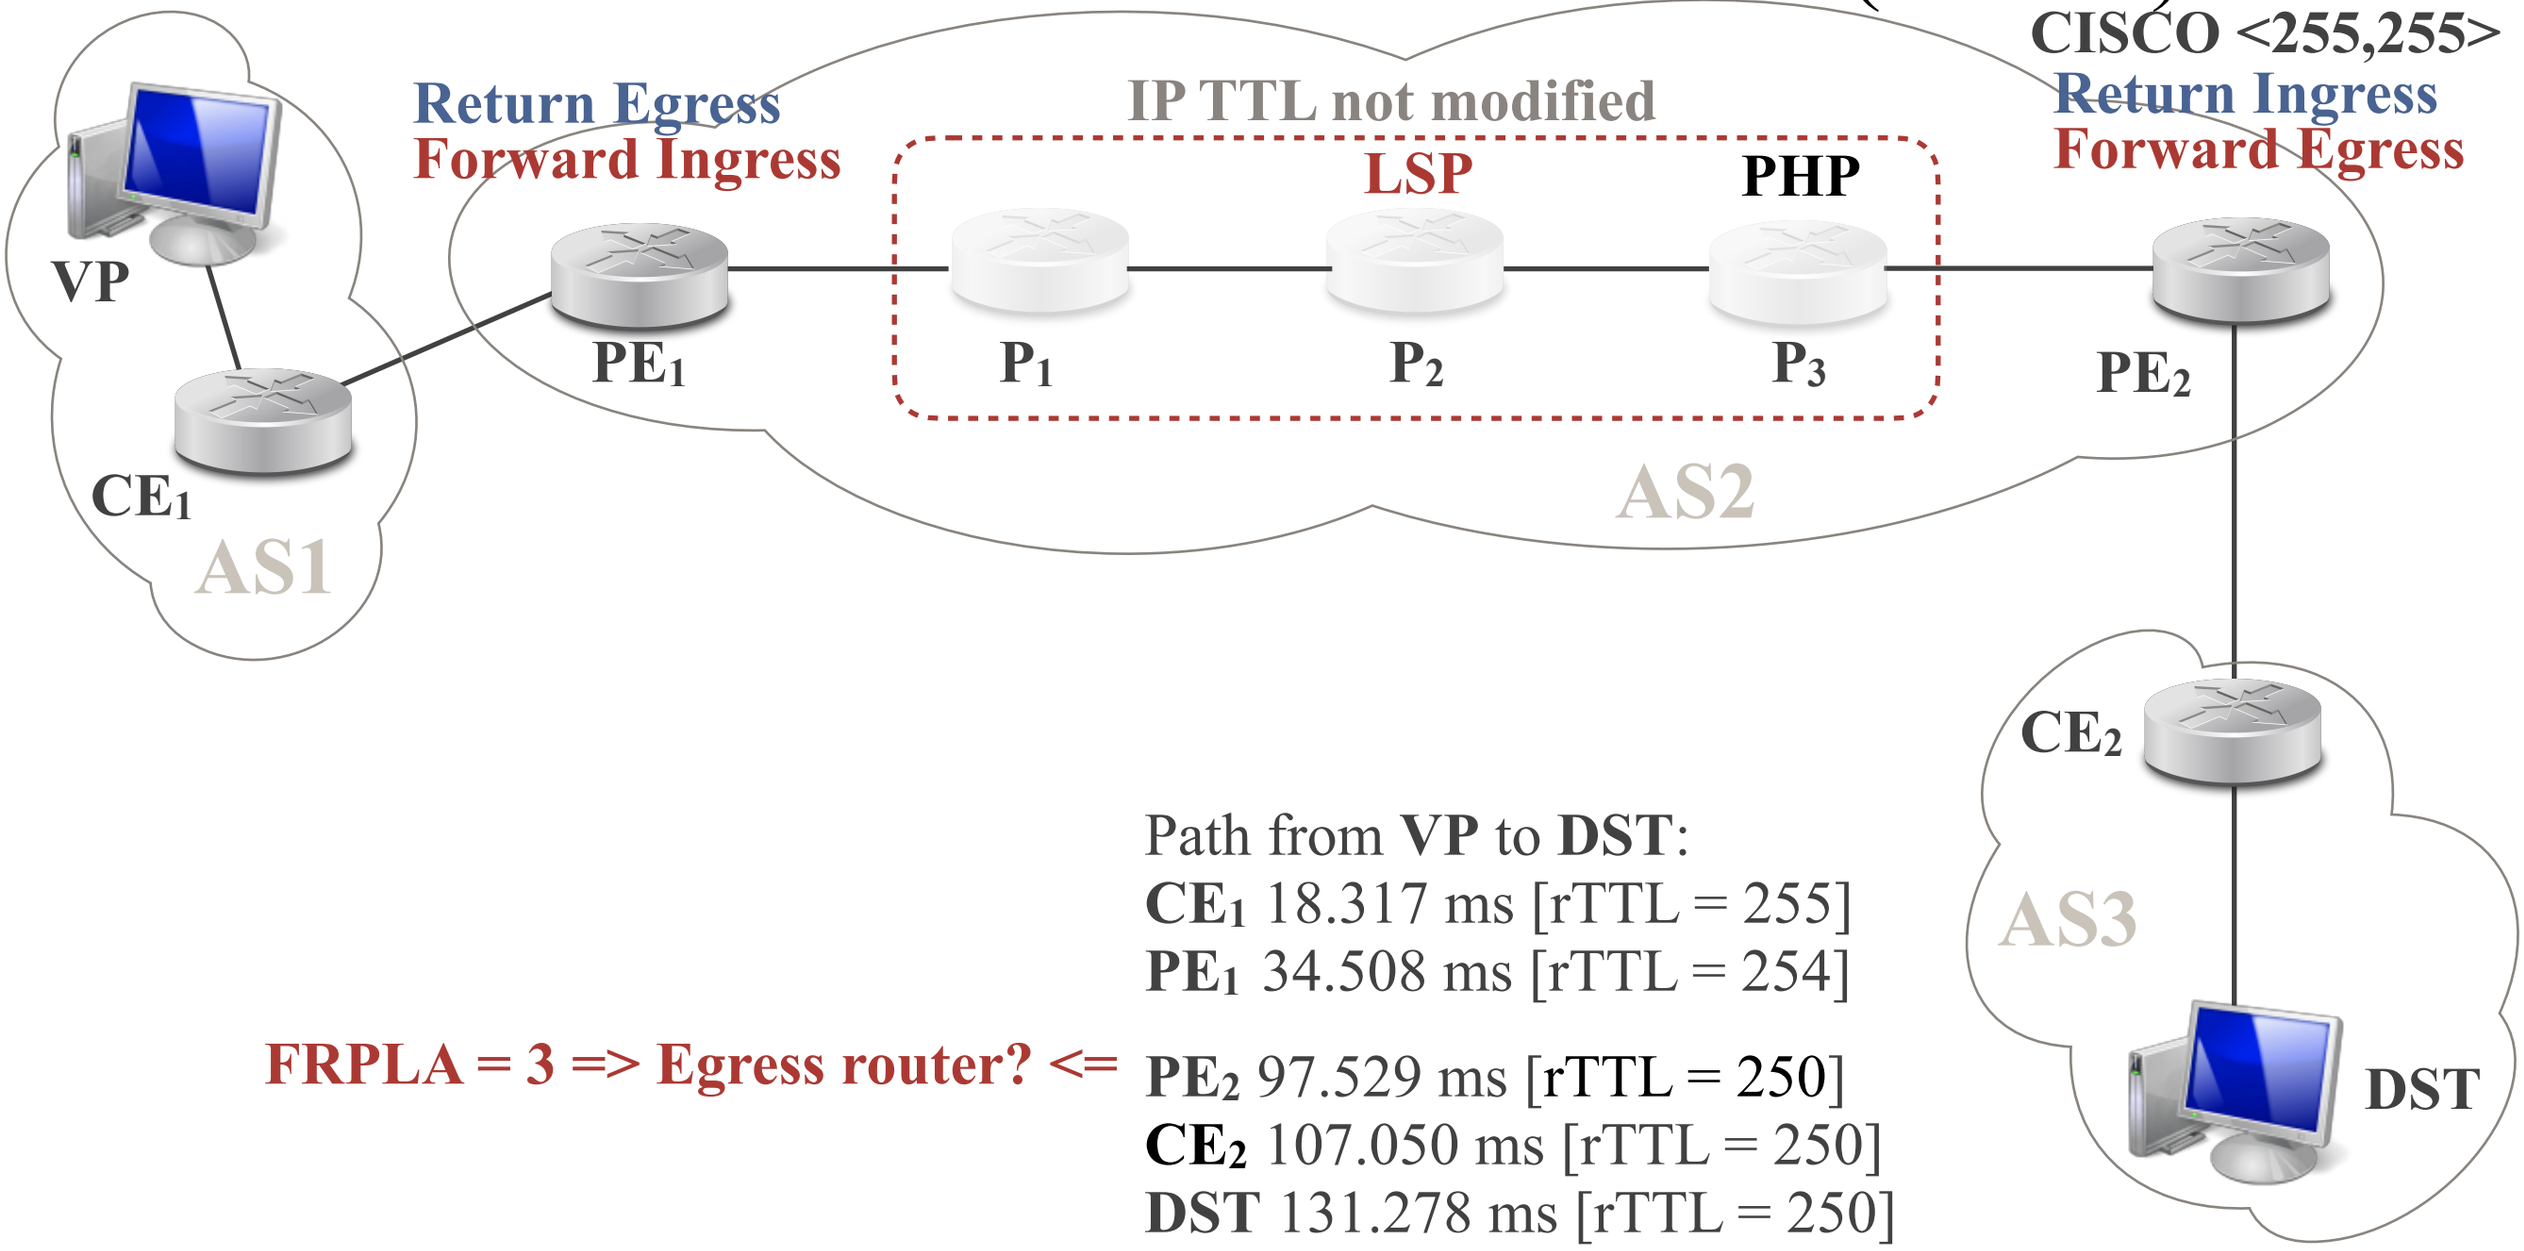

Path from **VP** to **DST**:
CE1 18.317 ms [rTTL = 255]
PE1 34.508 ms [rTTL = 254]

FRPLA = 3 => Egress router? <=

PE2 97.529 ms [rTTL = 250]
CE2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

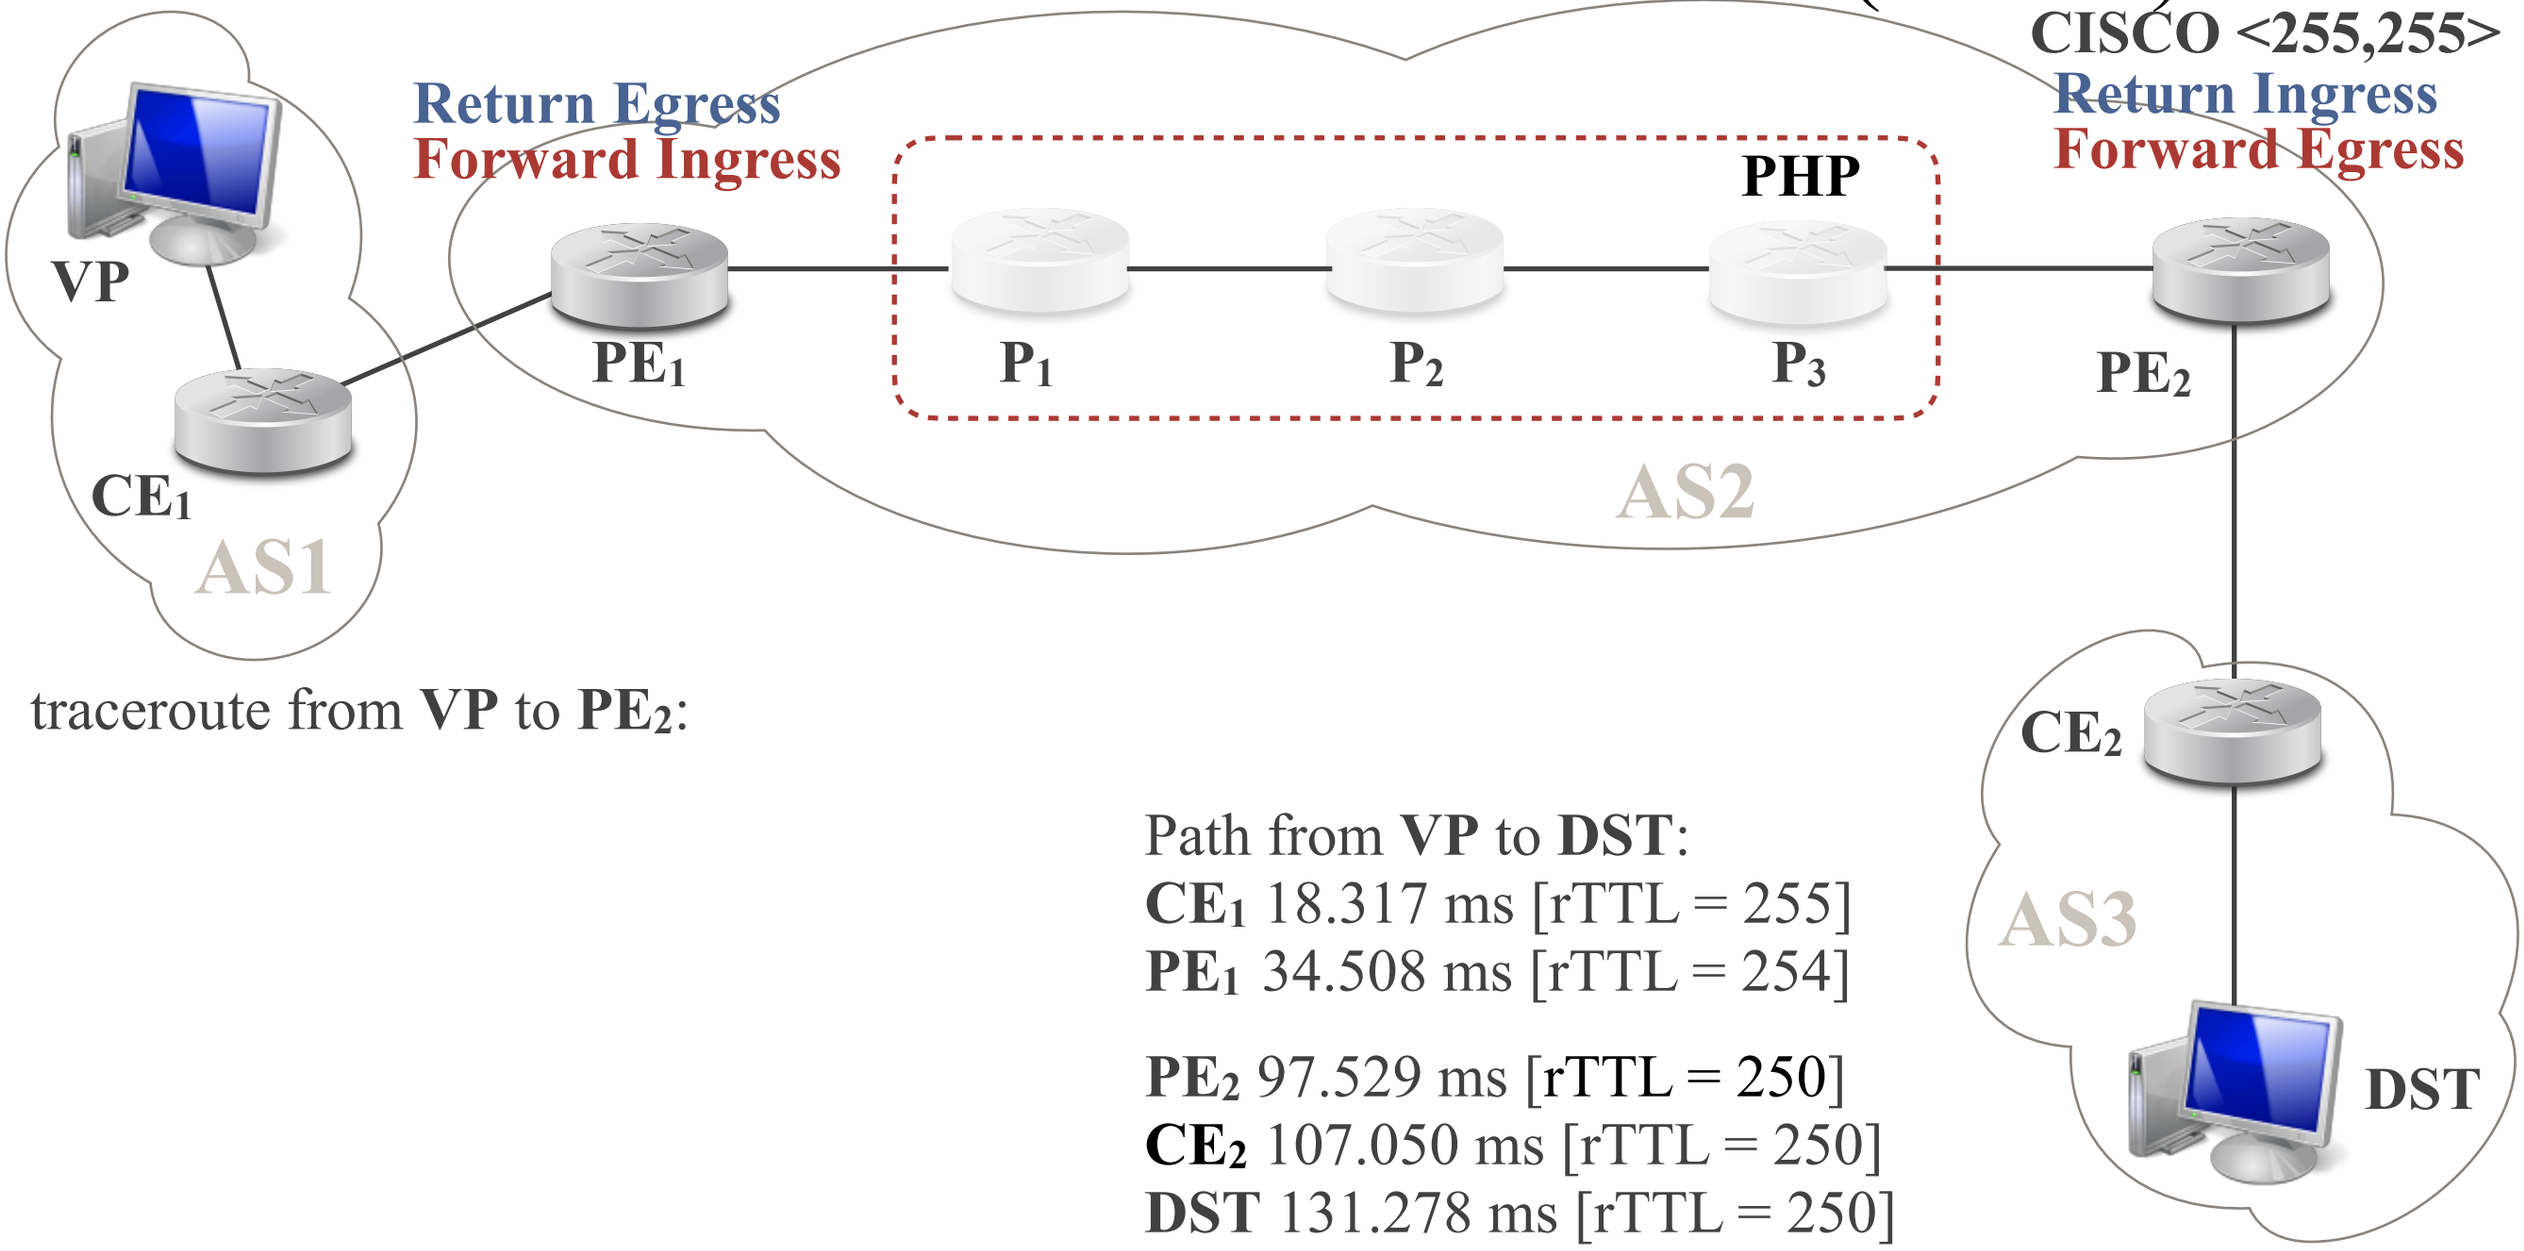

traceroute from **VP** to **PE$_2$**:

Path from **VP** to **DST**:
CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

PE$_2$ 97.529 ms [rTTL = 250]
CE$_2$ 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

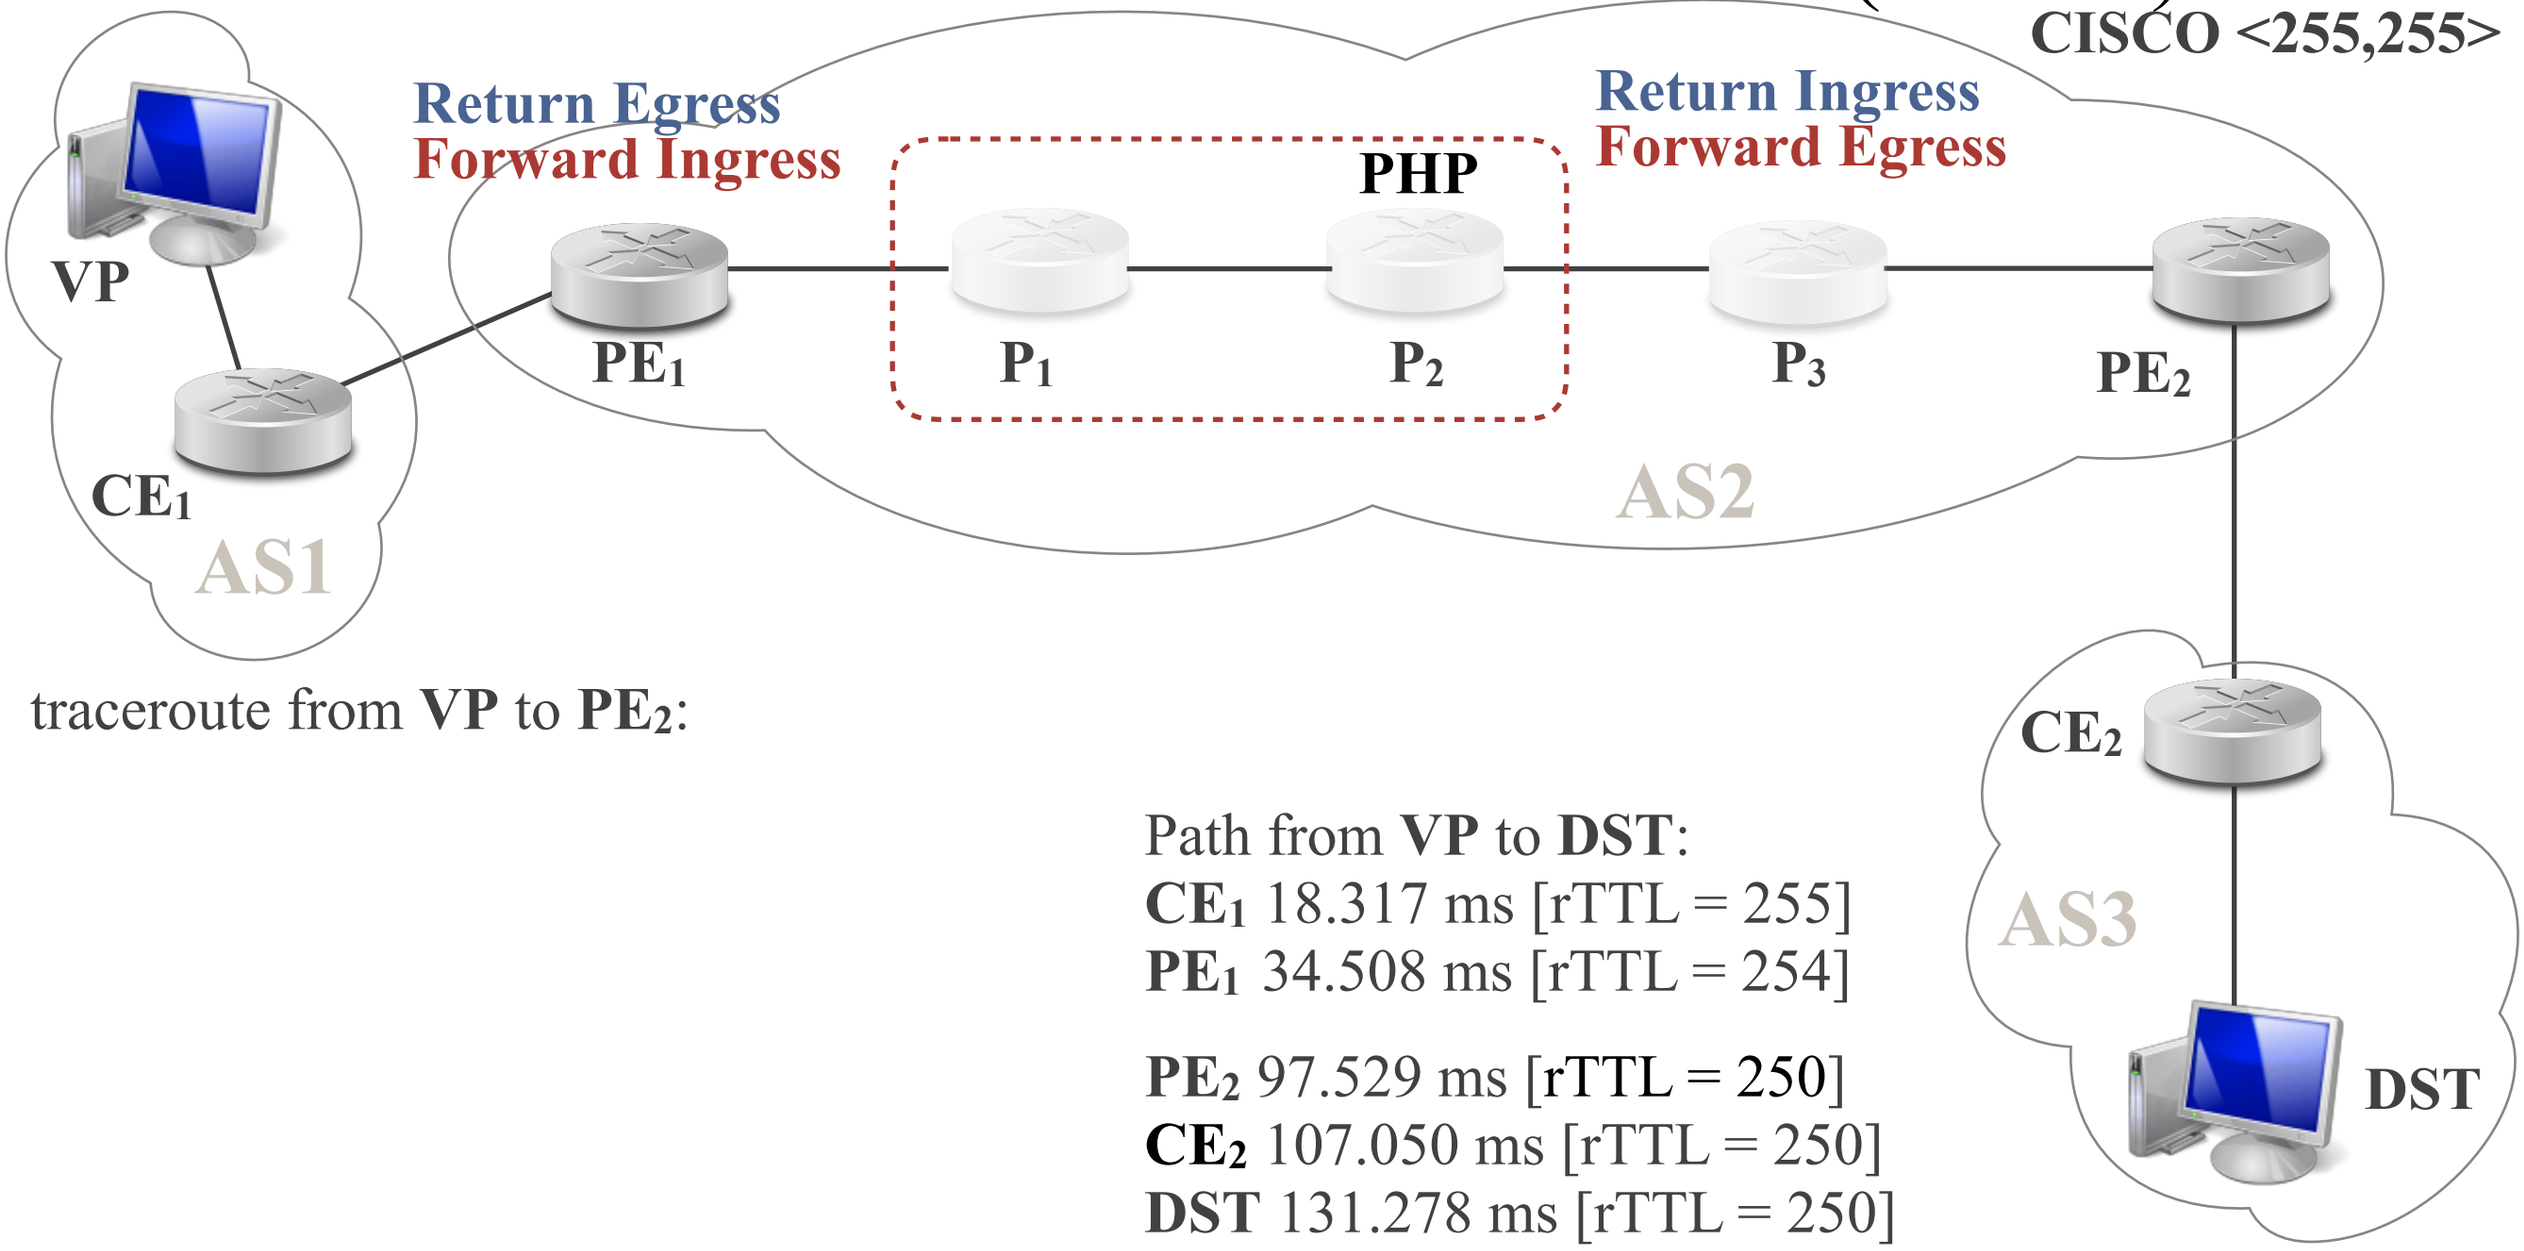

traceroute from **VP** to **PE$_2$**:

Path from **VP** to **DST**:
CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

PE$_2$ 97.529 ms [rTTL = 250]
CE$_2$ 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

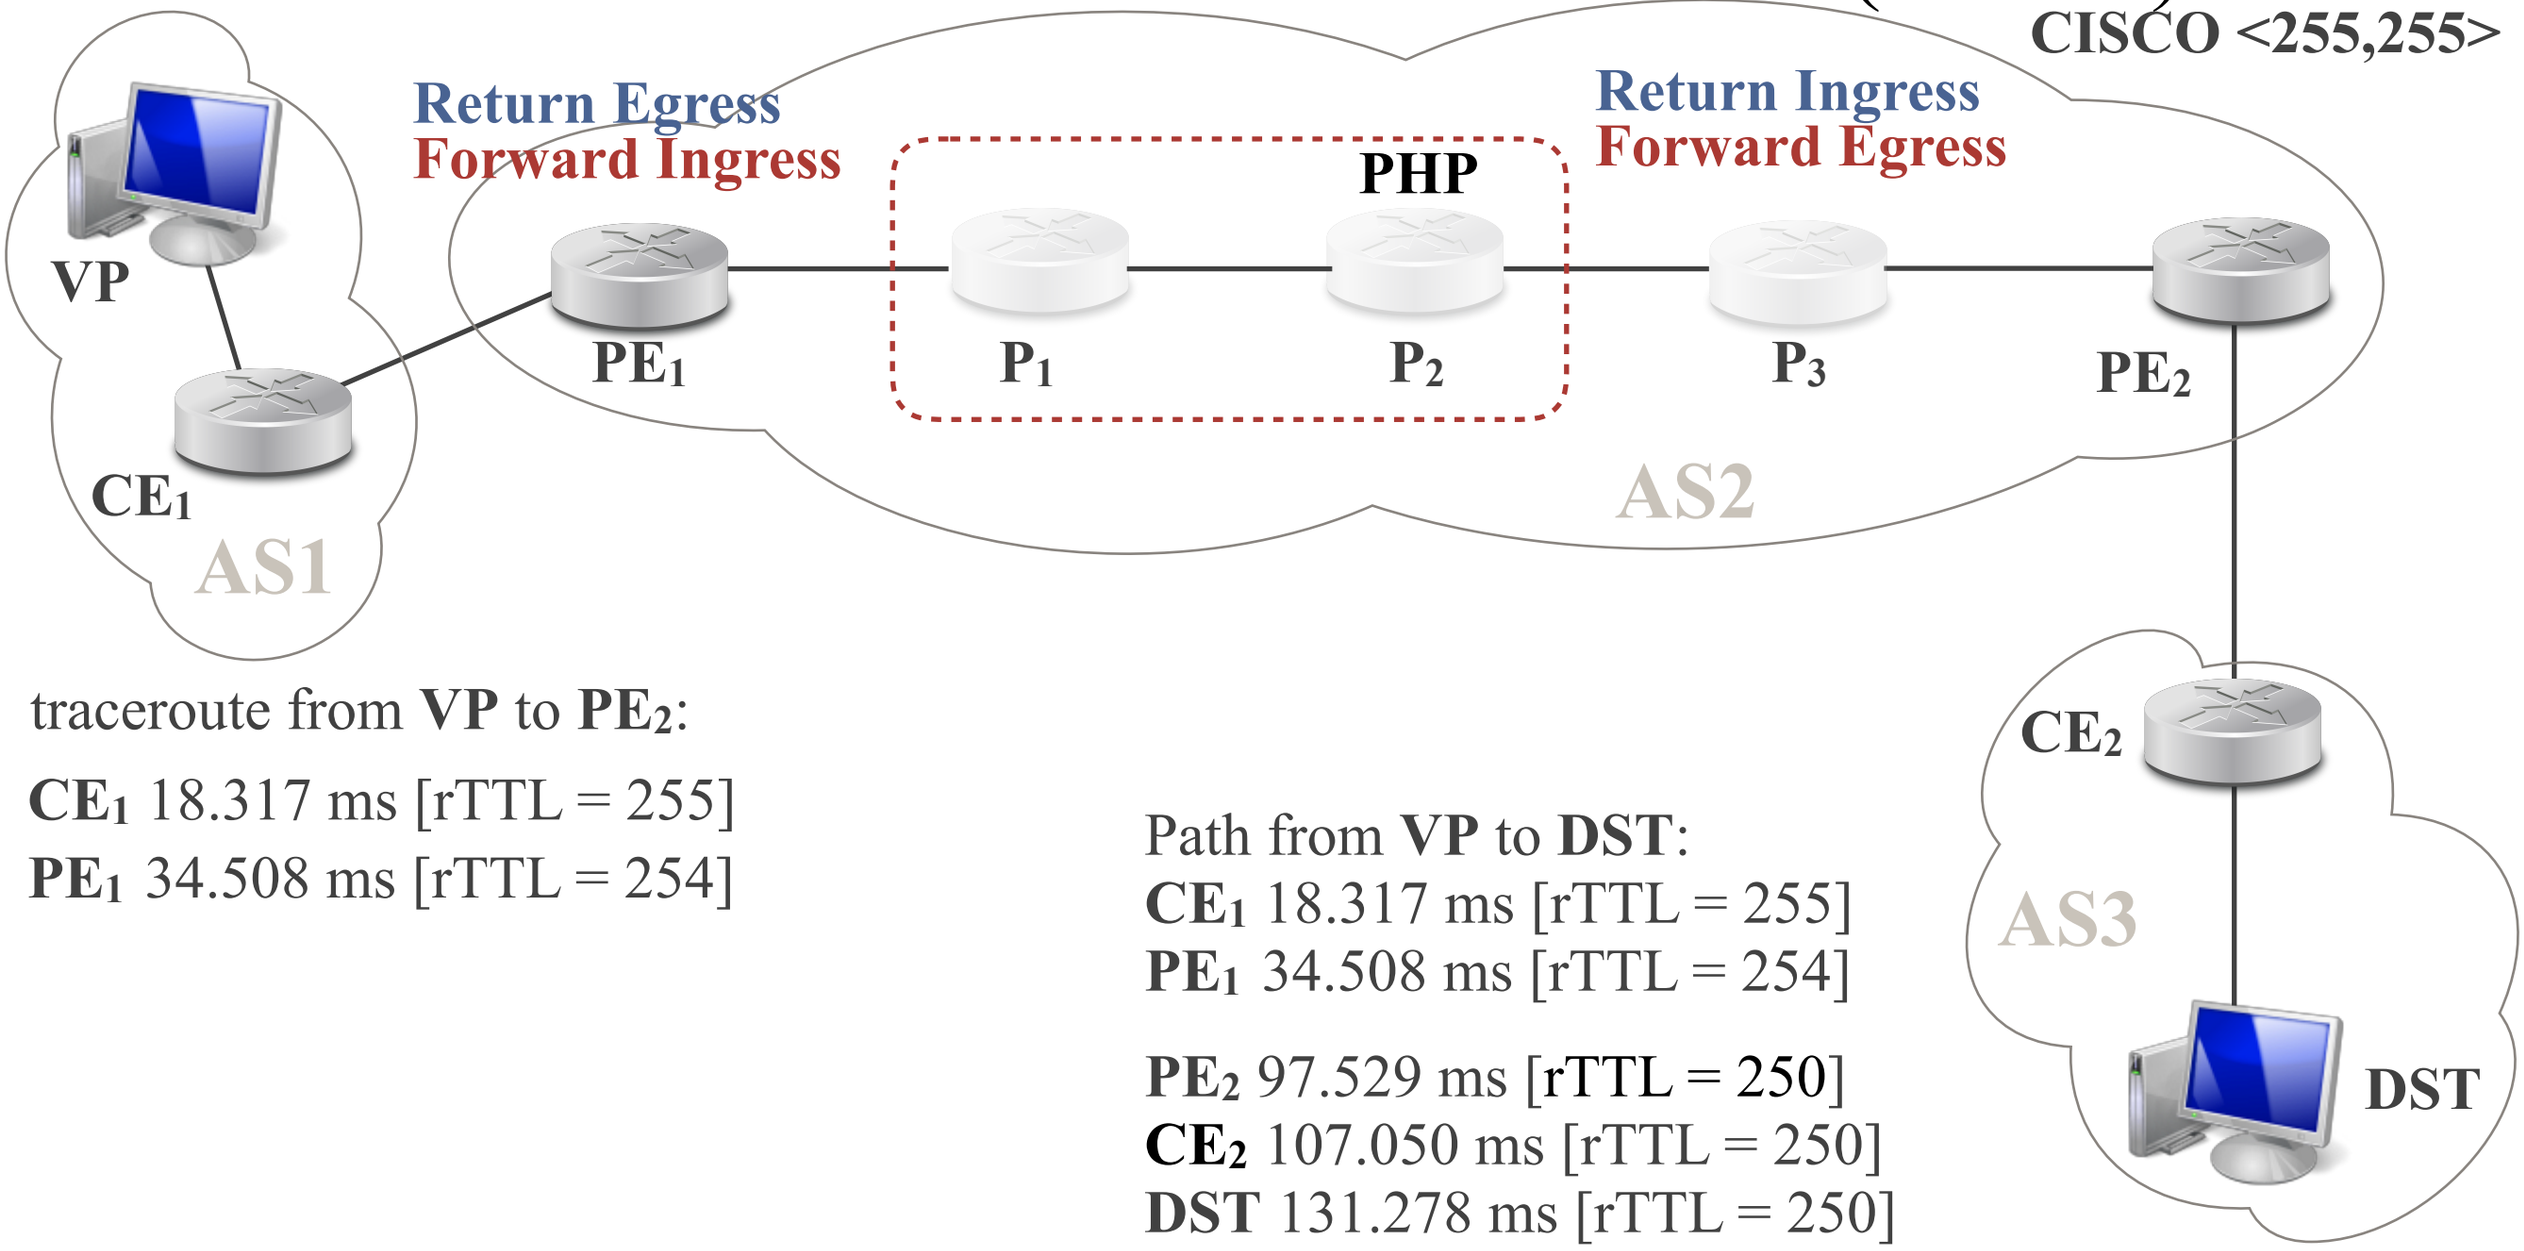

traceroute from **VP** to **PE$_2$**:

CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

Path from **VP** to **DST**:
CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

PE$_2$ 97.529 ms [rTTL = 250]
CE$_2$ 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

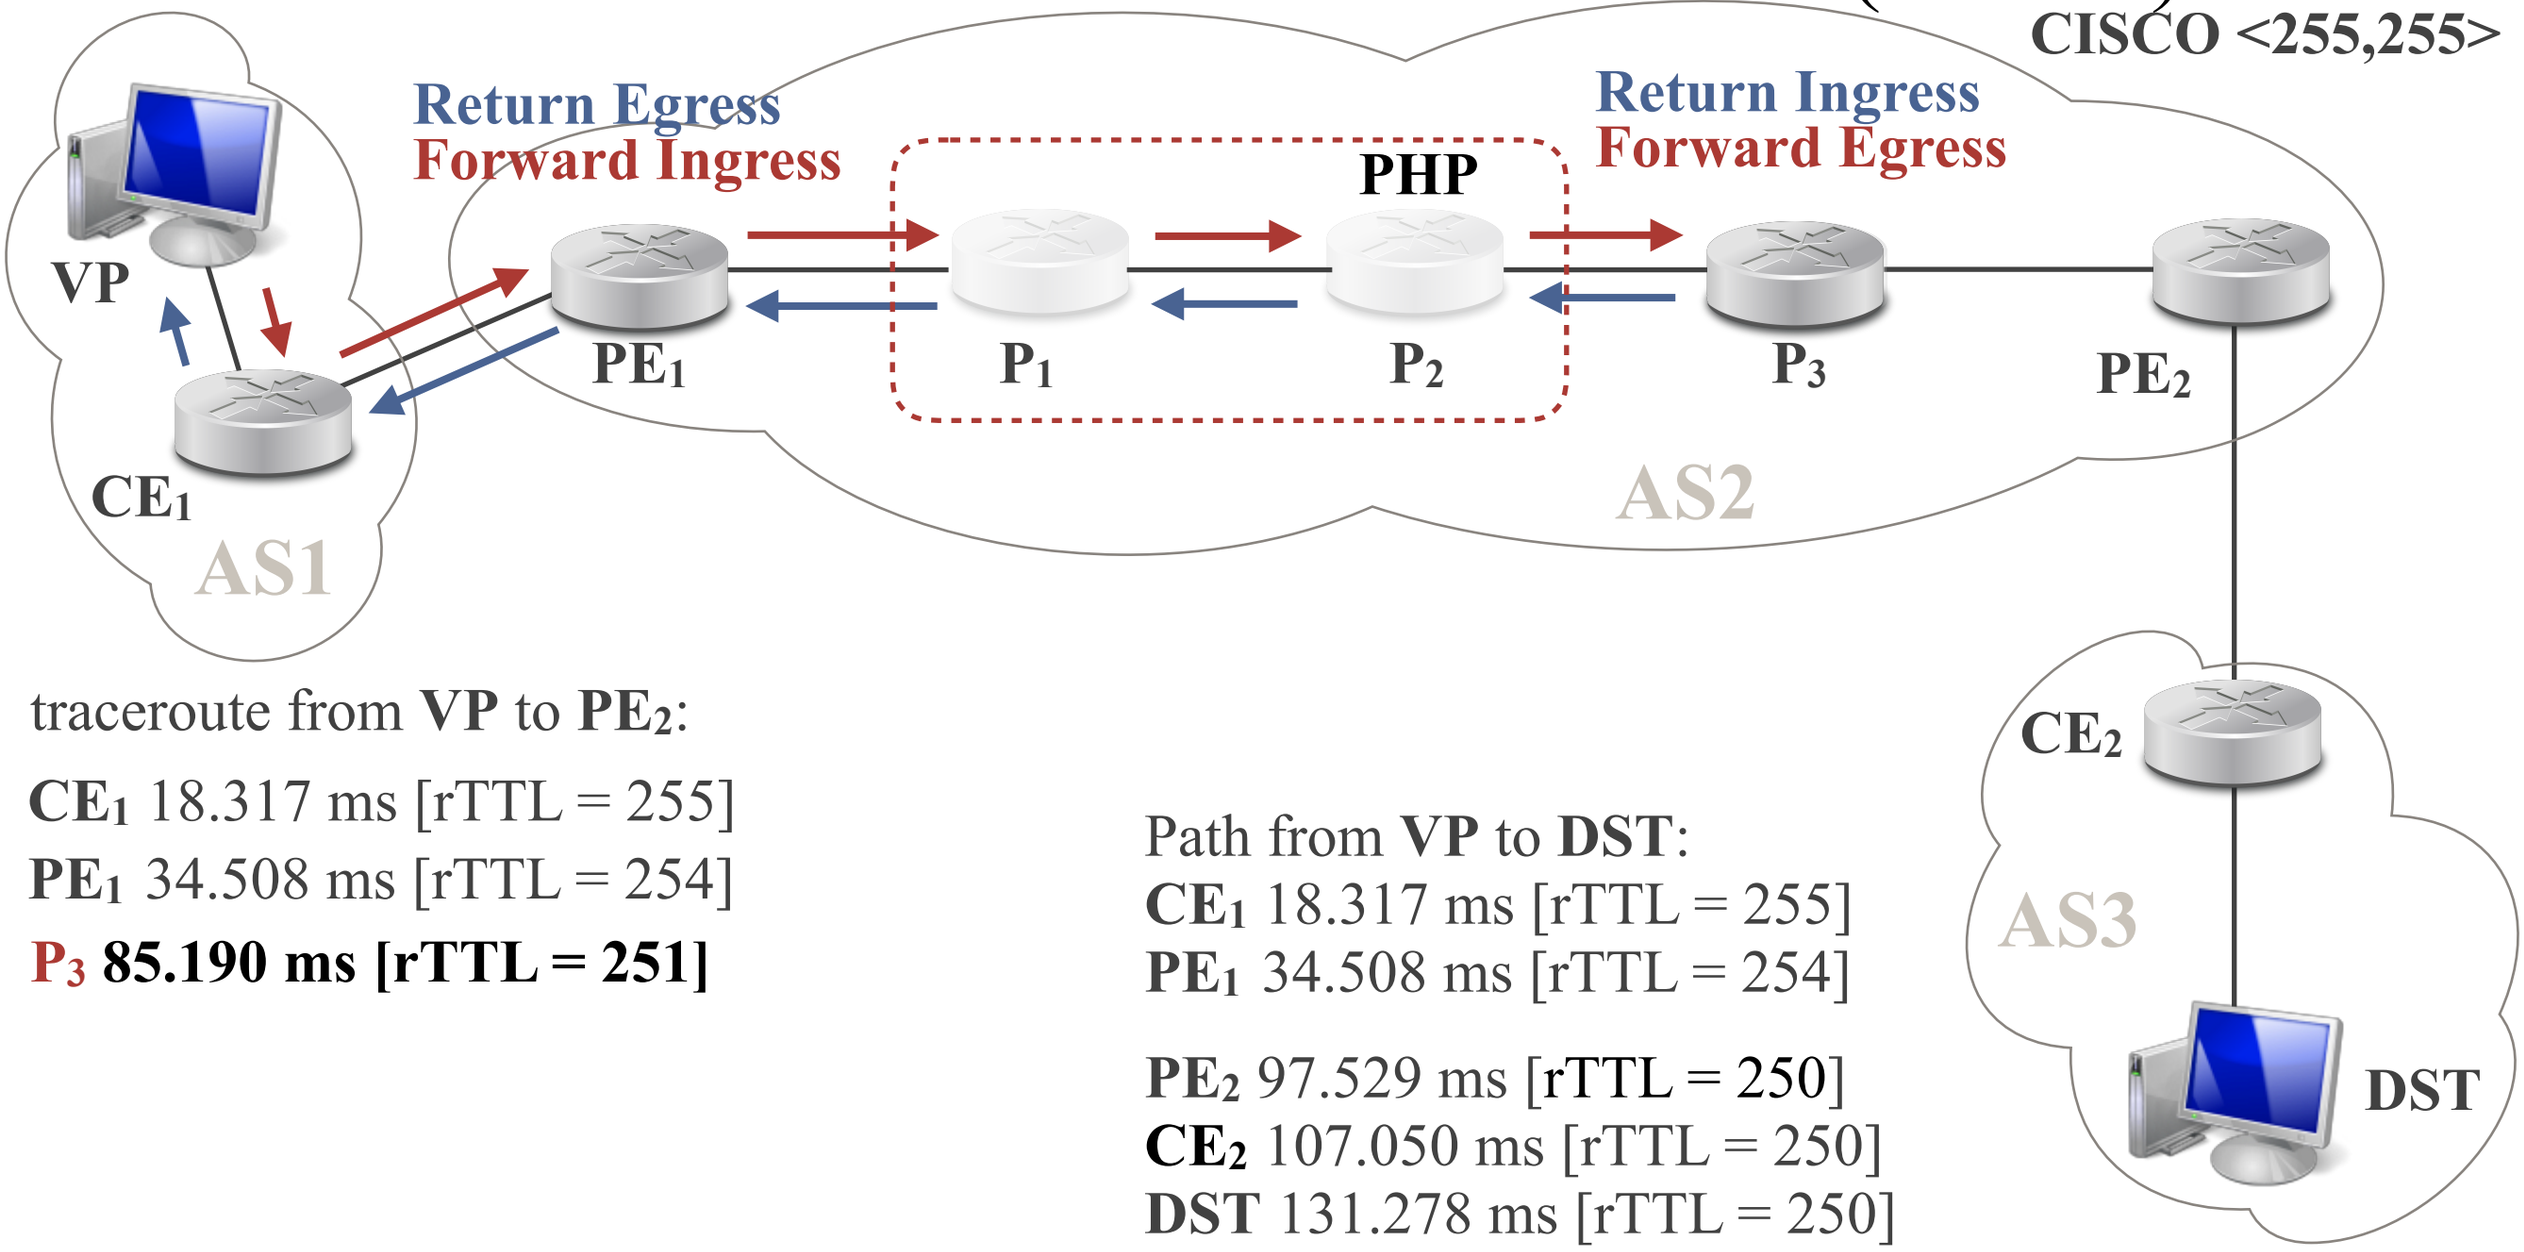

traceroute from **VP** to **PE$_2$**:

CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]
P$_3$ 85.190 ms [rTTL = 251]

Path from **VP** to **DST**:
CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

PE$_2$ 97.529 ms [rTTL = 250]
CE$_2$ 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

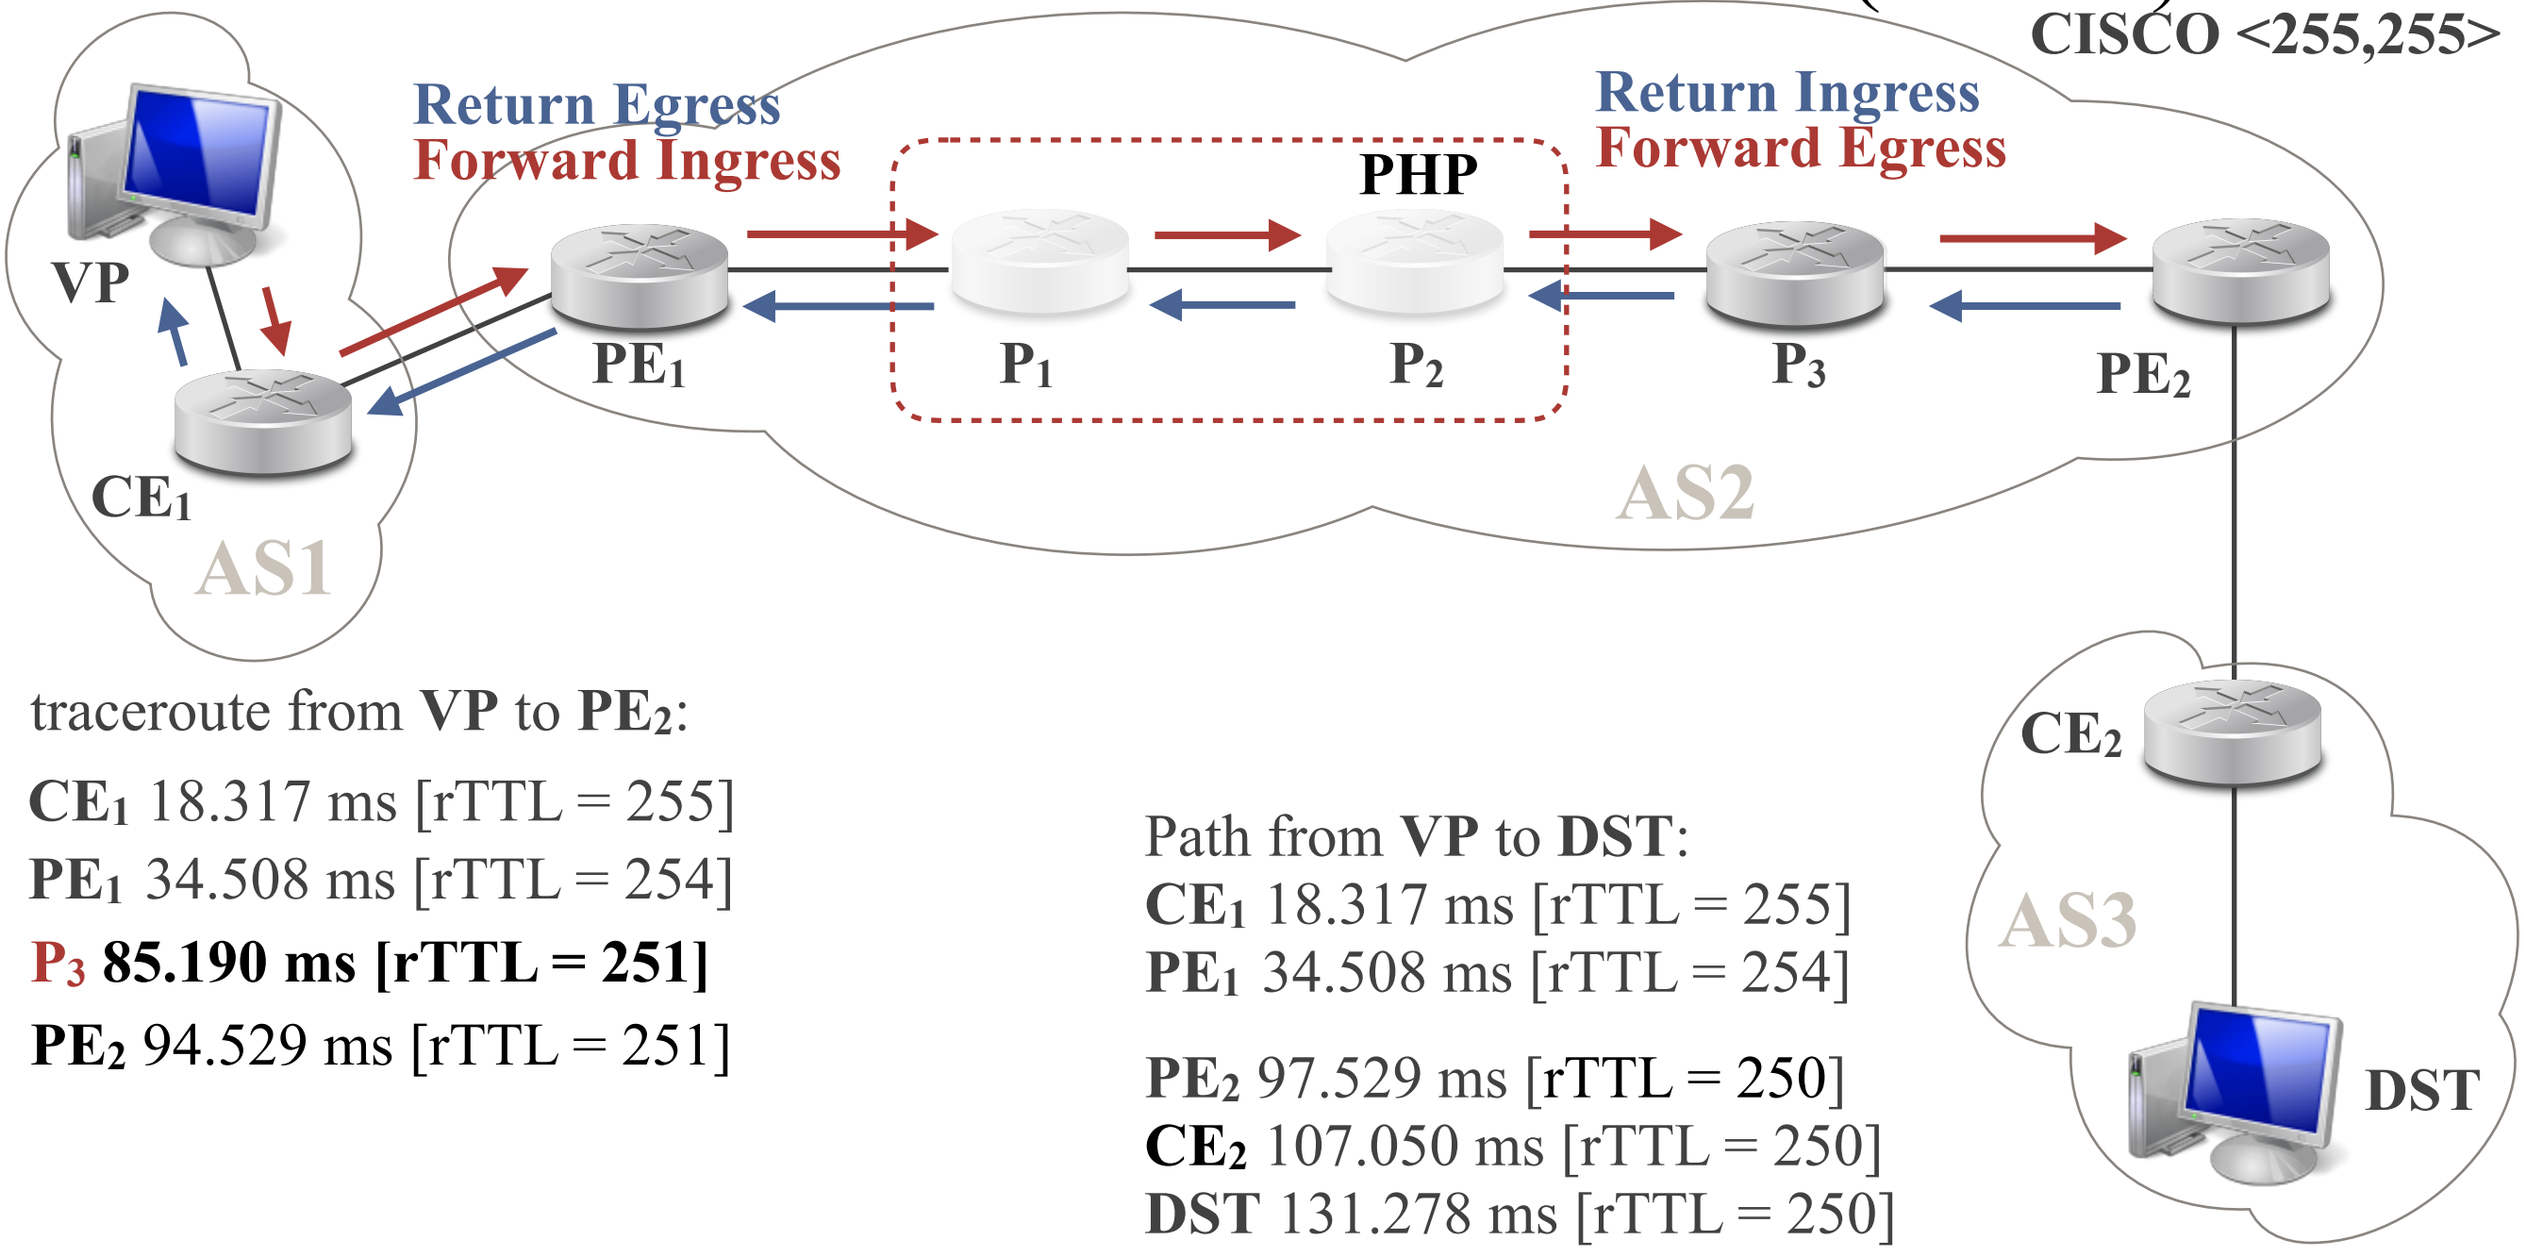

traceroute from **VP** to **PE$_2$**:

CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]
P$_3$ 85.190 ms [rTTL = 251]
PE$_2$ 94.529 ms [rTTL = 251]

Path from **VP** to **DST**:
CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

PE$_2$ 97.529 ms [rTTL = 250]
CE$_2$ 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

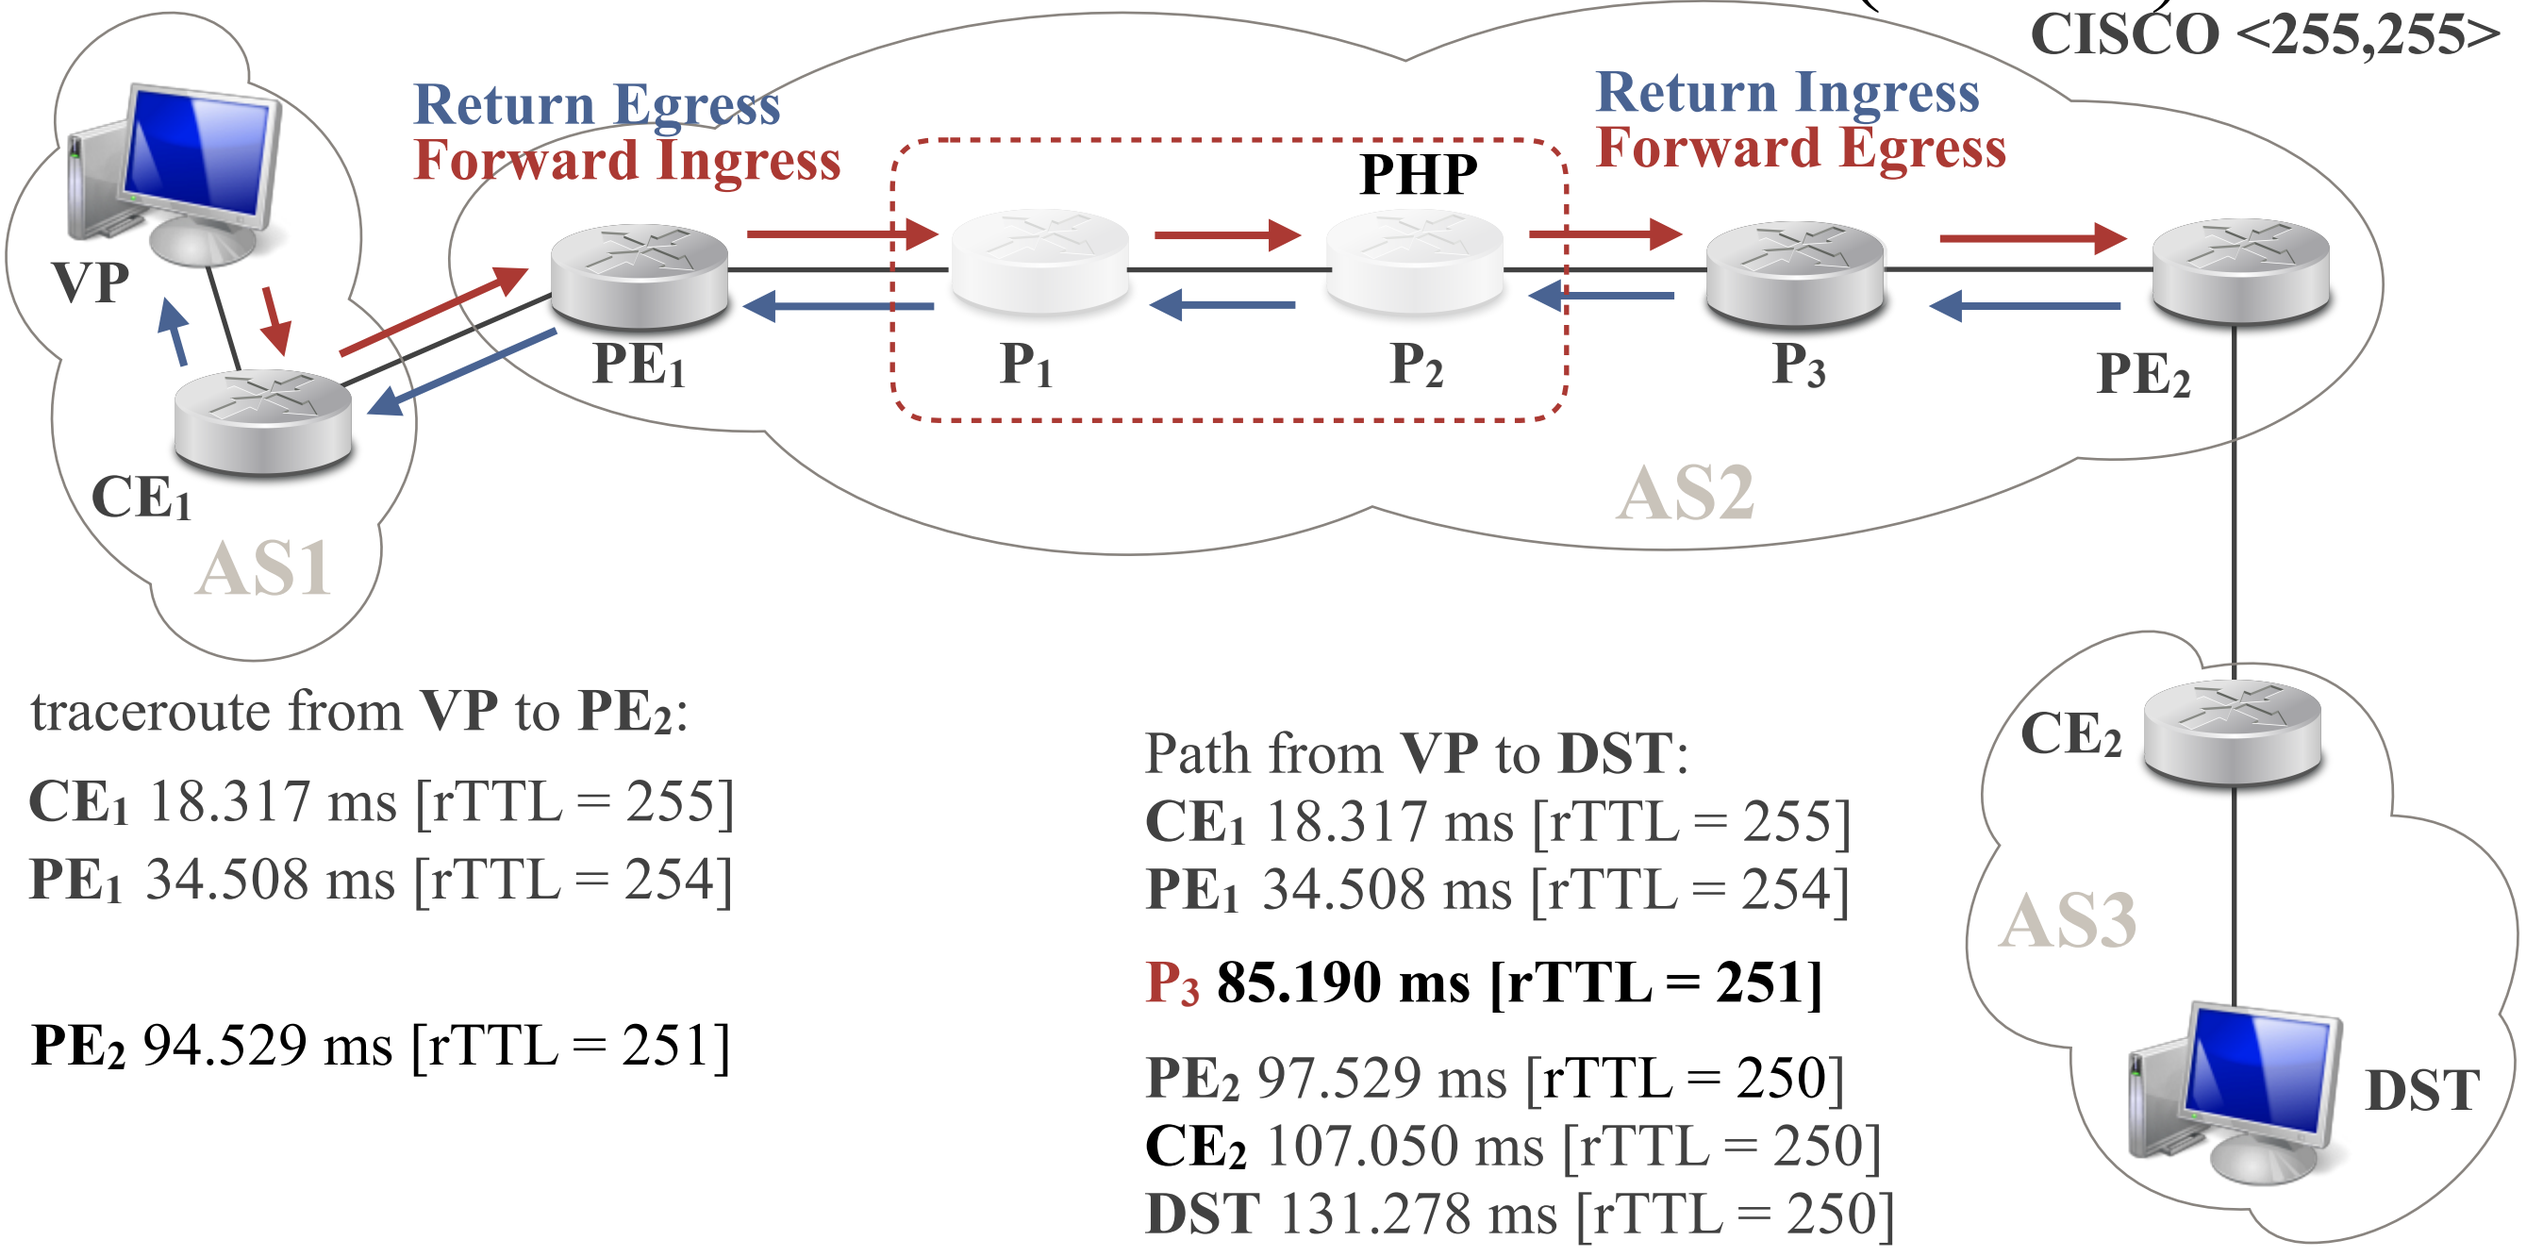

traceroute from **VP** to **PE_2**:

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

PE_2 94.529 ms [rTTL = 251]

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

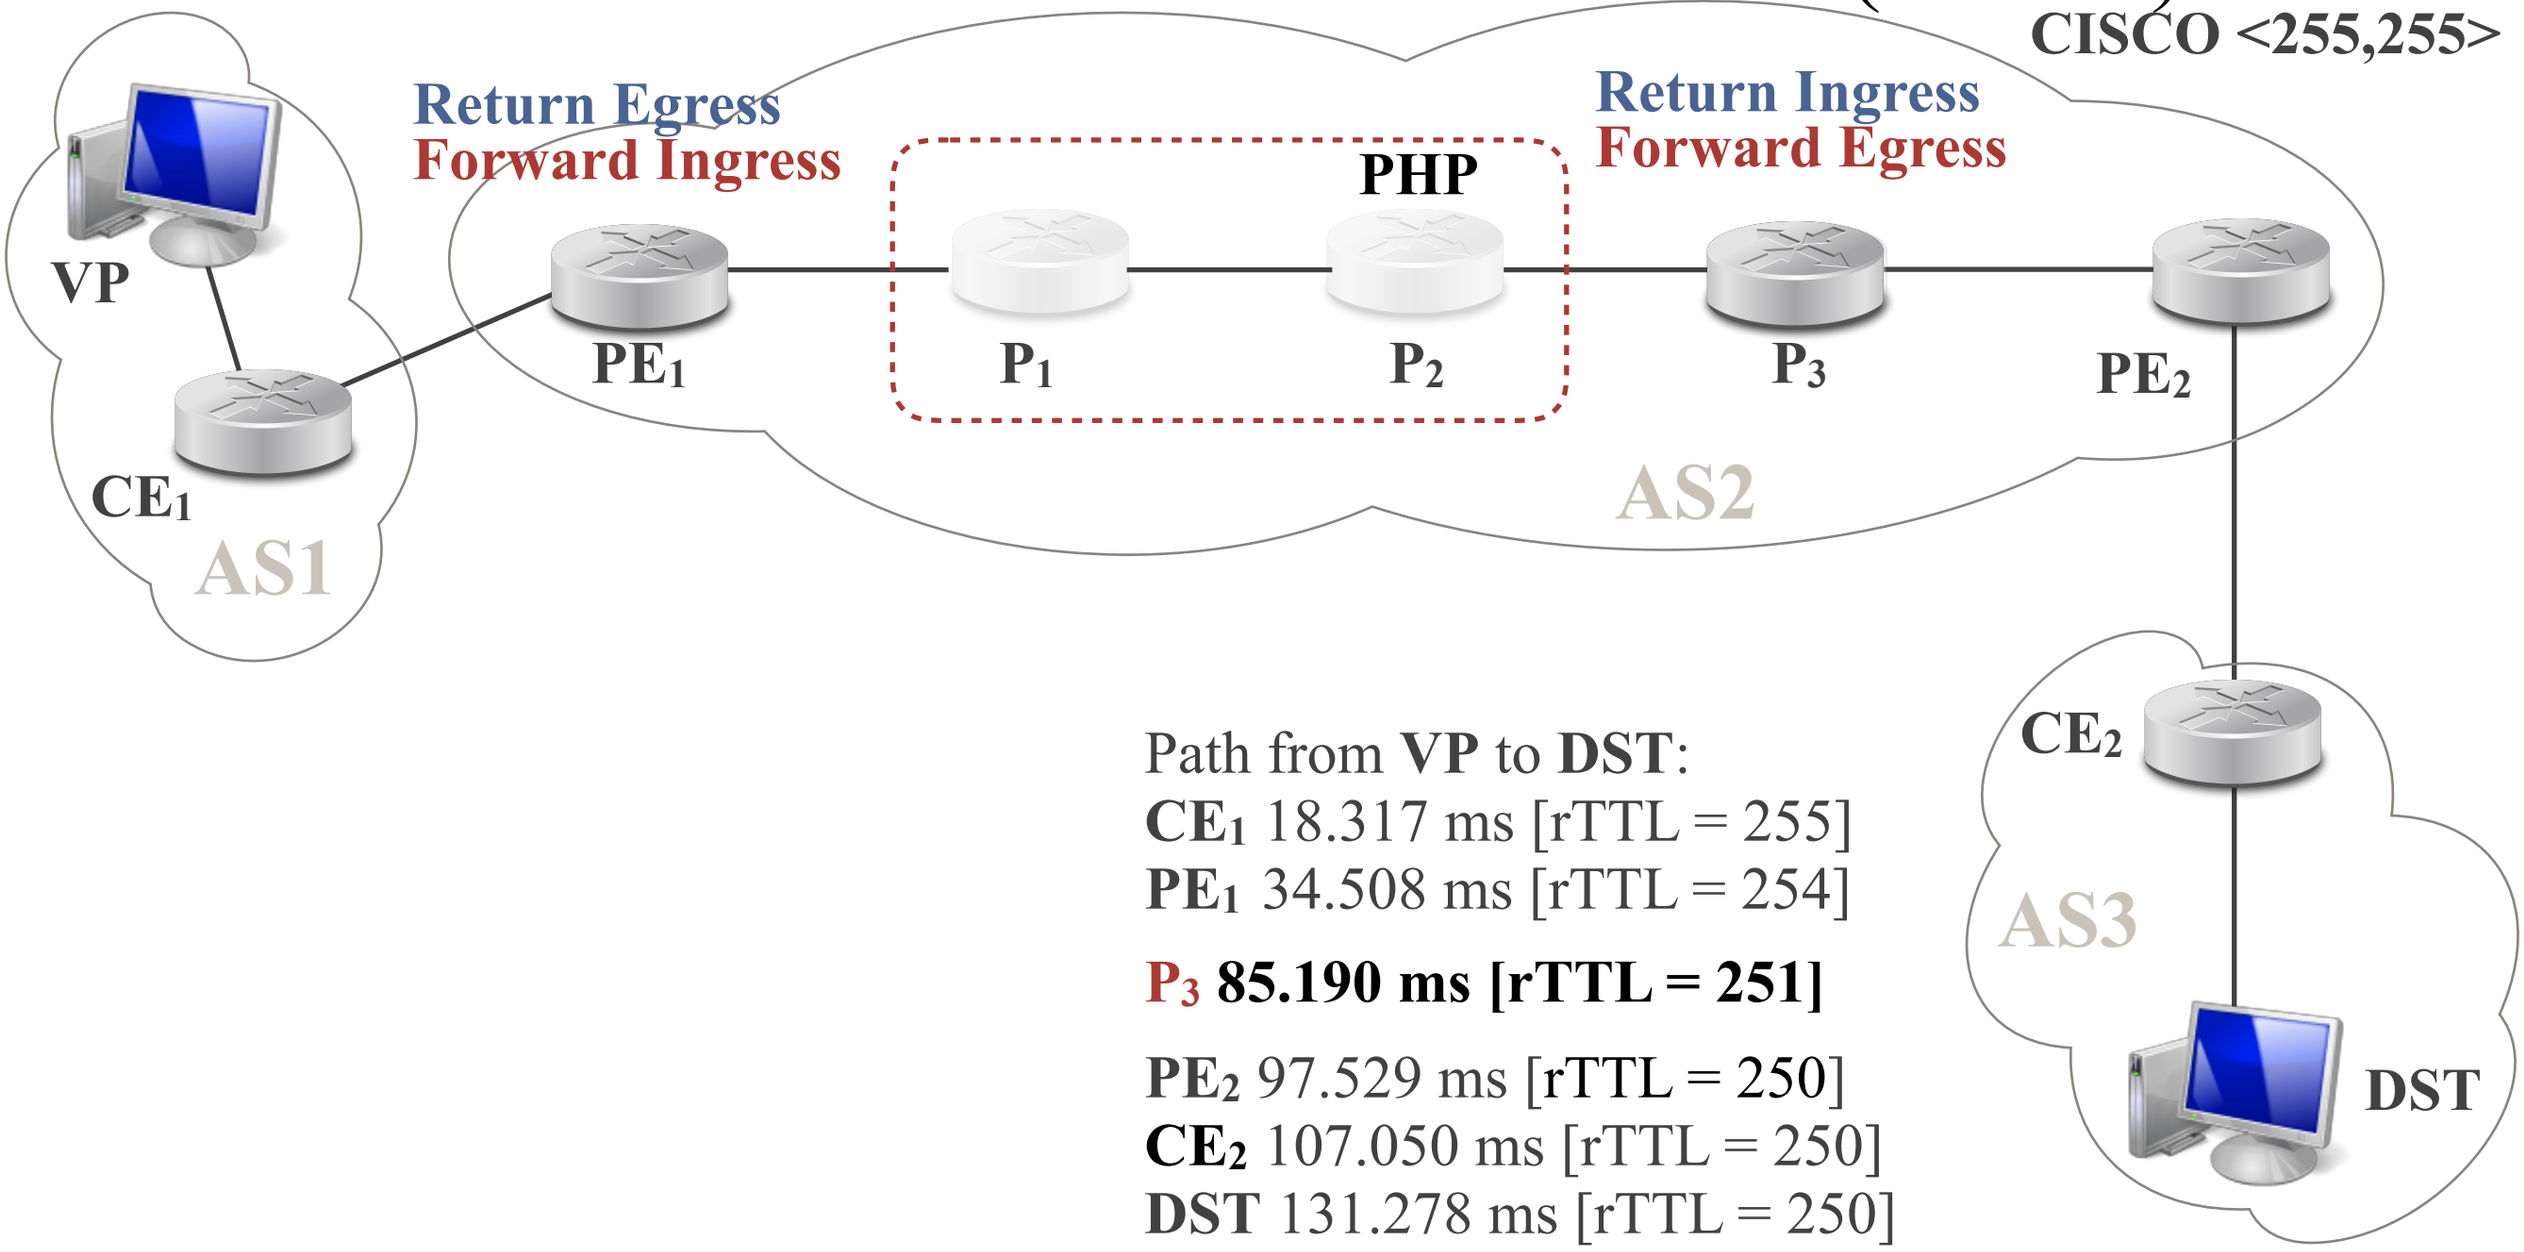

Path from **VP** to **DST**:
CE1 18.317 ms [rTTL = 255]
PE1 34.508 ms [rTTL = 254]

P3 85.190 ms [rTTL = 251]

PE2 97.529 ms [rTTL = 250]
CE2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

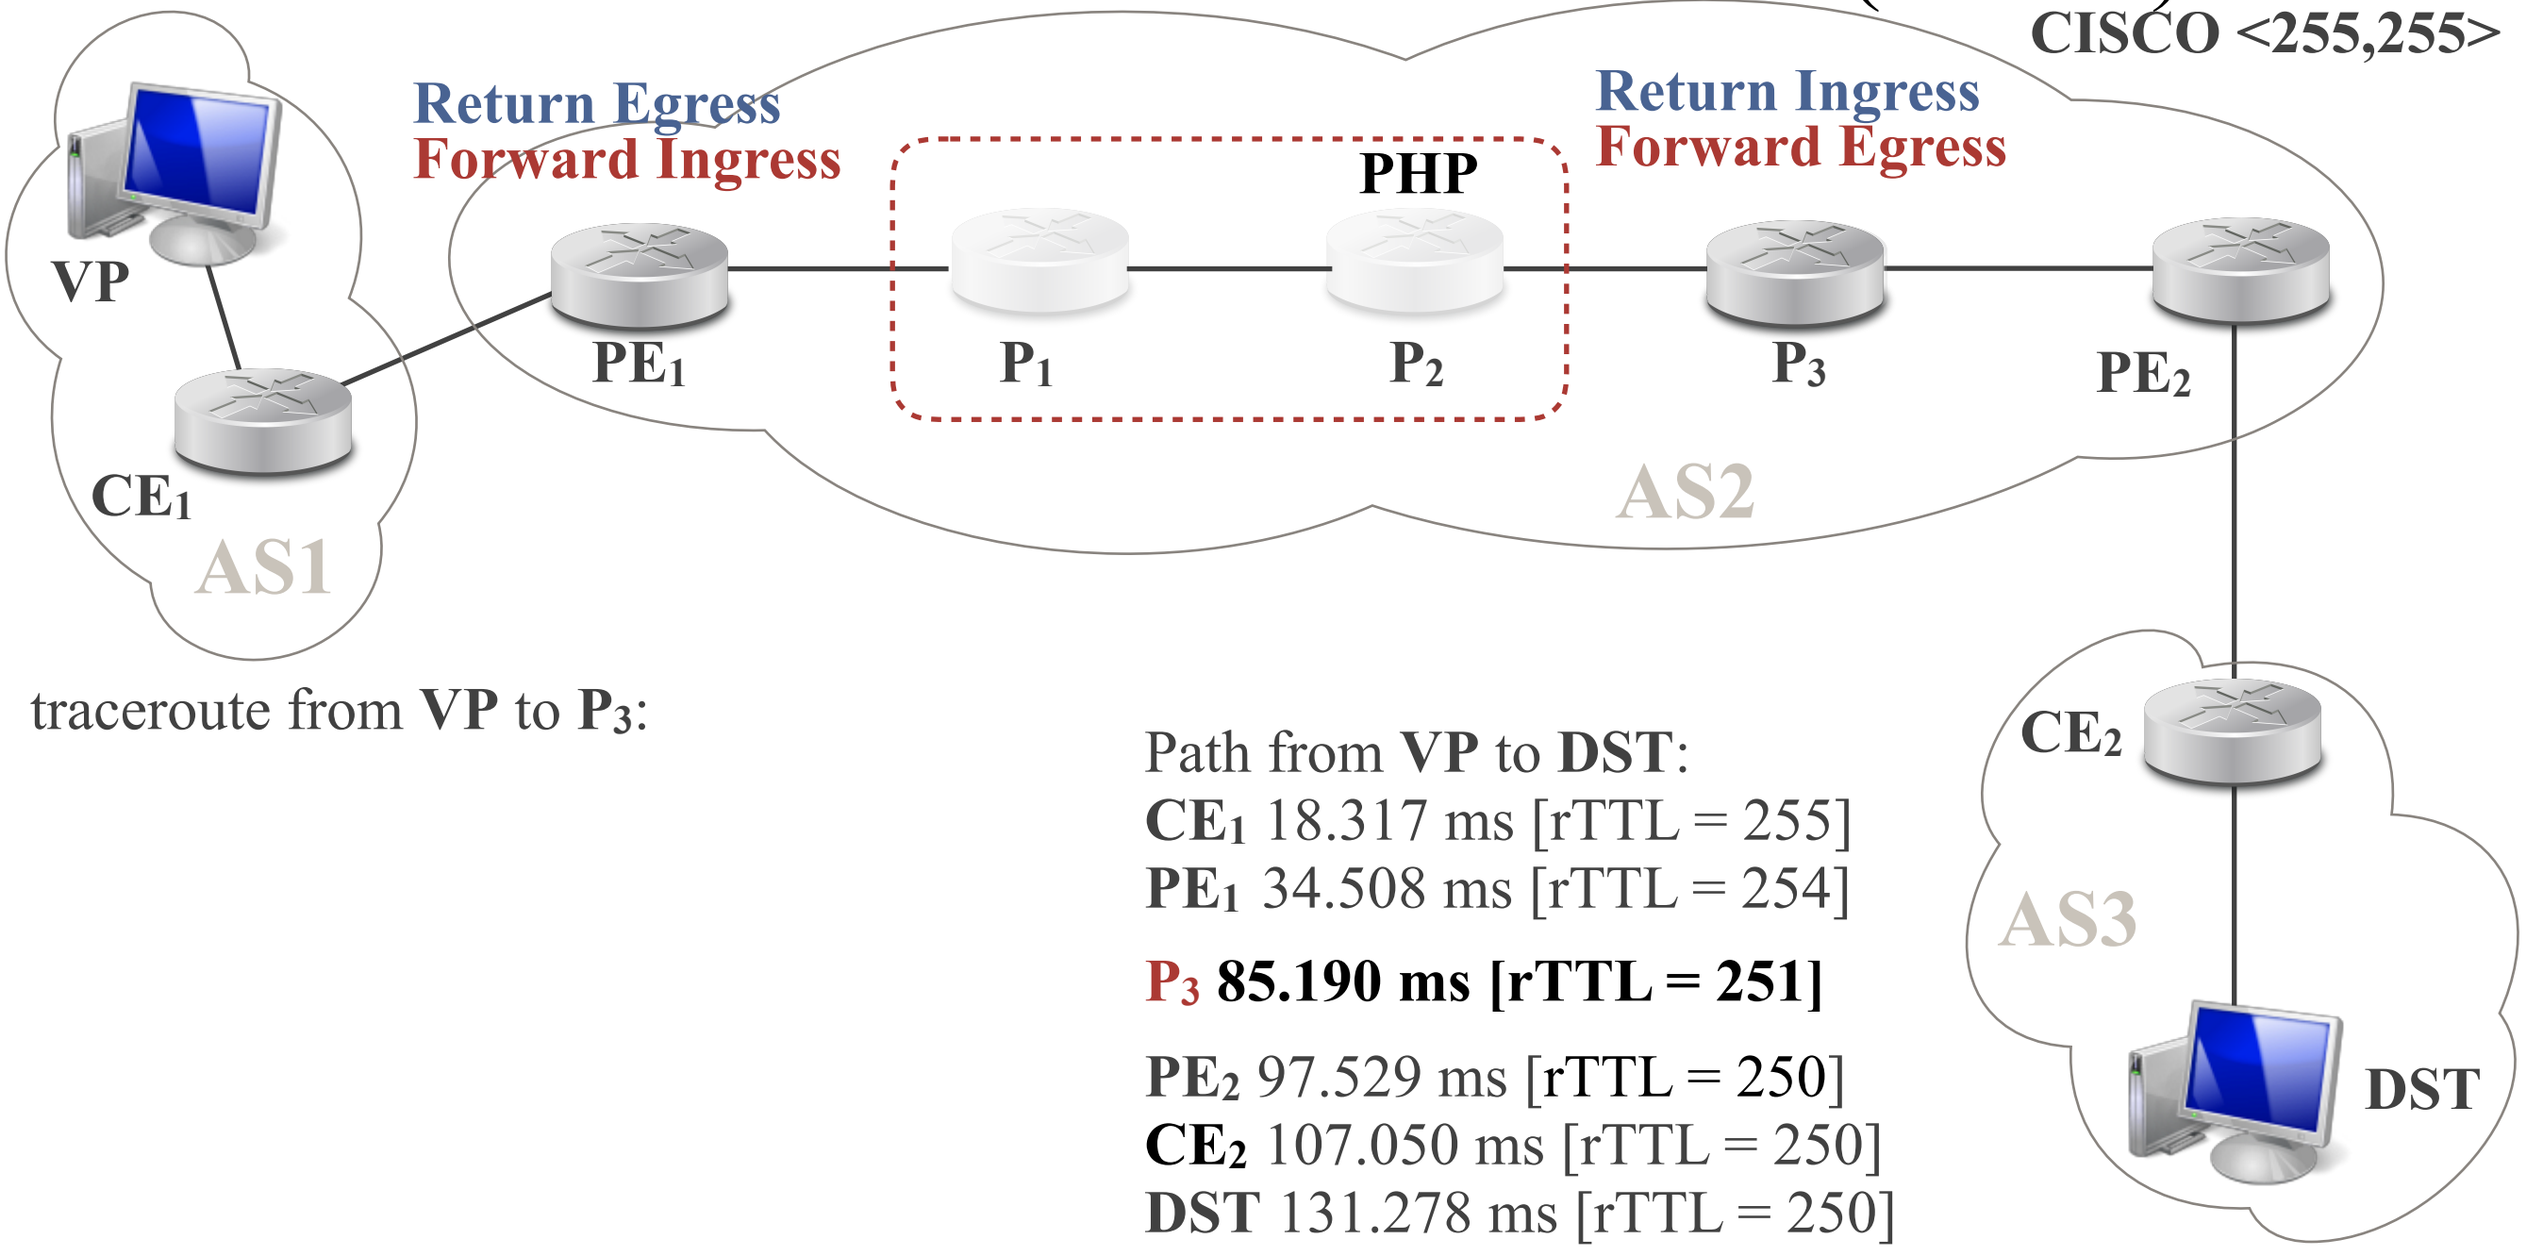

traceroute from **VP** to **P_3**:

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

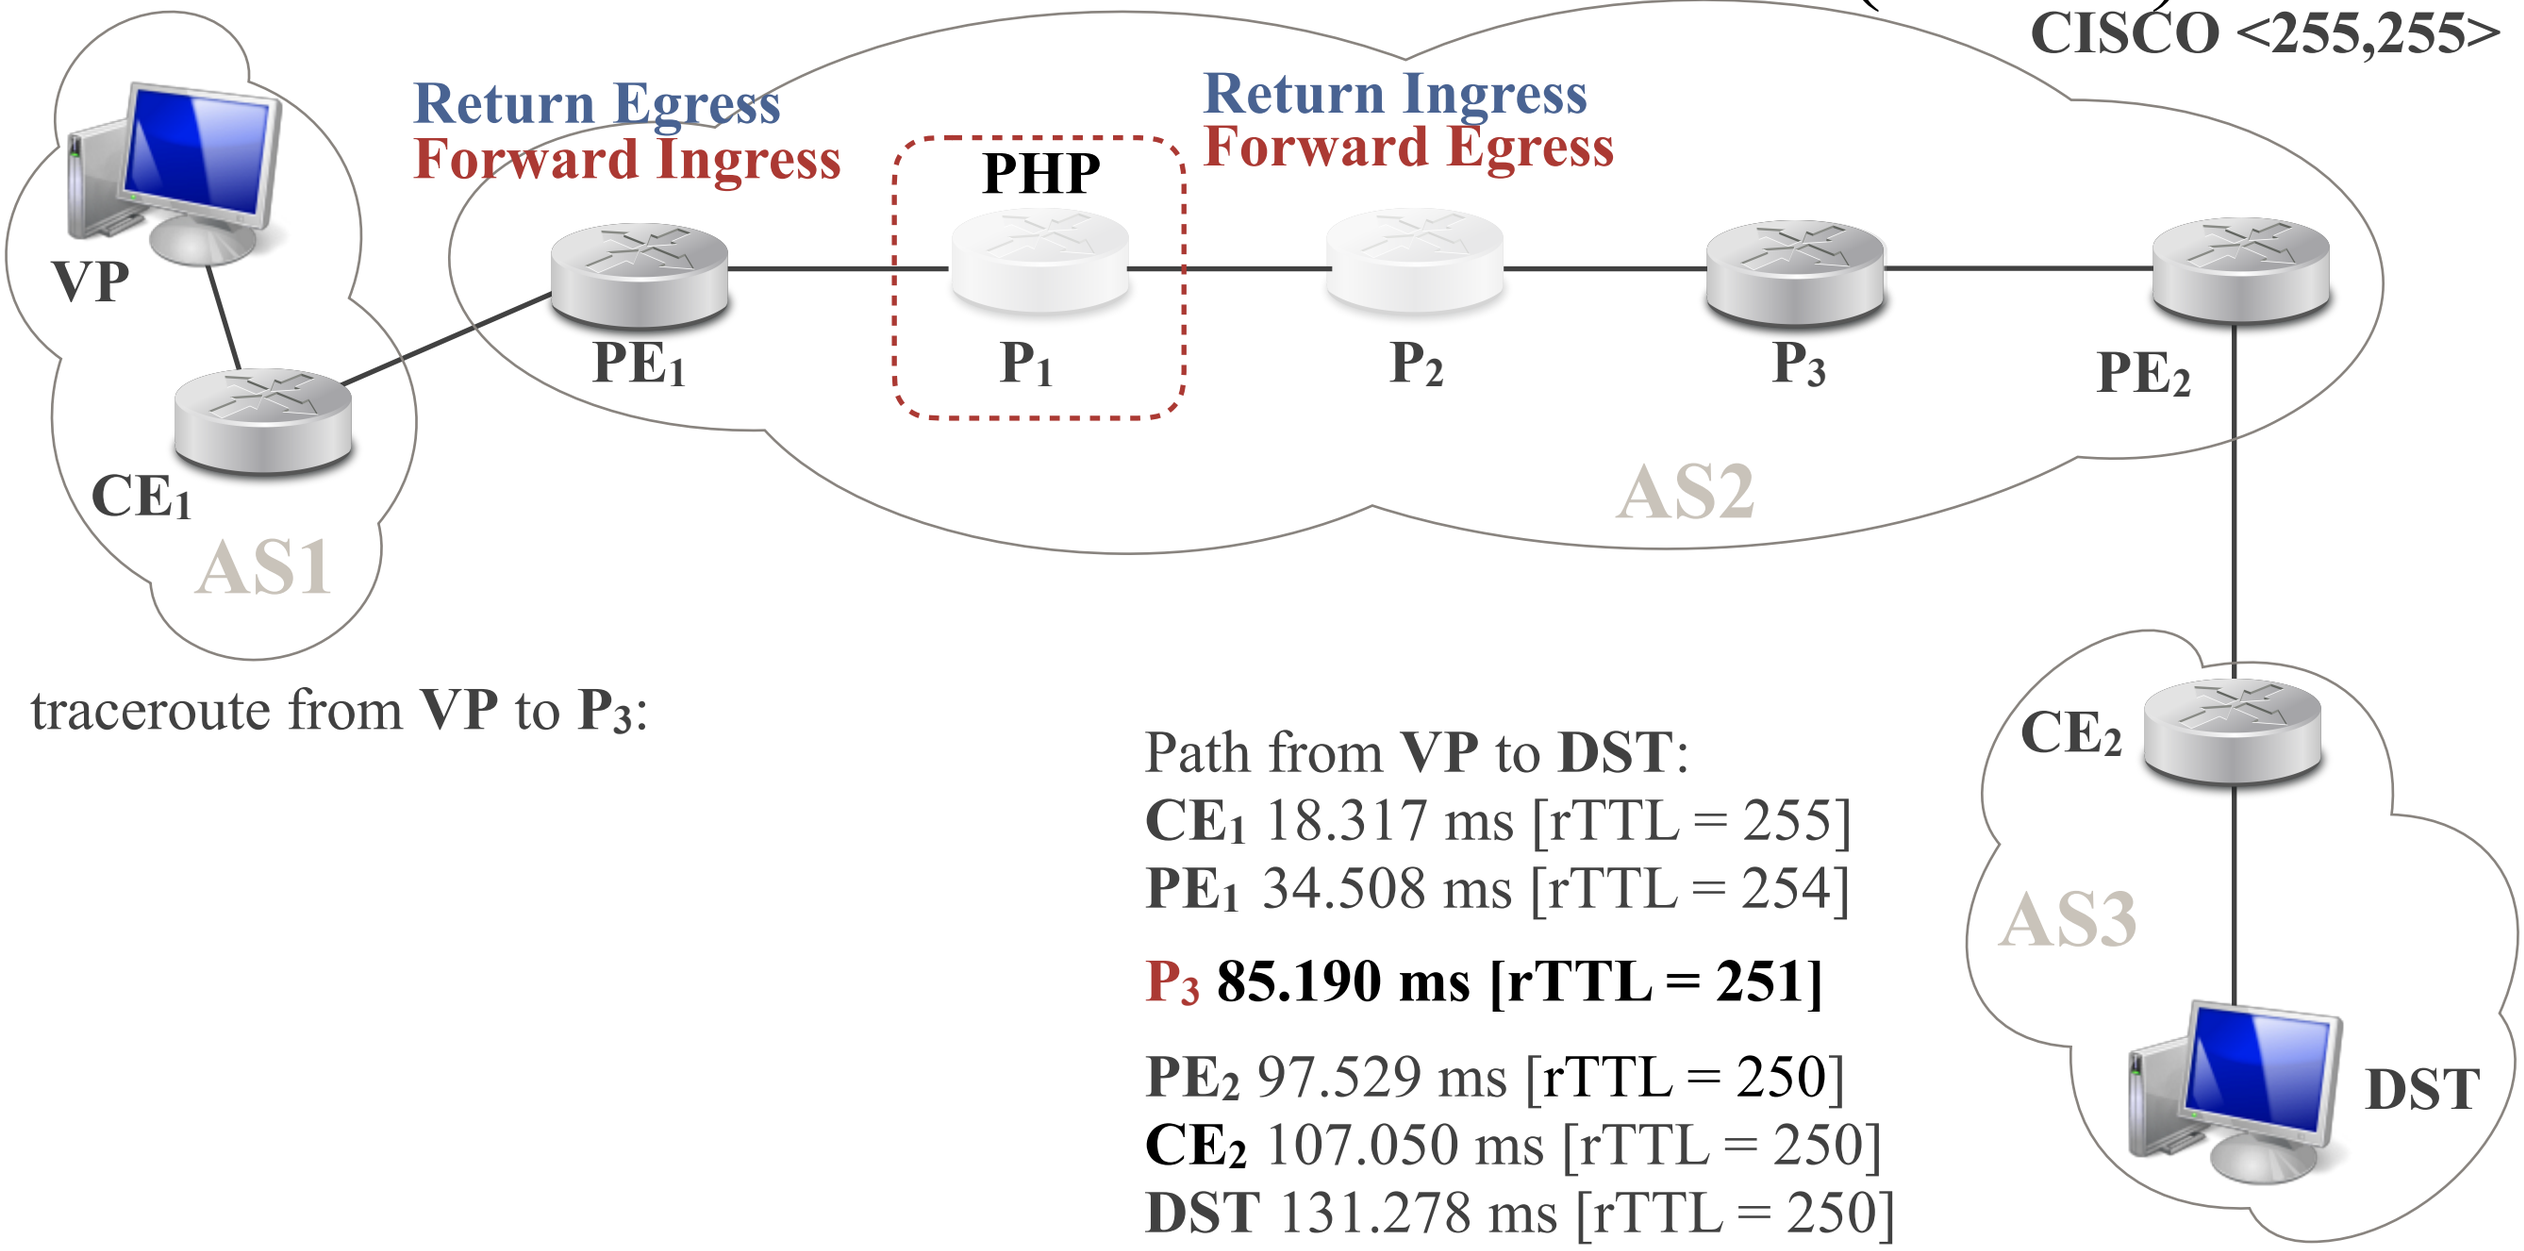

traceroute from **VP** to **P_3**:

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

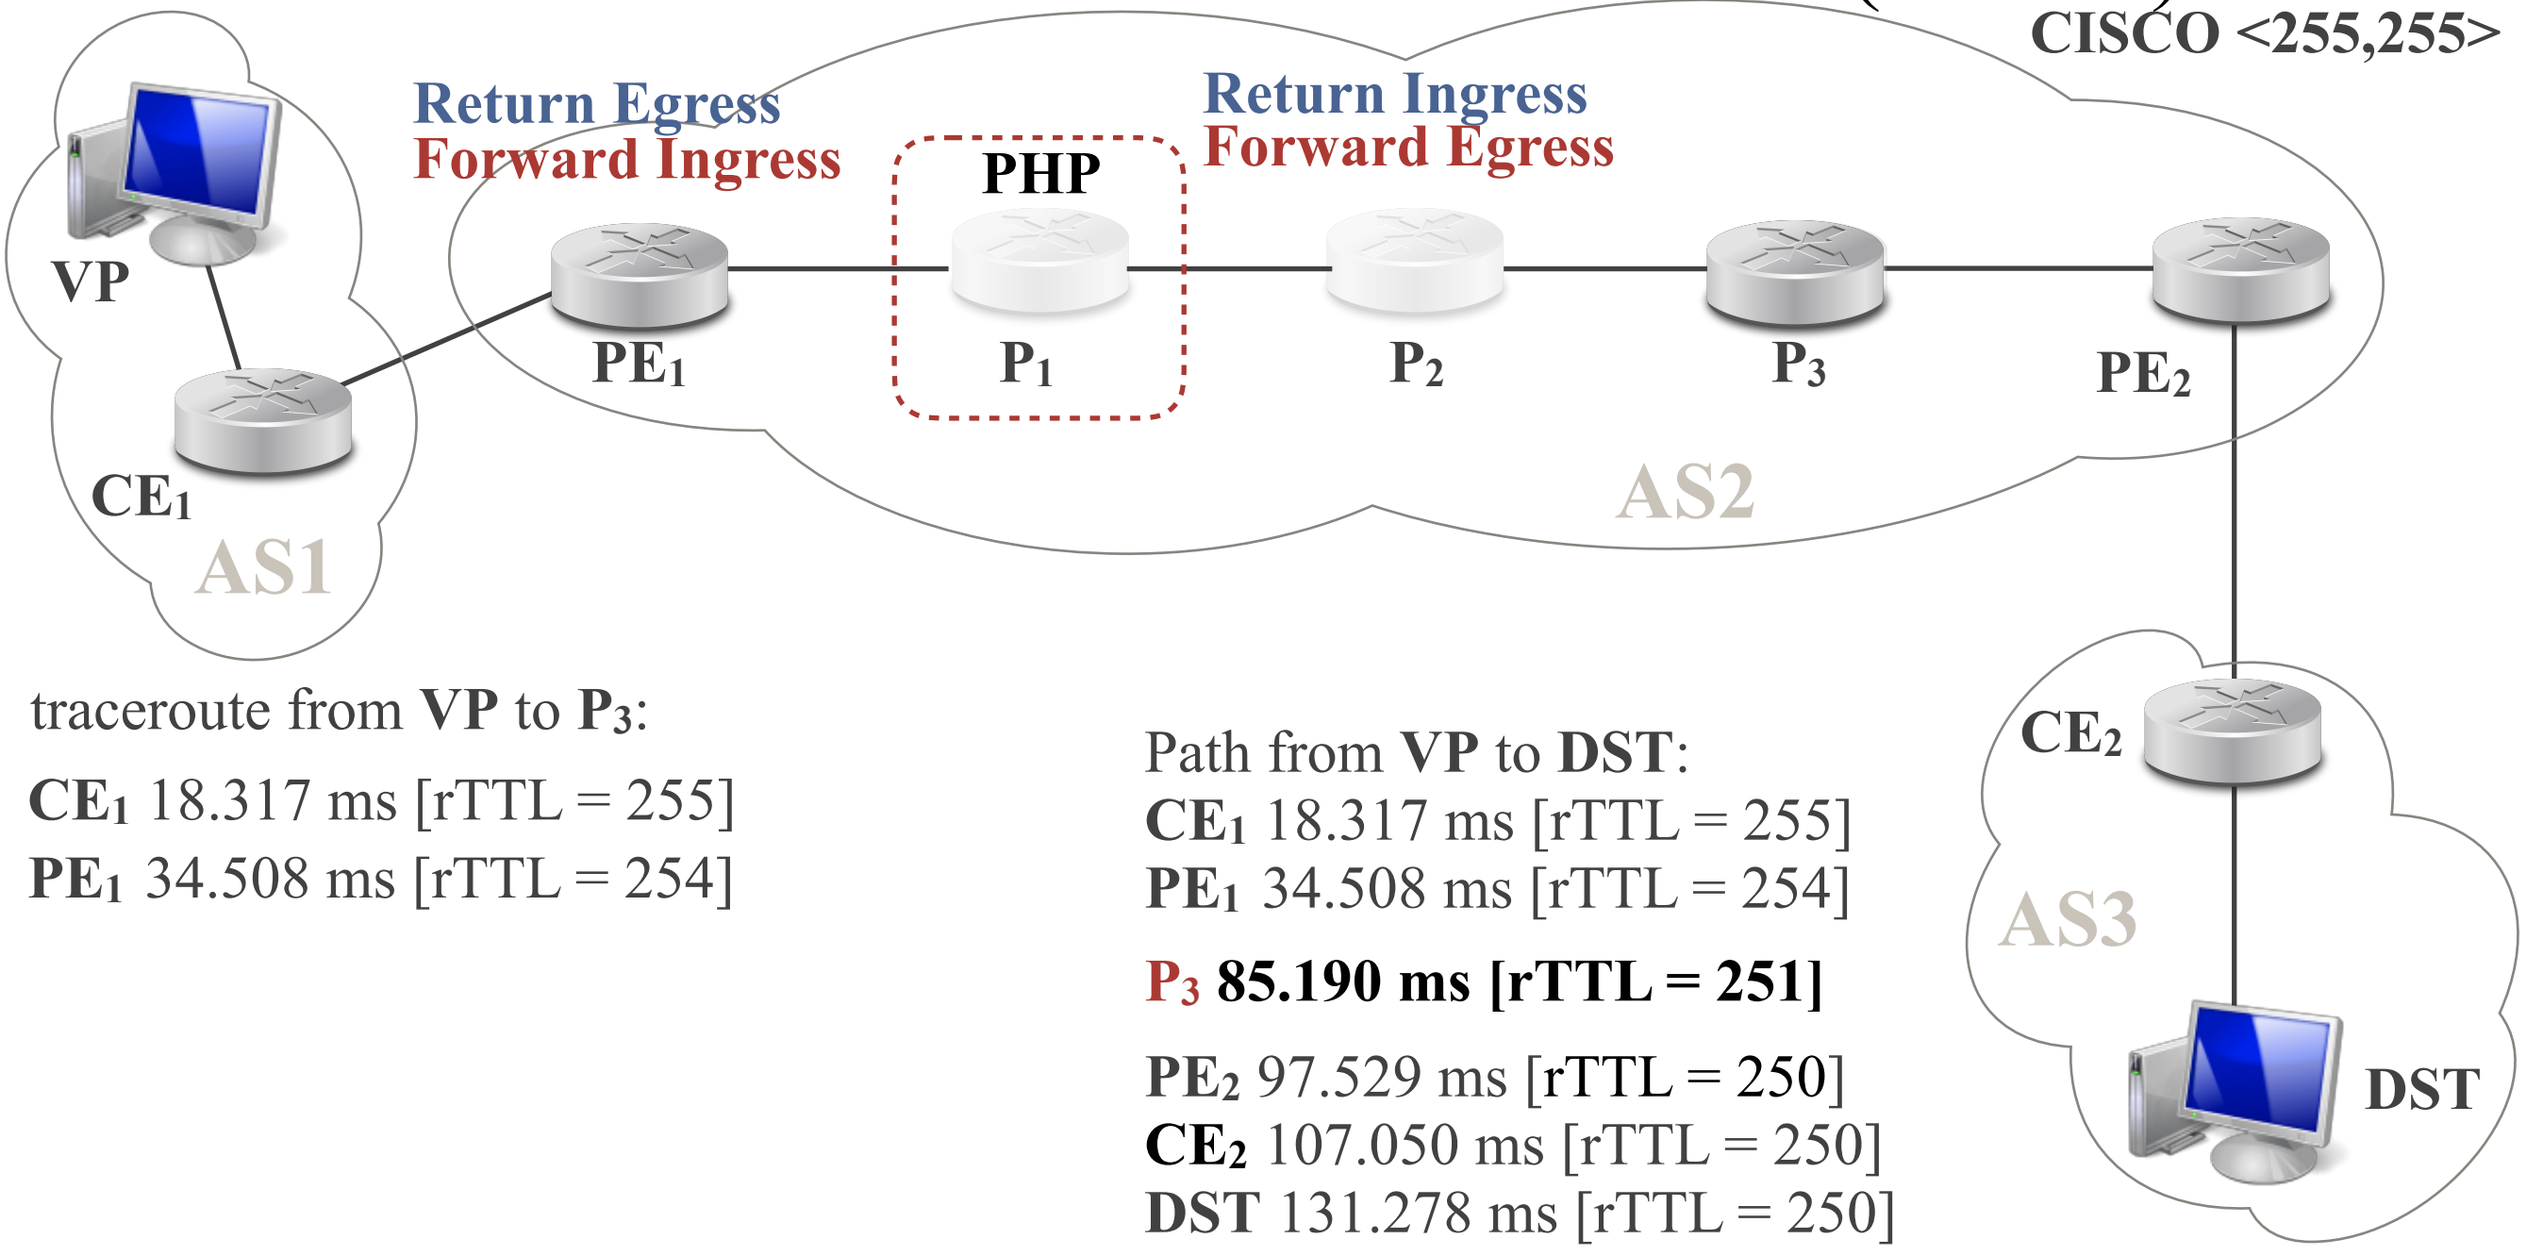

traceroute from **VP** to **P_3**:

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]
P_3 85.190 ms [rTTL = 251]
PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

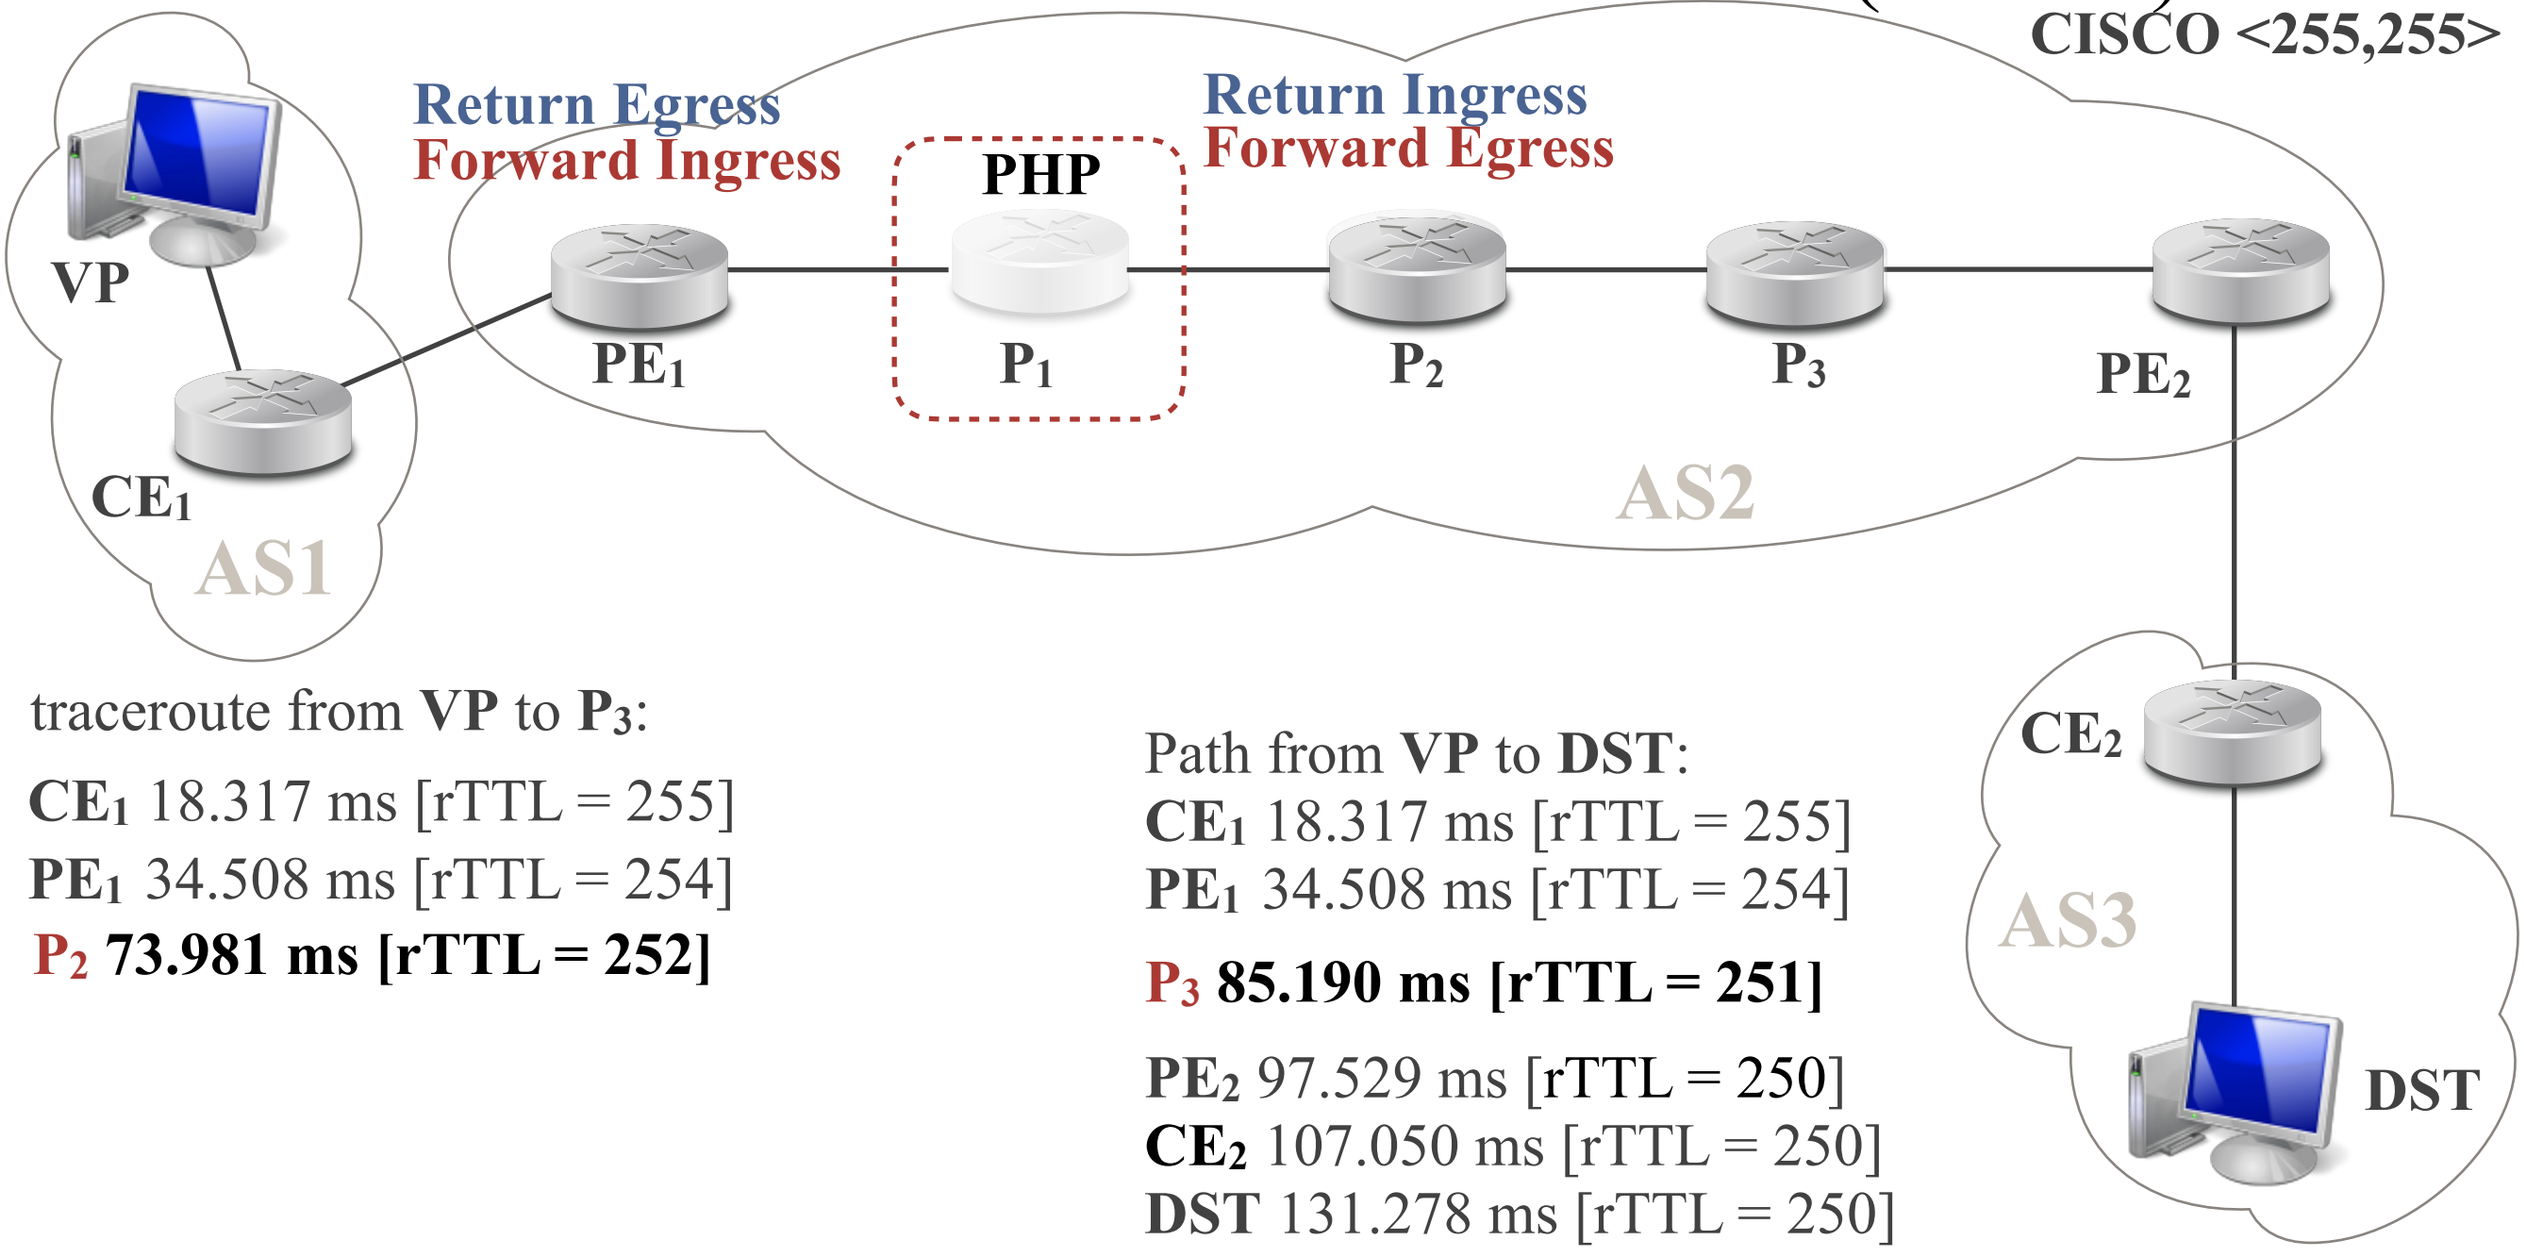

traceroute from **VP** to **P_3**:

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]
P_2 73.981 ms [rTTL = 252]

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]
P_3 85.190 ms [rTTL = 251]
PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

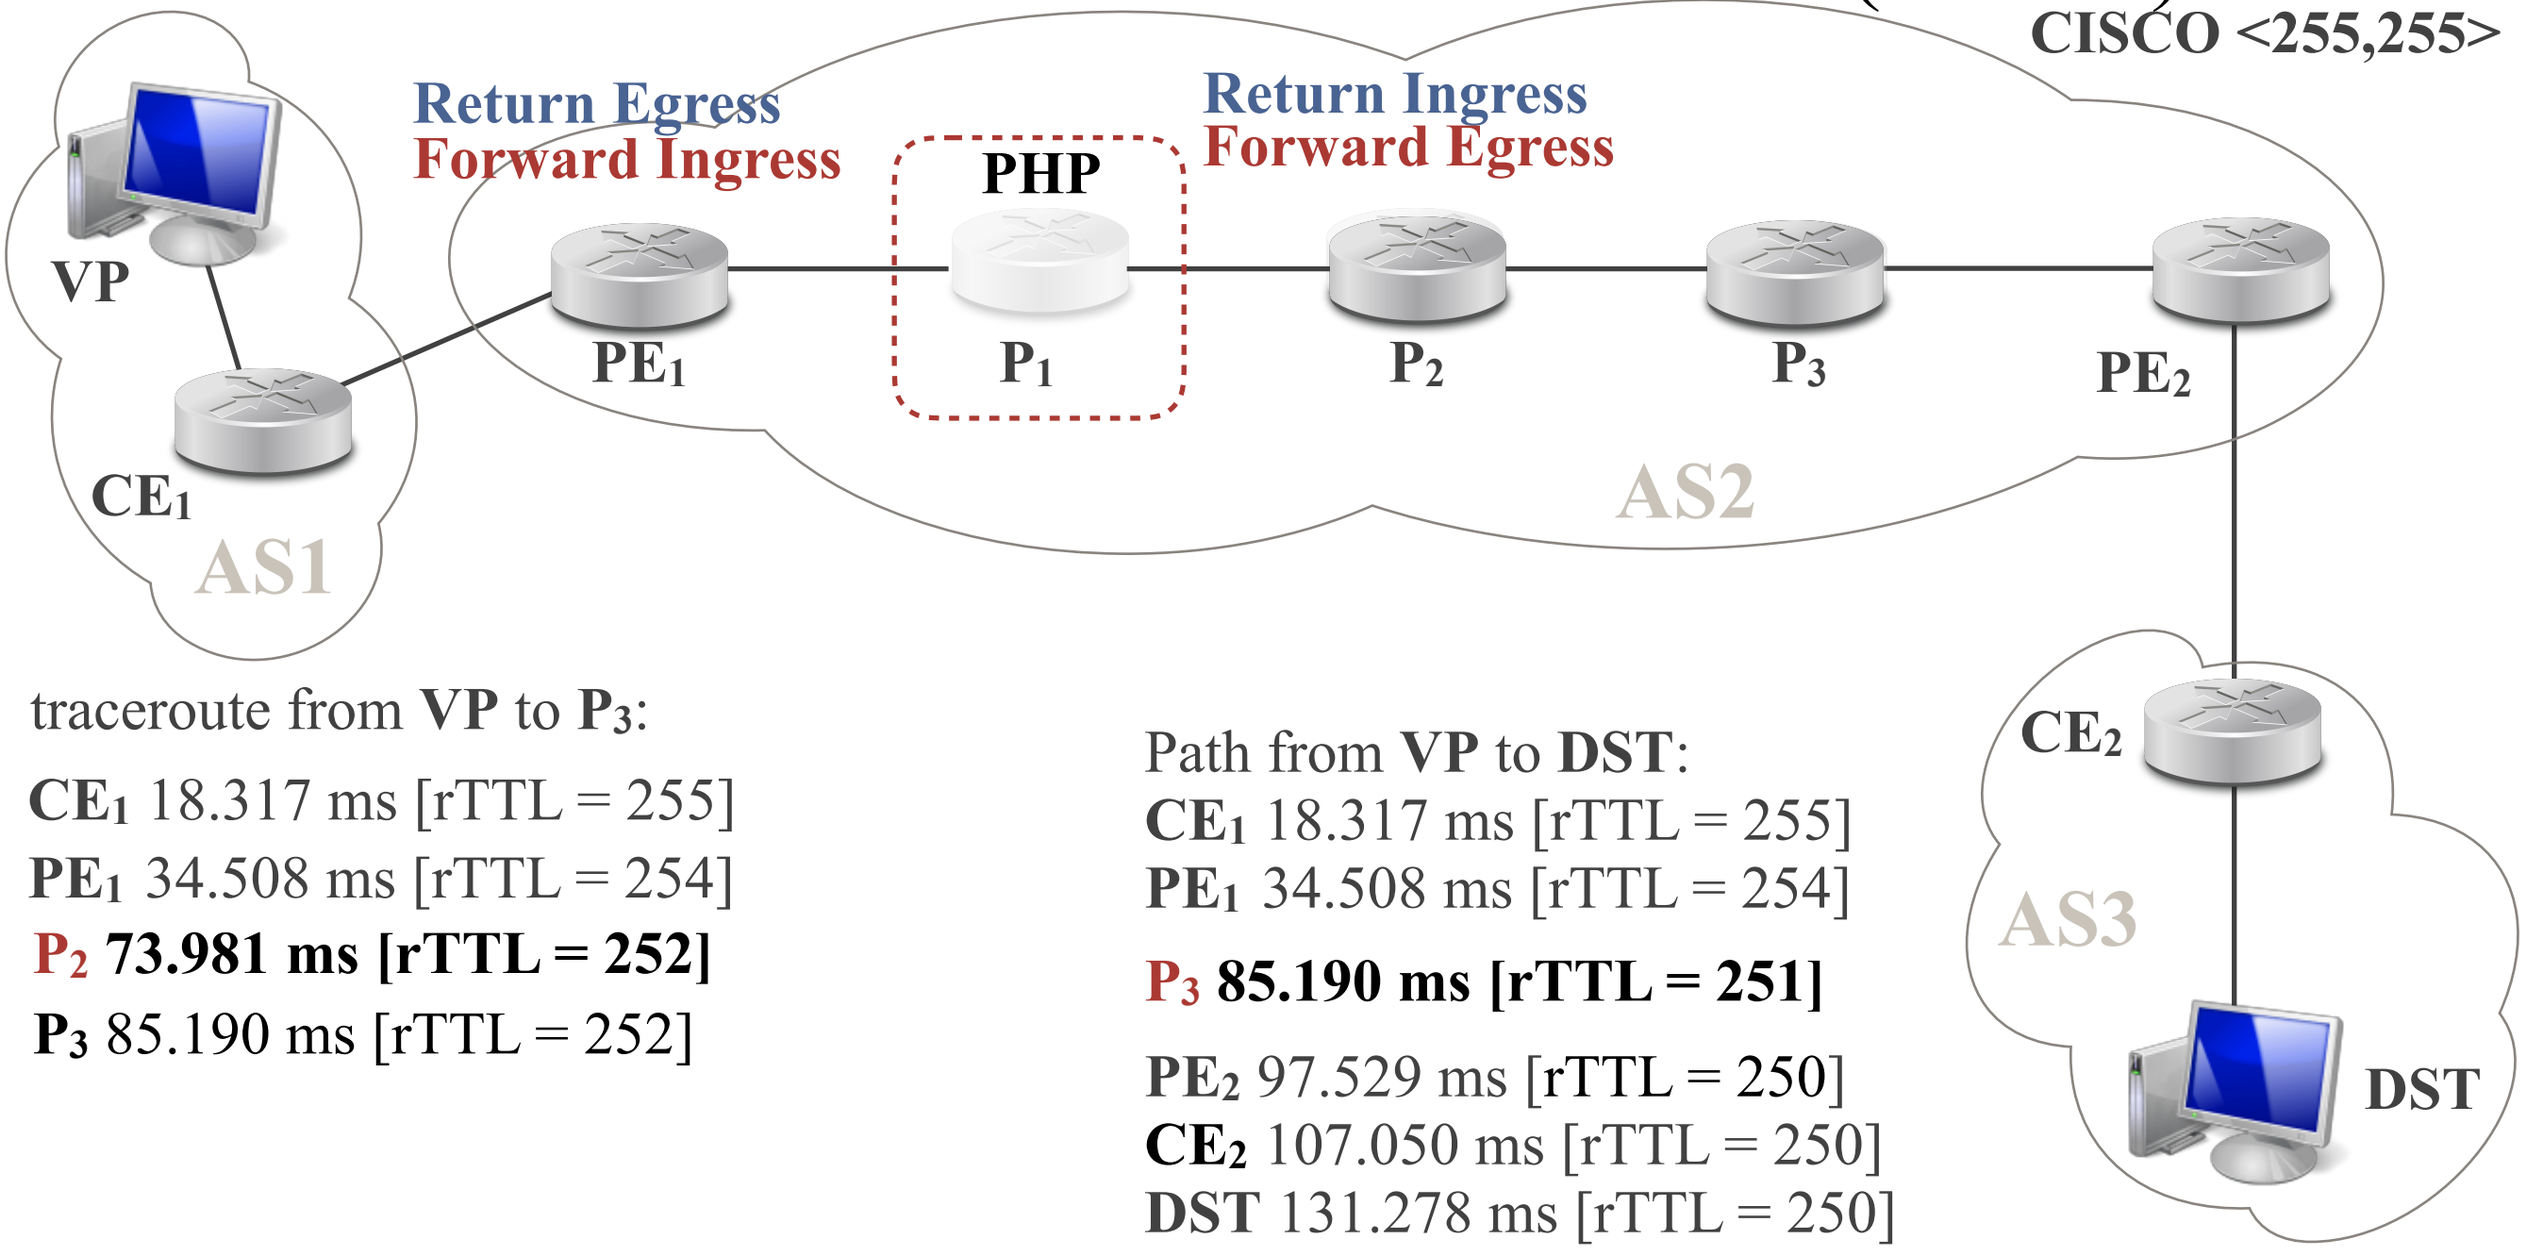

traceroute from **VP** to **P_3**:

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]
P_2 73.981 ms [rTTL = 252]
P_3 85.190 ms [rTTL = 252]

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]
P_3 85.190 ms [rTTL = 251]
PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

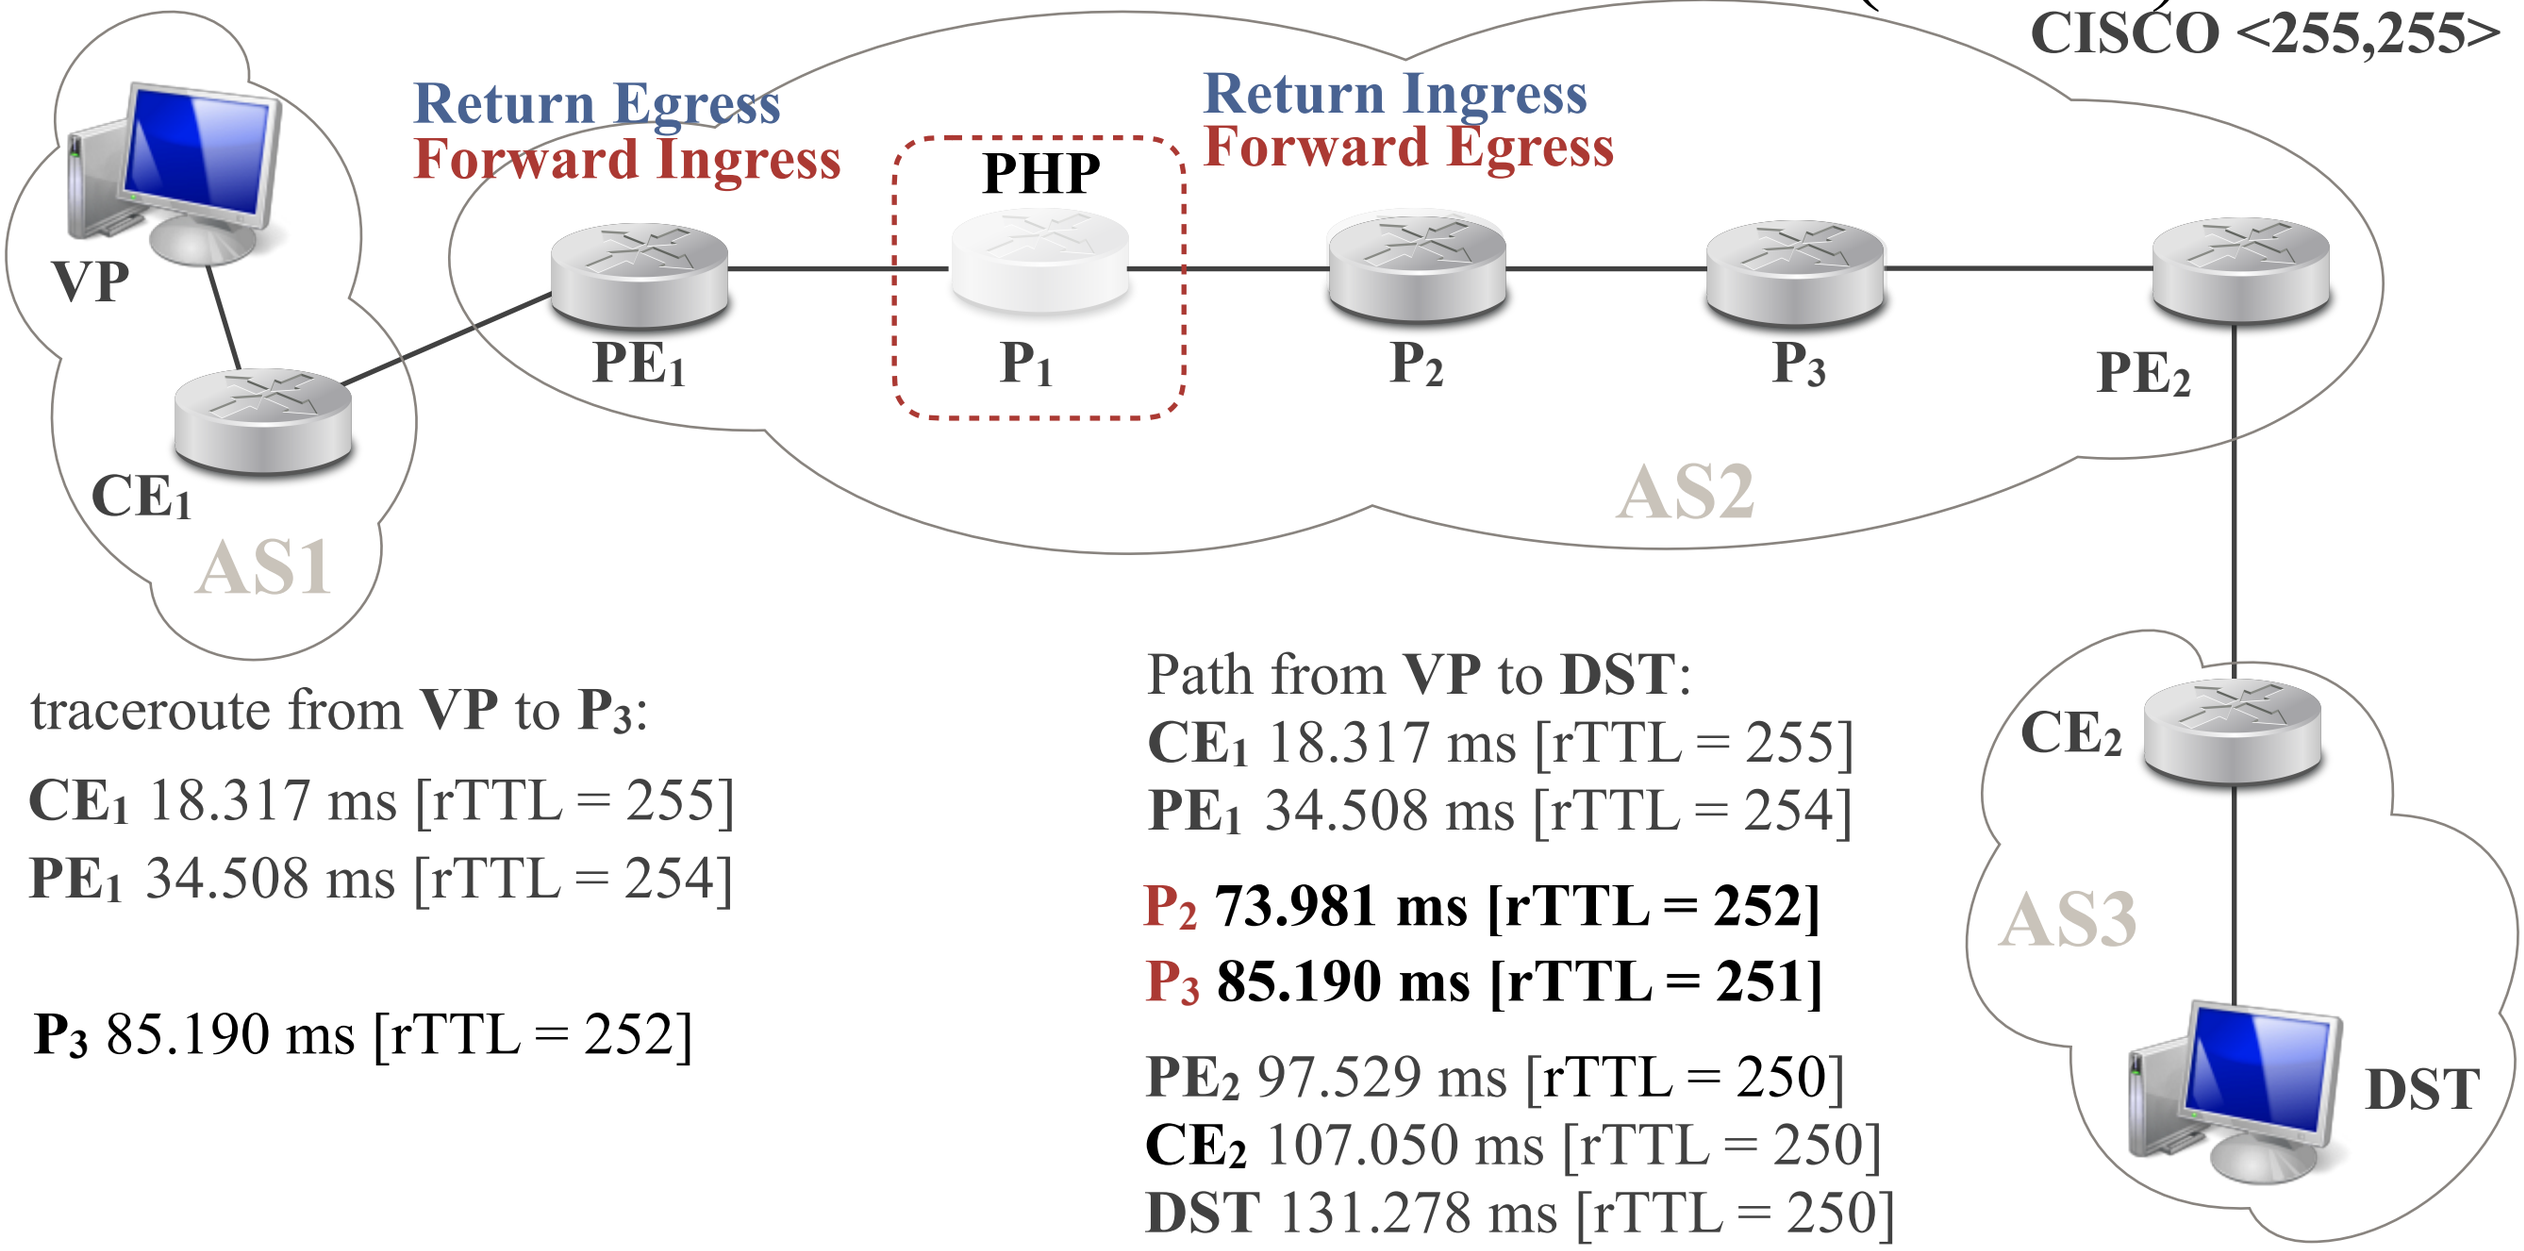

traceroute from **VP** to **P_3**:

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_3 85.190 ms [rTTL = 252]

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_2 73.981 ms [rTTL = 252]
P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

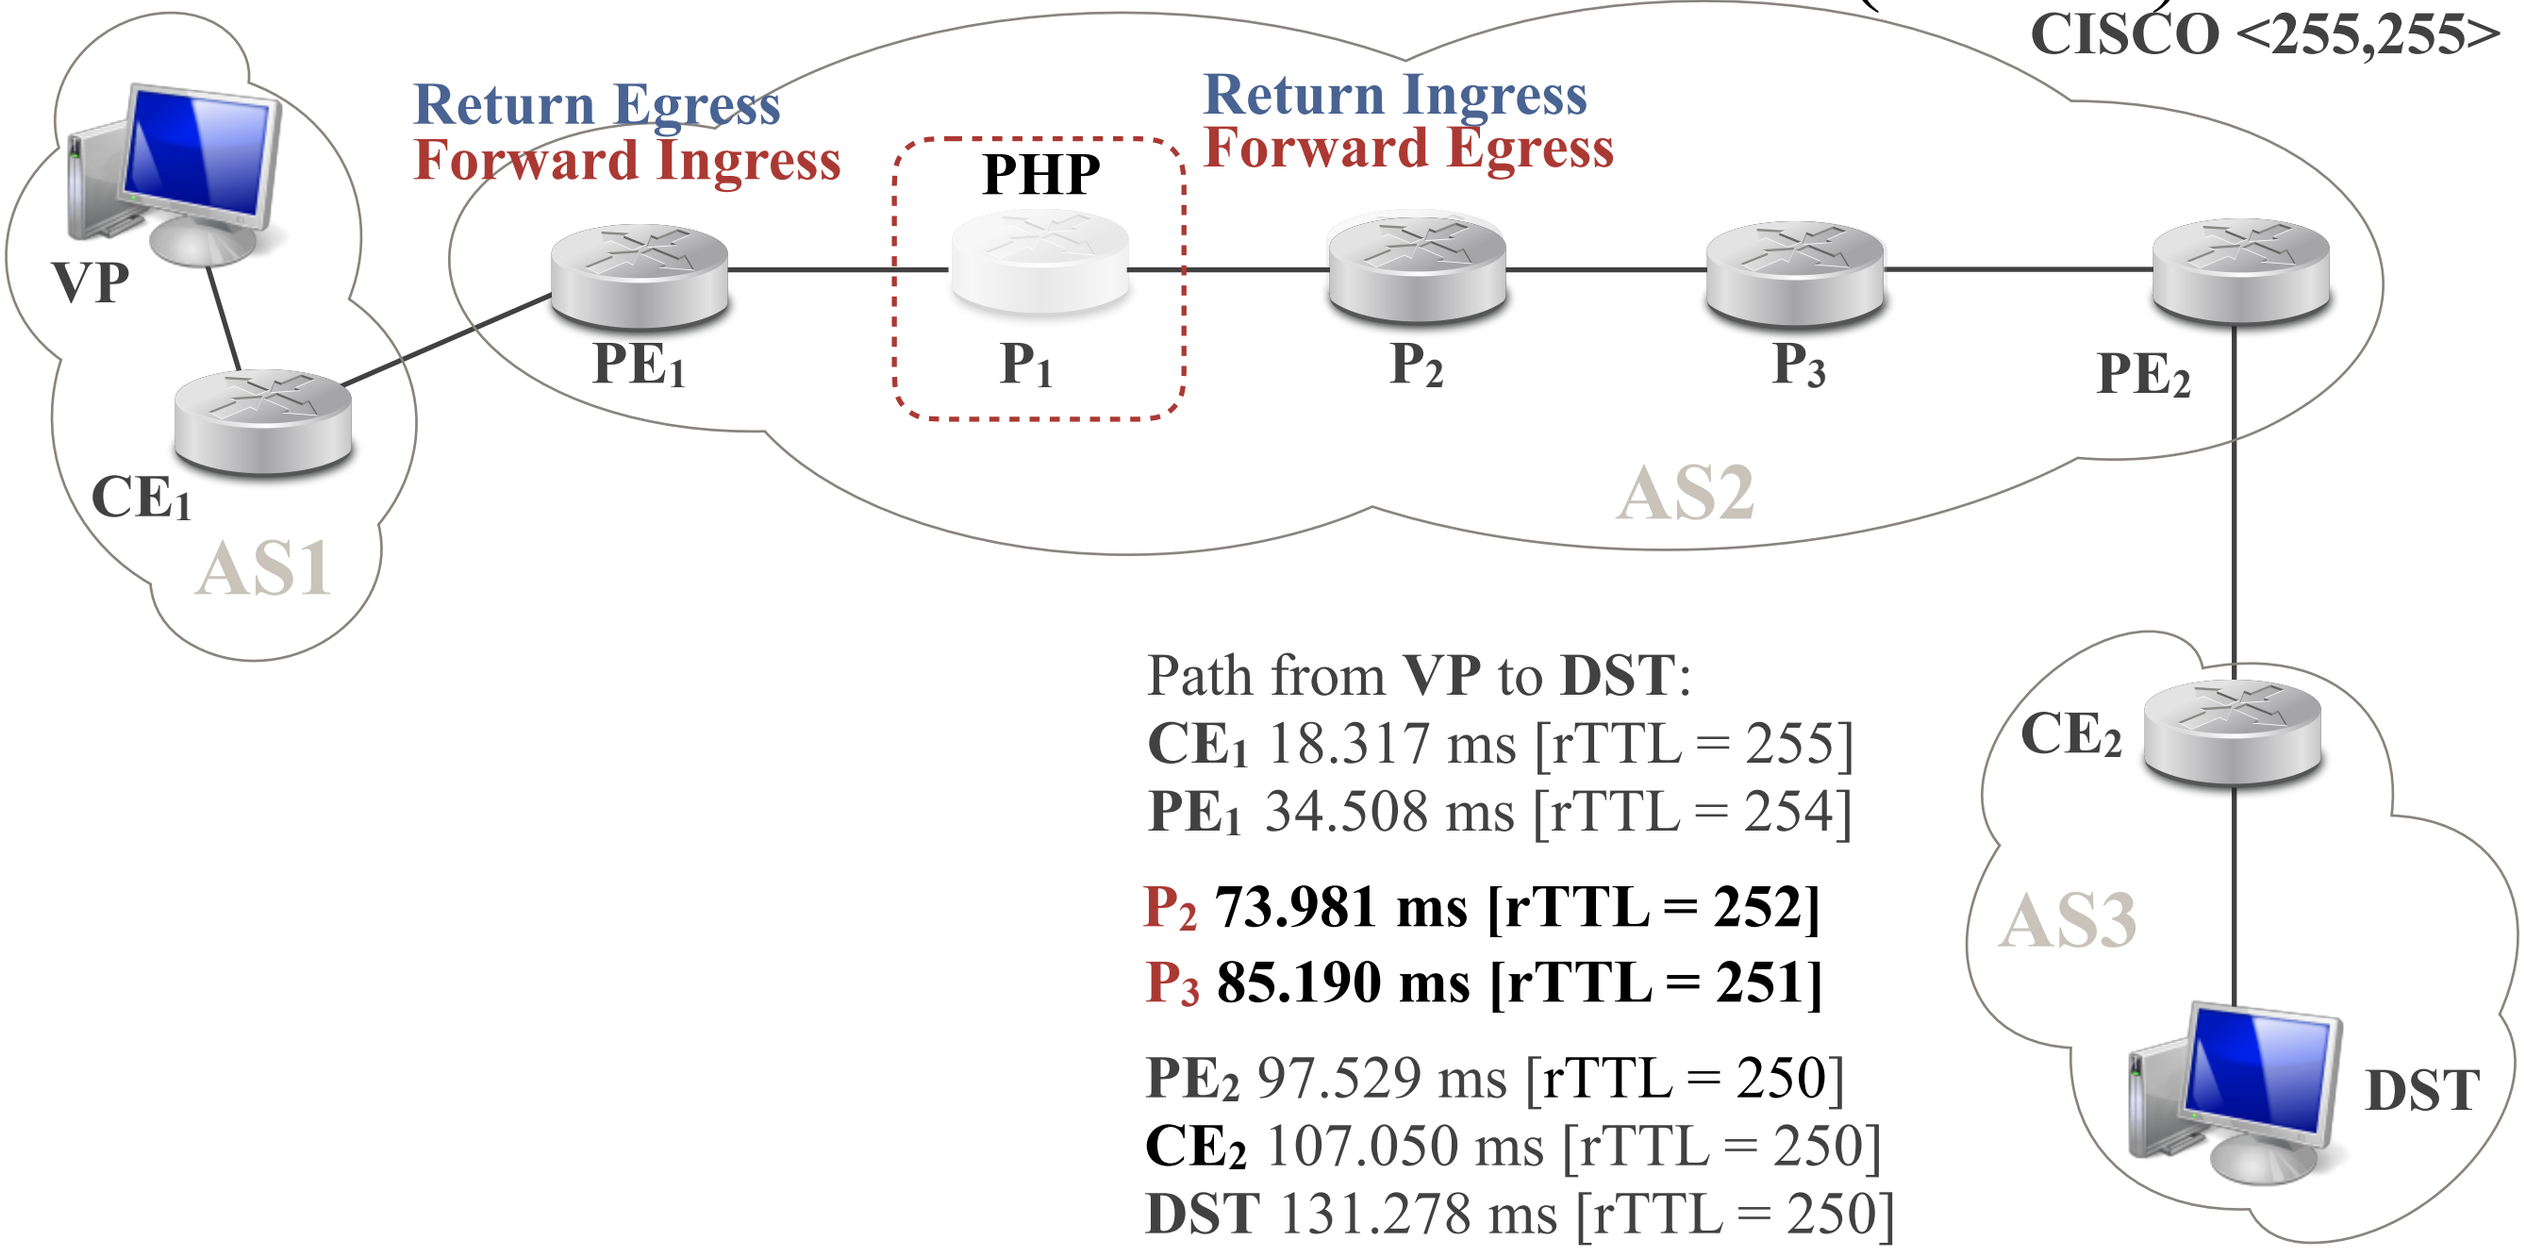

Path from **VP** to **DST**:
CE$_1$ 18.317 ms [rTTL = 255]
PE$_1$ 34.508 ms [rTTL = 254]

P$_2$ 73.981 ms [rTTL = 252]
P$_3$ 85.190 ms [rTTL = 251]

PE$_2$ 97.529 ms [rTTL = 250]
CE$_2$ 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

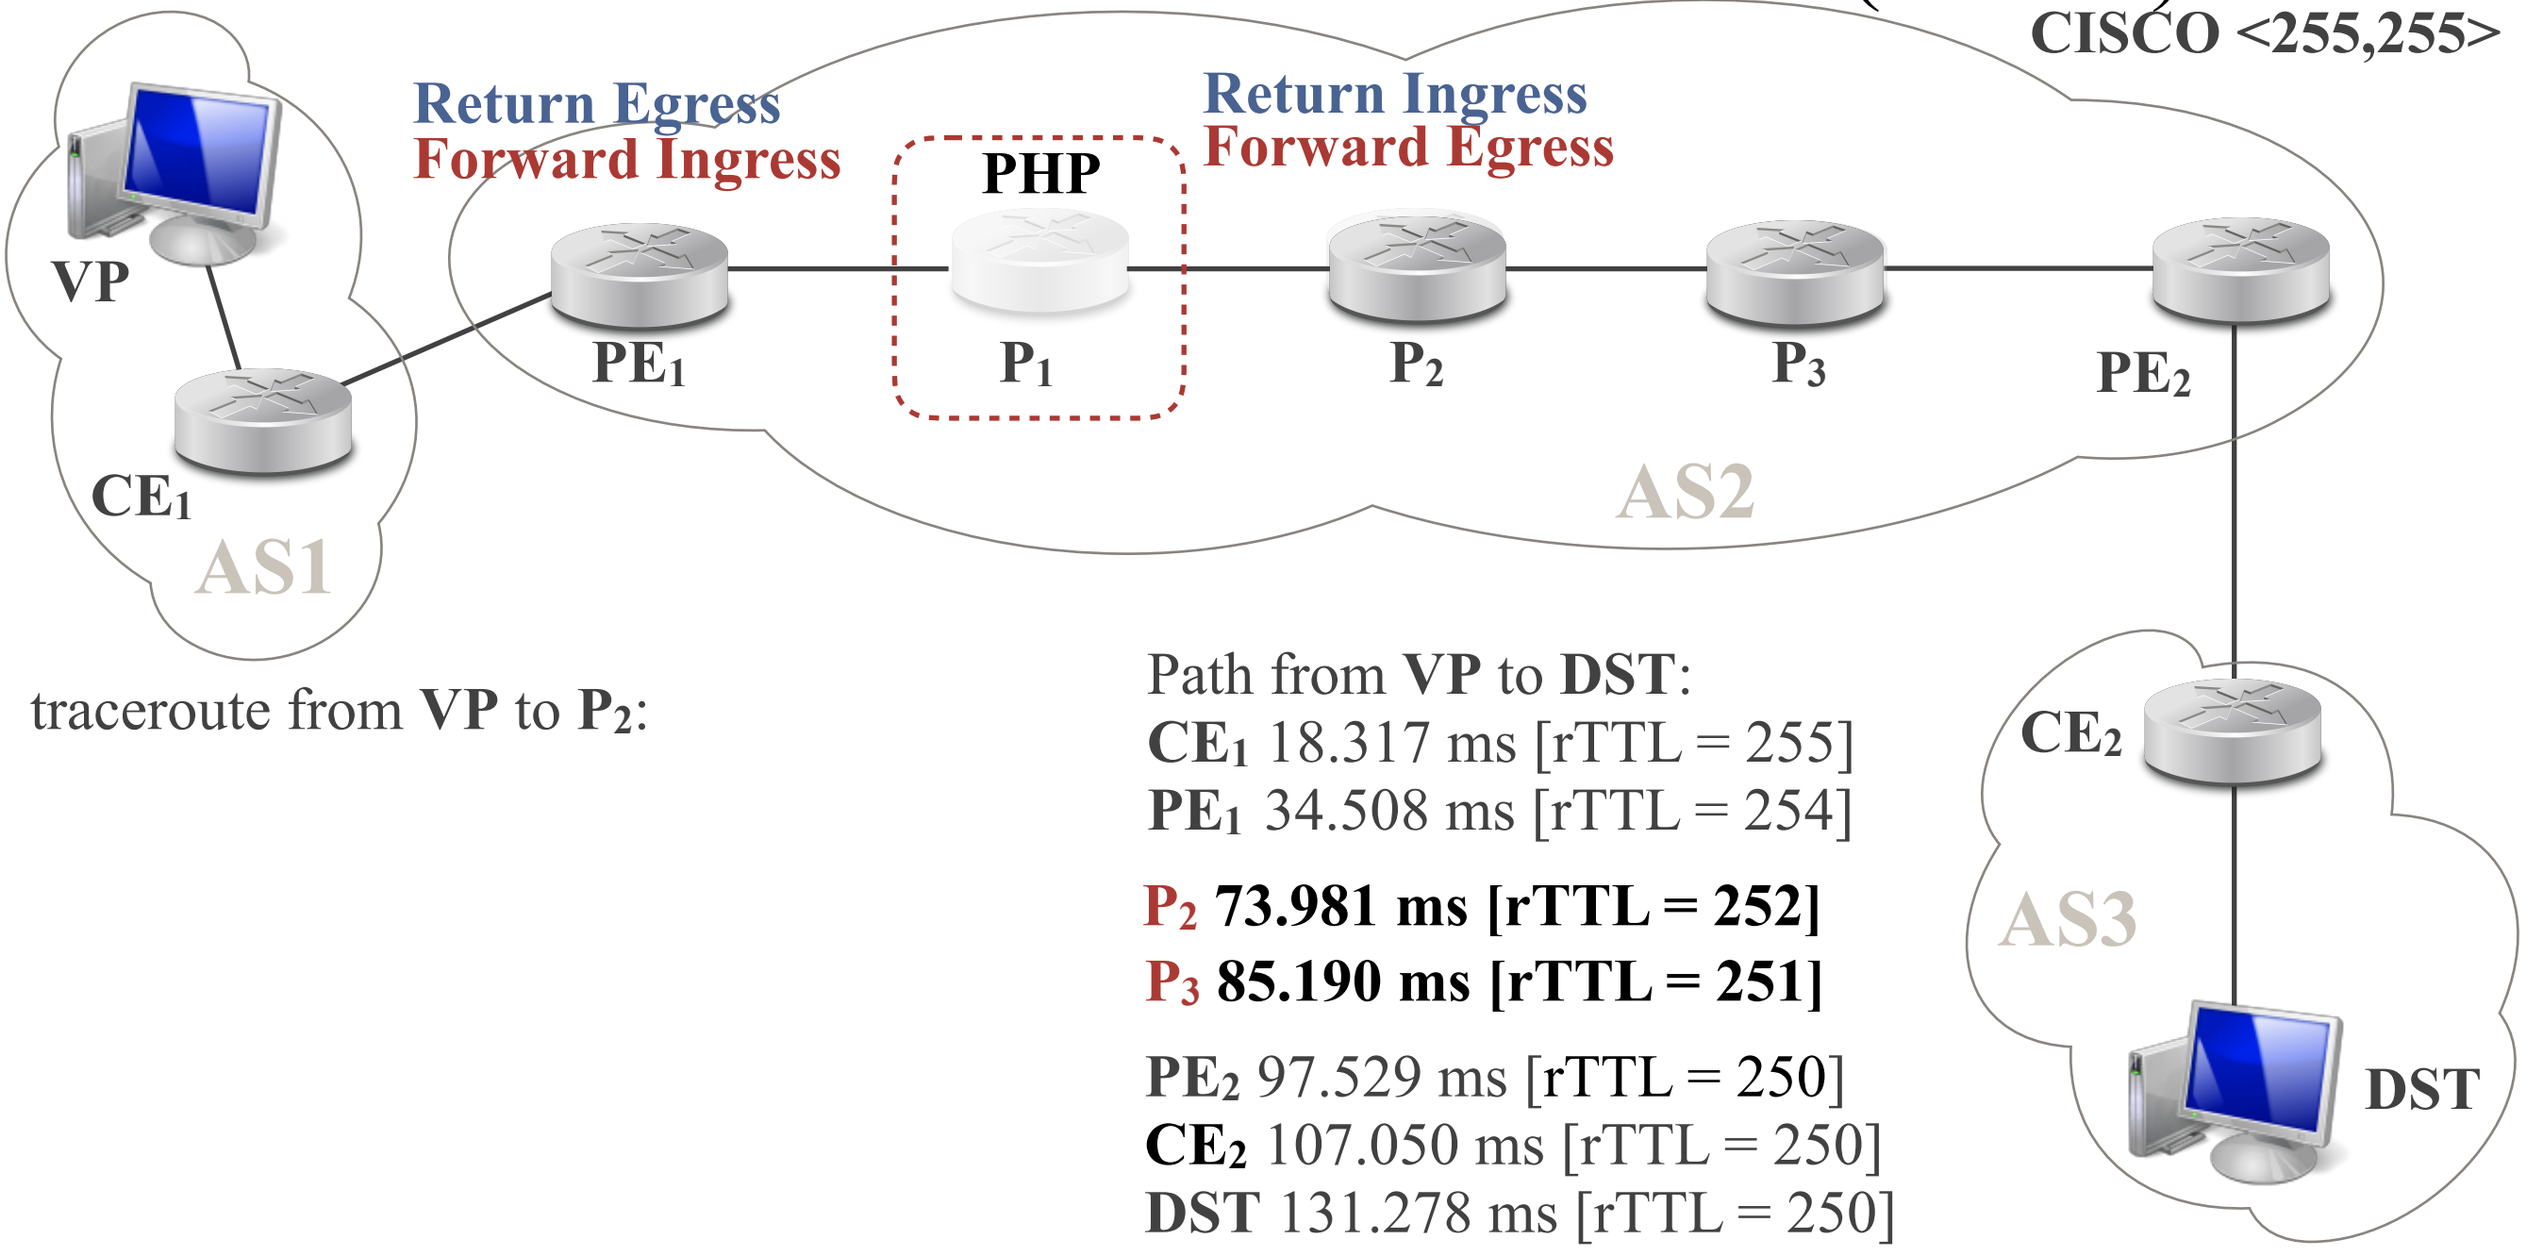

traceroute from **VP** to **P_2**:

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_2 73.981 ms [rTTL = 252]
P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

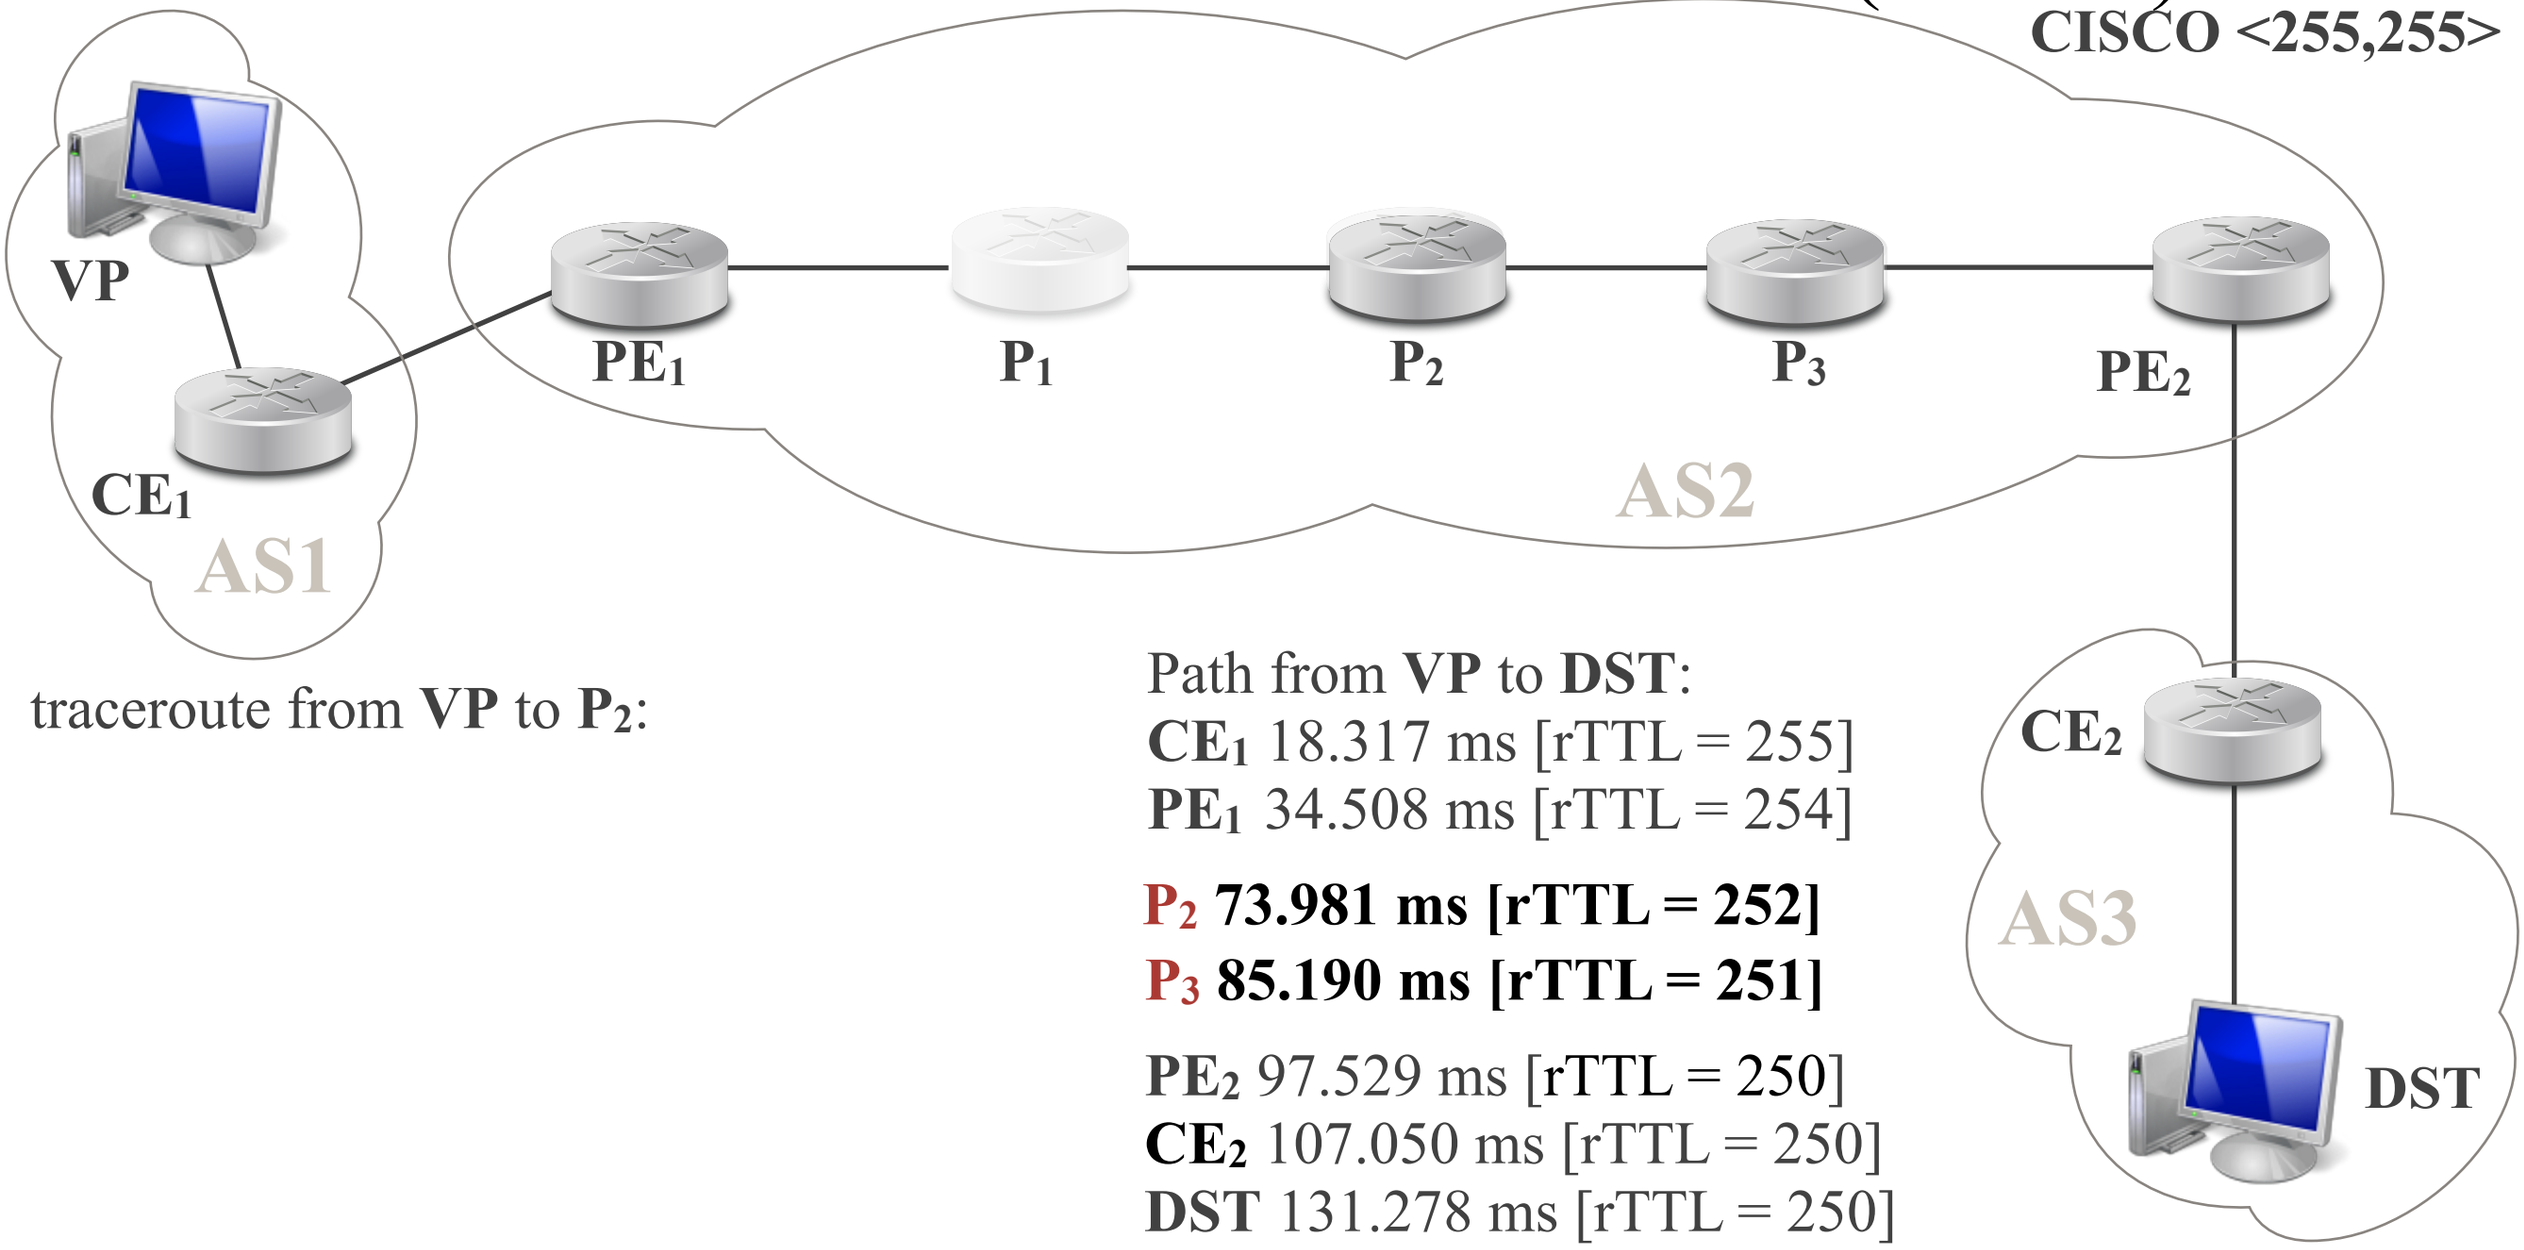

traceroute from **VP** to **P_2**:

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_2 73.981 ms [rTTL = 252]
P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

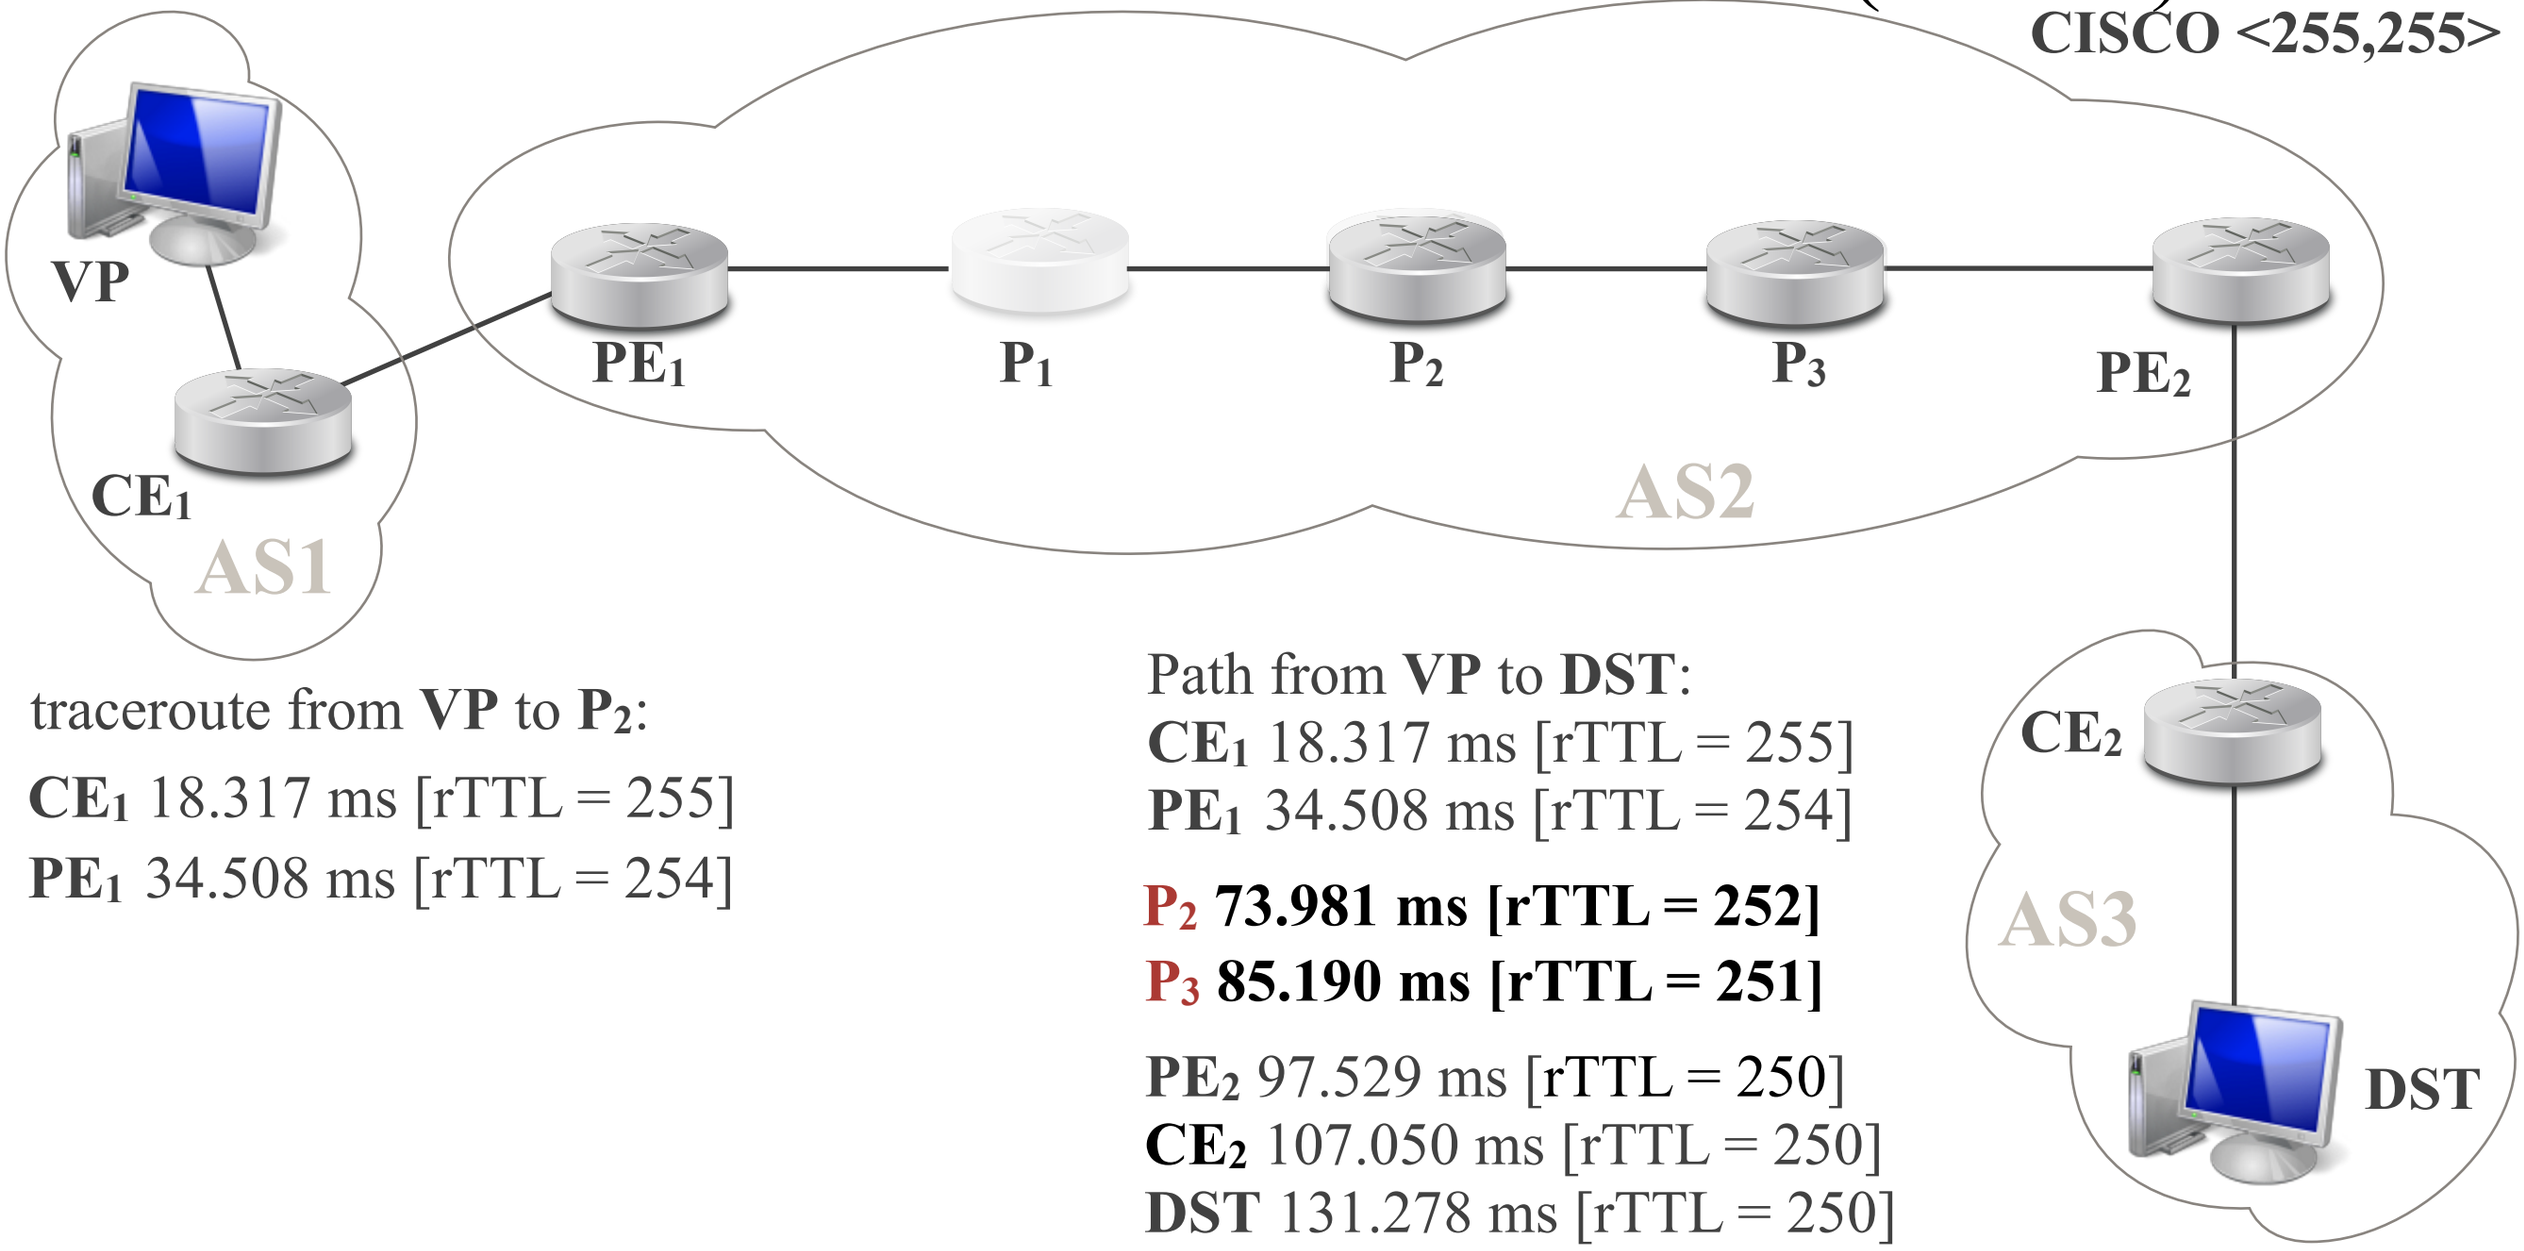

traceroute from **VP** to **P_2**:

CE_1 18.317 ms [rTTL = 255]

PE_1 34.508 ms [rTTL = 254]

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

CE_1 18.317 ms [rTTL = 255]

PE_1 34.508 ms [rTTL = 254]

P_2 73.981 ms [rTTL = 252]

P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]

CE_2 107.050 ms [rTTL = 250]

DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

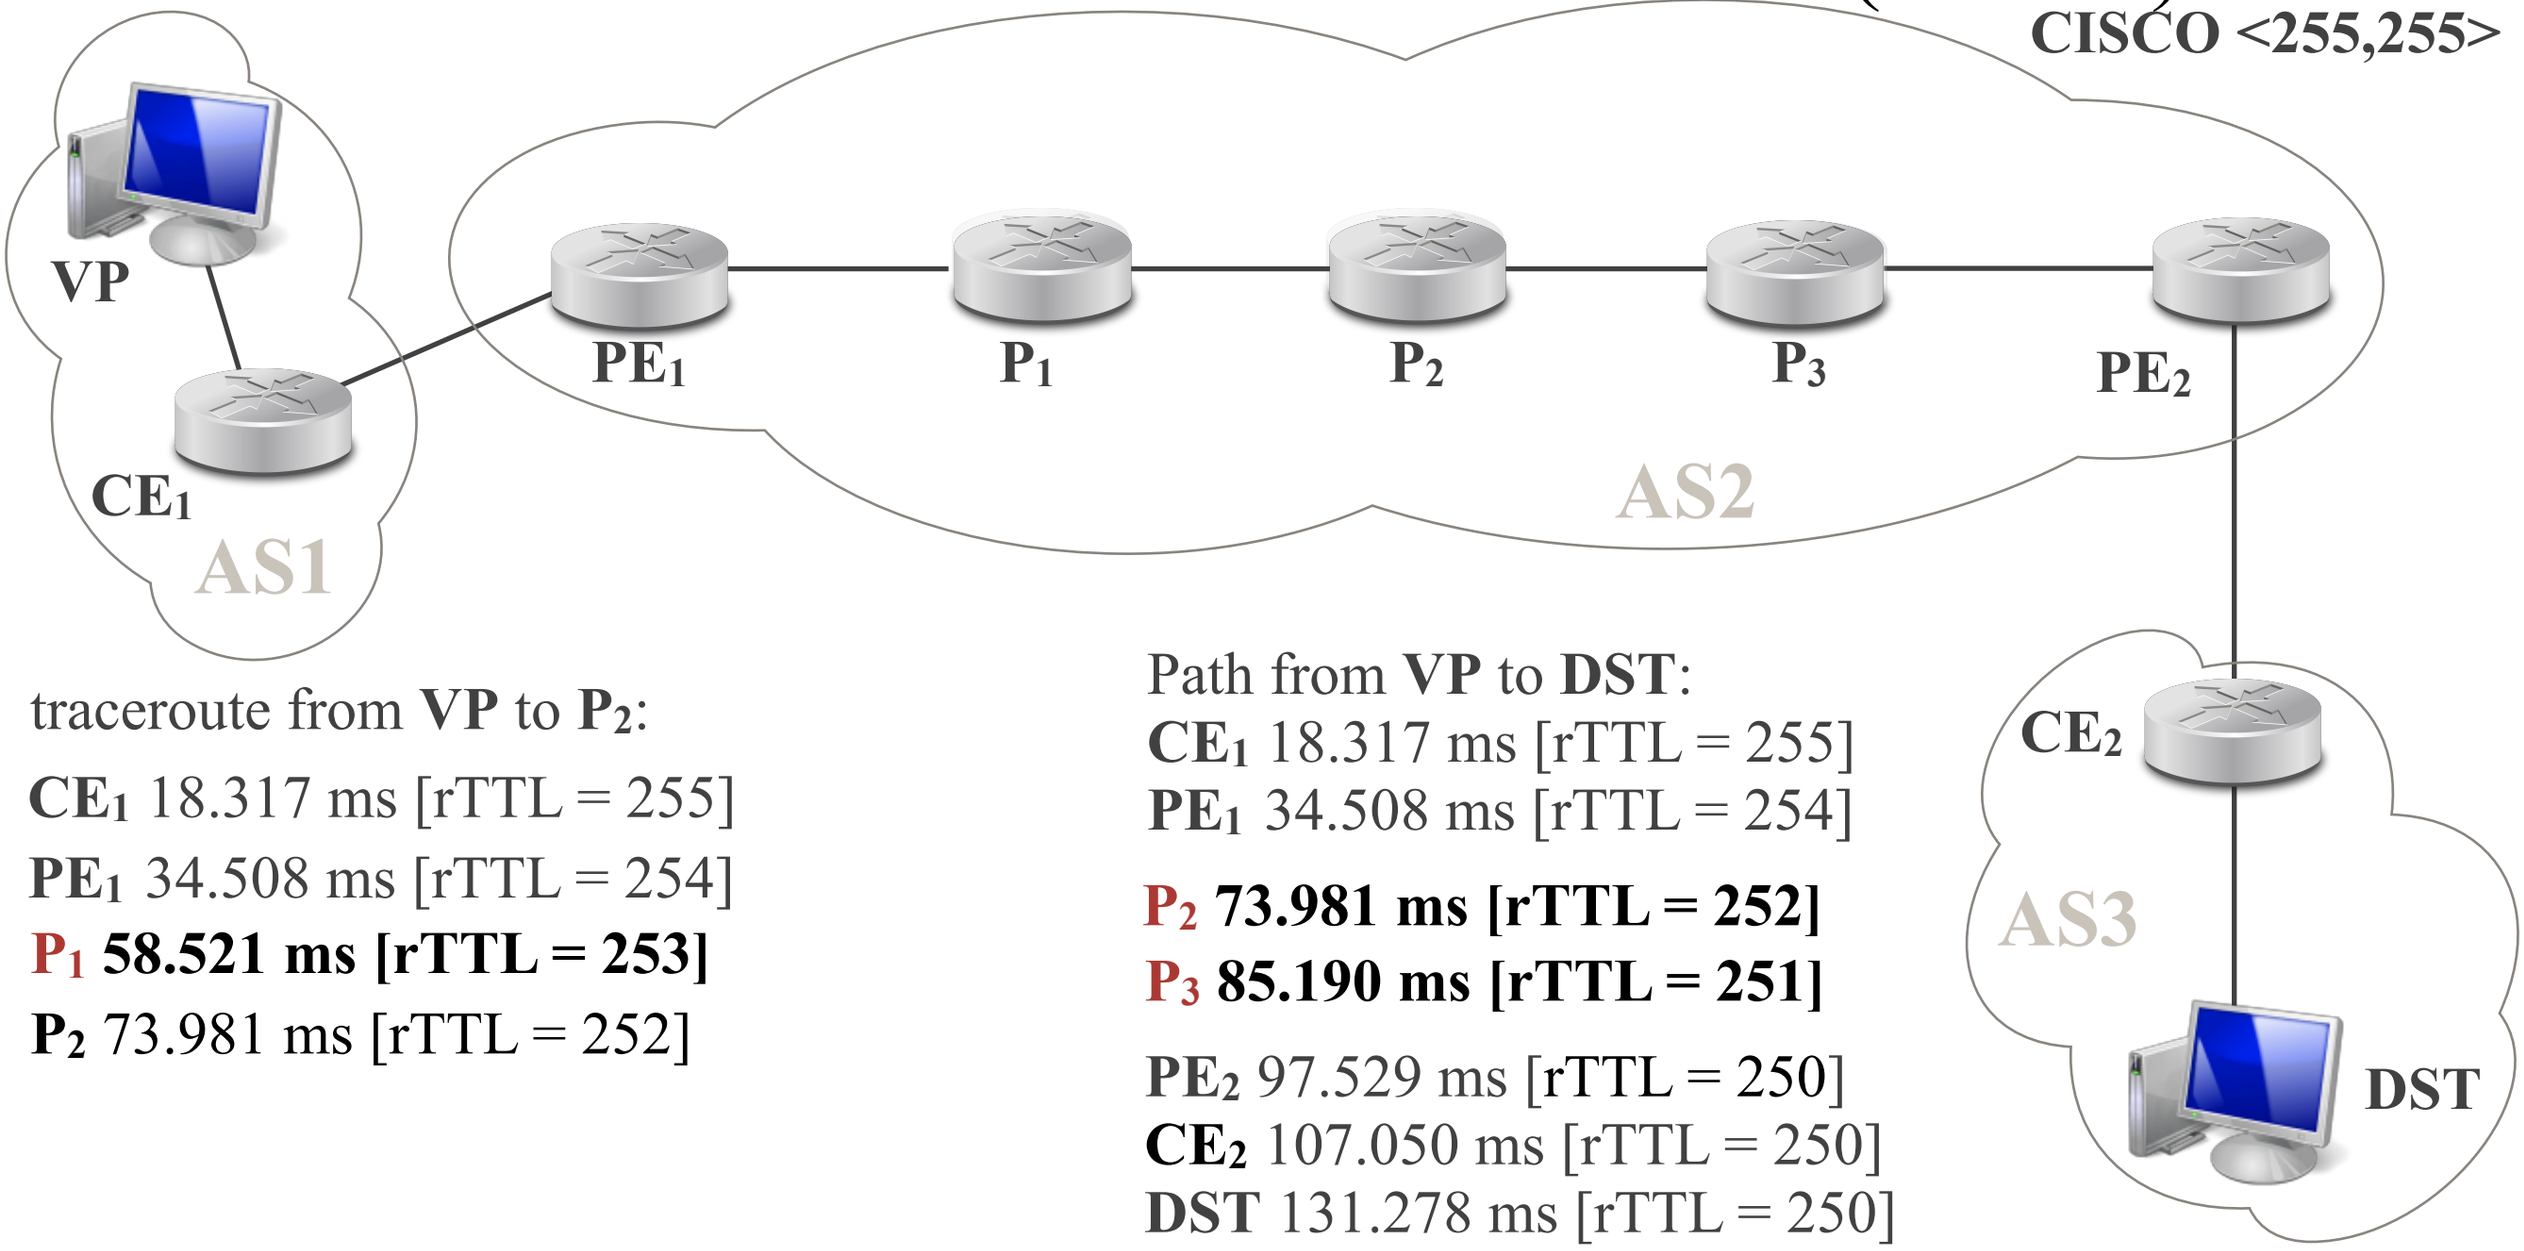

traceroute from **VP** to **P_2**:

CE_1 18.317 ms [rTTL = 255]

PE_1 34.508 ms [rTTL = 254]

P_1 58.521 ms [rTTL = 253]

P_2 73.981 ms [rTTL = 252]

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

CE_1 18.317 ms [rTTL = 255]

PE_1 34.508 ms [rTTL = 254]

P_2 73.981 ms [rTTL = 252]

P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]

CE_2 107.050 ms [rTTL = 250]

DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

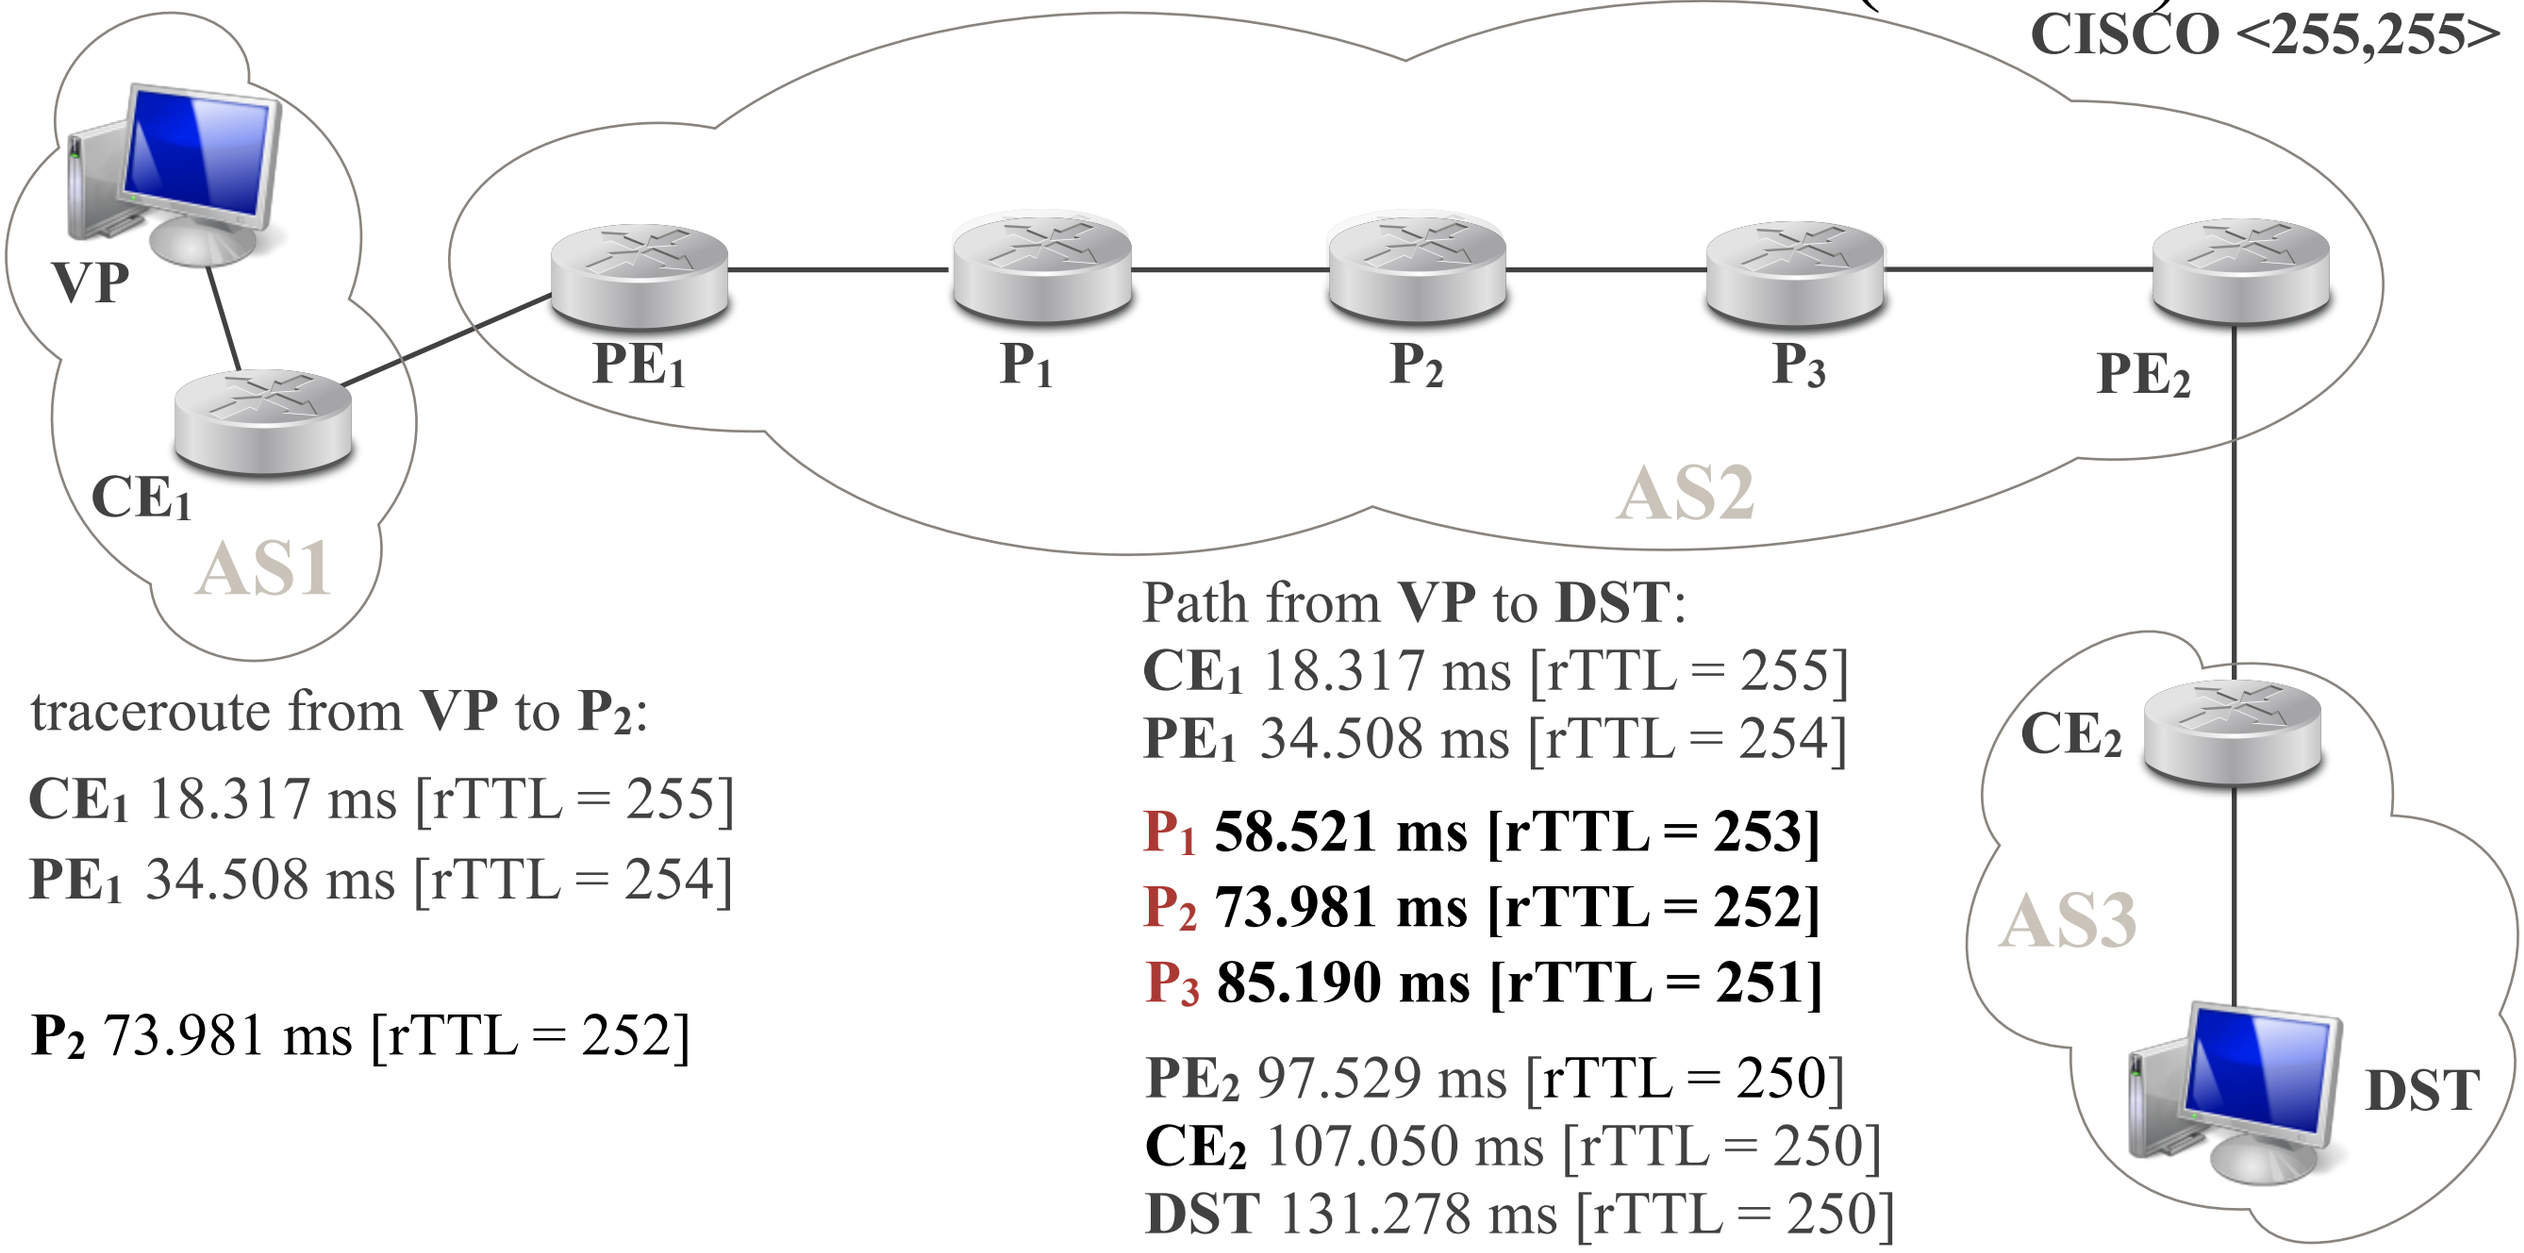

traceroute from **VP** to **P_2**:

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_2 73.981 ms [rTTL = 252]

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

CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_1 58.521 ms [rTTL = 253]
P_2 73.981 ms [rTTL = 252]
P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

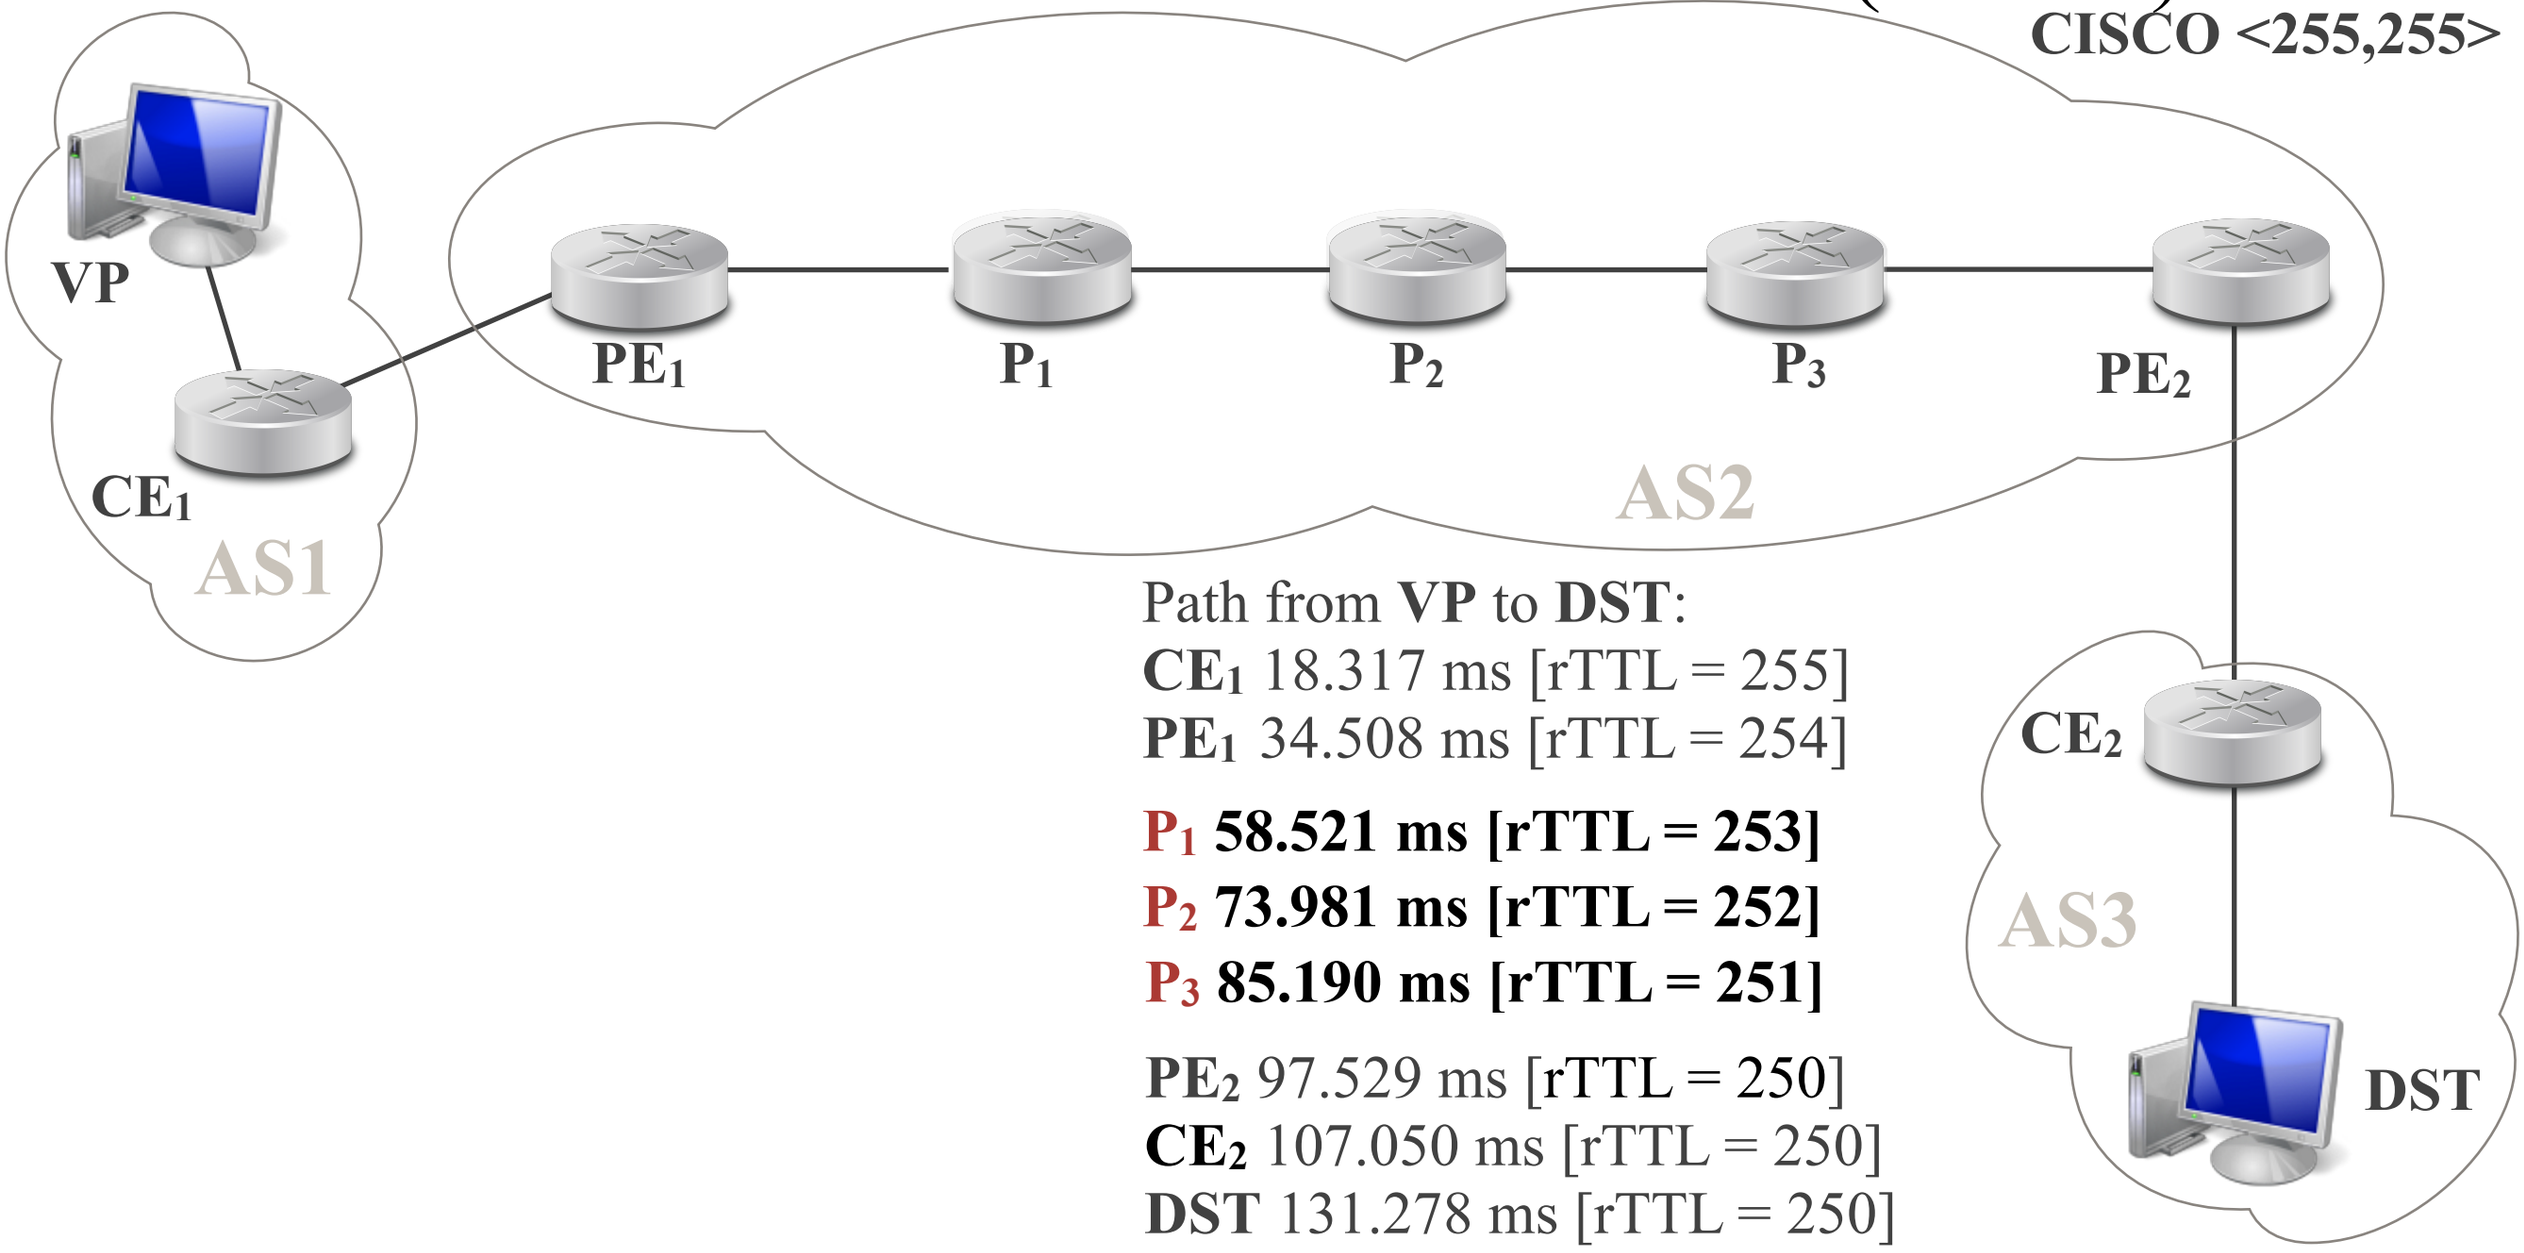

Path from **VP** to **DST**:
CE1 18.317 ms [rTTL = 255]
PE1 34.508 ms [rTTL = 254]

P1 58.521 ms [rTTL = 253]
P2 73.981 ms [rTTL = 252]
P3 85.190 ms [rTTL = 251]

PE2 97.529 ms [rTTL = 250]
CE2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

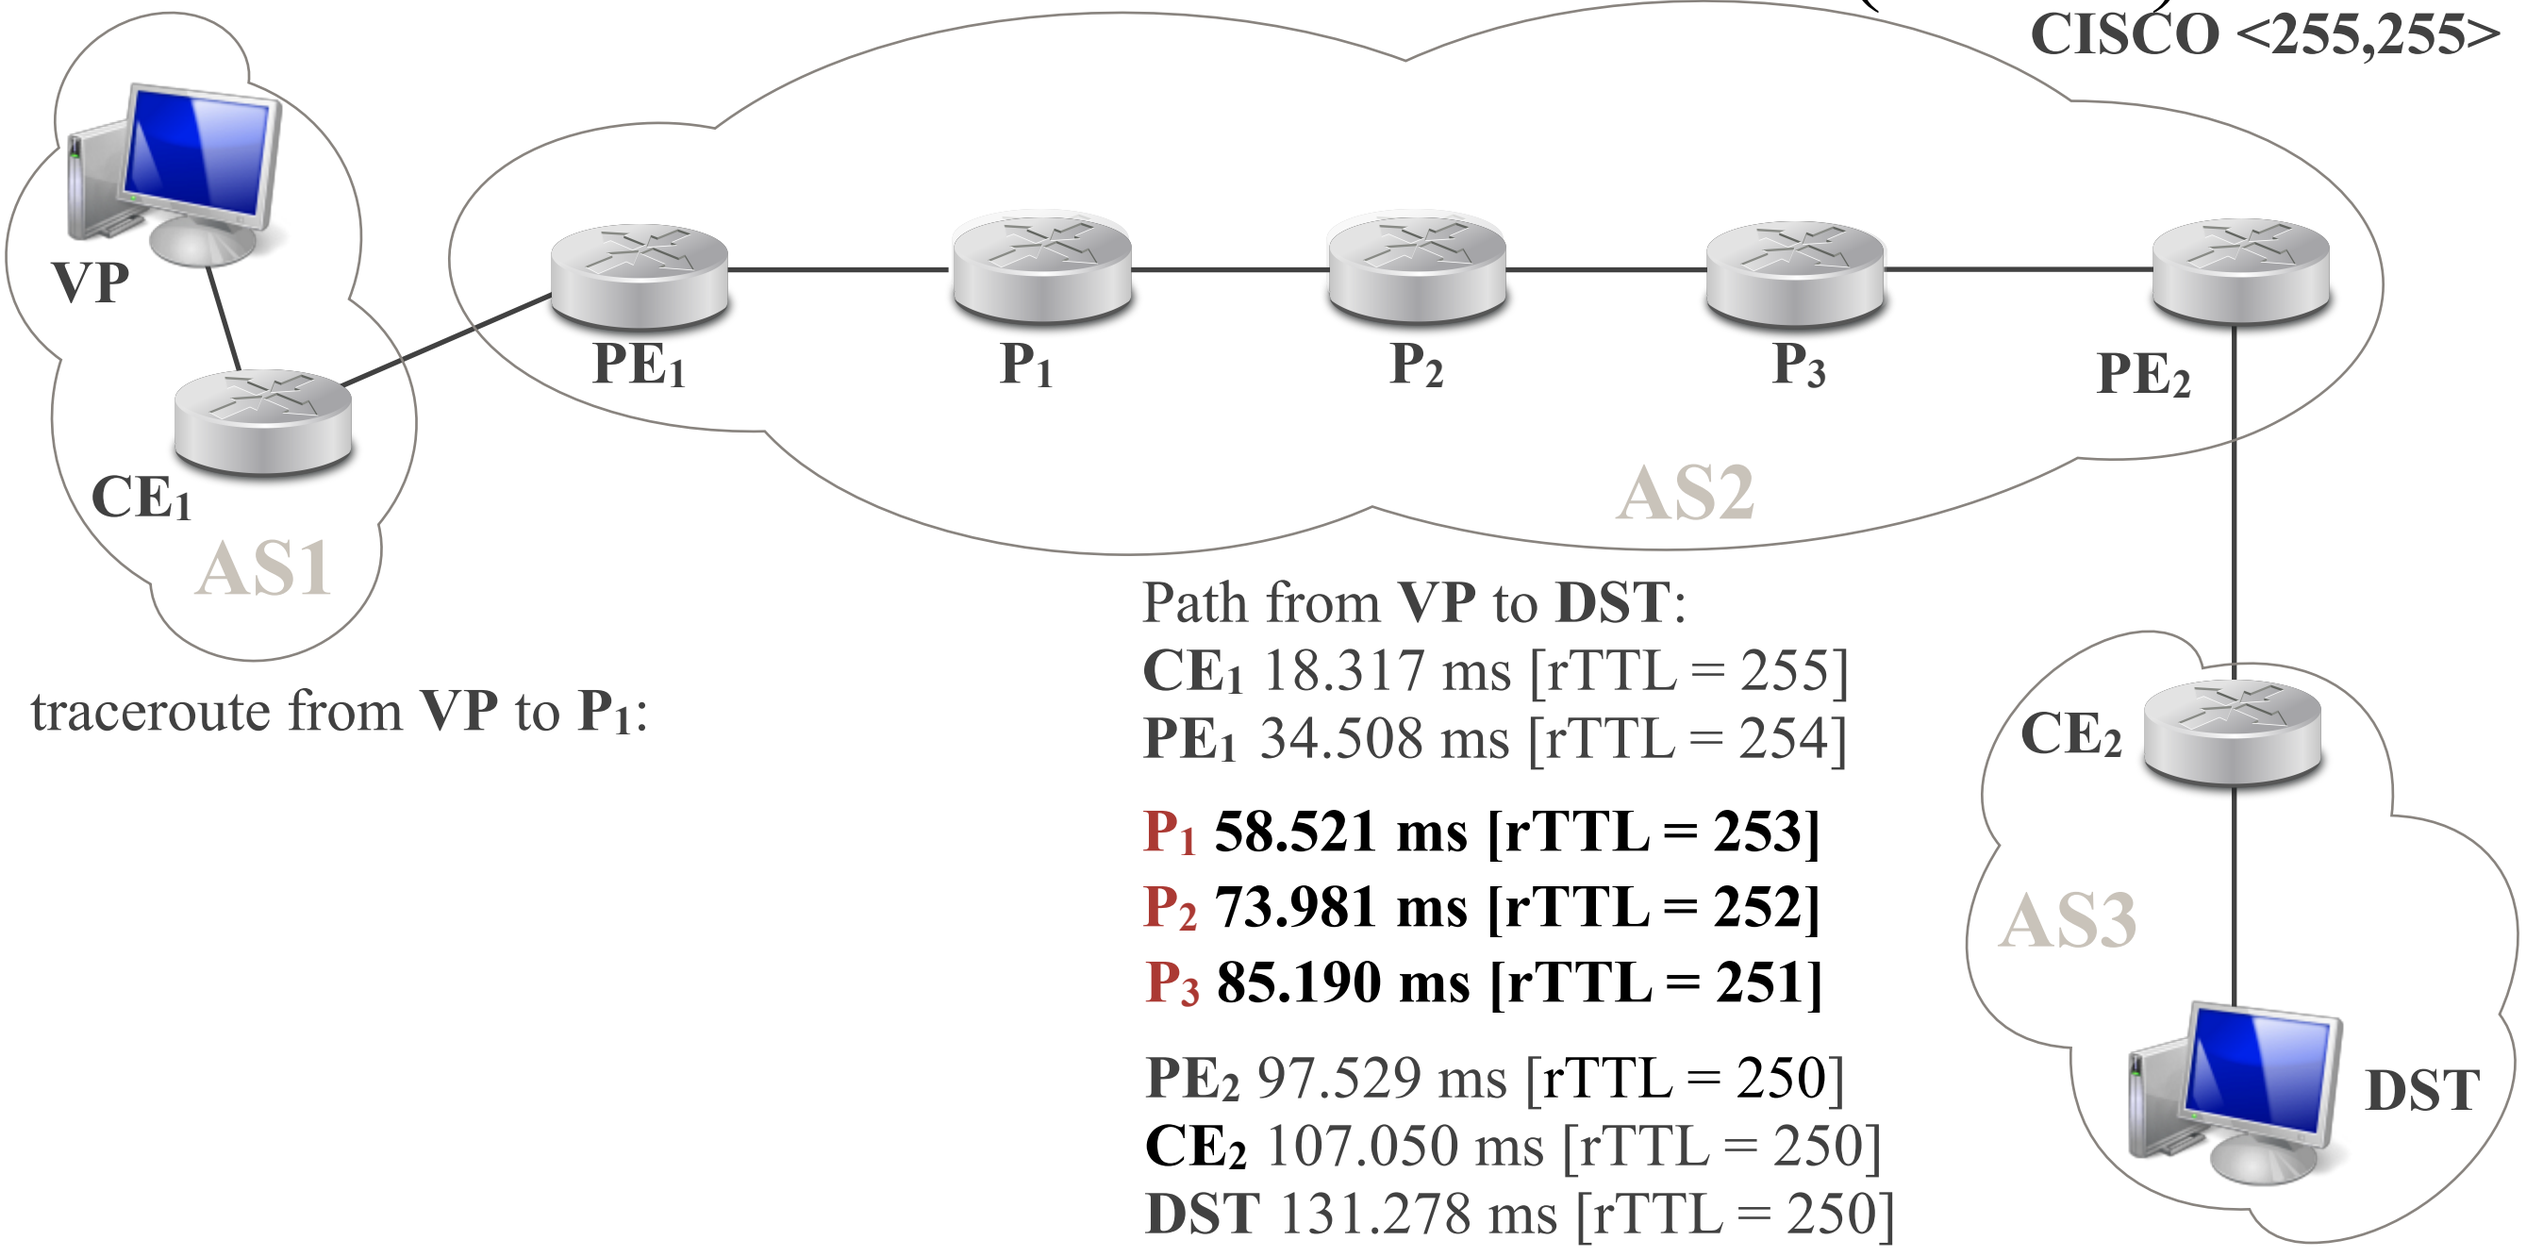

traceroute from **VP** to P_1:

Path from **VP** to **DST**:
CE_1 18.317 ms [rTTL = 255]
PE_1 34.508 ms [rTTL = 254]

P_1 58.521 ms [rTTL = 253]
P_2 73.981 ms [rTTL = 252]
P_3 85.190 ms [rTTL = 251]

PE_2 97.529 ms [rTTL = 250]
CE_2 107.050 ms [rTTL = 250]
DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

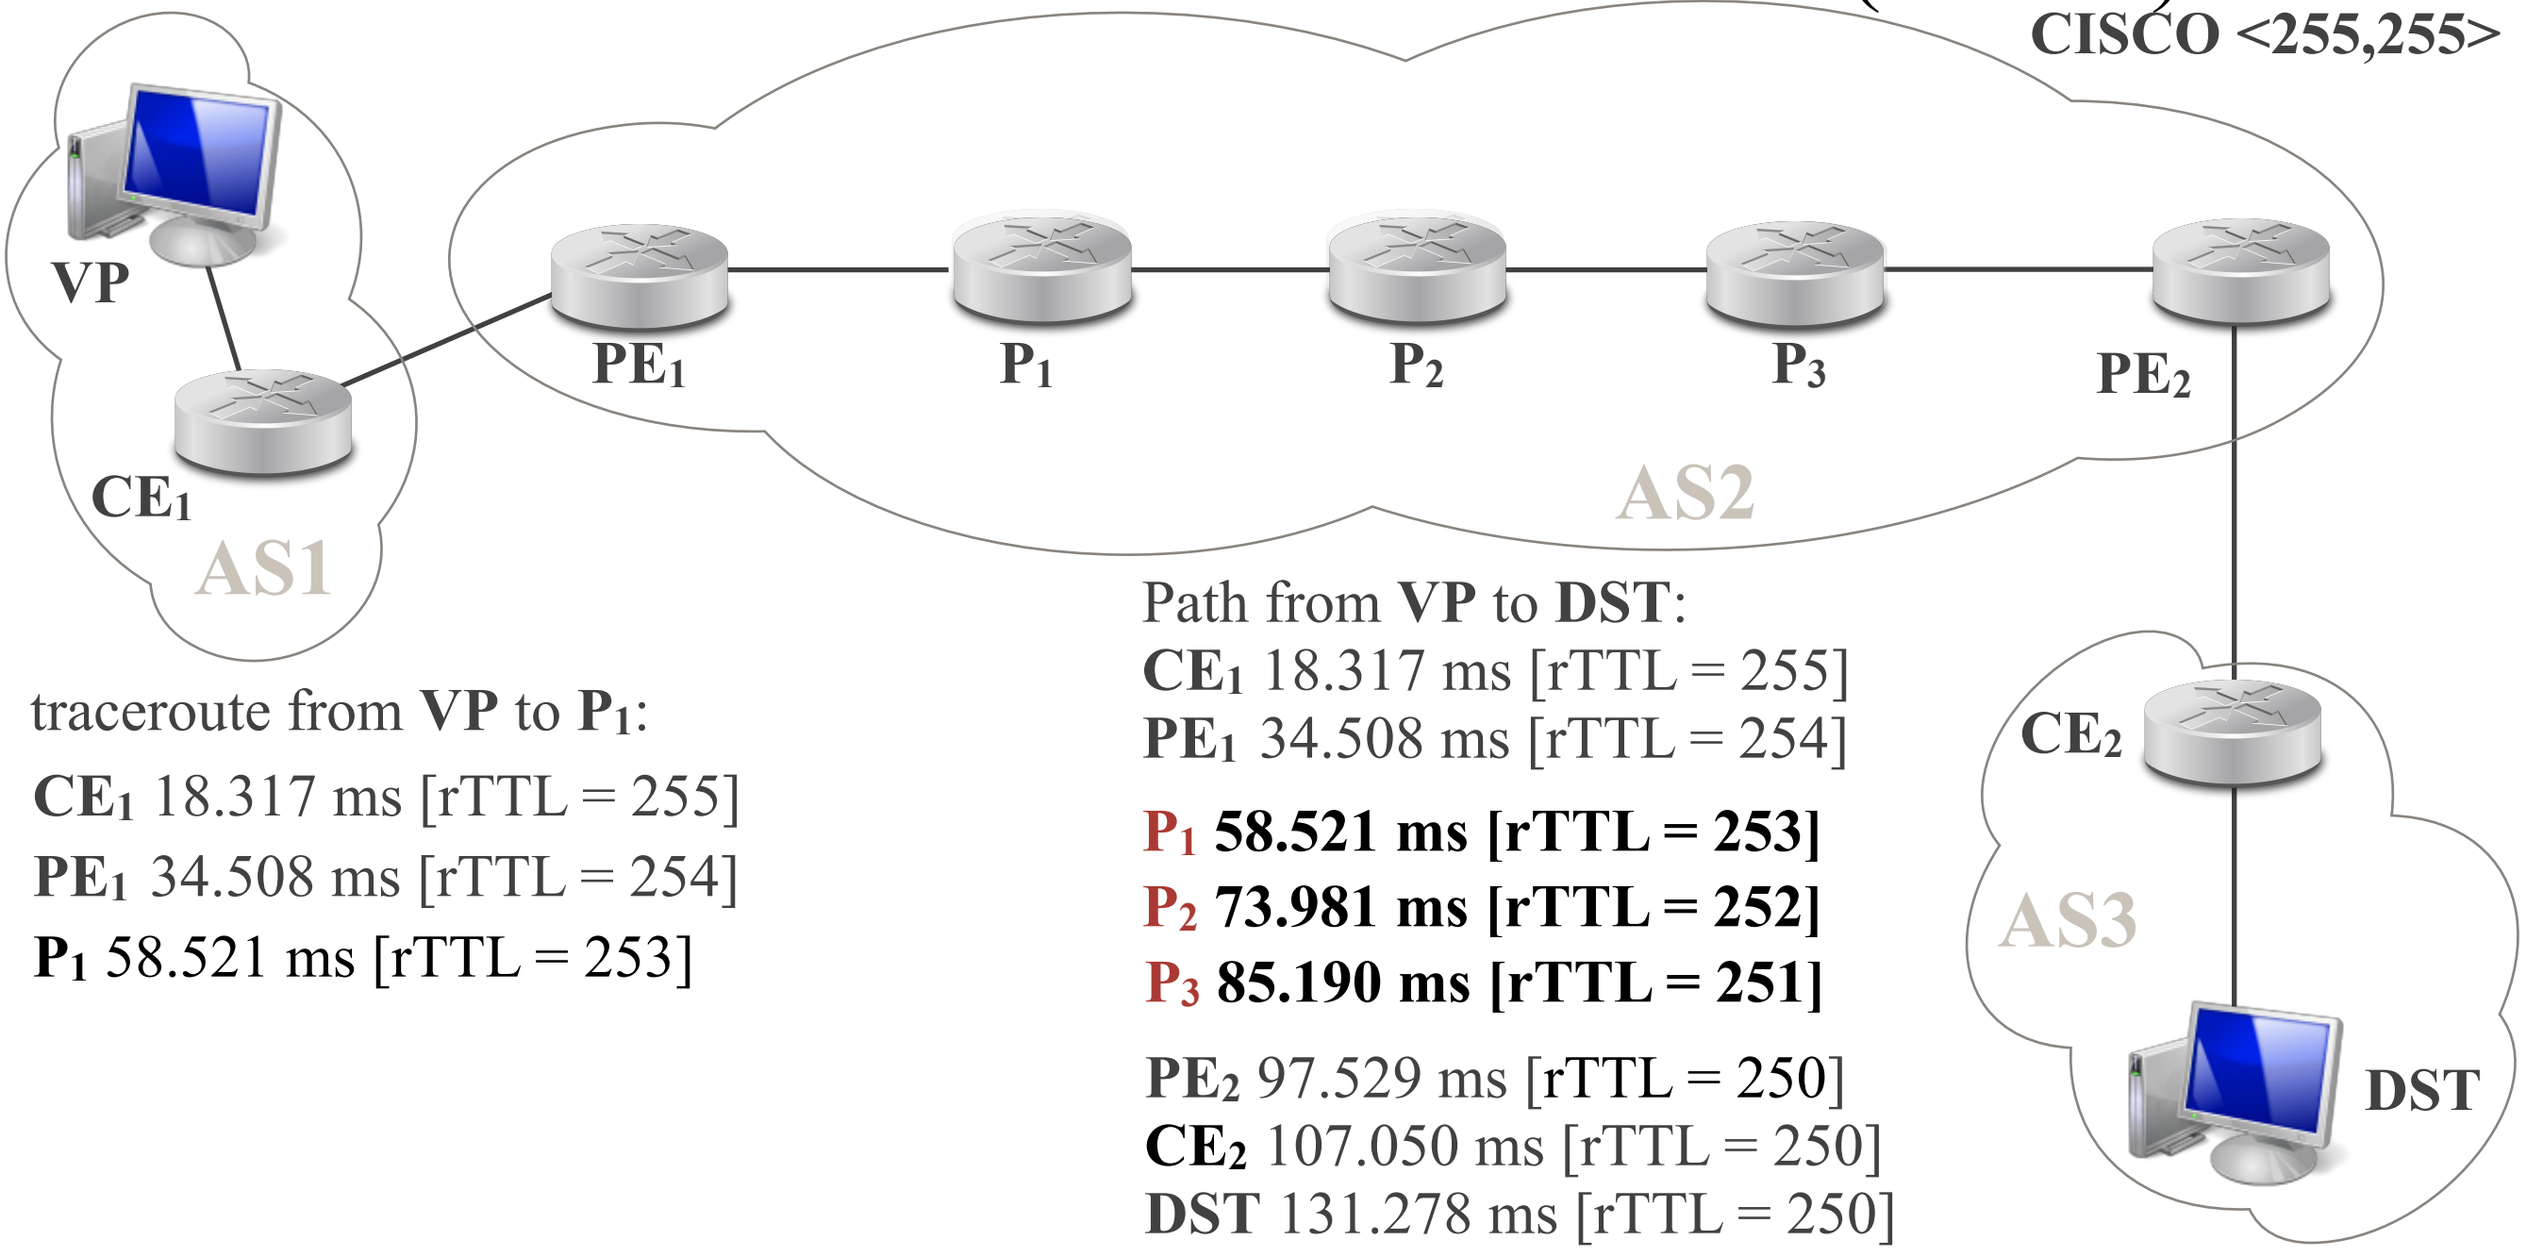

traceroute from **VP** to **P1**:

CE1 18.317 ms [rTTL = 255]

PE1 34.508 ms [rTTL = 254]

P1 58.521 ms [rTTL = 253]

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

CE1 18.317 ms [rTTL = 255]

PE1 34.508 ms [rTTL = 254]

P1 58.521 ms [rTTL = 253]

P2 73.981 ms [rTTL = 252]

P3 85.190 ms [rTTL = 251]

PE2 97.529 ms [rTTL = 250]

CE2 107.050 ms [rTTL = 250]

DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

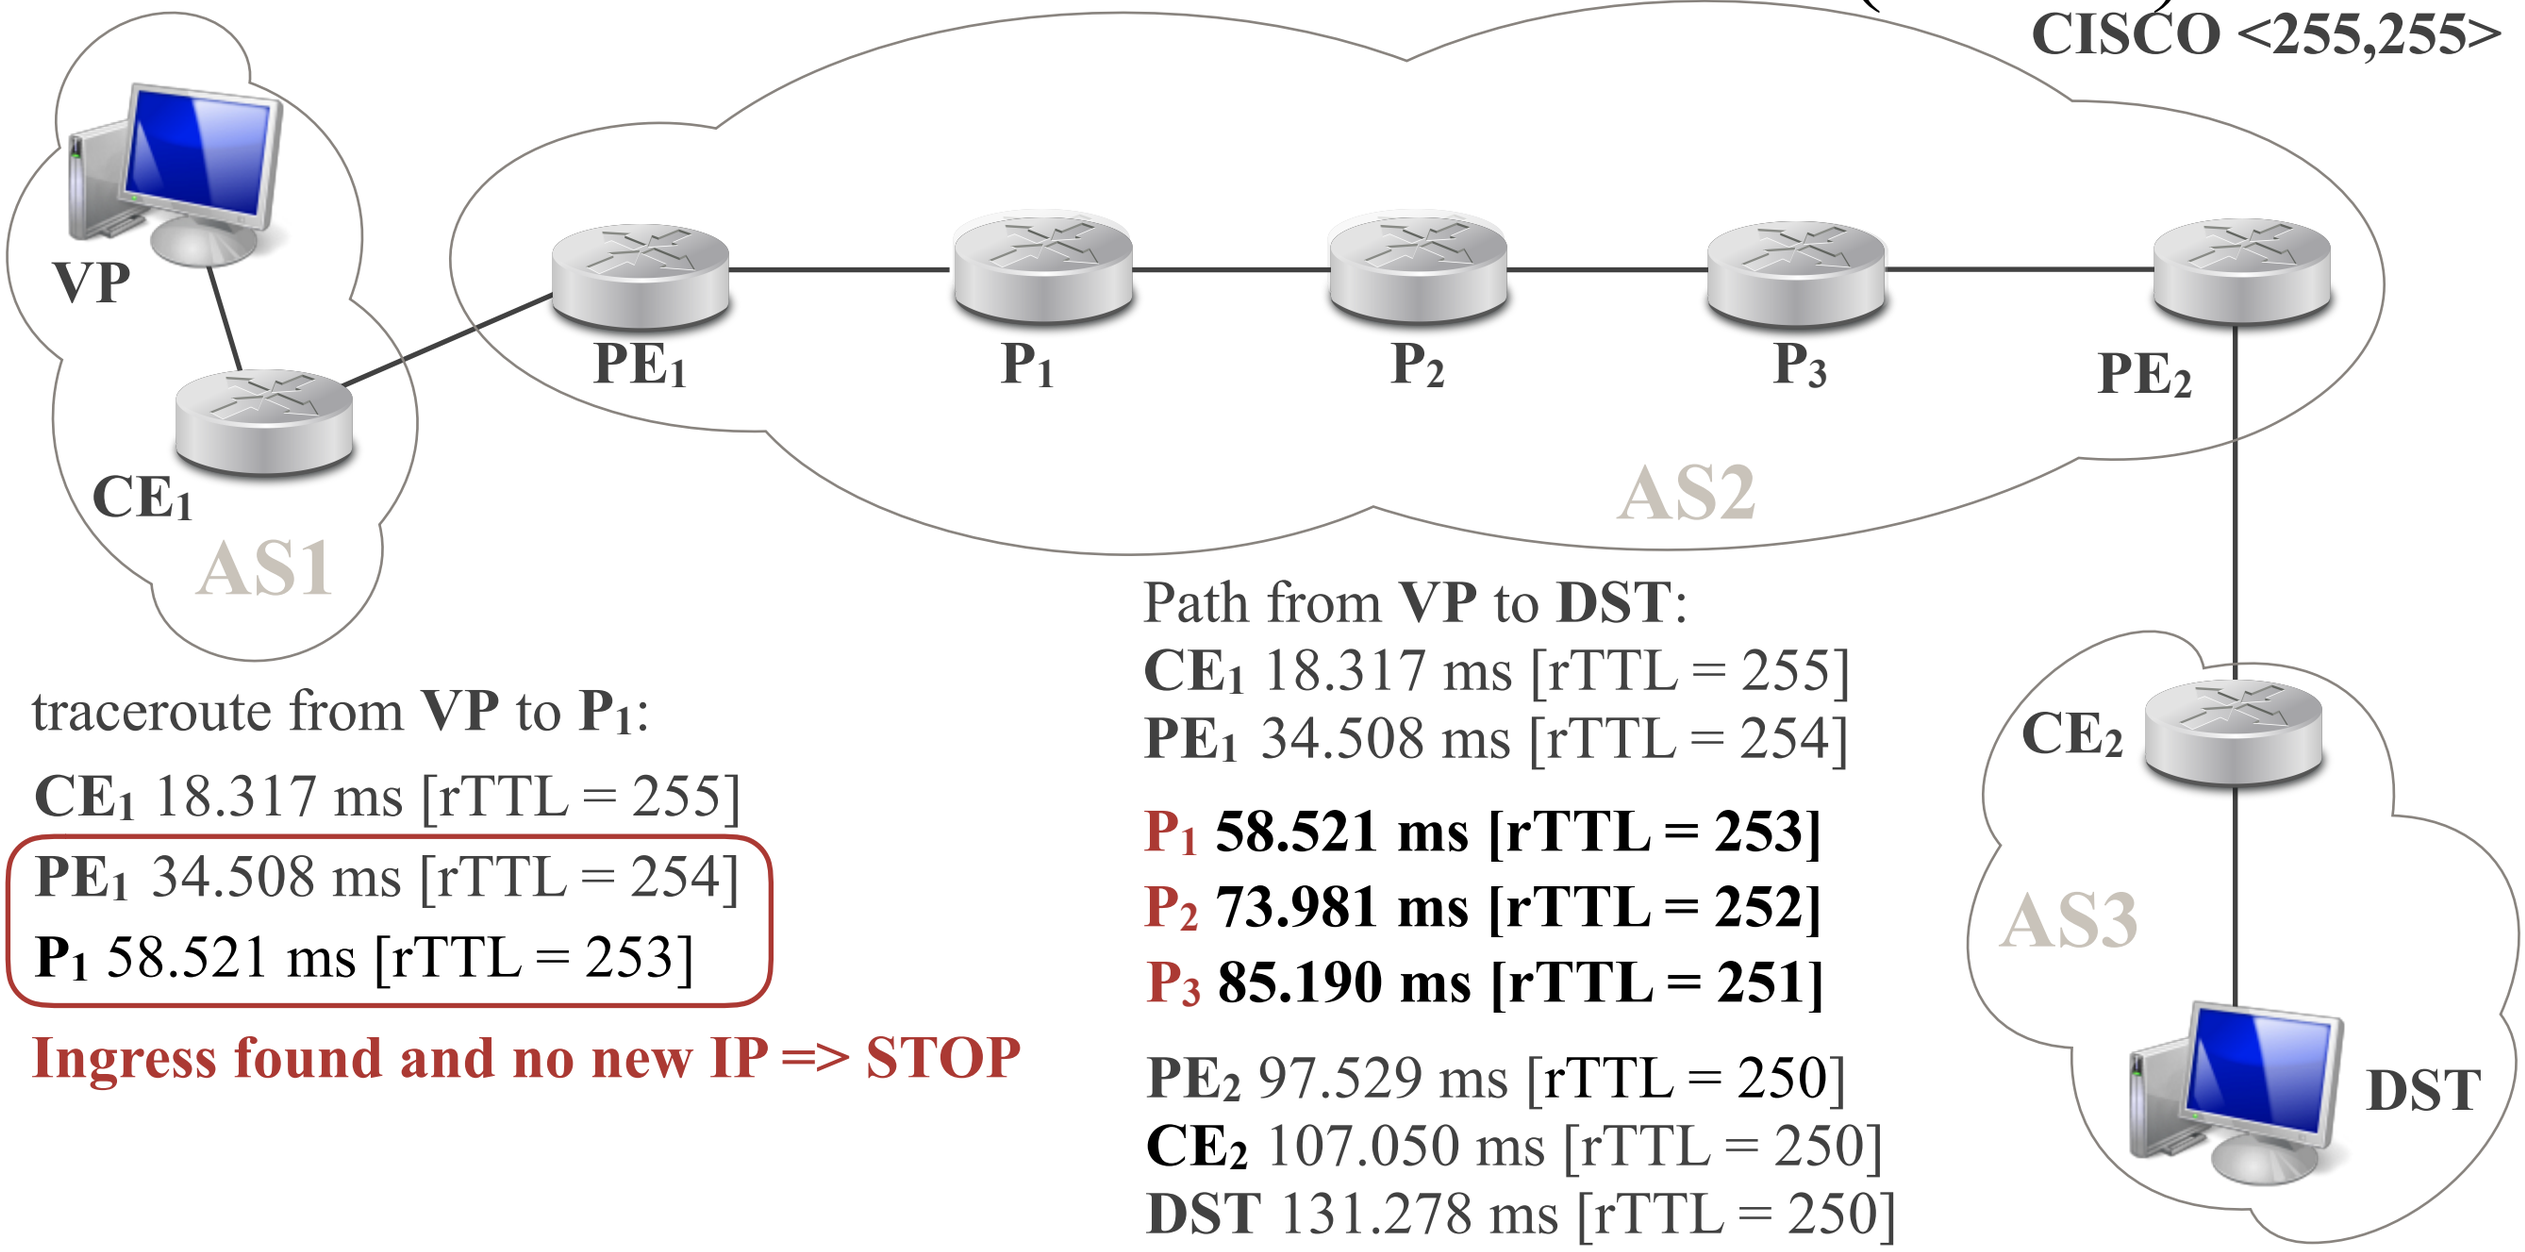

traceroute from **VP** to **P1**:

CE1 18.317 ms [rTTL = 255]

PE1 34.508 ms [rTTL = 254]

P1 58.521 ms [rTTL = 253]

Ingress found and no new IP => STOP

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

CE1 18.317 ms [rTTL = 255]

PE1 34.508 ms [rTTL = 254]

P1 58.521 ms [rTTL = 253]

P2 73.981 ms [rTTL = 252]

P3 85.190 ms [rTTL = 251]

PE2 97.529 ms [rTTL = 250]

CE2 107.050 ms [rTTL = 250]

DST 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (7)

- Backward Recursive Path Revelation (BRPR)

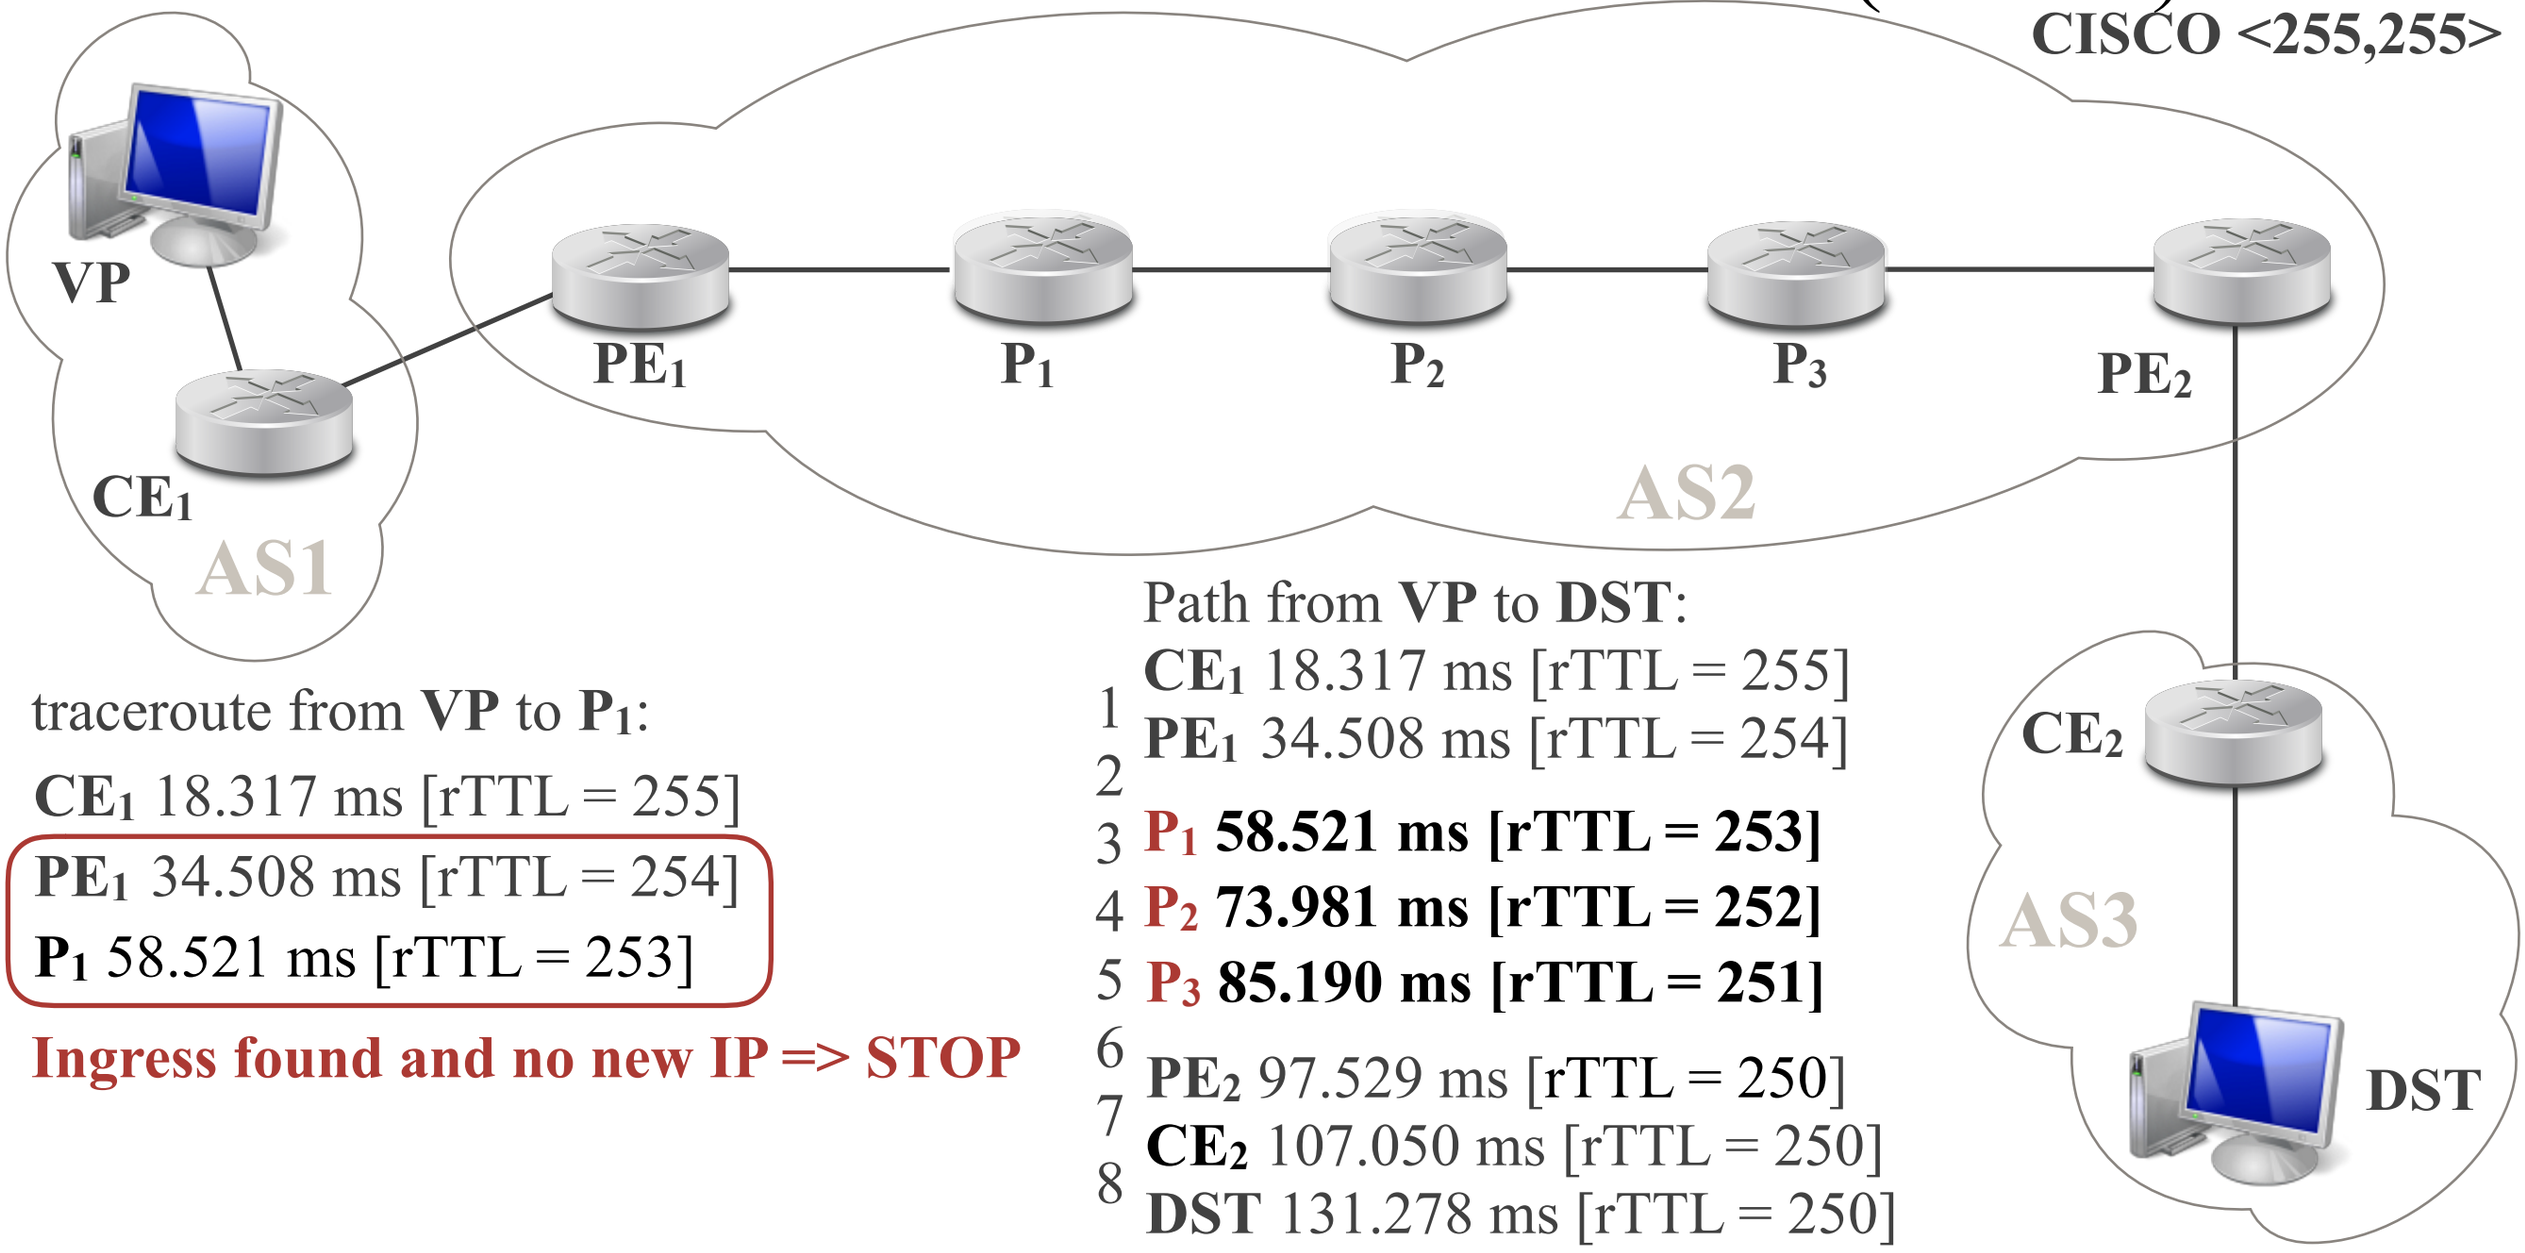

traceroute from **VP** to **P1**:

CE1 18.317 ms [rTTL = 255]

PE1 34.508 ms [rTTL = 254]

P1 58.521 ms [rTTL = 253]

Ingress found and no new IP => STOP

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

1 **CE1** 18.317 ms [rTTL = 255]

2 **PE1** 34.508 ms [rTTL = 254]

3 **P1 58.521 ms [rTTL = 253]**

4 **P2 73.981 ms [rTTL = 252]**

5 **P3 85.190 ms [rTTL = 251]**

6 **PE2** 97.529 ms [rTTL = 250]

7 **CE2** 107.050 ms [rTTL = 250]

8 **DST** 131.278 ms [rTTL = 250]

MPLS is used for internal traffic, with PHP enabled

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

Invisible Tunnels (8)

- What can we do with a duplicate IP address?

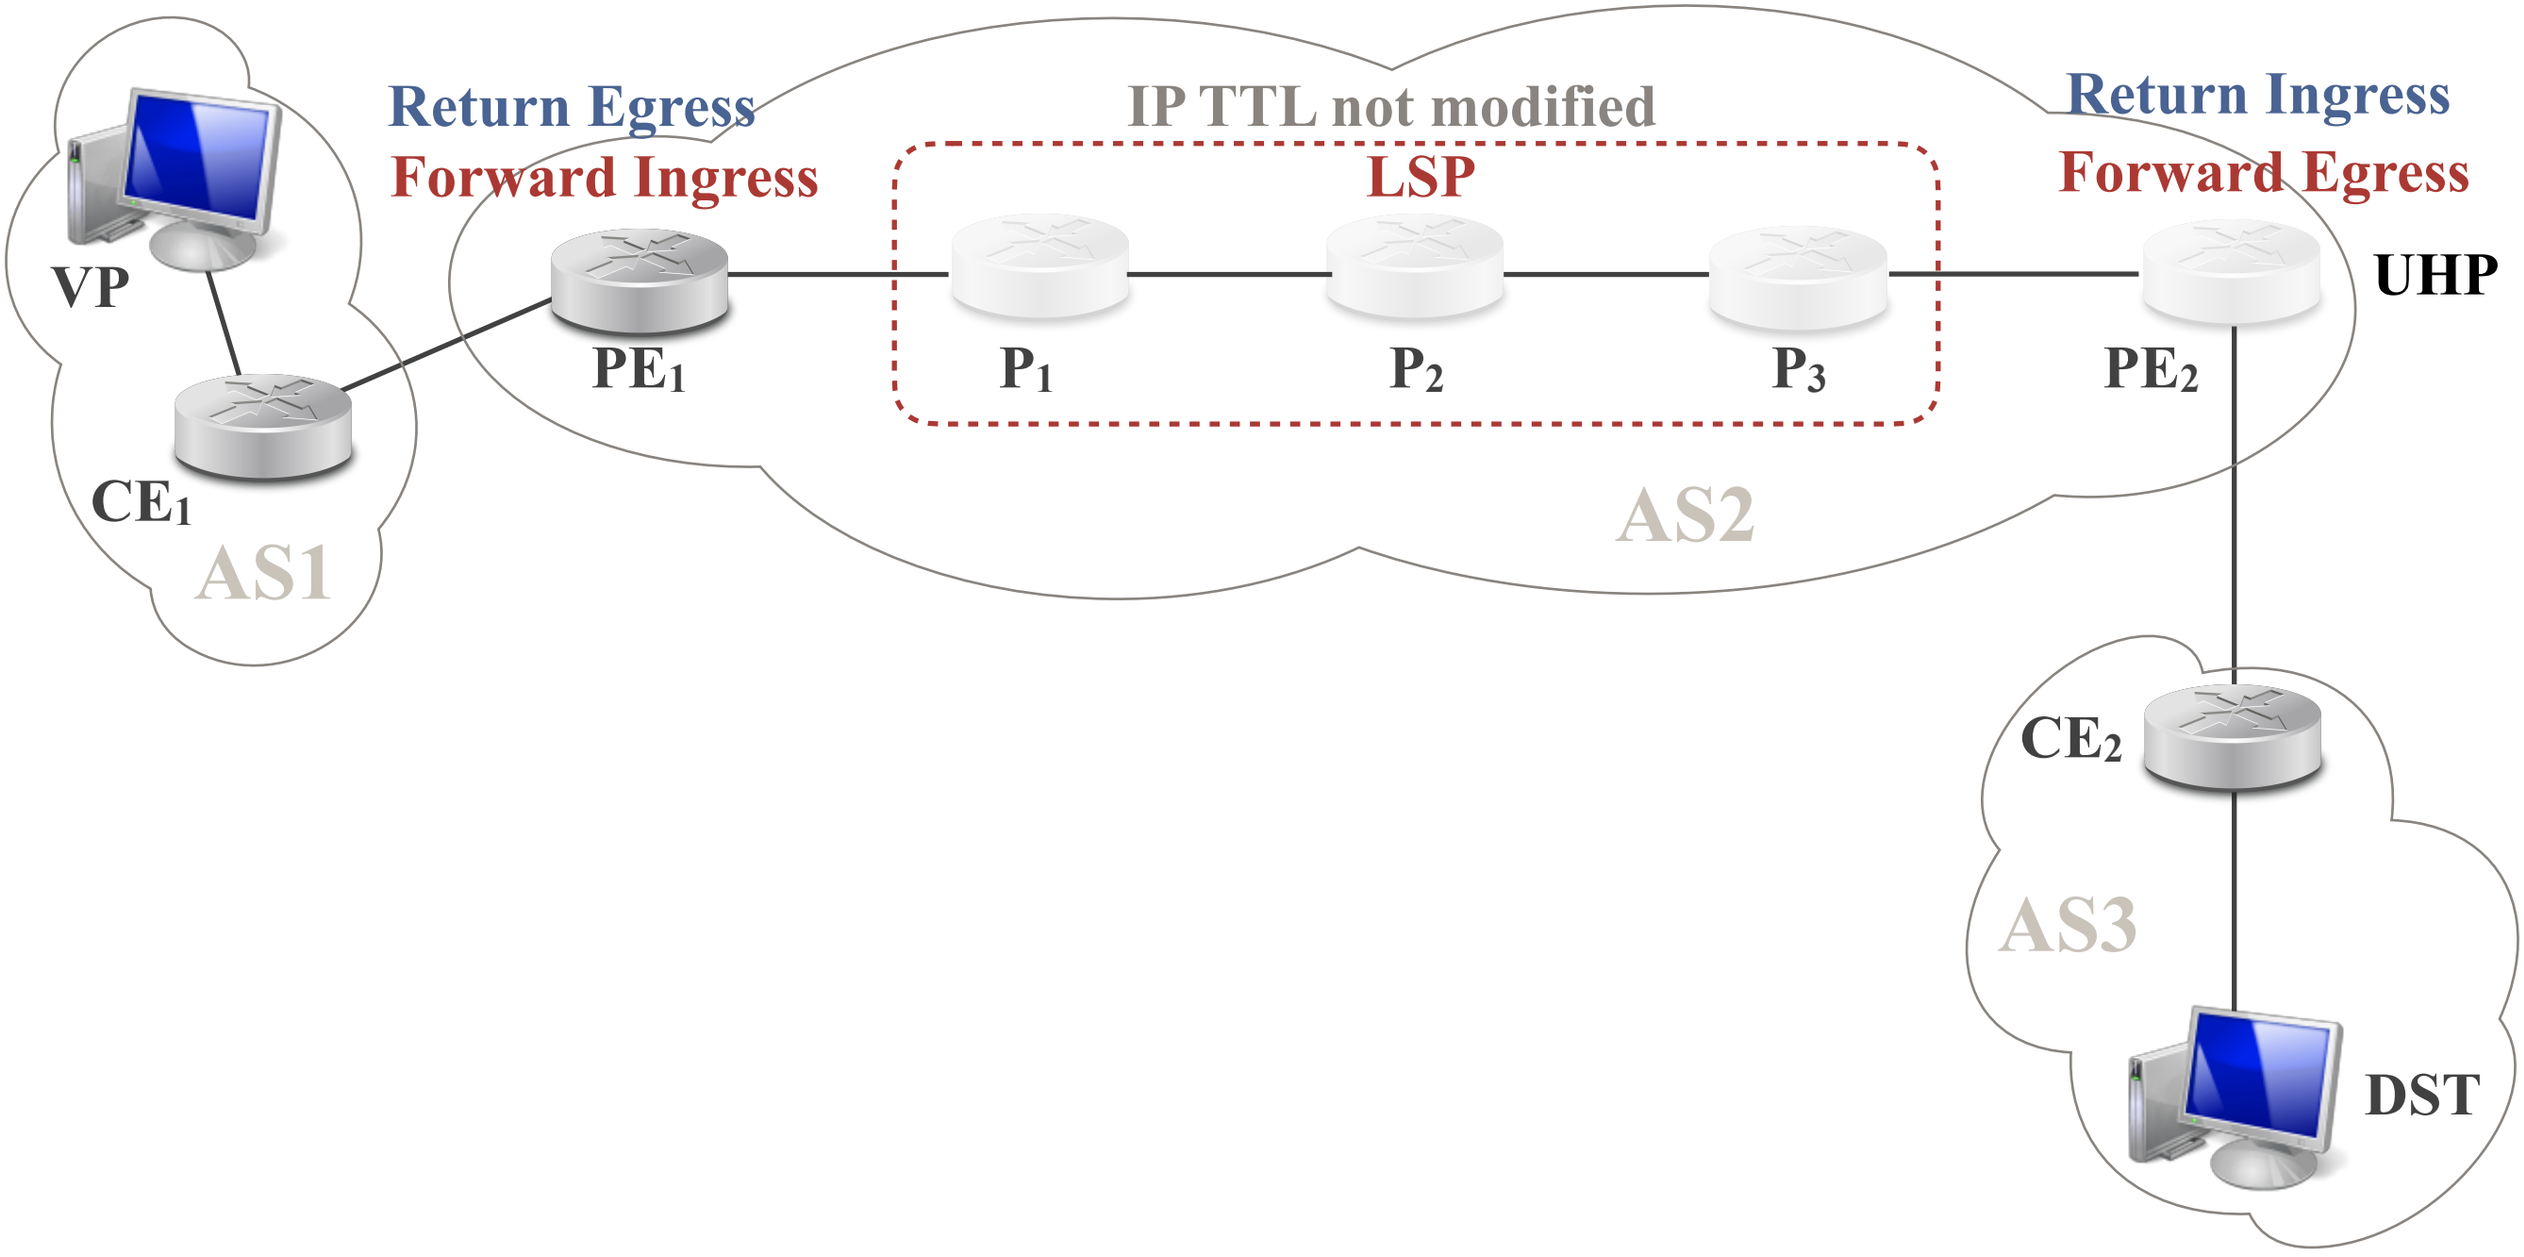

Invisible Tunnels (8)

- What can we do with a duplicate IP address?

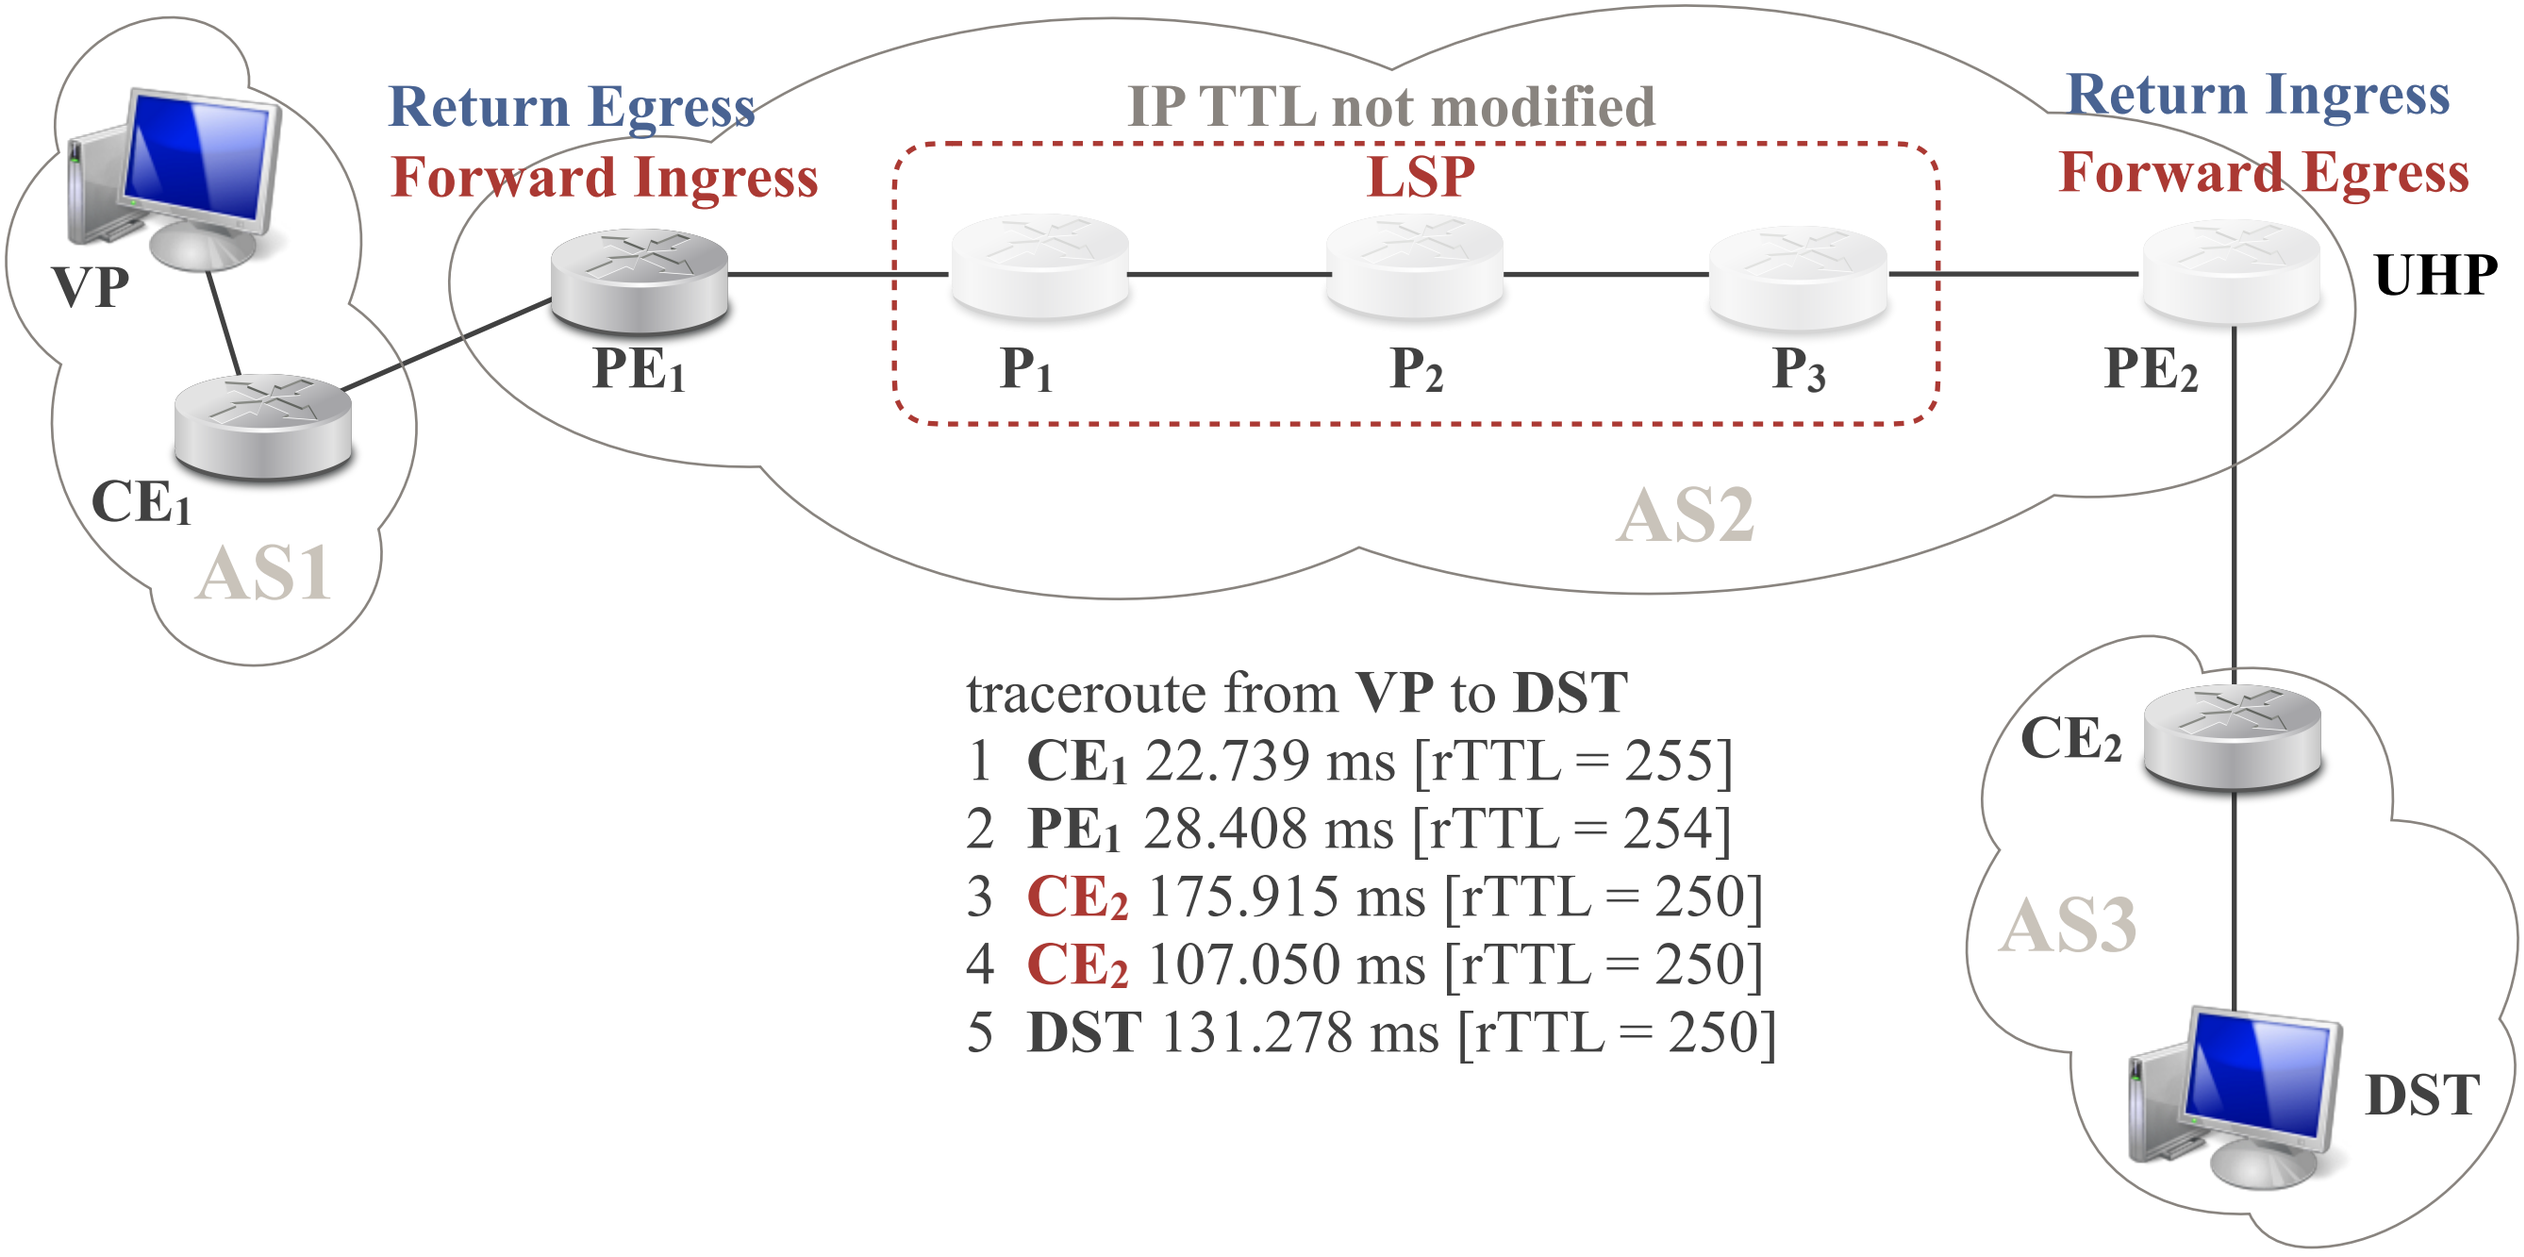

traceroute from **VP** to **DST**

1 **CE$_1$** 22.739 ms [rTTL = 255]
2 **PE$_1$** 28.408 ms [rTTL = 254]
3 **CE$_2$** 175.915 ms [rTTL = 250]
4 **CE$_2$** 107.050 ms [rTTL = 250]
5 **DST** 131.278 ms [rTTL = 250]

Invisible Tunnels (8)

- What can we do with a duplicate IP address?

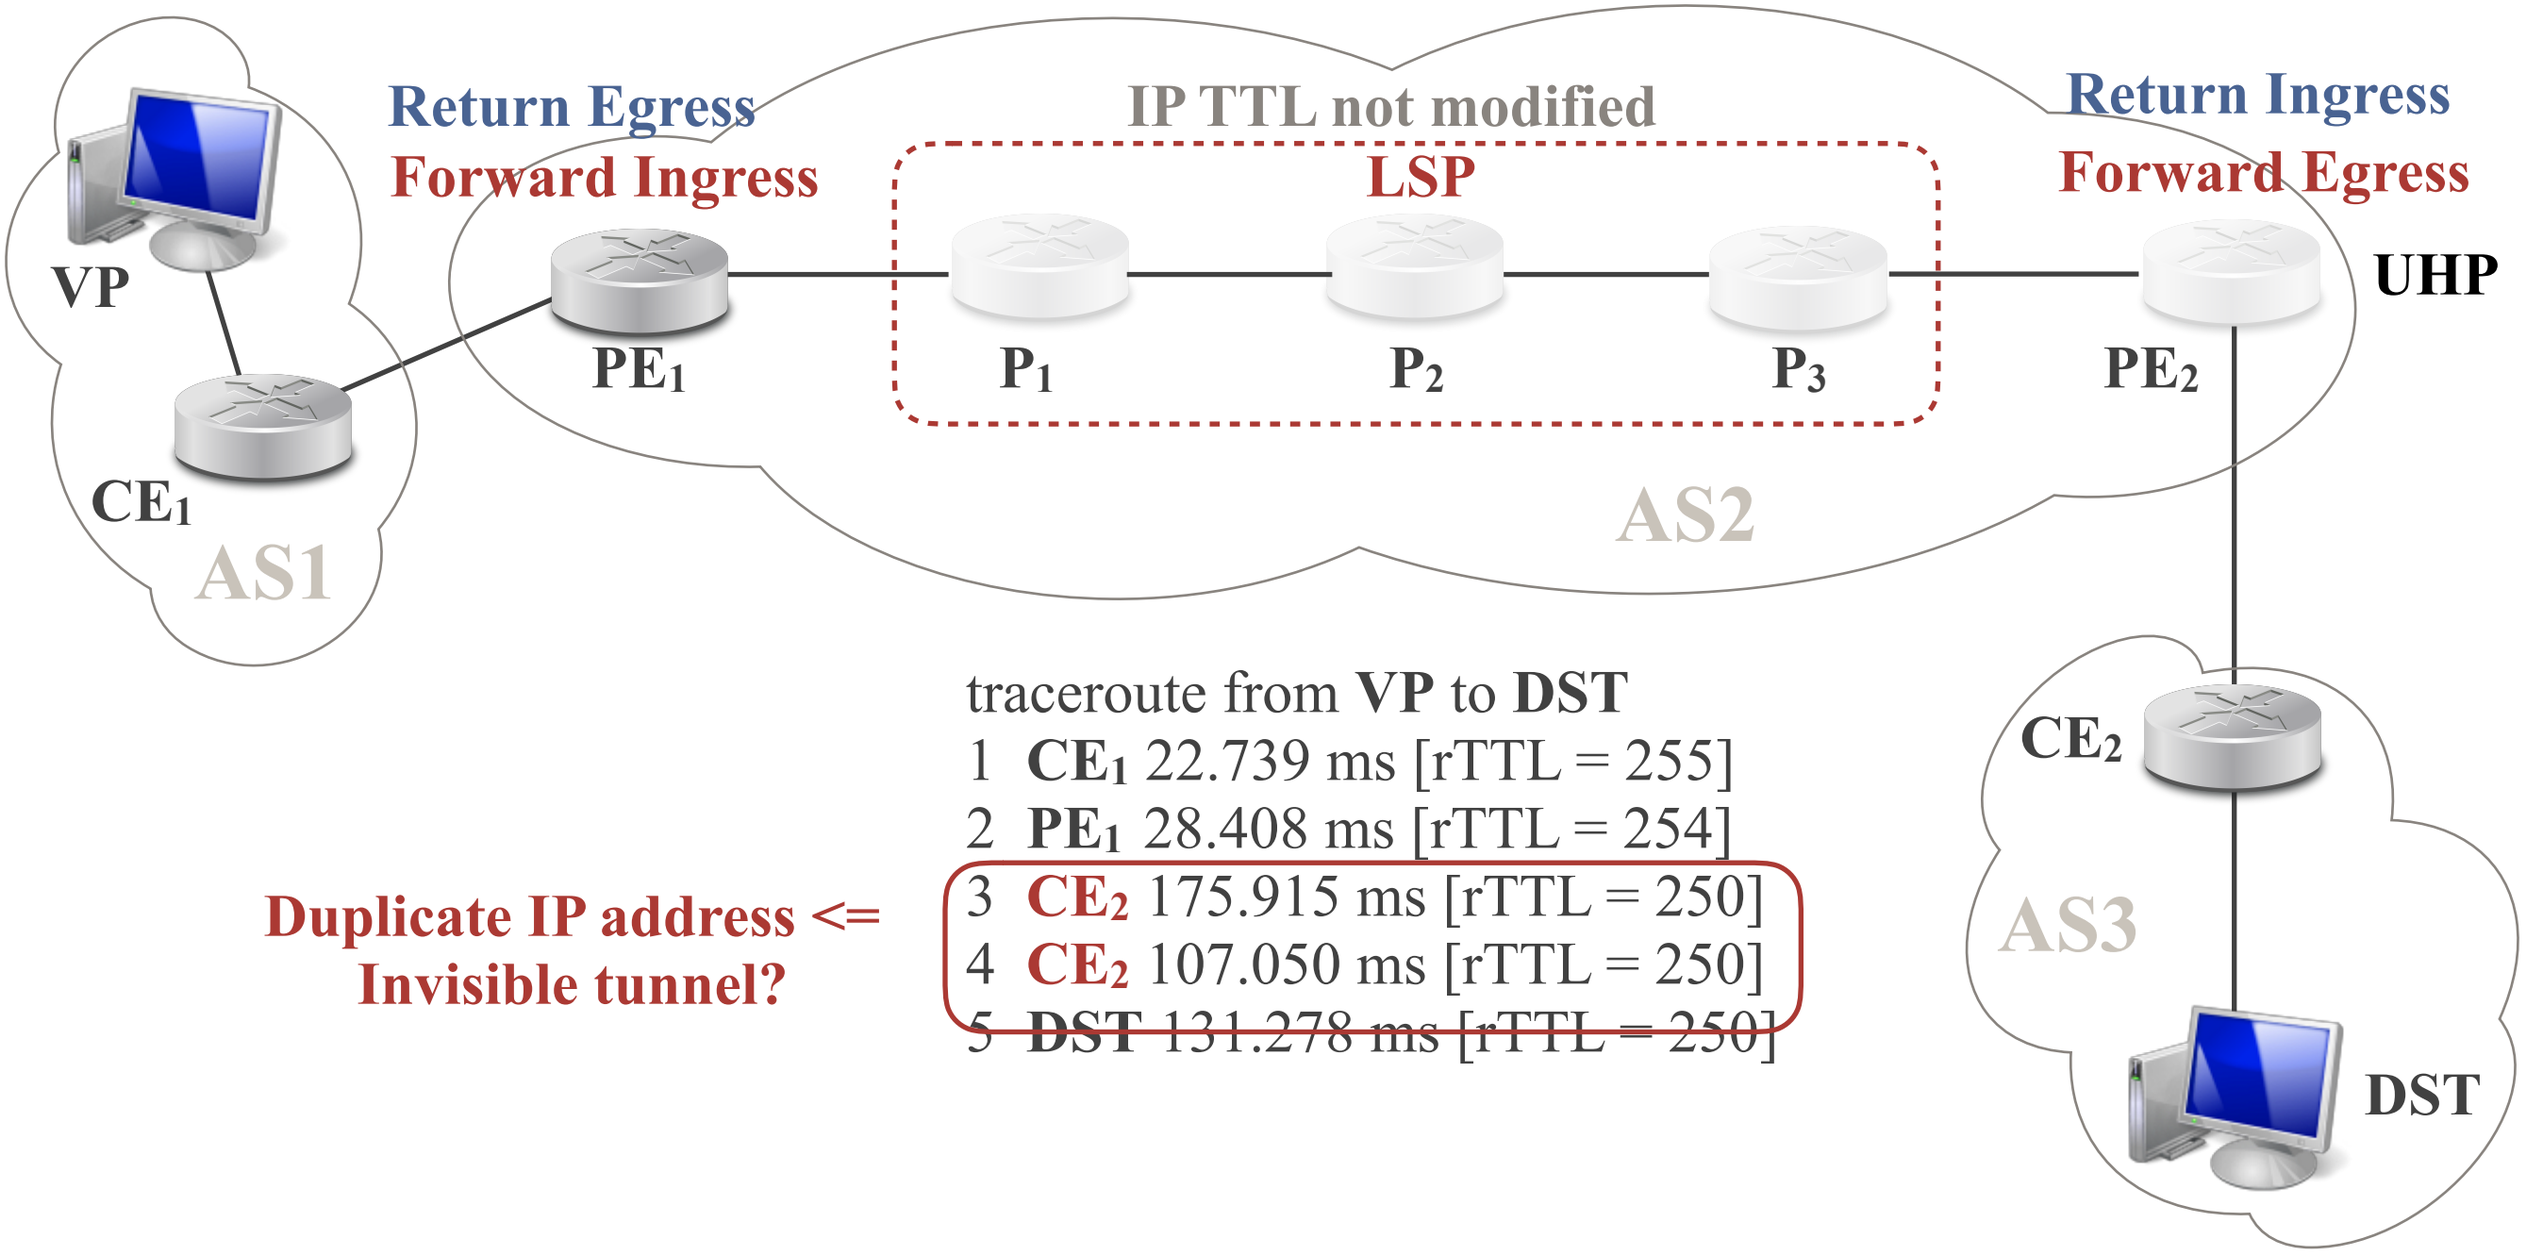

Duplicate IP address <=
Invisible tunnel?

traceroute from **VP** to **DST**
1 **CE1** 22.739 ms [rTTL = 255]
2 **PE1** 28.408 ms [rTTL = 254]
3 **CE2** 175.915 ms [rTTL = 250]
4 **CE2** 107.050 ms [rTTL = 250]
5 ~~**DST** 131.278 ms [rTTL = 250]~~

Invisible Tunnels (8)

- What can we do with a duplicate IP address?

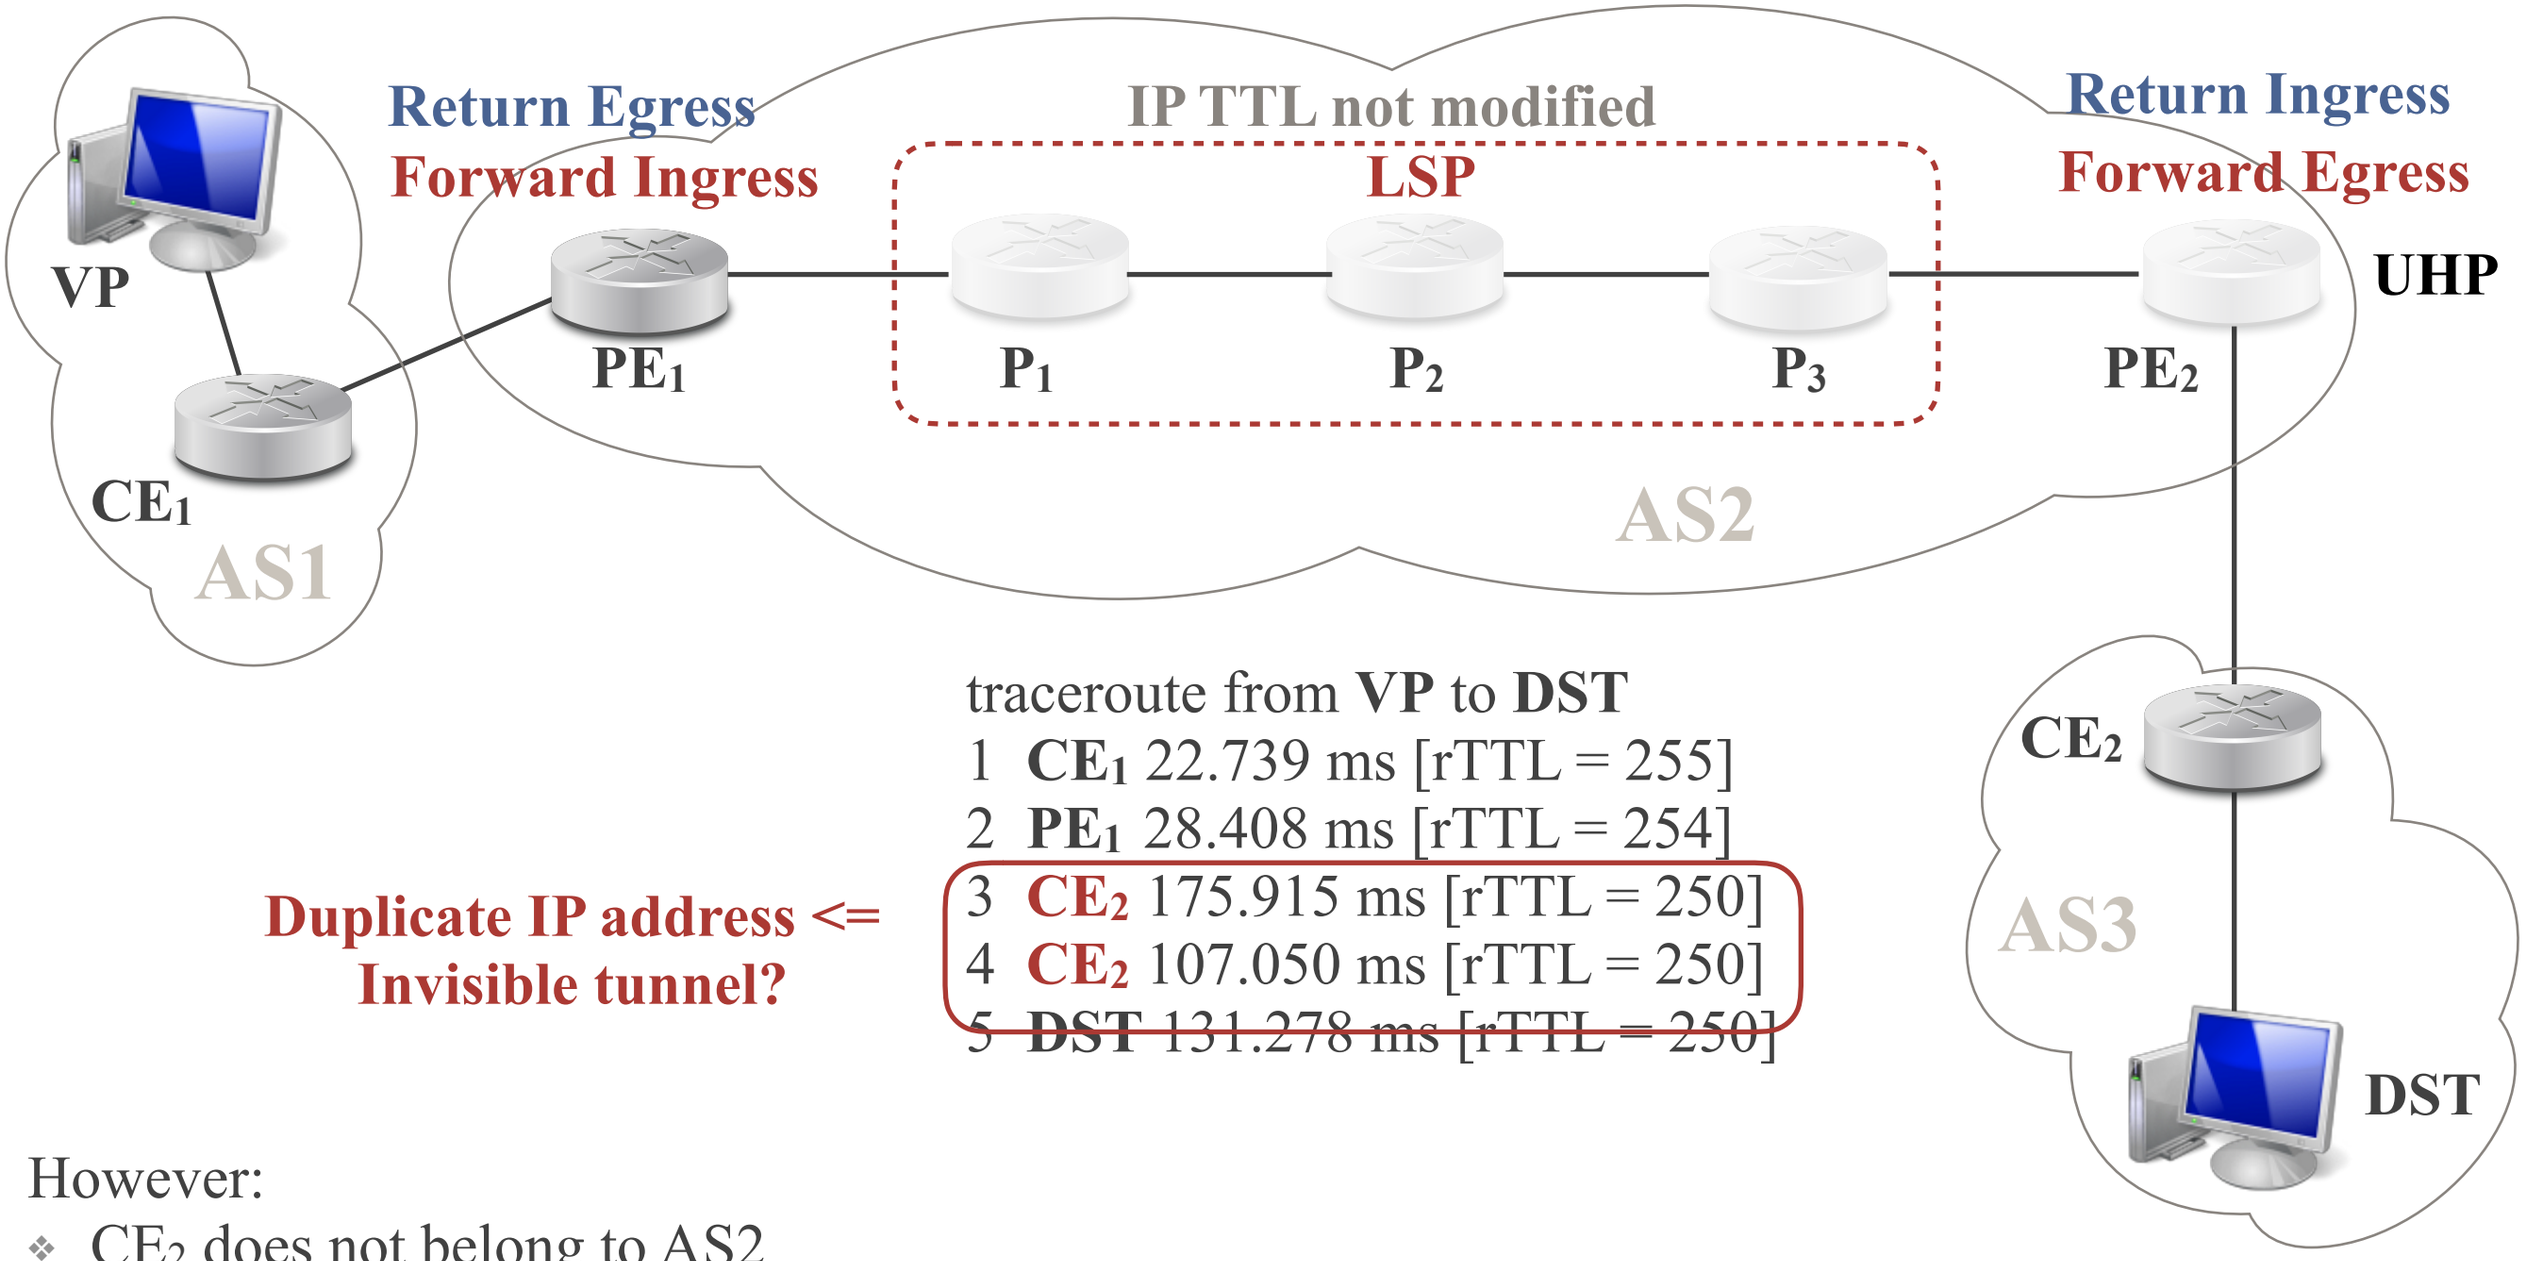

Duplicate IP address <=
Invisible tunnel?

traceroute from **VP** to **DST**
1 **CE1** 22.739 ms [rTTL = 255]
2 **PE1** 28.408 ms [rTTL = 254]
3 **CE2** 175.915 ms [rTTL = 250]
4 **CE2** 107.050 ms [rTTL = 250]
5 ~~**DST** 131.278 ms [rTTL = 250]~~

However:

- CE_2 does not belong to AS2
- **Probing CE_2 will not reveal any LSR! (not an internal prefix)**
- **Need to find a way to probe an internal prefix**

Invisible Tunnels (9)

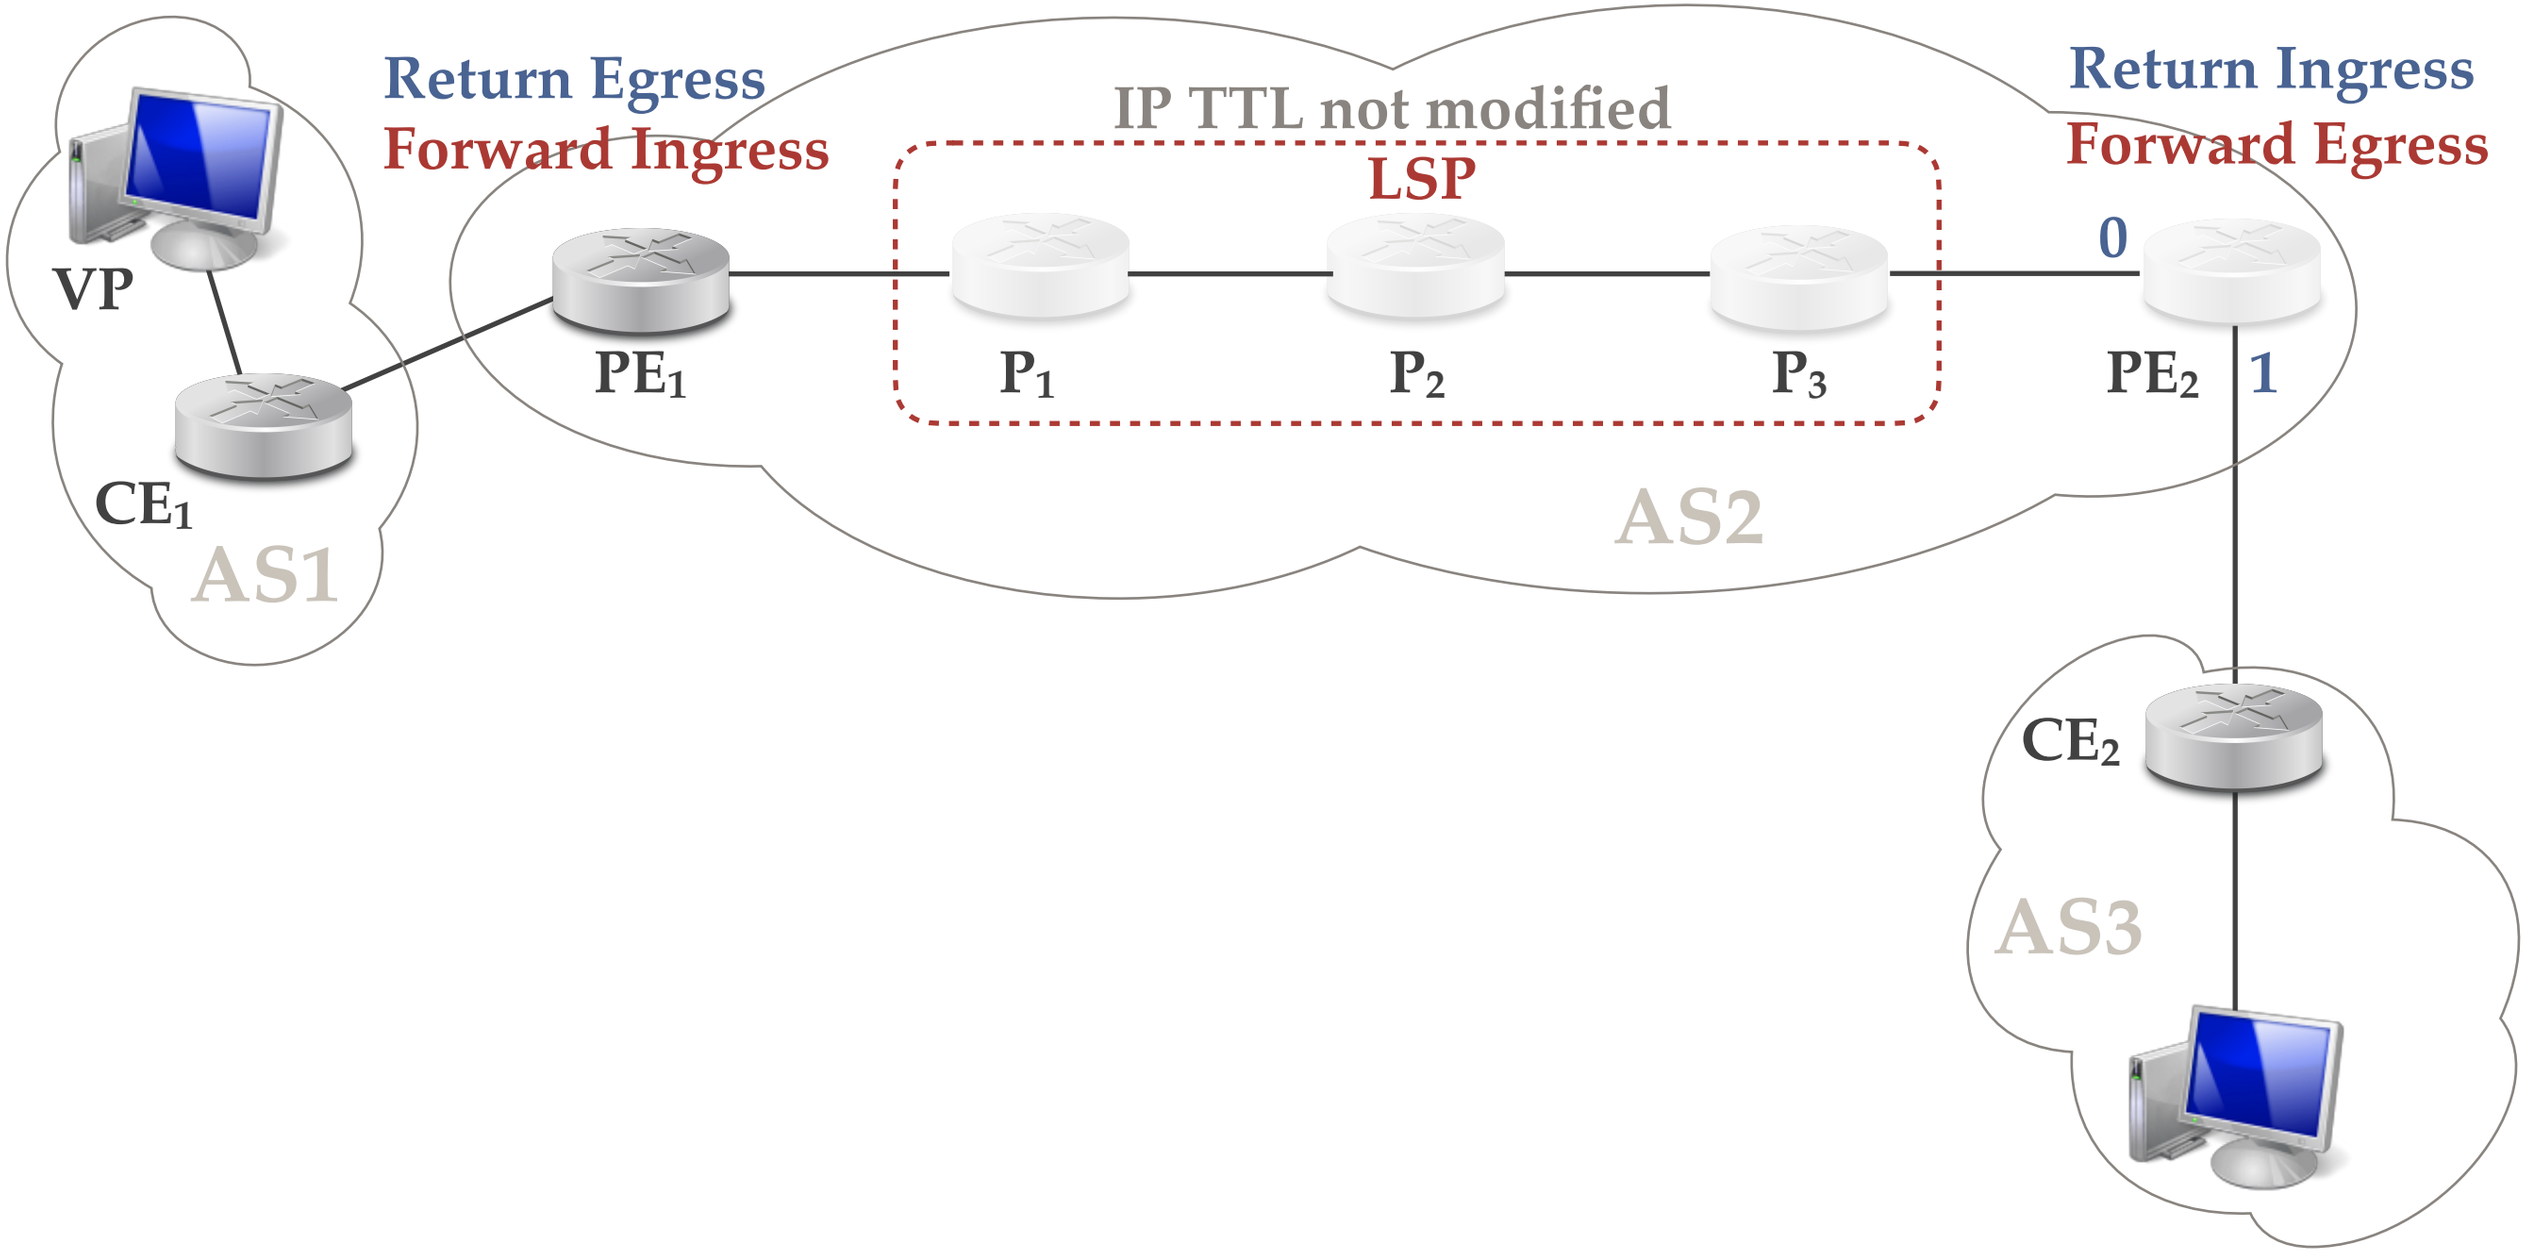

Invisible Tunnels (9)

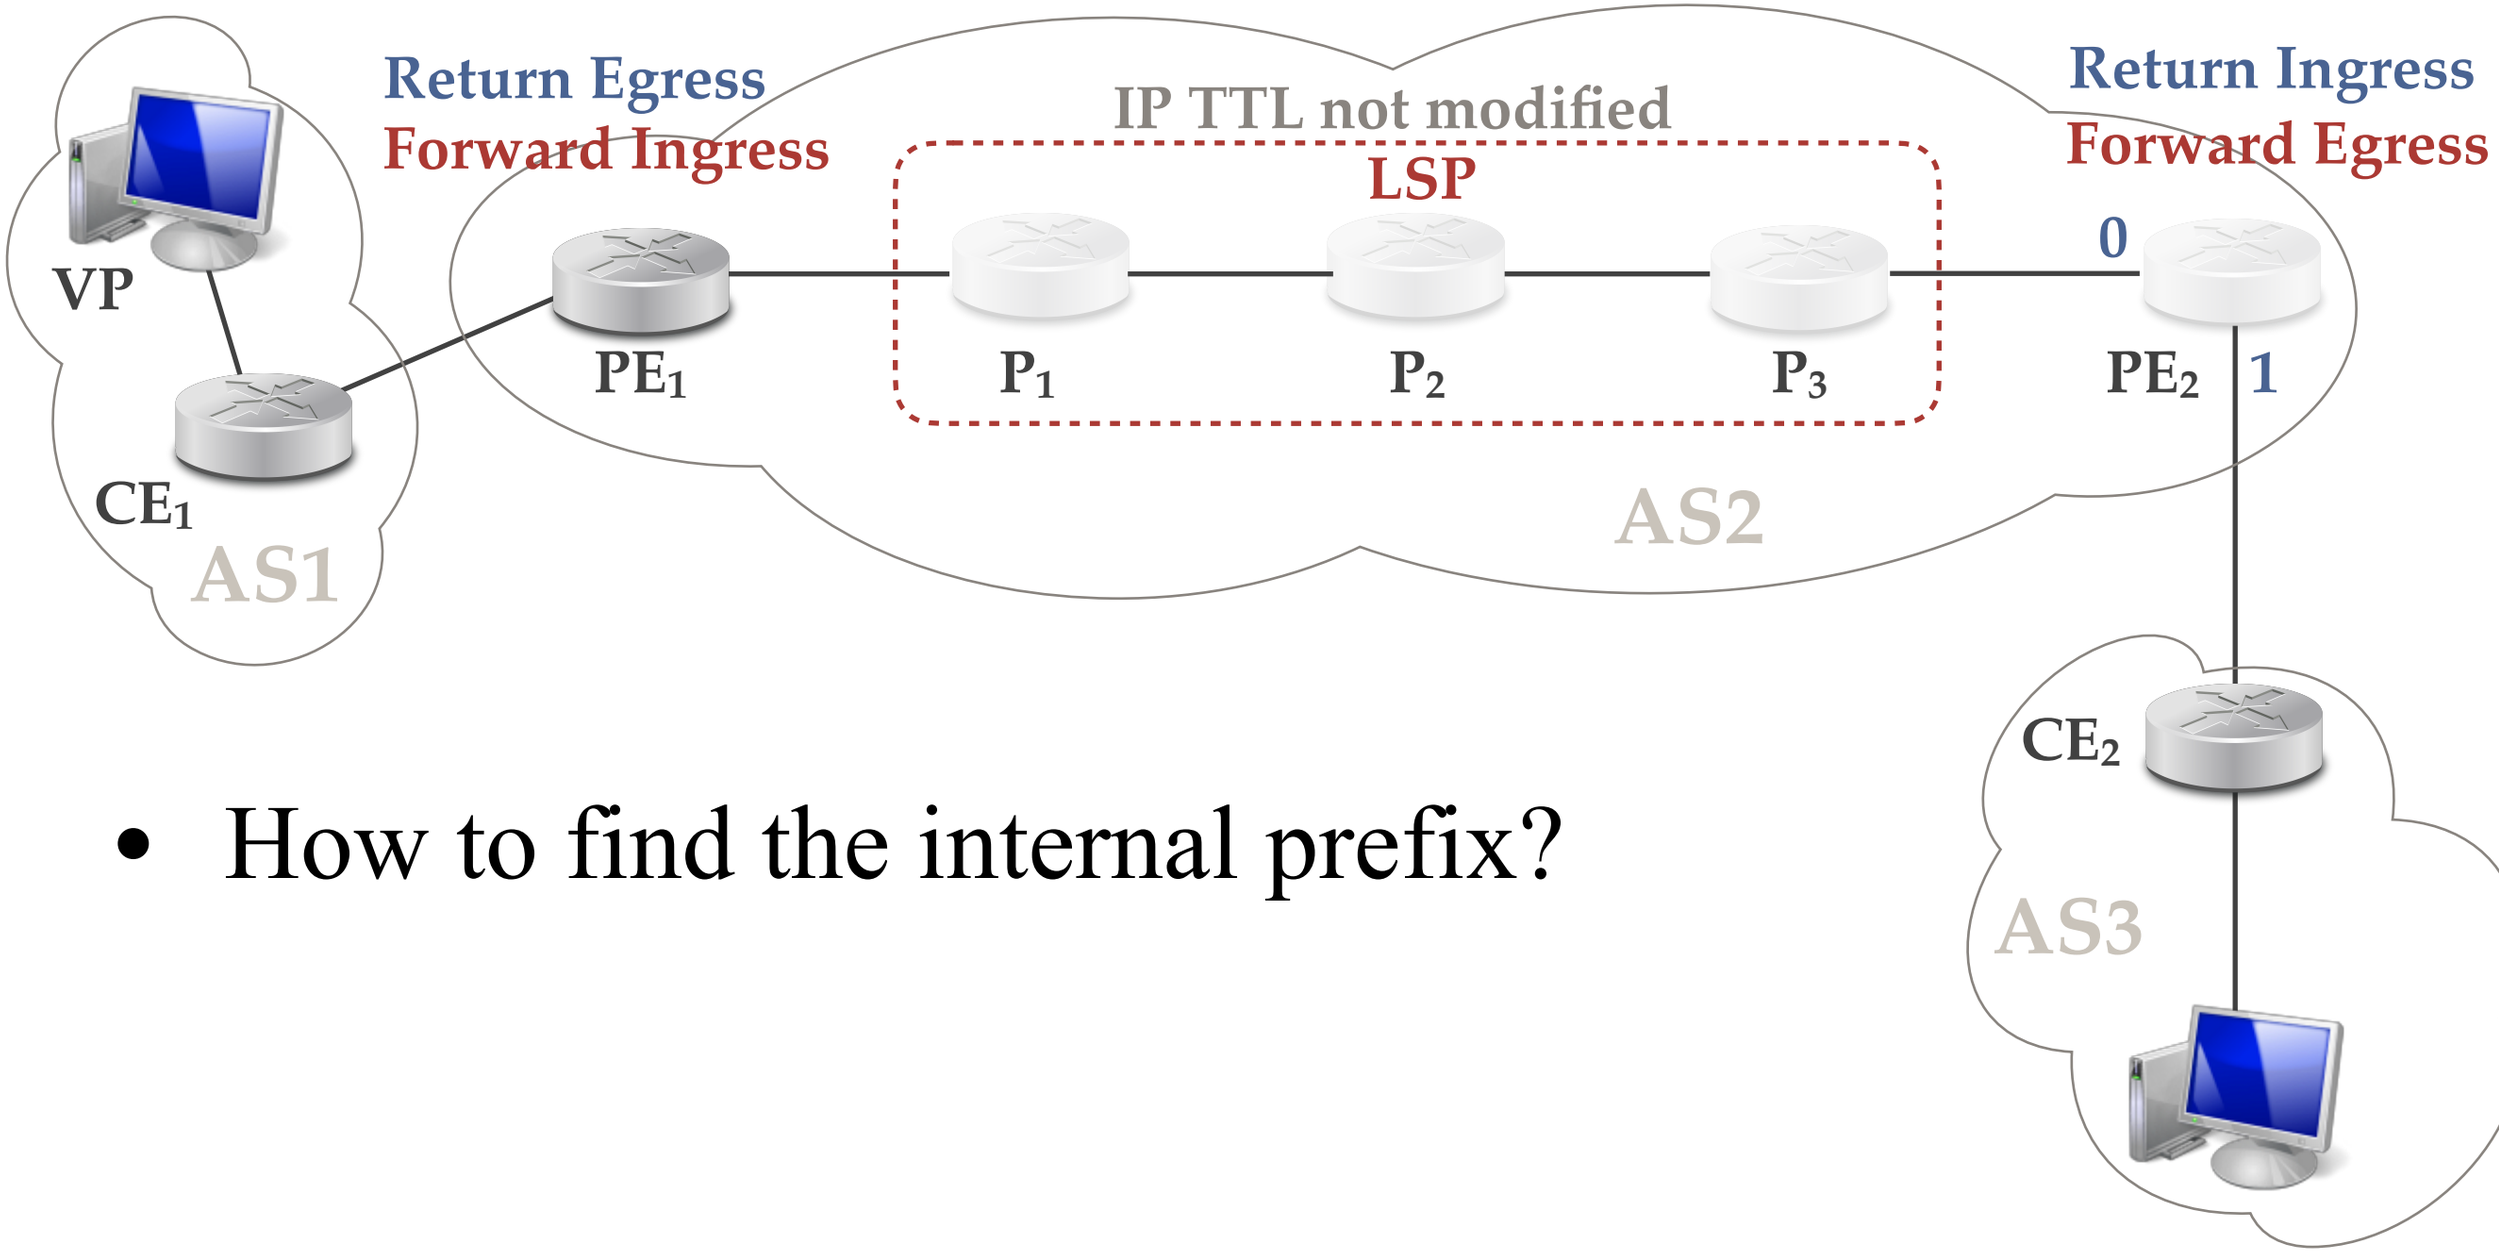

- How to find the internal prefix?

Invisible Tunnels (9)

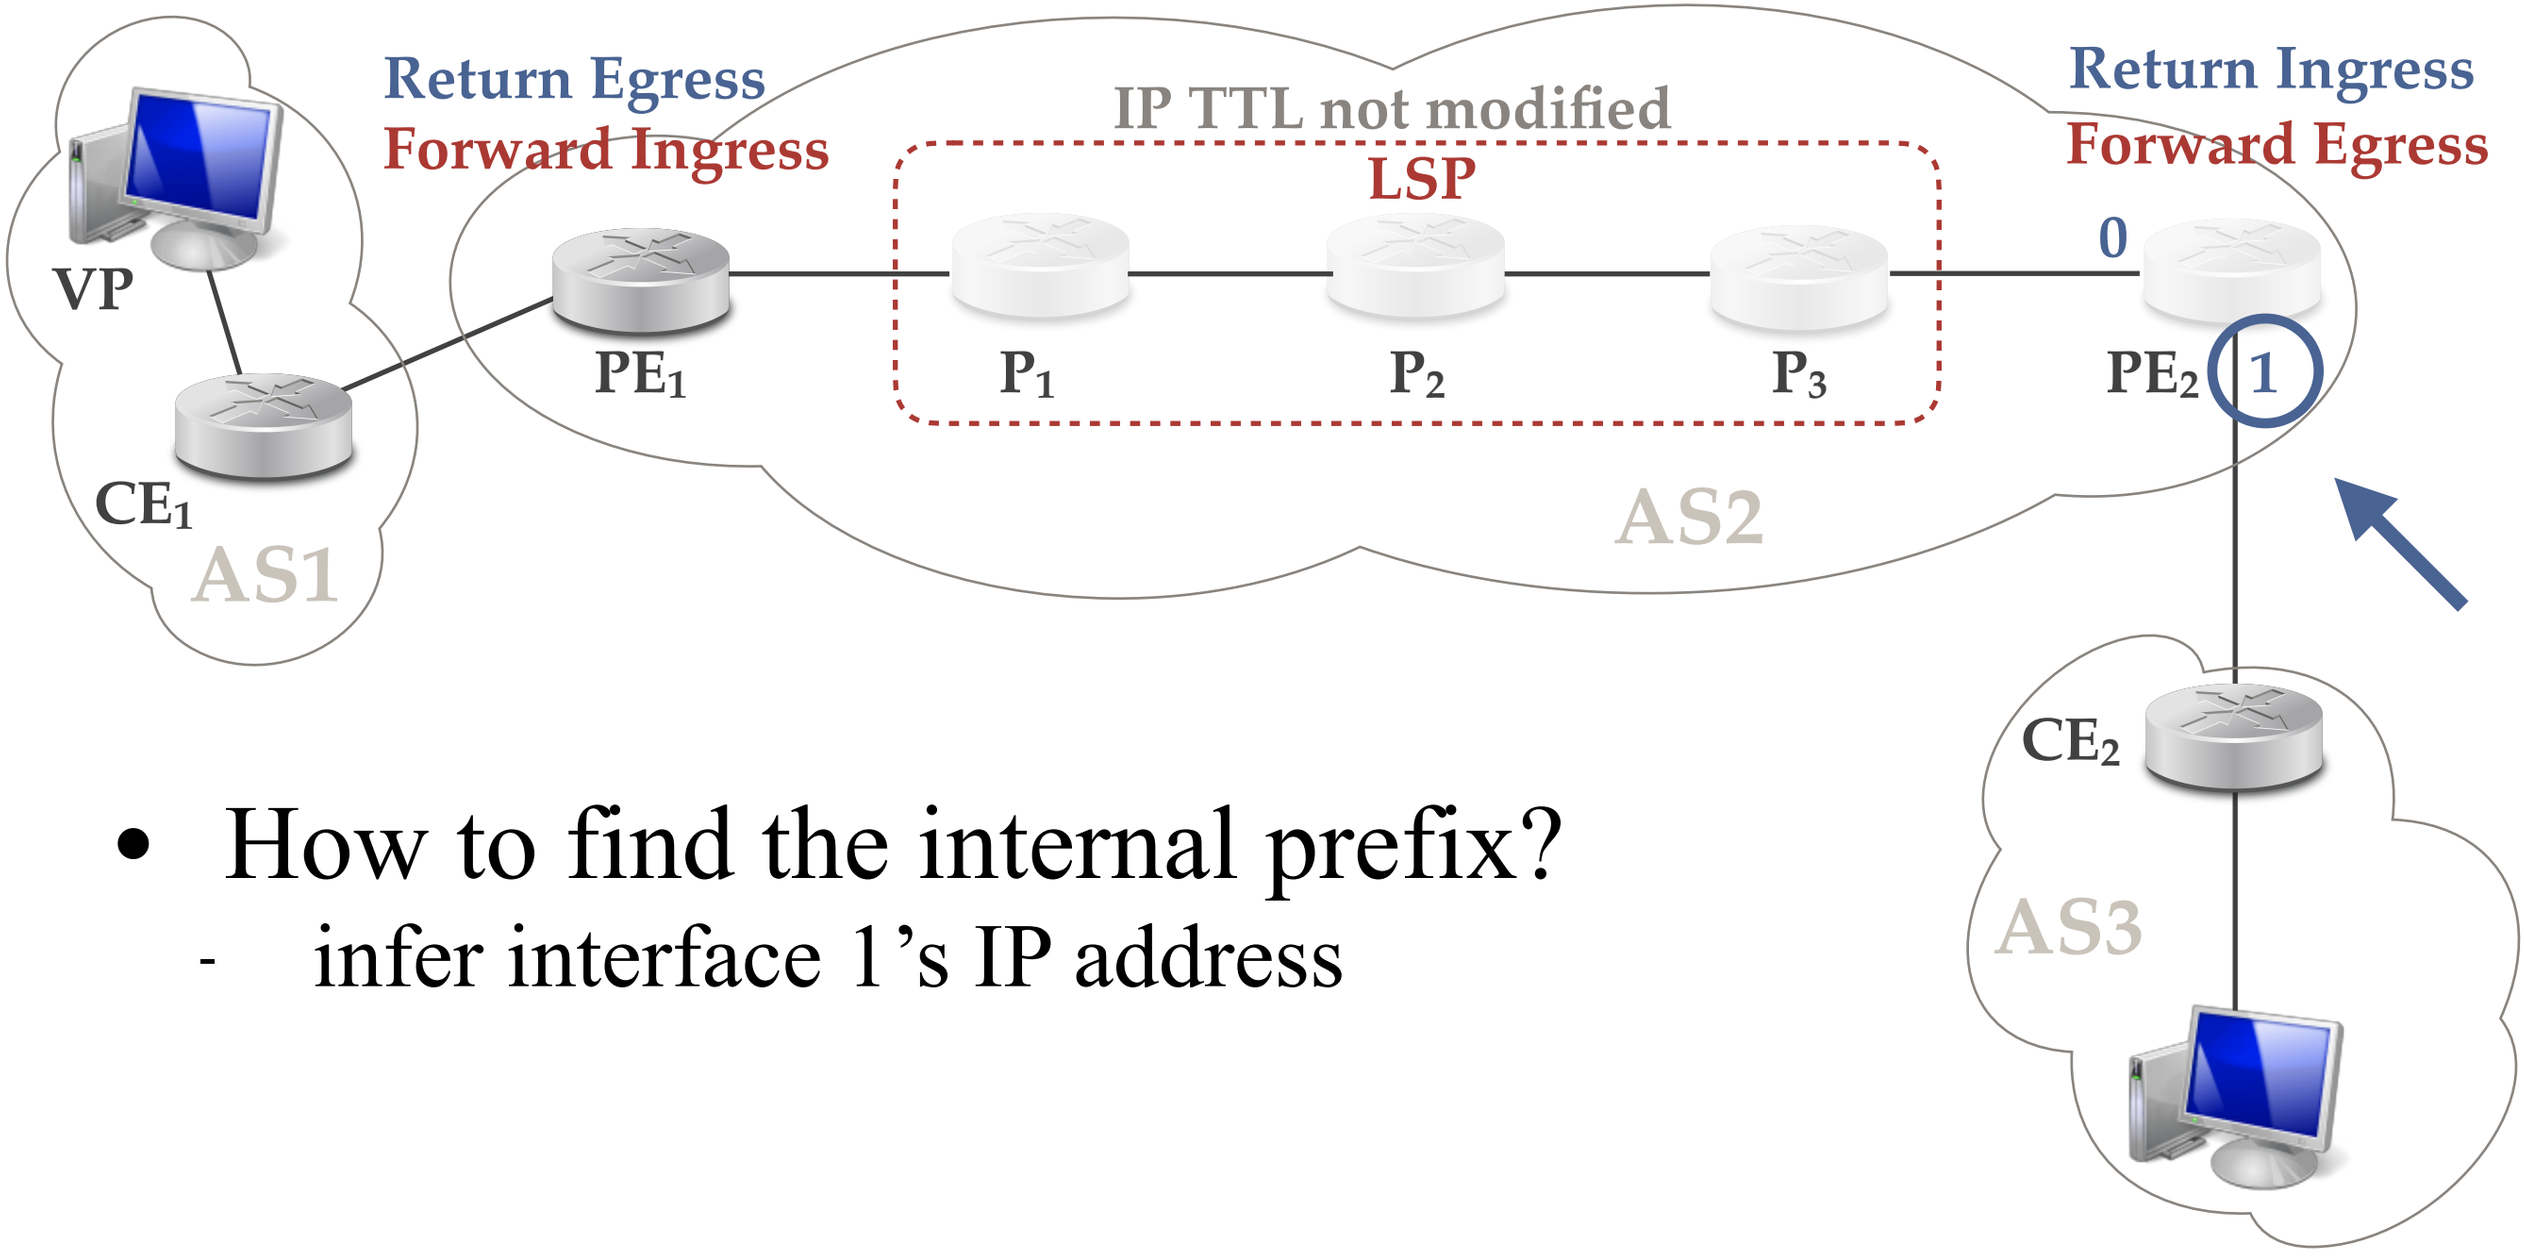

- How to find the internal prefix?
  - infer interface 1's IP address

Invisible Tunnels (9)

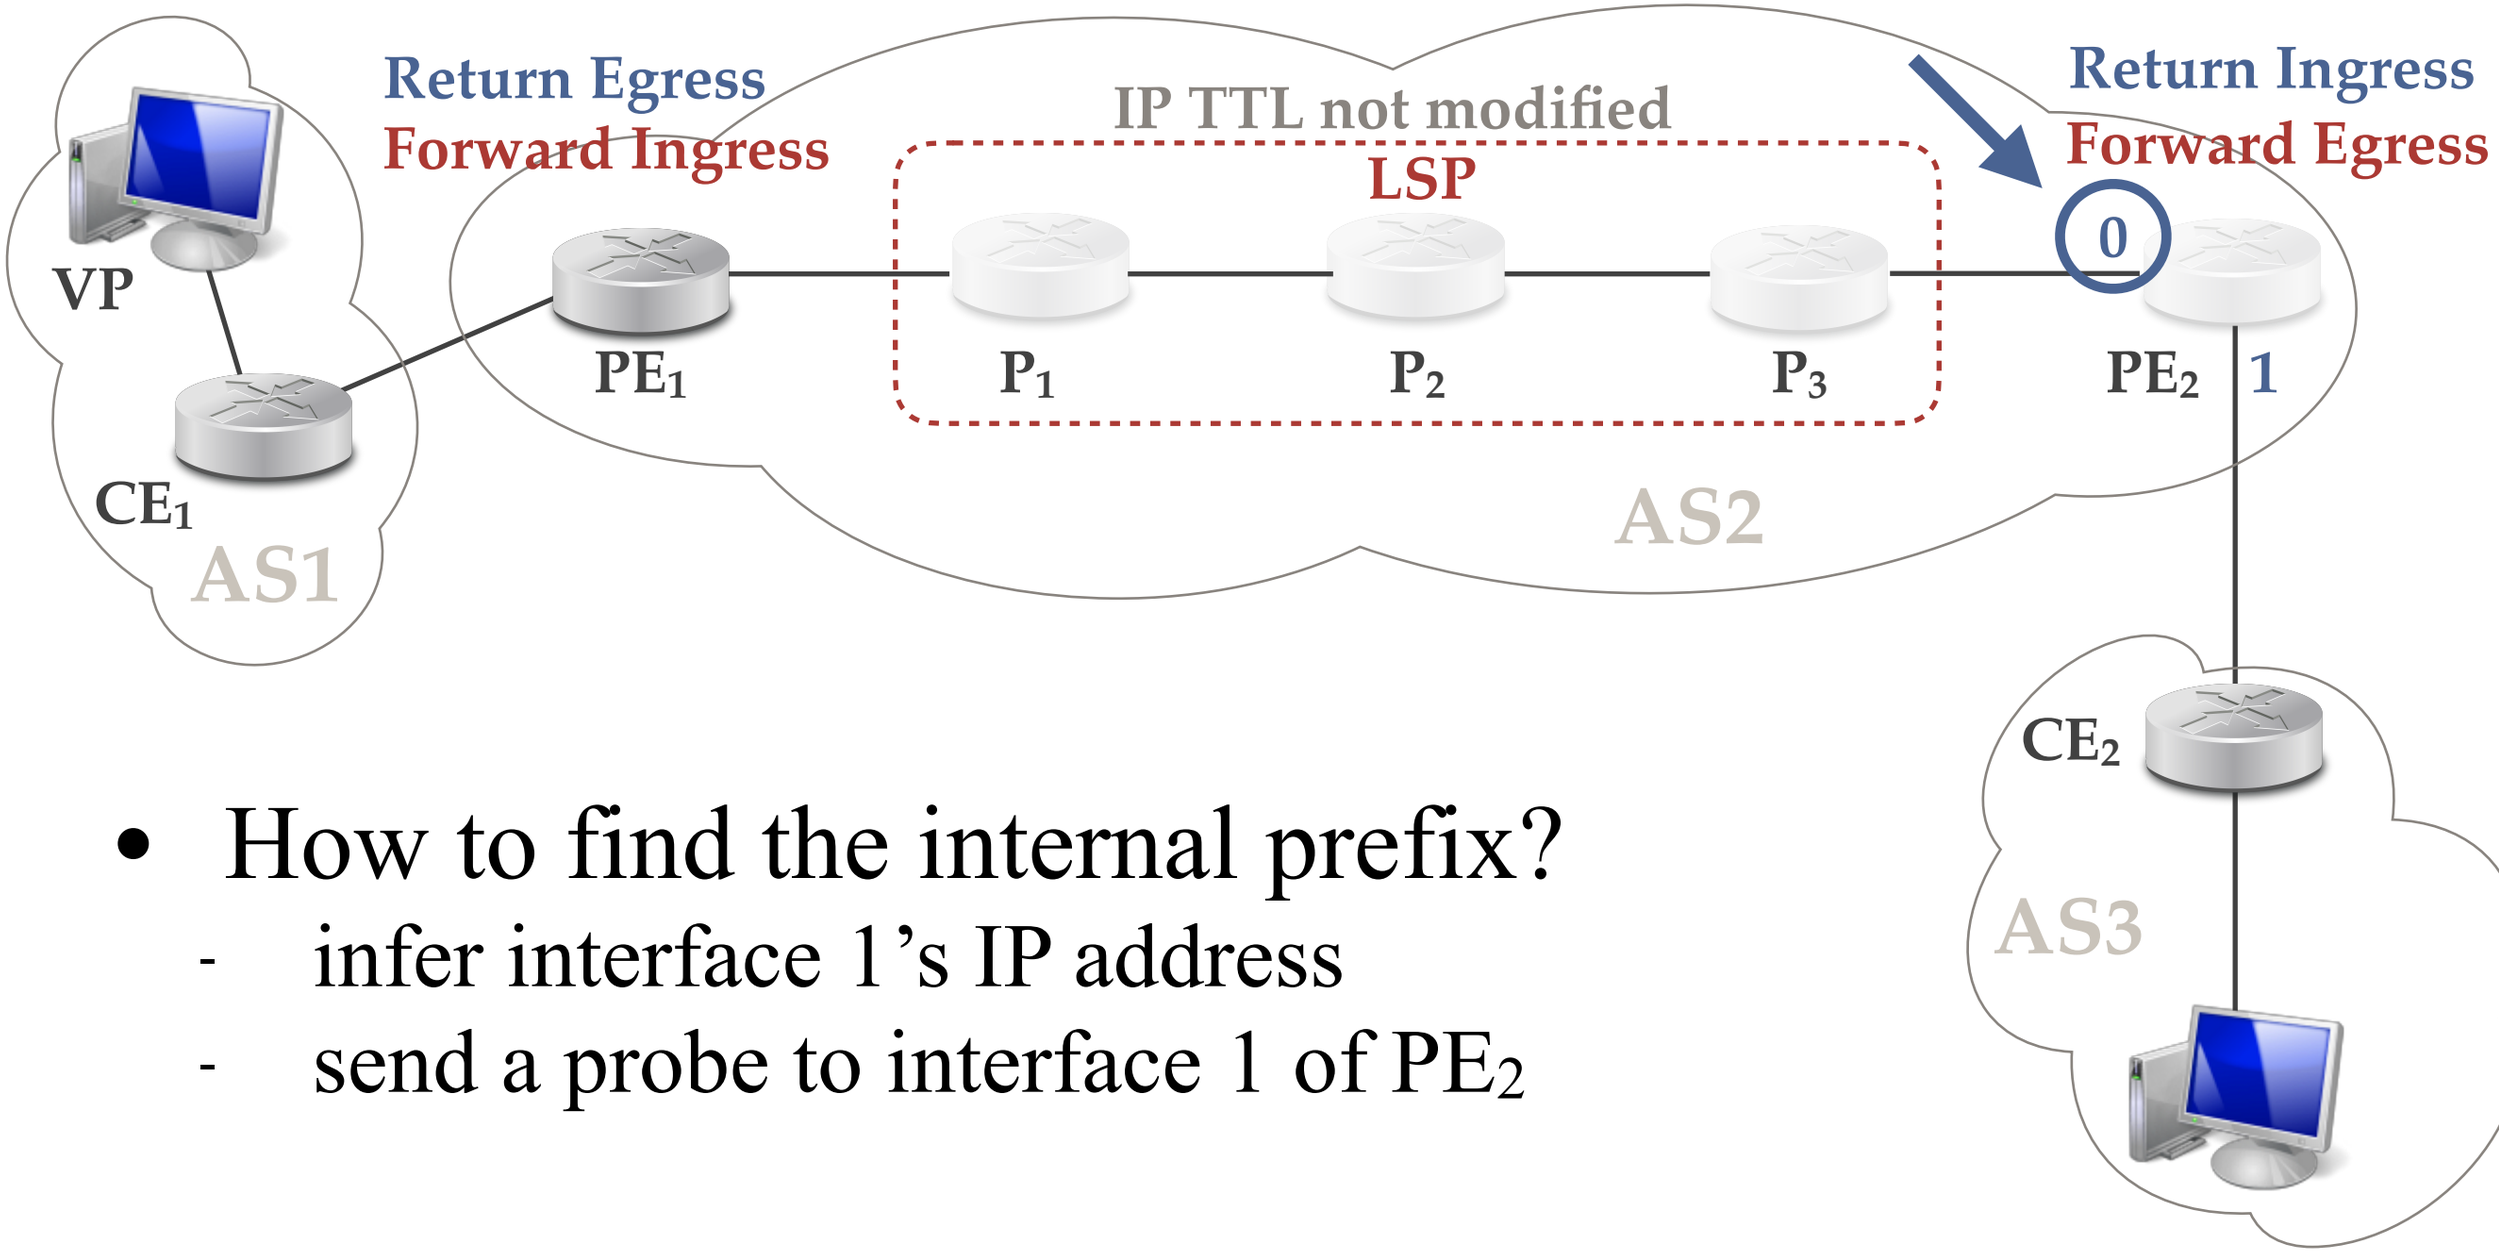

- How to find the internal prefix?
  - infer interface 1's IP address
  - send a probe to interface 1 of PE_2

Invisible Tunnels (9)

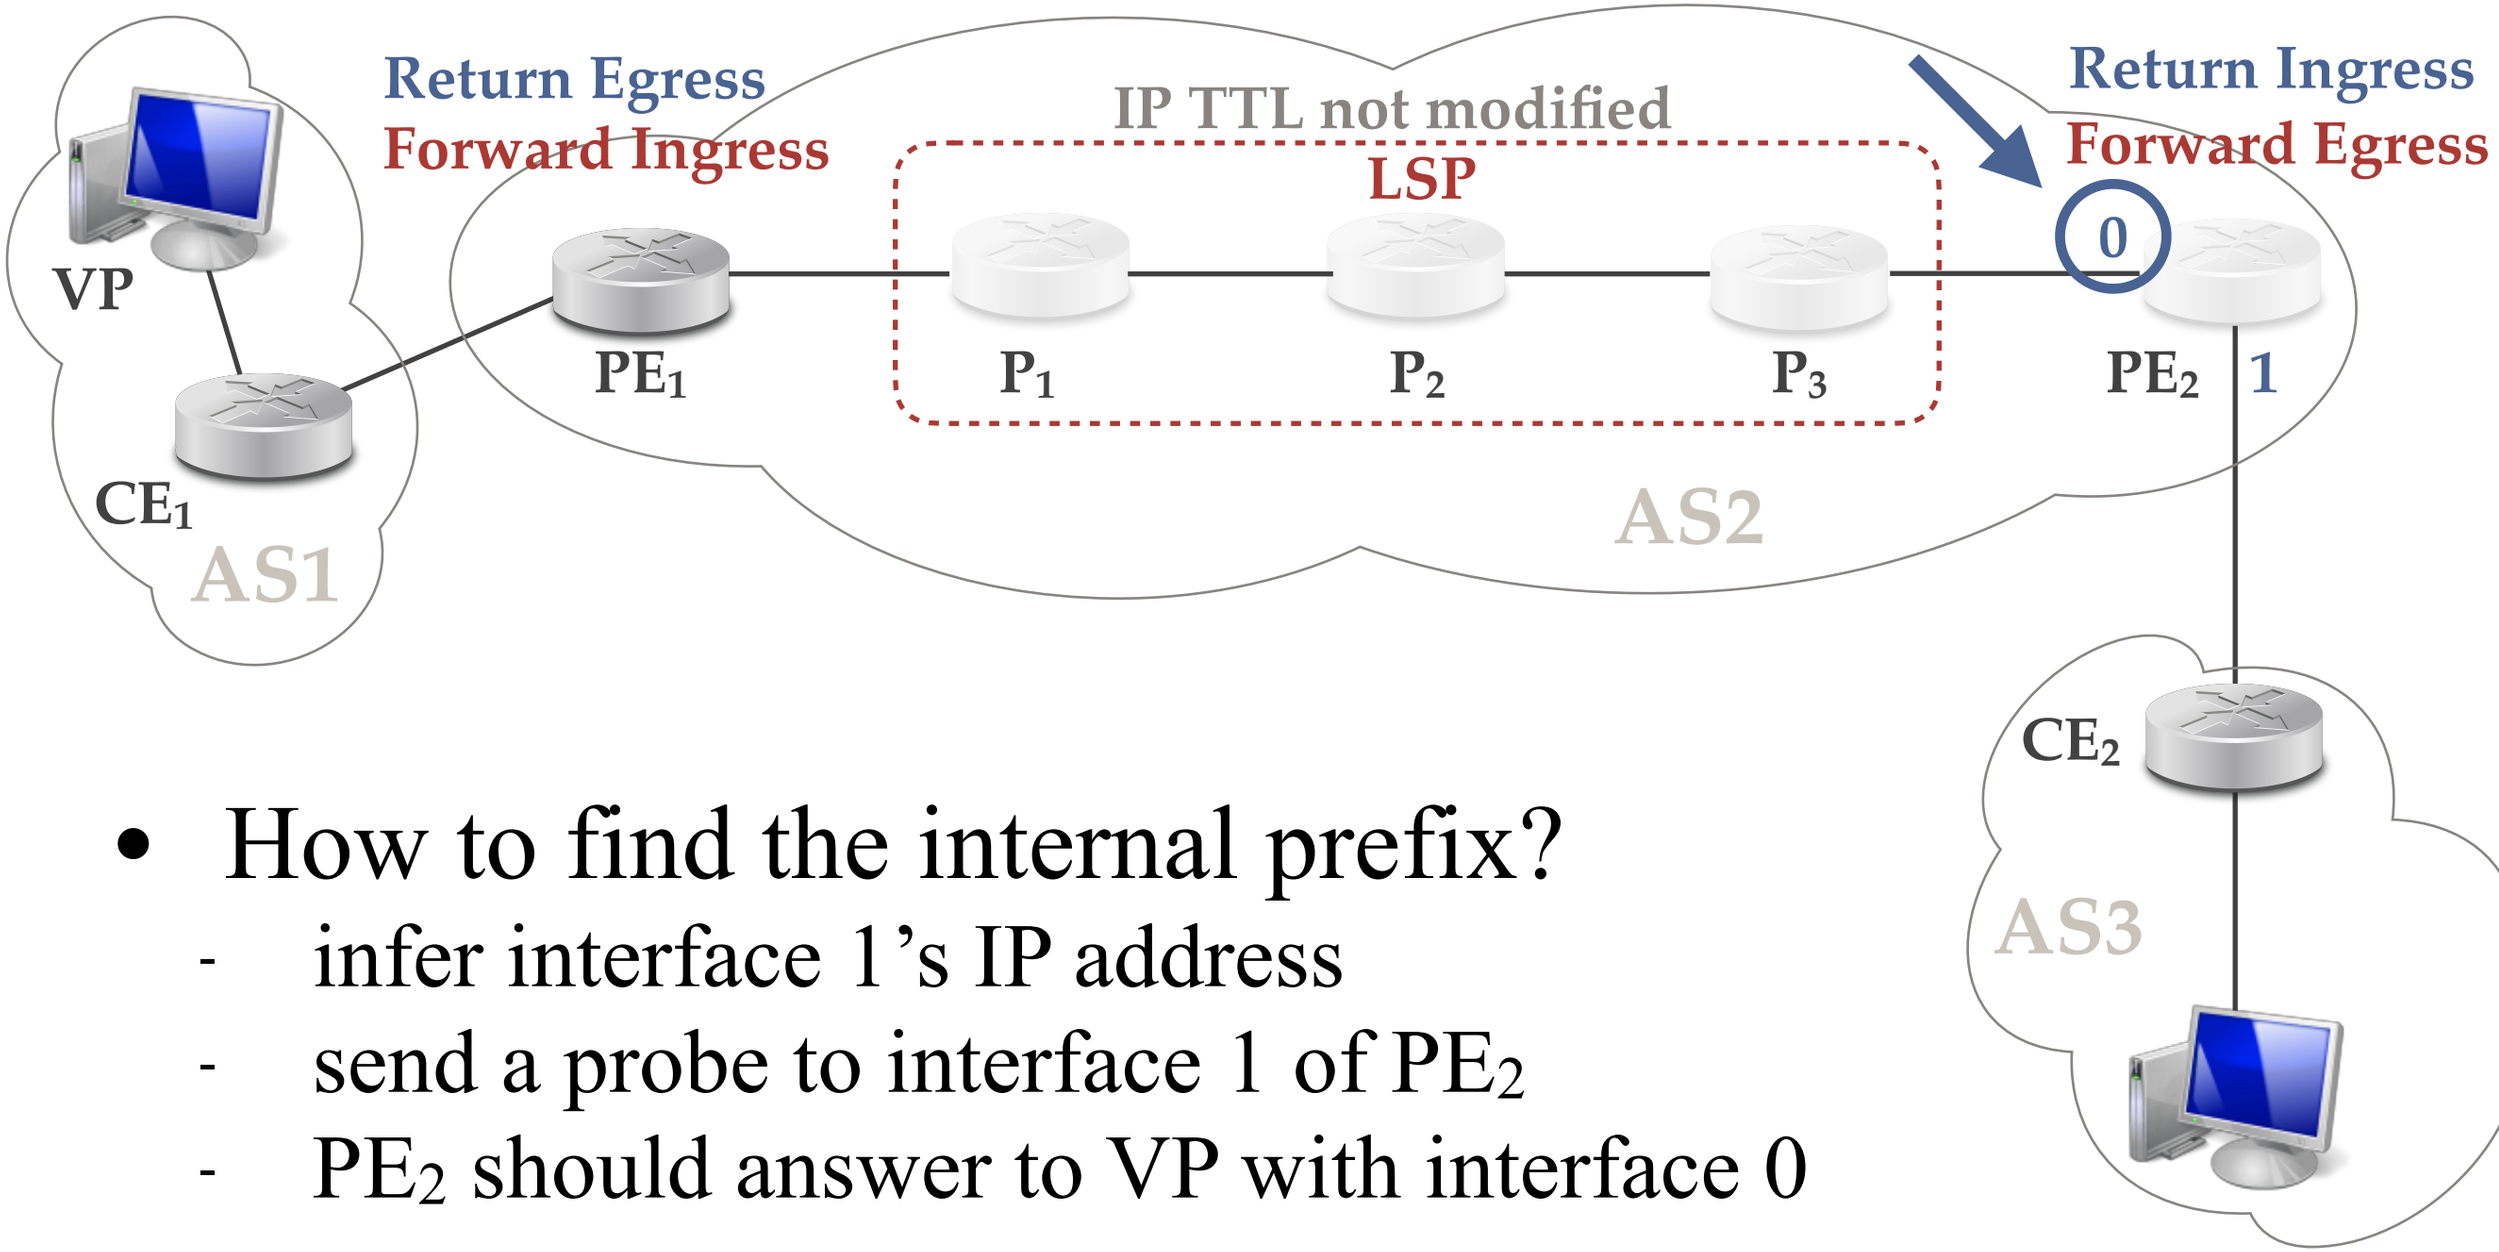

- How to find the internal prefix?
  - infer interface 1's IP address
  - send a probe to interface 1 of PE_2
  - PE_2 should answer to VP with interface 0

Invisible Tunnels (9)

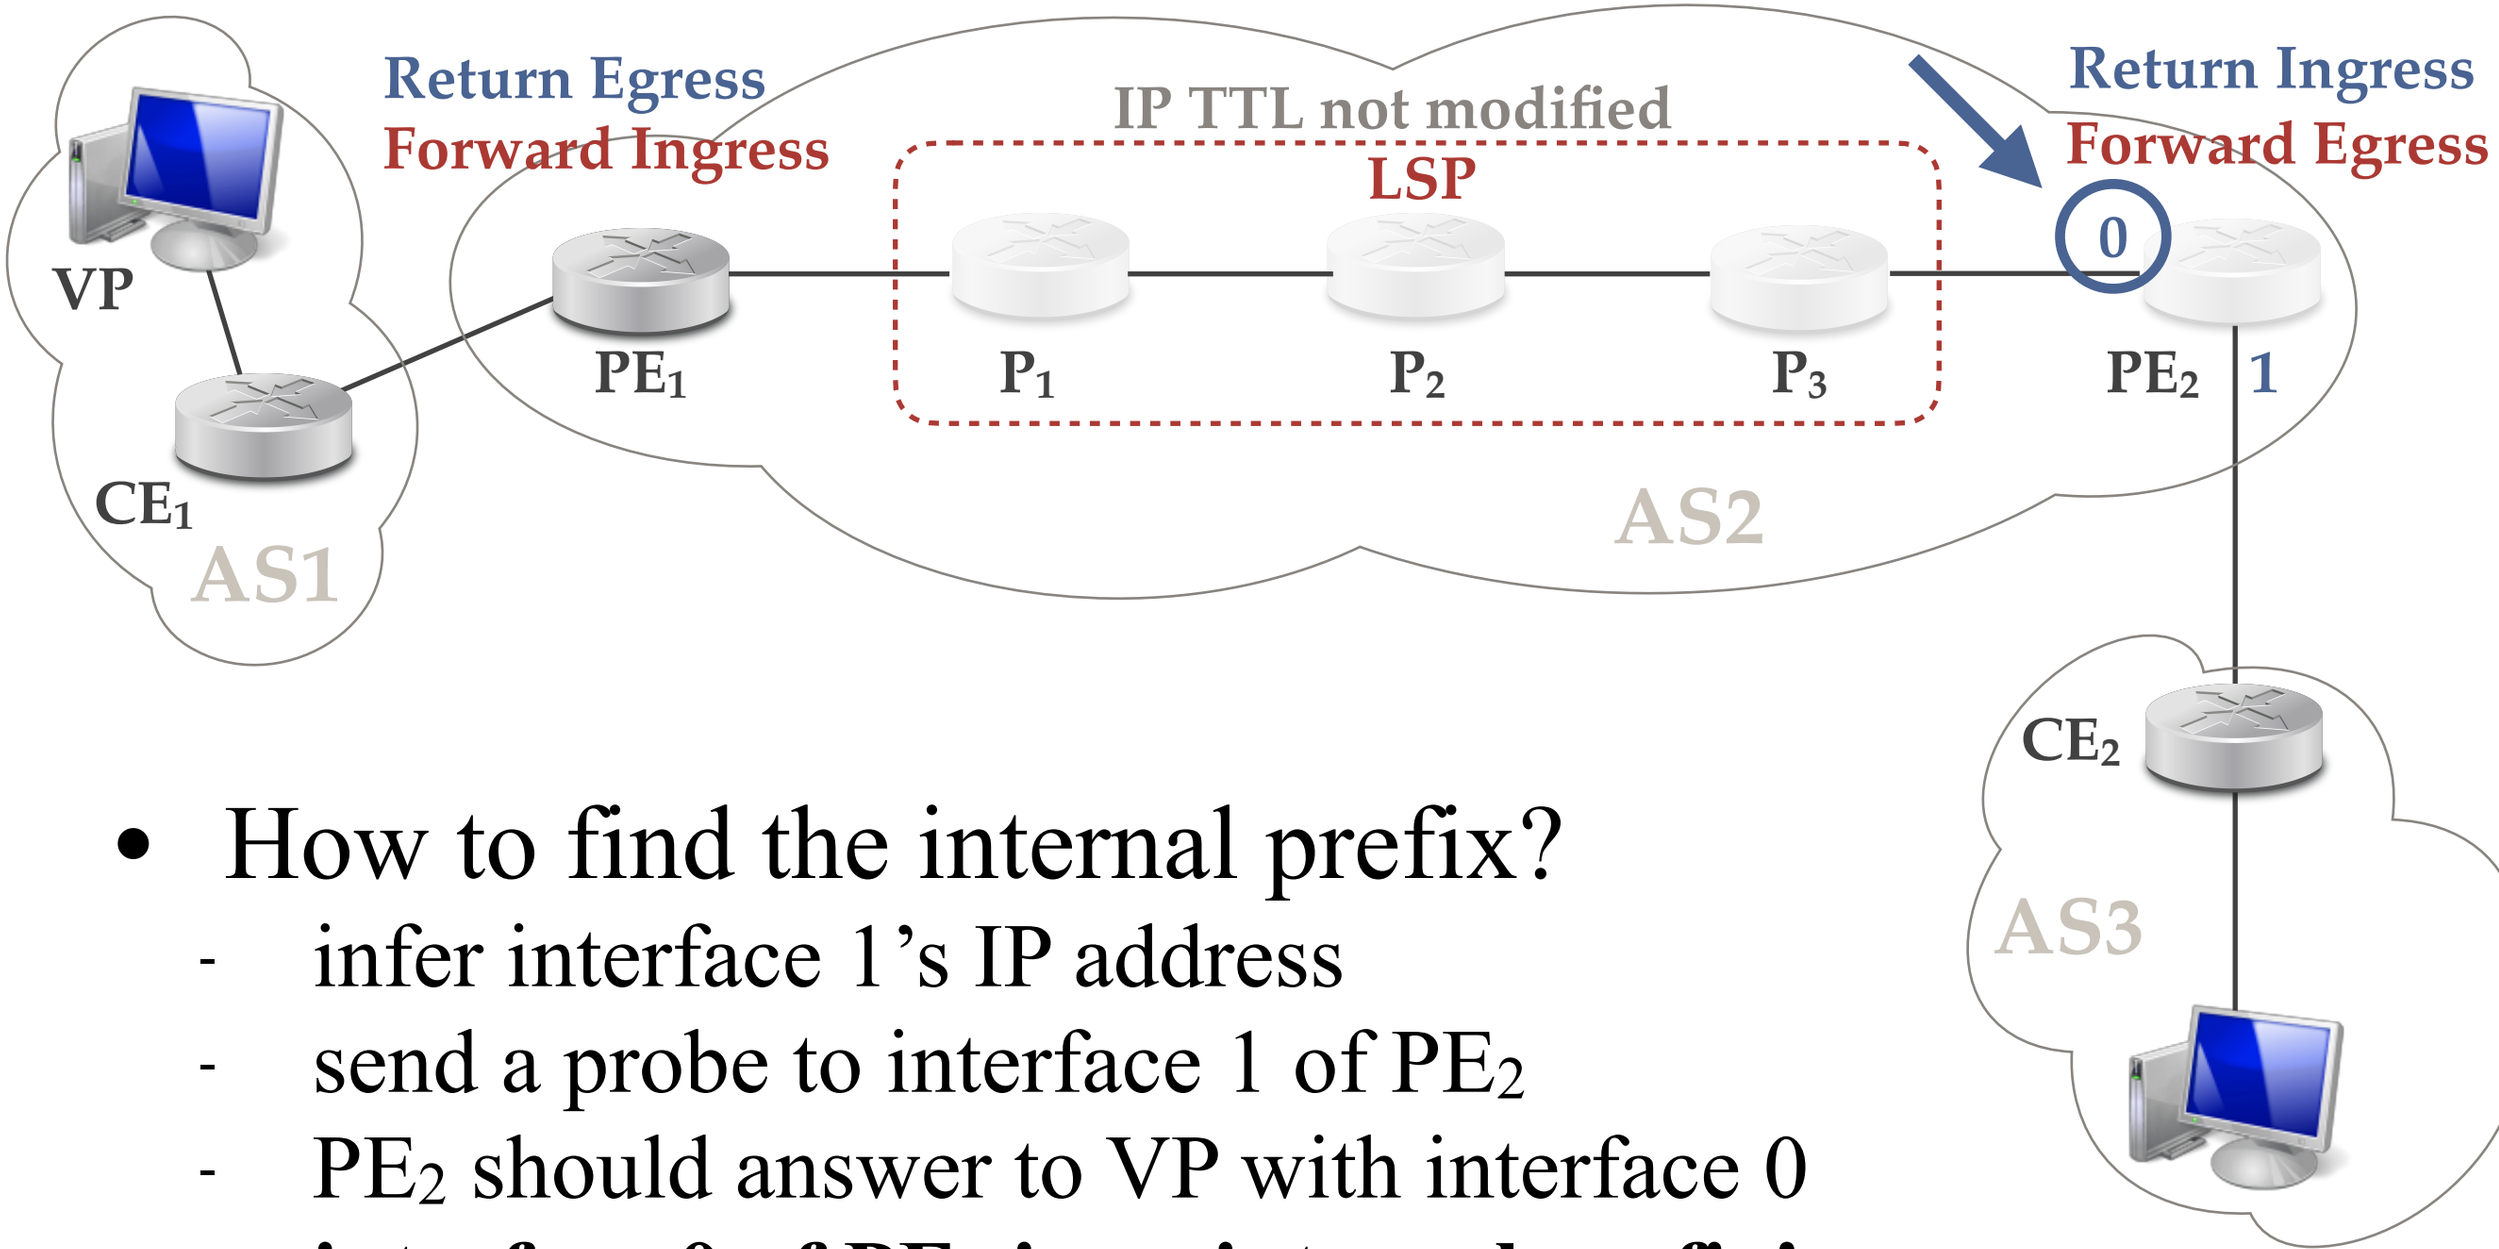

- How to find the internal prefix?
  - infer interface 1's IP address
  - send a probe to interface 1 of PE_2
  - PE_2 should answer to VP with interface 0
  - **interface 0 of PE_2 is an internal prefix!**

Summary (+*validation*)

Configurations	Pop	Cisco iOS15.2	Juniper VMX
P2P circuits	PHP	FRPLA, BRPR	RTLA, DPR
(e.g. LDP or	UHP	DUP_IP, BRPR ++	RTLA, DPR
RSVP-TE tunnels)		☑	☑
P2MP overlays	PHP	LSE-TTL, -	RTLA++, -
(e.g. VPRN: CsC or	UHP	LSE-TTL++, -	N/A
VPN BGP-MPLS)		☒	☒

Summary (+*validation*)

Configurations	Pop	Cisco iOS15.2	Juniper VMX
P2P circuits	PHP	FRPLA, BRPR	RTLA, DPR
(e.g. LDP or	UHP	DUP_IP, BRPR ++	RTLA, DPR
RSVP-TE tunnels)		☑	☑
P2MP overlays	PHP	LSE-TTL, -	RTLA++, -
(e.g. VPRN: CsC or	UHP	LSE-TTL++, -	N/A
VPN BGP-MPLS)		☒	☒

- Reproducibility analysis with GNS3

Summary (+*validation*)

Configurations	Pop	Cisco iOS15.2	Juniper VMX
P2P circuits	PHP	FRPLA, BRPR	RTLA, DPR
(e.g. LDP or	UHP	DUP_IP, BRPR ++	RTLA, DPR
RSVP-TE tunnels)		☑	☑
P2MP overlays	PHP	LSE-TTL, -	RTLA++, -
(e.g. VPRN: CsC or	UHP	LSE-TTL++, -	N/A
VPN BGP-MPLS)		☒	☒

- Reproducibility analysis with GNS3
  - What configuration?
    - P2P or P2MP

Summary (+*validation*)

Configurations	Pop	Cisco iOS15.2	Juniper VMX
P2P circuits	PHP	FRPLA, BRPR	RTLA, DPR
(e.g. LDP or	UHP	DUP_IP, BRPR ++	RTLA, DPR
RSVP-TE tunnels)		☑	☑
P2MP overlays	PHP	LSE-TTL, -	RTLA++, -
(e.g. VPRN: CsC or	UHP	LSE-TTL++, -	N/A
VPN BGP-MPLS)		☒	☒

- Reproducibility analysis with GNS3
  - What configuration?
    - P2P or P2MP
  - Popping function

Summary (+*validation*)

Configurations	Pop	Cisco iOS15.2	Juniper VMX
P2P circuits	PHP	FRPLA, BRPR	RTLA, DPR
(e.g. LDP or	UHP	DUP_IP, BRPR ++	RTLA, DPR
RSVP-TE tunnels)		☑	☑
P2MP overlays	PHP	LSE-TTL, -	RTLA++, -
(e.g. VPRN: CsC or	UHP	LSE-TTL++, -	N/A
VPN BGP-MPLS)		☒	☒

- Reproducibility analysis with GNS3
  - What configuration?
    - P2P or P2MP
  - Popping function
    - UHP or PHP

Calibration

- Quality of thresholds for triggering DPR/BRPR
  - FRPLA
  - RTLA

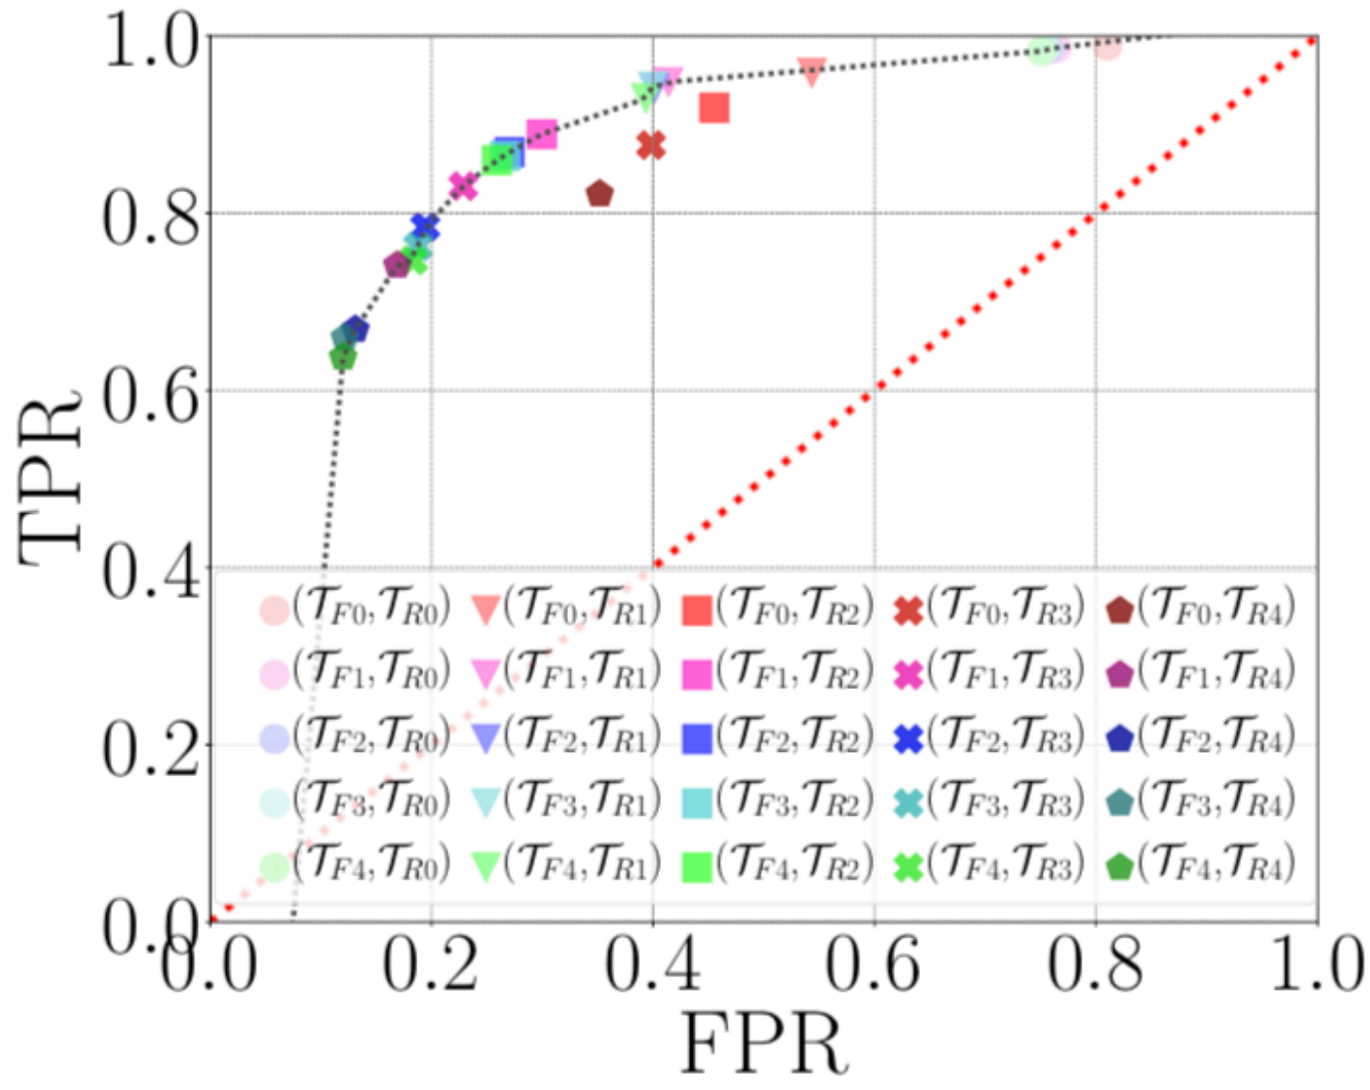

Results

- Fingerprinting, triggers and revelation techniques implemented in TNT
  - https://github.com/YvesVanaubel/TNT
- TNT deployed on Archipelago infrastructure
  - 28 vantage points
  - 2,800,000 destinations
  - impact of triggers investigated
    - ✓ not shown here
  - data collected on April 2018
- Dataset available
  - http://www.montefiore.ulg.ac.be/~bdonnet/mpls/data.html

Results (2)

- Tunnels discovered
  - per type
  - per revelation technique

Tunnel Type	Indicator/Trigger	Revelation Technique DPR	BRPR	1Hop_Lsp	Mix	# Tunnels
Explicit	LSE headers	-	-	-	-	150,036
Implicit	qTTL	-	-	-	-	2,689
	Uturn	-	-	-	-	7,216
Opaque	LSE-TTL	22	17	43	-	3,346
Invisible PHP	Rtla	11,268	1,191	2,595	279	15,333
	Frpla	5,903	2,555	3,260	1,012	12,730
Invisible UHP	Dup_Ip	1,609	1,531	686	296	4,122
Total		18,802	5,294	6,584	1,587	195,525

Results (3)

- Impact of invisible tunnels on
  - degree distribution
  - delay

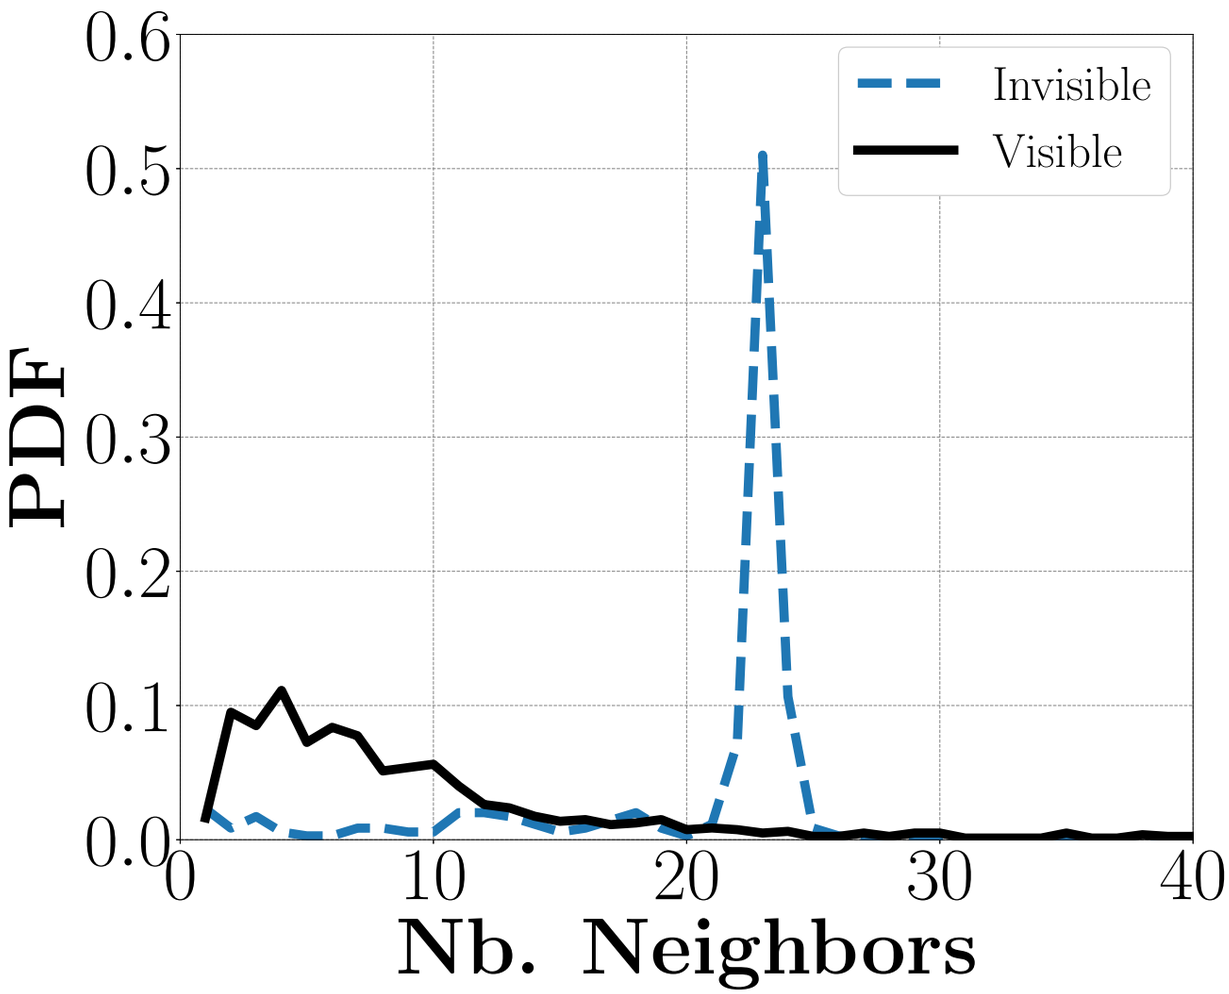

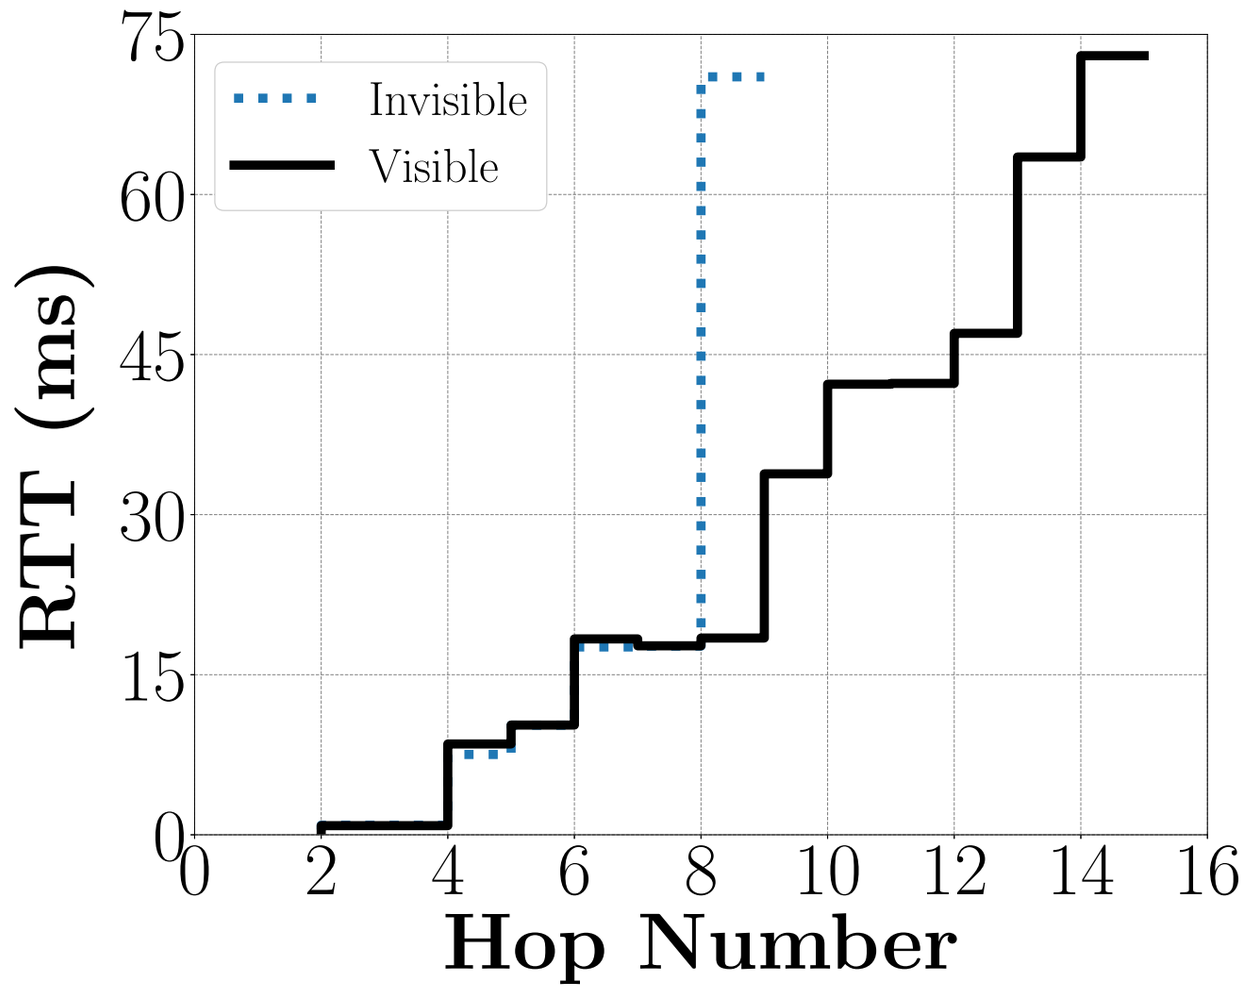

Agenda

- Motivations
- Network Fingerprinting
- MPLS Background
- TNT and MPLS Invisible Tunnels
- Conclusion

Conclusion

- TNT
  - allows for revealing all MPLS tunnels
  - implemented within scamper
- Might be the basis for a "European Network Observatory"
  - useful for security purpose

Let’s try?

- traceroute -P icmp -f <X> [-n] -q 1 130.79.245.185

Results 1

// 17 robinet.u-strasbg.fr (130.79.91.202) 132.660 ms // ma machine
18 ce1-in.u-strasbg.fr (130.79.245.179) 53.334 ms
19 ingress-pe1-in.u-strasbg.fr (130.79.245.181) 80.554 ms
20 egress-pe2-in.u-strasbg.fr (130.79.245.173) 180.629 ms // <- next target
21 ce2-in-dup.u-strasbg.fr (130.79.245.187) 160.719 ms
22 target-ce3-lo.u-strasbg.fr (130.79.245.185) 190.704 ms

18 ce1-in.u-strasbg.fr (130.79.245.179) 53.436 ms
19 ingress-pe1-in.u-strasbg.fr (130.79.245.181) 79.704 ms
20 ph-p3-in.u-strasbg.fr (130.79.245.171) 150.658 ms
21 egress-pe2-in.u-strasbg.fr (130.79.245.173) 157.074 ms

18 ce1-in.u-strasbg.fr (130.79.245.179) 60.061 ms
19 ingress-pe1-in.u-strasbg.fr (130.79.245.181) 81.464 ms
20 p2-in.u-strasbg.fr (130.79.245.169) 118.925 ms
21 ph-p3-in.u-strasbg.fr (130.79.245.171) 115.040 ms

18 ce1-in.u-strasbg.fr (130.79.245.179) 57.176 ms
19 ingress-pe1-in.u-strasbg.fr (130.79.245.181) 80.675 ms
20 p1-in.u-strasbg.fr (130.79.245.167) 110.718 ms
21 p2-in.u-strasbg.fr (130.79.245.169) 95.650 ms

18 ce1-in.u-strasbg.fr (130.79.245.179) 52.486 ms
19 ingress-pe1-in.u-strasbg.fr (130.79.245.181) 69.836 ms
20 p1-in.u-strasbg.fr (130.79.245.167) 99.689 ms

Result 2 : TNT

9 robinet.u-strasbg.fr (130.79.91.202) <56,*> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0] [PING Timeout]
10 130.79.245.179 (130.79.245.179) <246,246> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0]
11 130.79.245.181 (130.79.245.181) <245,245> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0]

FRPLA | Length estimation : 3 | Revealed : 3 (difference : 0)
11.1 [REVEALED] 130.79.245.167 (130.79.245.167) <244,244> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0] - step 2
11.2 [REVEALED] 130.79.245.169 (130.79.245.169) <243,243> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0] - step 1
11.3 [REVEALED] 130.79.245.171 (130.79.245.171) <242,242> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0] - step 0

12 130.79.245.173 (130.79.245.173) <241,241> [frpla = 3][qttl = 1][uturn = 0][meta = 3, 0, 0] // <- CIBLE
13 130.79.245.187 (130.79.245.187) <241,241> [frpla = 2][qttl = 1][uturn = 0][meta = -1, 0, 0]
14 130.79.245.189 (130.79.245.189) <241,241> [frpla = 0][qttl = 1][uturn = 0][meta = 0, 0, 0]