

**UNIVERSITÀ DEGLI STUDI ROMA TRE**  
Dipartimento di Informatica e Automazione

# **Monitoring the Status of MPLS VPN and VPLS Based on BGP Signaling Information**

Giuseppe Di Battista  
Massimo Rimondini  
Giorgio Sadolfo



IEEE/IFIP NOMS 2012  
18/04/2012



UNIVERSITÀ DEGLI STUDI ROMA TRE  
Dipartimento di Informatica e Automazione

# Monitoring the Status of **MPLS** **VPN** and **VPLS** Based on BGP Signaling Information

Giuseppe Di Battista  
Massimo Rimondini  
Giorgio Sadolfo



IEEE/IFIP NOMS 2012  
18/04/2012



UNIV  
Dipar

**MPLS VPN**

RC  
utor

**VPLS**

# Monitoring the Status of MPLS VPN and VPLS Based on BGP Signaling Information

Giuseppe Di Battista  
Massimo Rimondini  
Giorgio Sadolfo



IEEE/IFIP NOMS 2012  
18/04/2012

# About MPLS VPNs/VPLS

## Monitoring the Status of MPLS VPN and VPLS Based on BGP Signaling Information

Giuseppe Di Battista  
Massimo Rimondini  
Giorgio Sadolfo



IEEE/IFIP NOMS 2012  
18/04/2012

# About MPLS VPNs/VPLS

**Customer**

**Customer**

# About MPLS VPNs/VPLS

**Customer  
site**

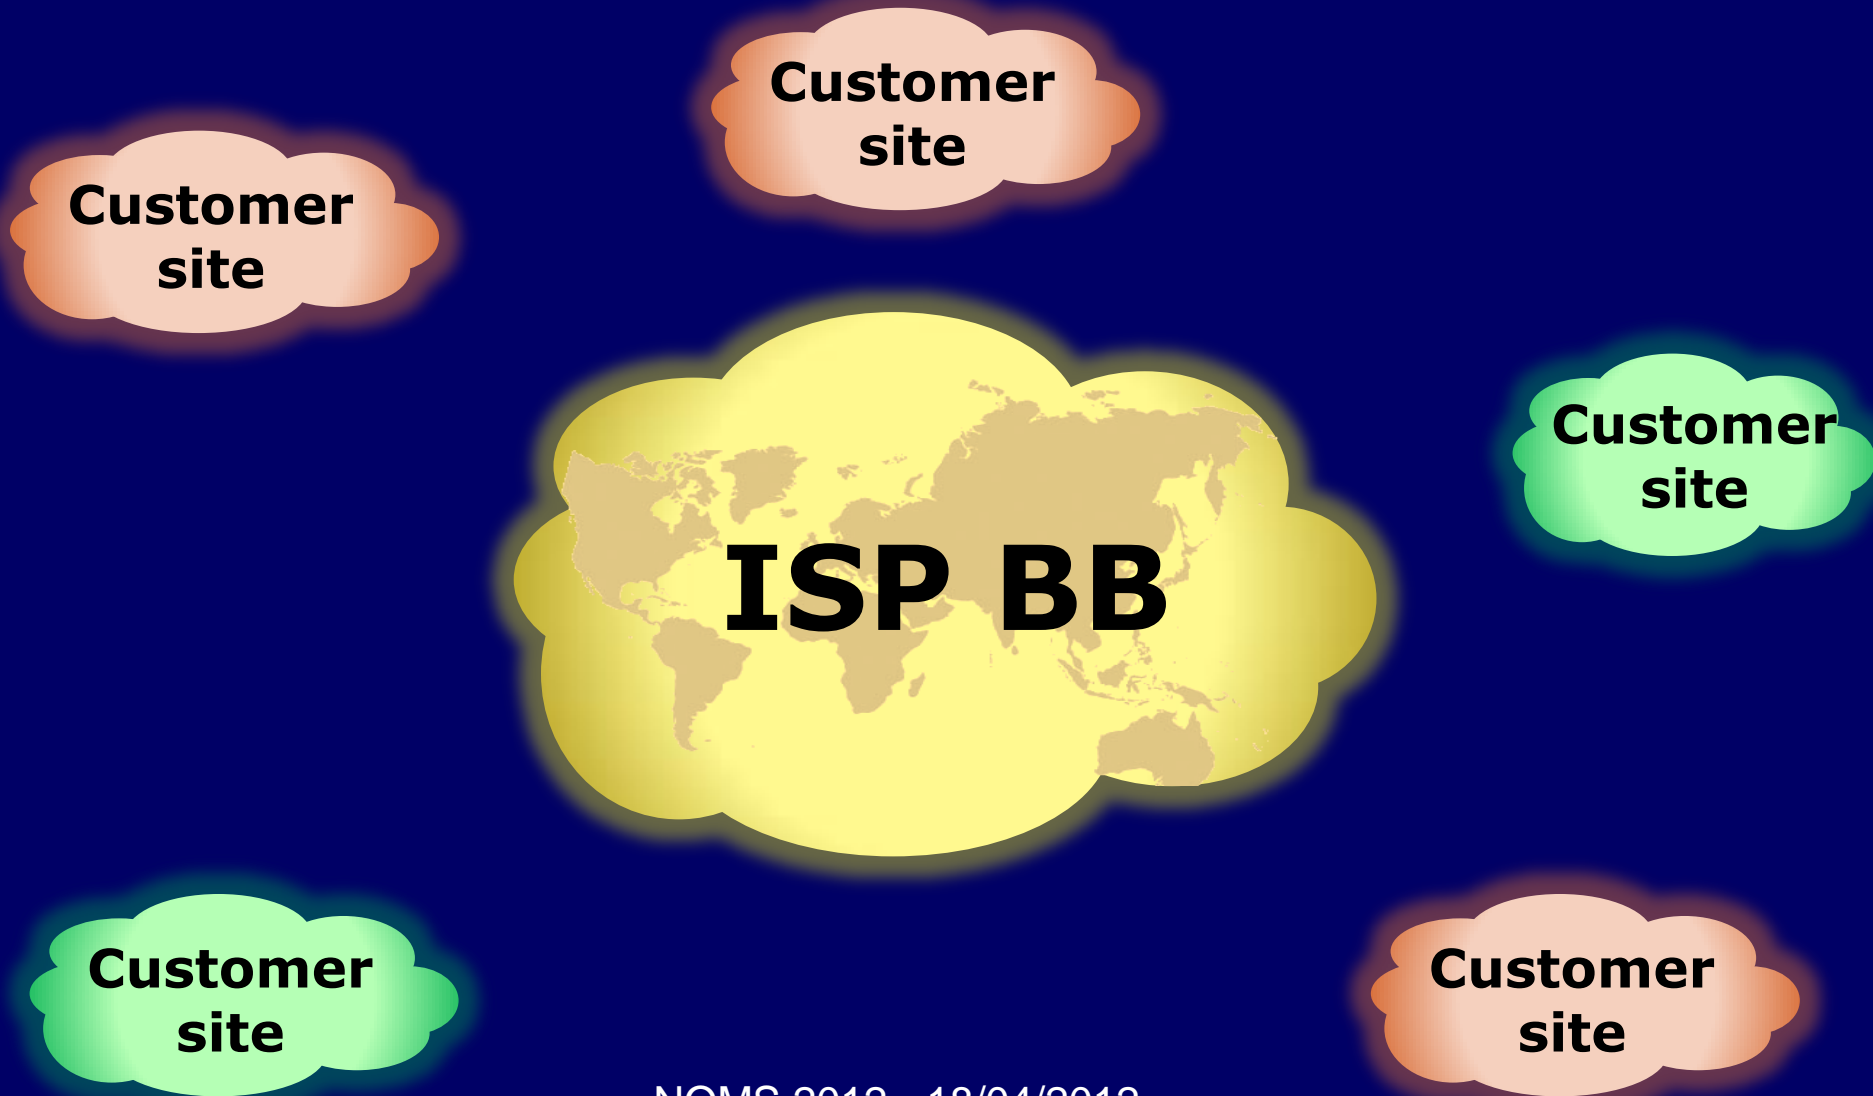
**Customer  
site**

**Customer  
site**

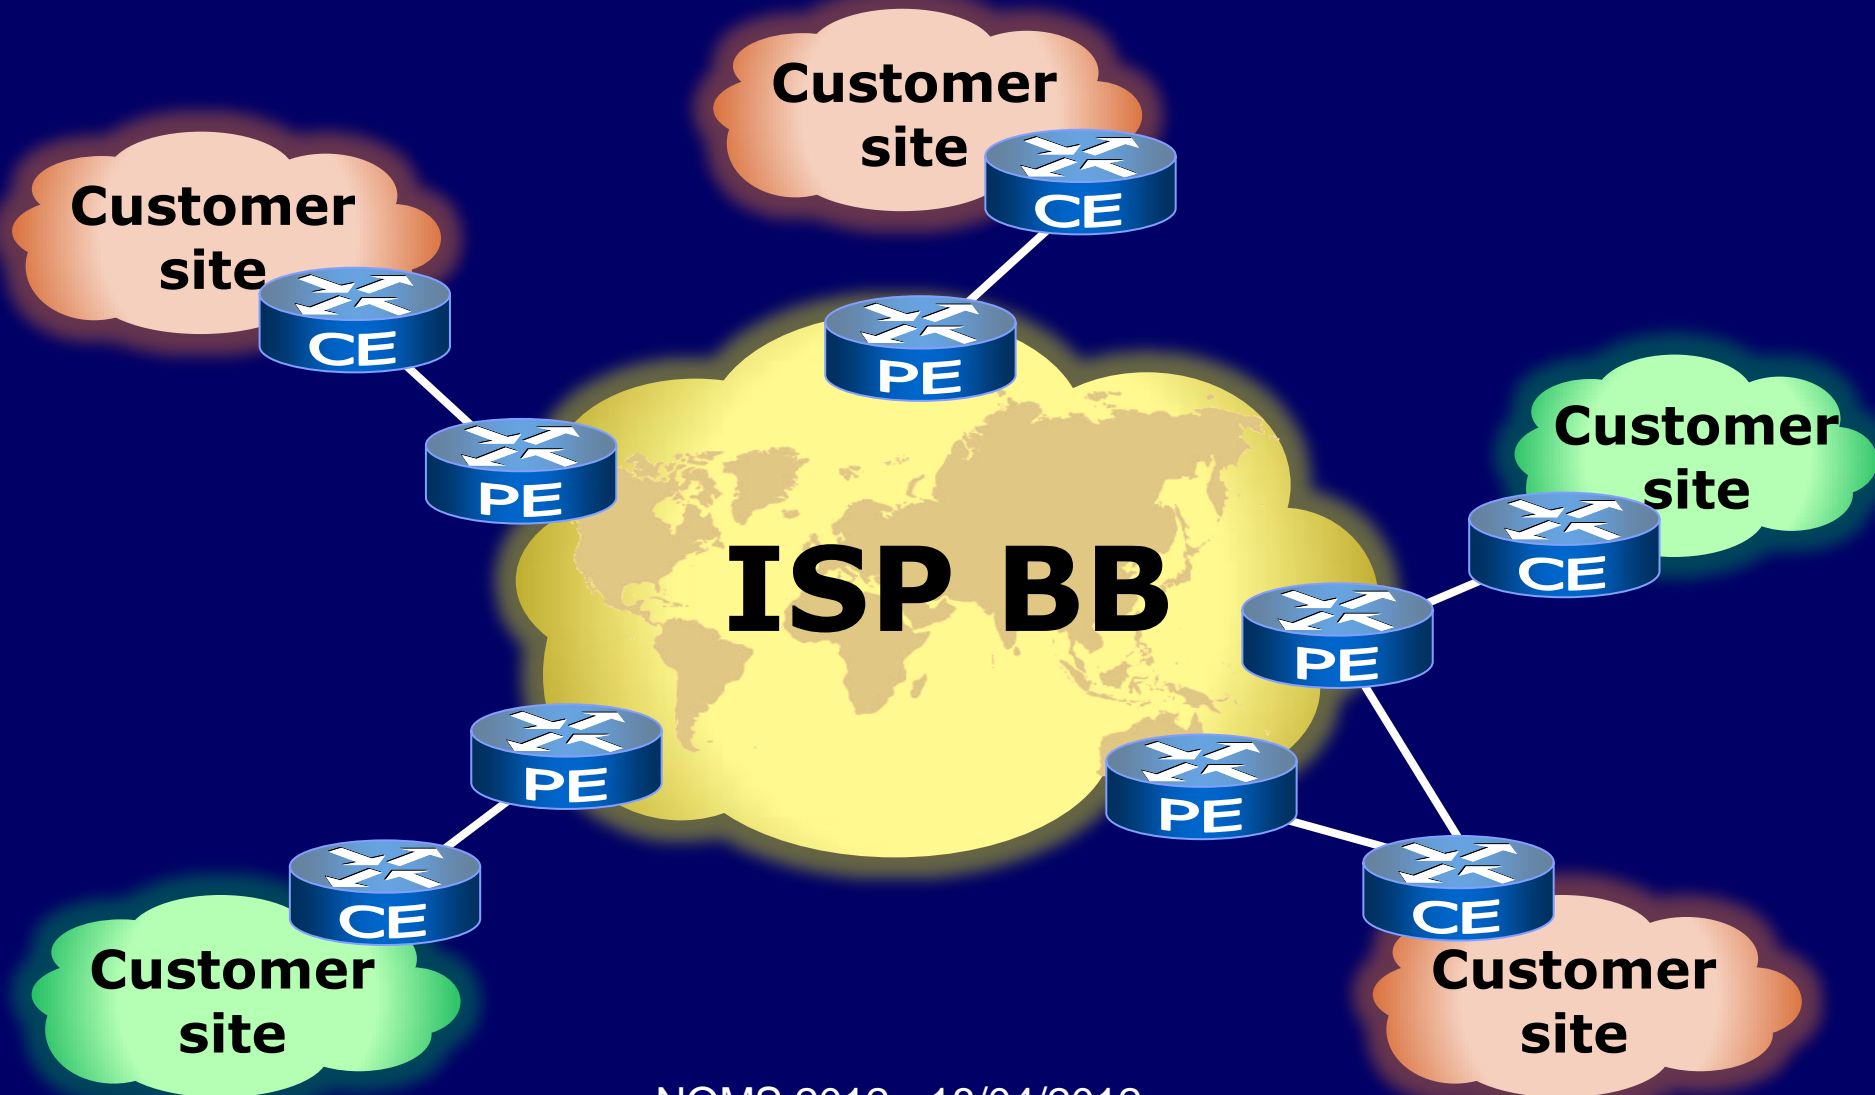
**Customer  
site**

**Customer  
site**

# About MPLS VPNs/VPLS




# About MPLS VPNs/VPLS






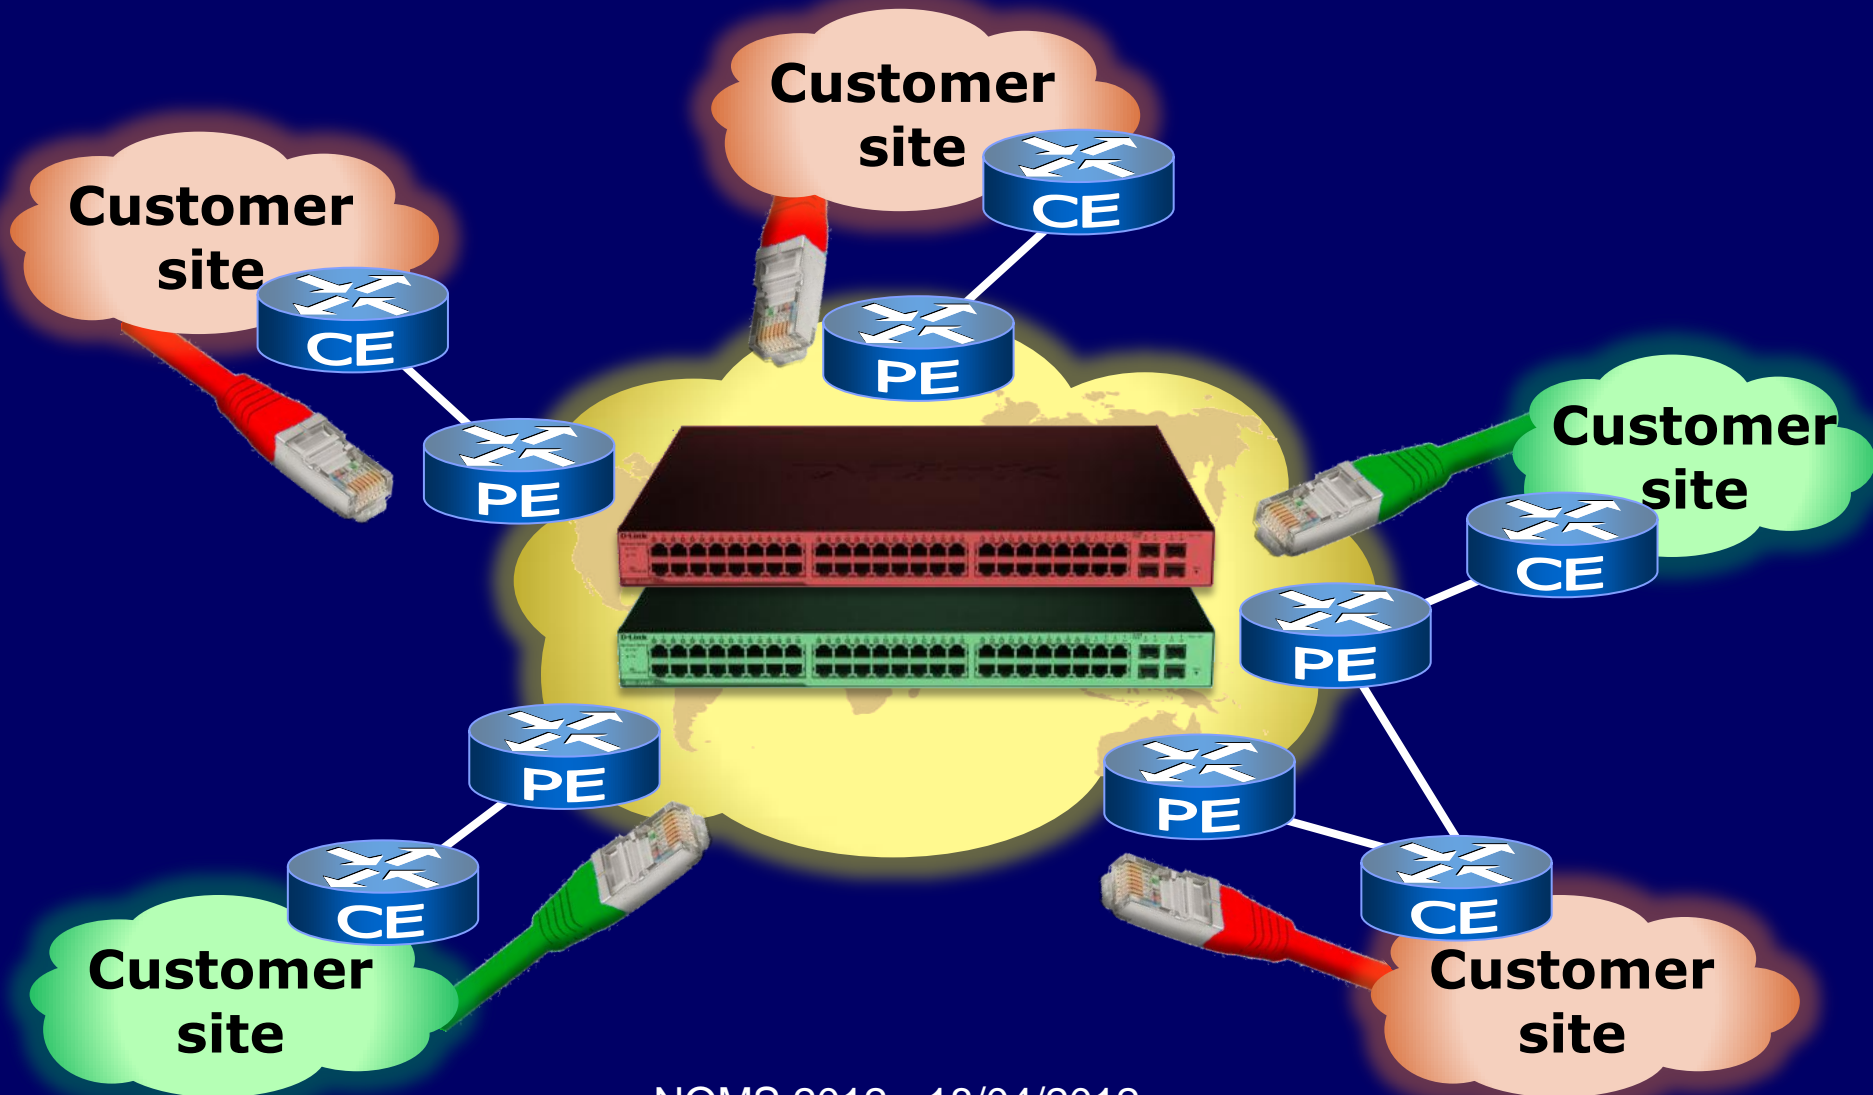
# About MPLS VPNs/VPLS




# About MPLS VPNs/VPLS



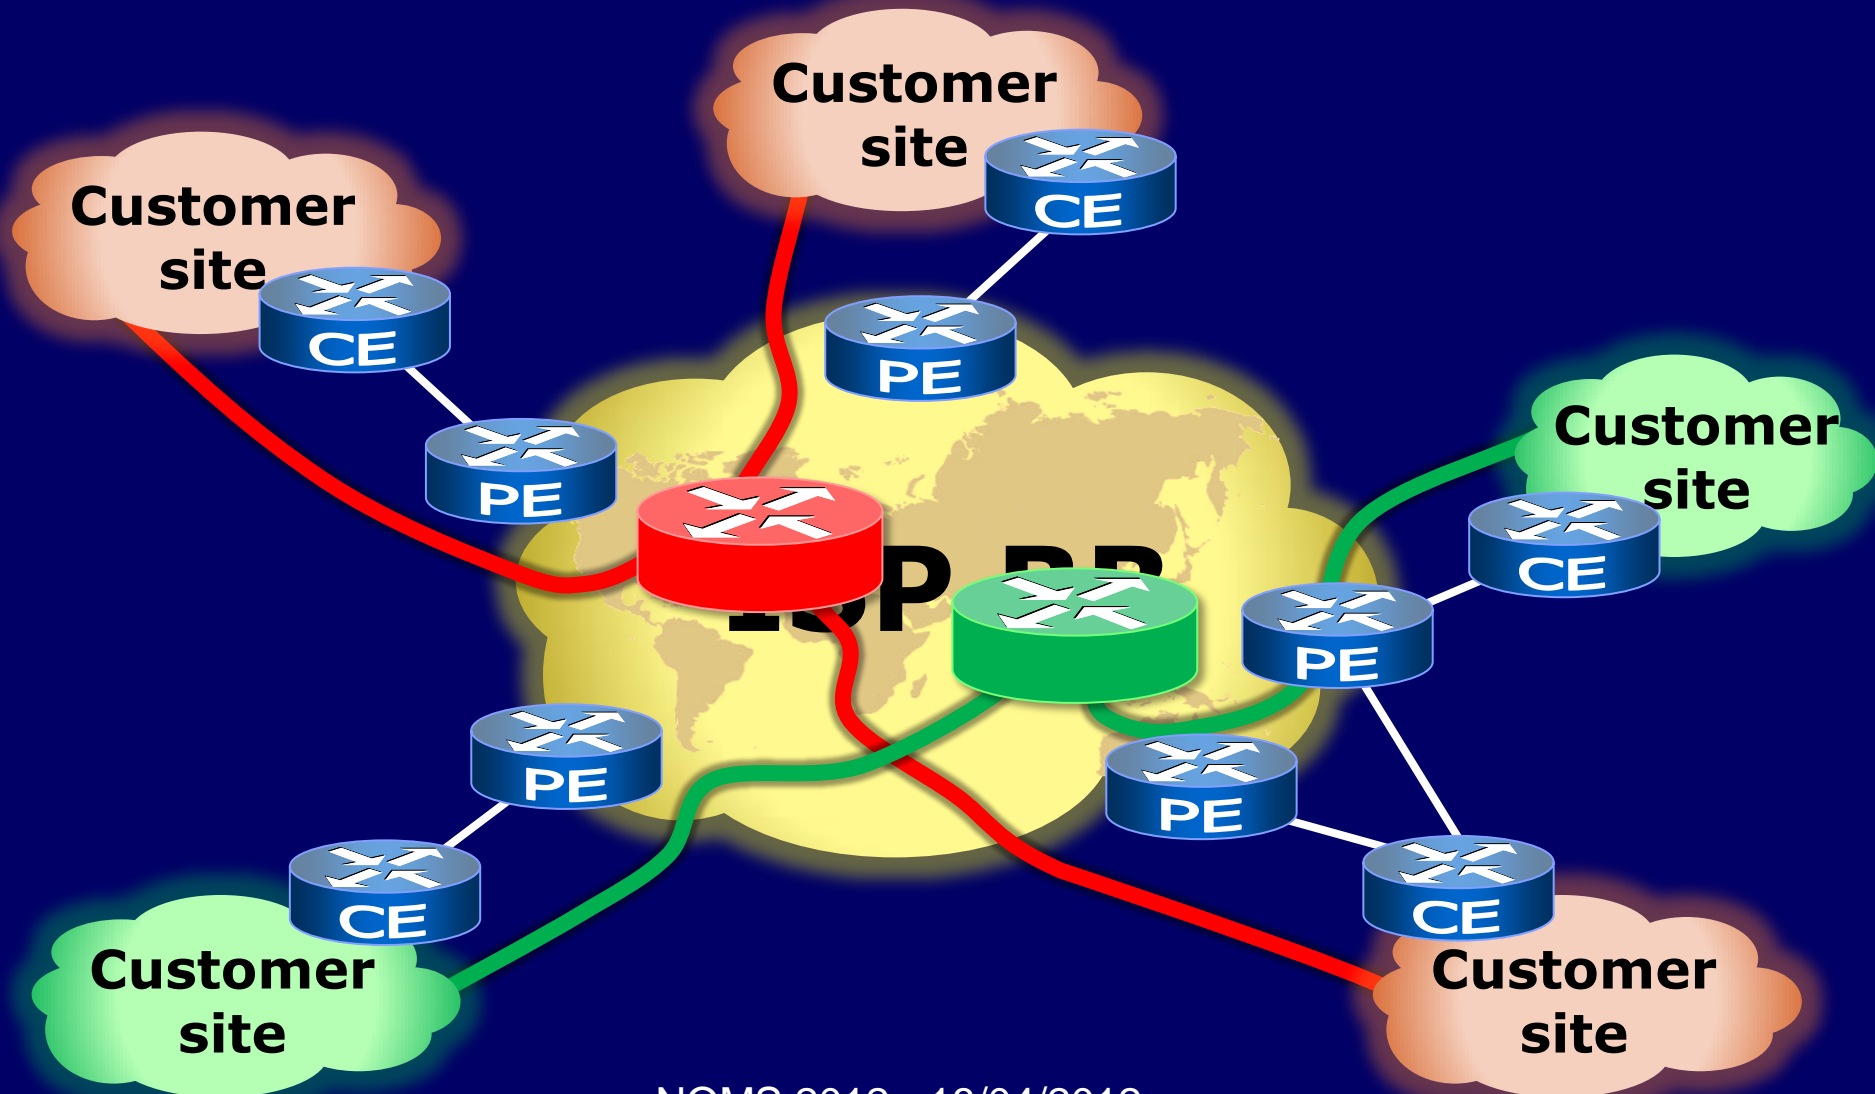
# About MPLS VPNs/VPLS




# About MPLS VPNs/VPLS




# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



The background of the slide features a network diagram. It shows several 'Customer site' clouds, each containing a 'CE' (Customer Edge) router. These CE routers are connected to a central 'PE' (Provider Edge) router. The connections are represented by lines, some of which are highlighted in red and green. The diagram illustrates the architecture of an MPLS VPN or VPLS, where customer sites are isolated from each other but can communicate through the provider's network.

Global Crossing **Level(3)** 's EtherSphere™





# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



Service Inventory | Service Design | Monitoring | Administration  
Inventory and Connection Manager | Deployment Flow Manager | Device Console

You Are Here: Service Inventory > Inventory and Connection Manager > Customers > CPE Devices

**Edi1 CPE Device**

Device Name: enc131.cisco.com  
Site Name: c131  
Customer Name: cust  
Management Type: Managed  
Wildcard Pre-shared Key:

IP Address Ranges

Showing 1-10 of 11 records

#	Name	IP Address	IP Address Type	Encapsulation	Description	IPsec	Firewall	NAT	GoS Candidate
1	Ethernet0/0	192.168.129.88/30	STATIC	ETHERNET	Line to enr3 (192.168.129.88) ! DON'T MODIFY or REMOVE!	Private	Inside	Inside	None
2	Ethernet1/1	192.168.130.33/30	STATIC	ETHERNET		Public	Outside	Outside	None
3	Loopback0	192.168.115.115/32	STATIC	UNKNOWN		Public	None	None	None
4	Ethernet1/0		STATIC	ETHERNET		None	None	None	None
5	Ethernet1/2		STATIC	ETHERNET		None	None	None	None
6	Ethernet1/2	10.13.13.18/30	STATIC	ETHERNET		None	None	None	None
7	Loopback2	12.12.12.11/32	STATIC	UNKNOWN	IPSec Secured Tunnel Endpoint ! DON'T MODIFY or REMOVE!	None	None	None	None
8	Tunnel33		STATIC	UNKNOWN		None	None	None	None
9	Tunnel34	158.50.43.3/24	STATIC	UNKNOWN		None	None	None	None
10	Tunnel37	158.50.73.3/24	STATIC	UNKNOWN		None	None	None	None

Rows per page: 10 | Page 1 of 2 | Go to page: 00 | Save | Cancel

# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS



JUNIPER  
NETWORKS

CISCO



192.168.0.4



# About MPLS VPNs/VPLS



**192.168.0.4**



# About MPLS VPNs/VPLS



**192.168.0.4**



# About MPLS VPNs/VPLS



# About MPLS VPNs/VPLS






# **State of the Art (in MPLS/VPLS monitoring)**



# State of the Art (in MPLS/VPLS monitoring)



# State of the Art (in MPLS/VPLS monitoring)



control plane  
monitoring

# State of the Art (in MPLS/VPLS monitoring)



MPLS *and* VPLS  
control plane  
monitoring



# State of the Art (in MPLS/VPLS monitoring)

industry

research

technology


# State of the Art (in MPLS/VPLS monitoring)

industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)



industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)



IP Solution Center



Service Aware Manager



Service Activator Solution  
for VPN Services



Tivoli Network Manager



VPN Explorer

industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)

industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)

◆ Routing convergence

◆ Scalability

◆ Monitoring

industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)

## ◆ Routing convergence

- D. Pei, J. Van der Merwe. BGP Convergence in Virtual Private Networks. Proc. IMC, 2006.

## ◆ Scalability

- C. Kim, A. Gerber, C. Lund, D. Pei, S. Sen. Scalable VPN Routing via Relaying. Proc. SIGMETRICS, 2008.

## ◆ Monitoring

- M. K. Thottan, G. K. Swanson, M. Cancone, T. K. Ho, J. Ren, S. Paul. SEQUIN: An SNMP-based MPLS Network Monitoring System. Bell Labs Technical Journal 8(1), 95–111, 2003.

# State of the Art (in MPLS/VPLS monitoring)

industry

research

technology



# State of the Art (in MPLS/VPLS monitoring)

- ◆ SNMP
- ◆ TIBCO Rendezvous Message Transport
- ◆ Oracle DBMS
- ◆ RCP, RSH
- ◆ Telnet, SSH
- ◆ TFTP, FTP

industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)

- ◆ SNMP
- ◆ TIBCO Rendezvous Message Transport
- ◆ Oracle DBMS
- ◆ RCP, RSH
- ◆ Telnet, SSH
- ◆ TFTP, FTP

industry

research

technology

# State of the Art (in MPLS/VPLS monitoring)

- ◆ SNMP
- ◆ TIBCO Rendezvous Message Transport
- ◆ Oracle DBMS
- ◆ RCP, RSH
- ◆ Telnet, SSH
- ◆ TFTP, FTP

industry

research

technology

# Our Contributions



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring + provisioning



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring
  - Observation of the network status





# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring
  - Observation of effects of network events



# Our Contributions

- ✦ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring
  - Observation of effects of network events
    - Reconfigurations
    - Failures



# Our Contributions

- ✦ MPLS VPN/VPLS monitoring methodology
    - Focus on monitoring
    - Observation of effects of network events
      - Reconfigurations
      - Failures
- ✓ Exhaustive analysis of observable effects



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring
  - Observation of effects of network events
    - Reconfigurations ✓ Exhaustive analysis of observable effects
    - Failures
  - Instant snapshot of device states



# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations
  - Failures
- (Almost) Instant snapshot of device states

✓ Exhaustive analysis of observable effects



# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations
  - Failures
- (Almost) Instant snapshot of device states + history



# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations
  - Failures
- (Almost) Instant snapshot of device states + history
- Additional technologies required



# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations      ✓ Exhaustive analysis of observable effects
  - Failures
- (Almost) Instant snapshot of device states + history
- Standard technologies (BGP)





# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations ✓ Exhaustive analysis of observable effects
  - Failures
- (Almost) Instant snapshot of device states + history
- Standard technologies (BGP)
- Requires access to devices



# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations ✓ Exhaustive analysis of observable effects
  - Failures
- (Almost) Instant snapshot of device states + history
- Standard technologies (BGP)
- Unobtrusive



# Our Contributions

## ◆ MPLS VPN/VPLS monitoring methodology

- Focus on monitoring
- Observation of effects of network events
  - Reconfigurations ✓ Exhaustive analysis of observable effects
  - Failures
- (Almost) Instant snapshot of device states + history
- Standard technologies (BGP)
- Unobtrusive
- Graphical visualization of VPN states



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring
  - Observation of effects of network events
    - Reconfigurations ✓ Exhaustive analysis of observable effects
    - Failures
  - (Almost) Instant snapshot of device states + history
  - Standard technologies (BGP)
  - Unobtrusive
  - Graphical visualization of VPN states
- ◆ Extensive discussion on **scalability** vs **visibility** of (the effects of) network events



# Our Contributions

- ◆ MPLS VPN/VPLS monitoring methodology
  - Focus on monitoring
  - Observation of effects of network events
    - Reconfigurations ✓ Exhaustive analysis of observable effects
    - Failures
  - (Almost) Instant snapshot of device states + history
  - Standard technologies (BGP)
  - Unobtrusive
  - Graphical visualization of VPN states
- ◆ Extensive discussion on **scalability** vs **visibility** of (the effects of) network events
- ◆ Architecture, prototype, experimentation in Junosphere



# Our Contributions

- (Almost) Instant snapshot of device states & history
- Standard technologies (BGP)
- Unobtrusive
- Graphical visualization of VPN states
- ◆ Extensive discussion on **scalability** vs **visibility** (the effects of) network events
- ◆ Architecture, prototype, experimentation in Junosphere
- ◆ Discovery of a subtle anomaly in the routing software, confirmed by Juniper





# Methodology




# Methodology



Collect signaling messages



# Methodology



1 Collect signaling messages

2 Reconstruct visibility of VPNs at PEs

# Methodology



1 Collect signaling messages

2 Reconstruct visibility of VPNs at PEs

3 Visualize VPN states

# Methodology





Collect signaling messages

Reconstruct visibility of VPNs at PEs

Visualize VPN states





# Methodology

## 1. Collection

**Approach**

**Drawback(s)**

# Methodology

## 1. Collection

### Approach

Monitor network traffic

### Drawback(s)

Undetermined in absence of traffic

# Methodology

## 1. Collection

### Approach

### Drawback(s)

Monitor network traffic

Undetermined in absence of traffic

Inject network traffic

Intrusive; hard to tune

# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions may apply

# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions may apply
Watch router states	Same as above + untimely



# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions may apply
Watch router states	Same as above + untimely
Notifications (e.g., SNMP)	Additional technologies required

# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions may apply
Watch router states	Same as above + untimely
Notifications (e.g., SNMP)	Additional technologies required
*	Limited visibility of the effect of a configuration

# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions • Actual propagation of information • Routing decisions @ PEs
Watch router states	Stateful
Notifications (e.g., SNMP)	Access required
*	Limited visibility of the effect of a configuration

# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions may apply
Watch router states	Same as above + untimely
Notifications (e.g., SNMP)	Additional technologies required
*	Limited visibility of the effect of a configuration
Monitor signaling messages	N/A

# Methodology

## 1. Collection

Approach	Drawback(s)
Monitor network traffic	Undetermined in absence of traffic
Inject network traffic	Intrusive; hard to tune
Watch router configurations	Intrusive; access restrictions may apply
Watch router states	Same as above + untimely
Notifications (e.g., SNMP)	Additional technologies required
*	Limited visibility of the effect of a configuration
Monitor signaling messages	N/A

# Methodology

## 1. Collection

### Approach

### Drawback(s)

Monitor network traffic

Undetermined in absence of traffic

Inject network traffic

Hard to tune

Watch router configuration

Access restrictions

Watch router state

Same as above + untimely

Notifications (e.g., SNMP)

Additional technologies required

\*

Limited visibility of the effect of a configuration

Monitor signaling messages

N/A



# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	<b>Autodiscovery</b>	<b>Signaling</b>	<b>Vendor</b>
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	<b>Autodiscovery</b>	<b>Signaling</b>	<b>Vendor</b>
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

BGP-based VPLS  
Autodiscovery



# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	<b>Autodiscovery</b>	<b>Signaling</b>	<b>Vendor</b>
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

BGP-based VPLS  
Autodiscovery

LDP-BGP VPLS  
Interworking

# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	<b>Autodiscovery</b>	<b>Signaling</b>	<b>Vendor</b>	
RFC 4762 (Kompella)	N/A	LDP	Cisco	BGP-based VPLS Autodiscovery  LDP-BGP VPLS Interworking
RFC 4761	BGP	BGP	Juniper	

# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	Autodiscovery	Signaling	Vendor
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

BGP-based VPLS  
Autodiscovery

LDP-BGP VPLS  
Interworking

# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	Autodiscovery	Signaling	Vendor
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

BGP-based VPLS  
Autodiscovery

LDP-BGP VPLS  
Interworking

BGP is also...

# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	Autodiscovery	Signaling	Vendor
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

BGP-based VPLS  
Autodiscovery

LDP-BGP VPLS  
Interworking

BGP is also...  
easy to set up  
scalable  
policy-aware

# Methodology

## 1. Collection

- ◆ VPN signaling
  - MPLS: BGP
  - VPLS:

	Autodiscovery	Signaling	Vendor
RFC 4762 (Kompella)	N/A	LDP	Cisco
RFC 4761	BGP	BGP	Juniper

BGP-based VPLS  
Autodiscovery


LDP-BGP VPLS  
Interworking



BGP is also...  
easy to set up  
scalable  
policy-aware

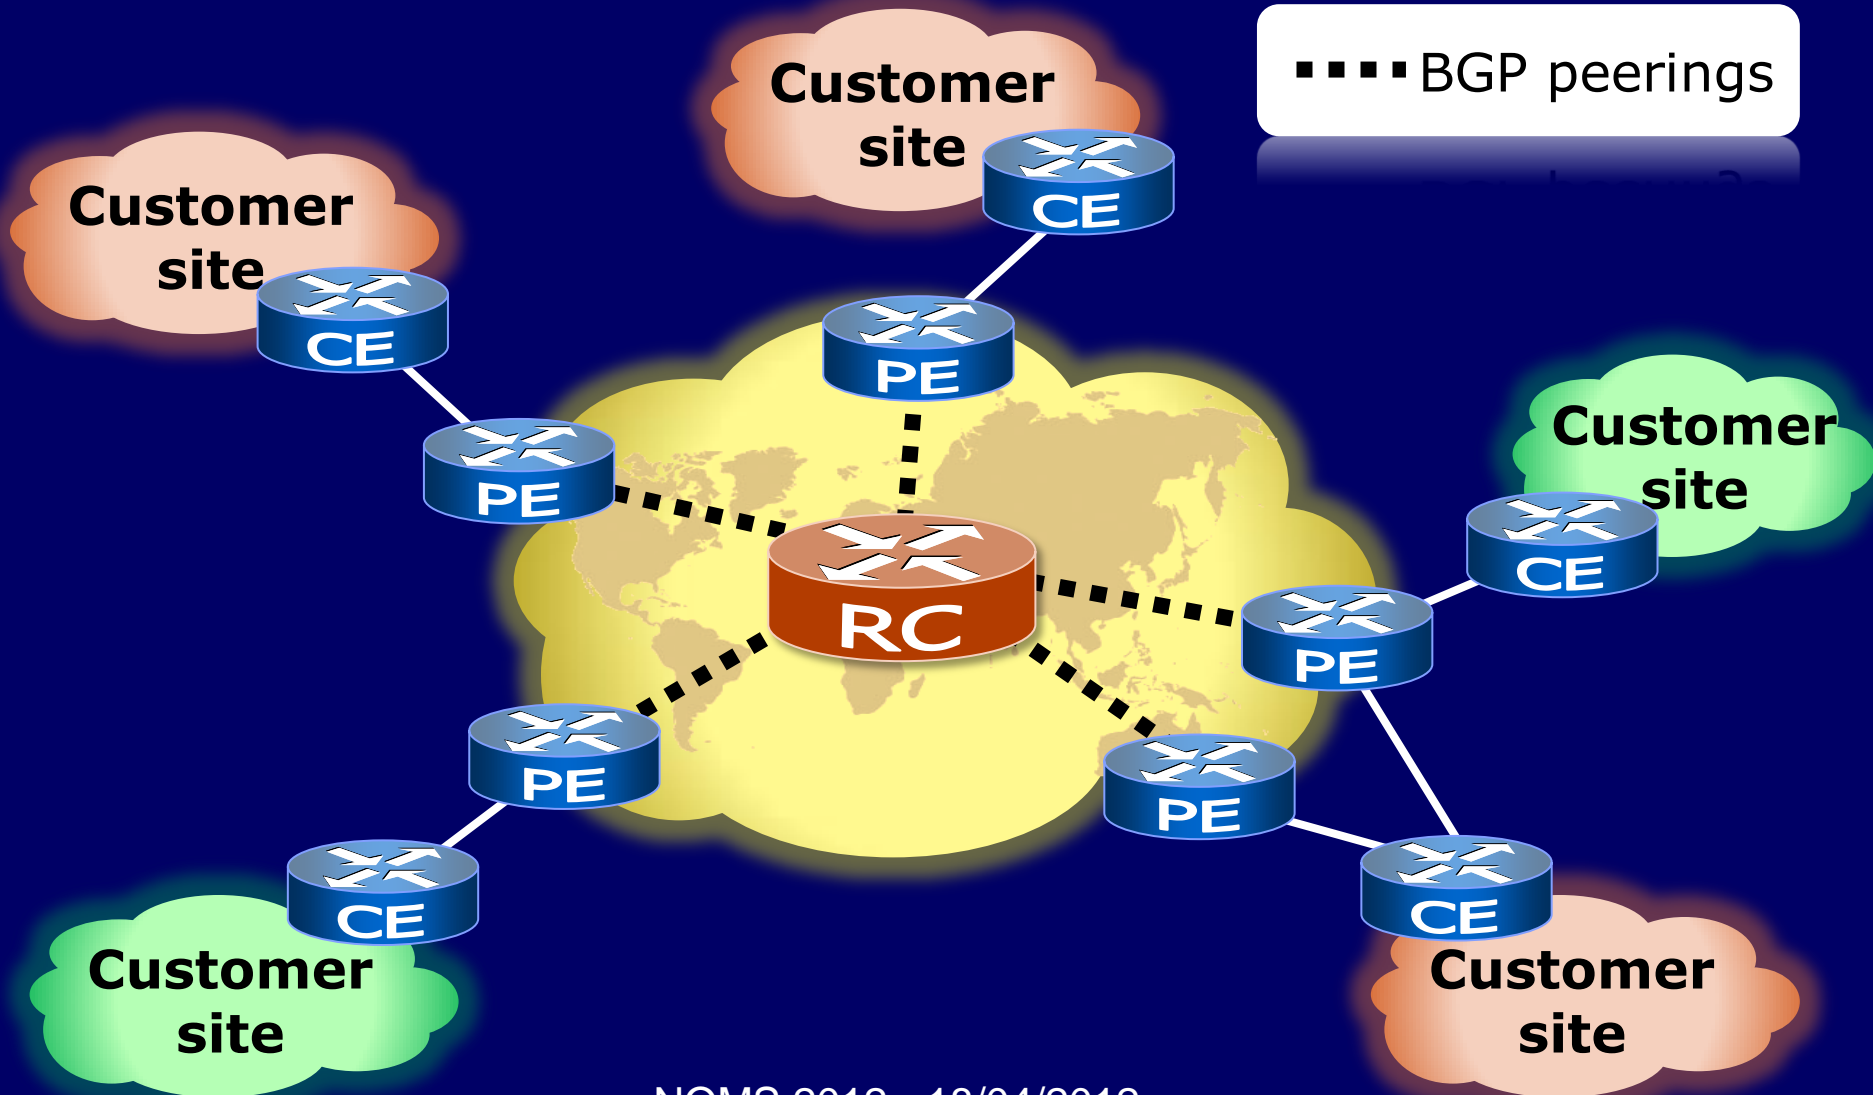
# Methodology

## 1. Collection



# Methodology

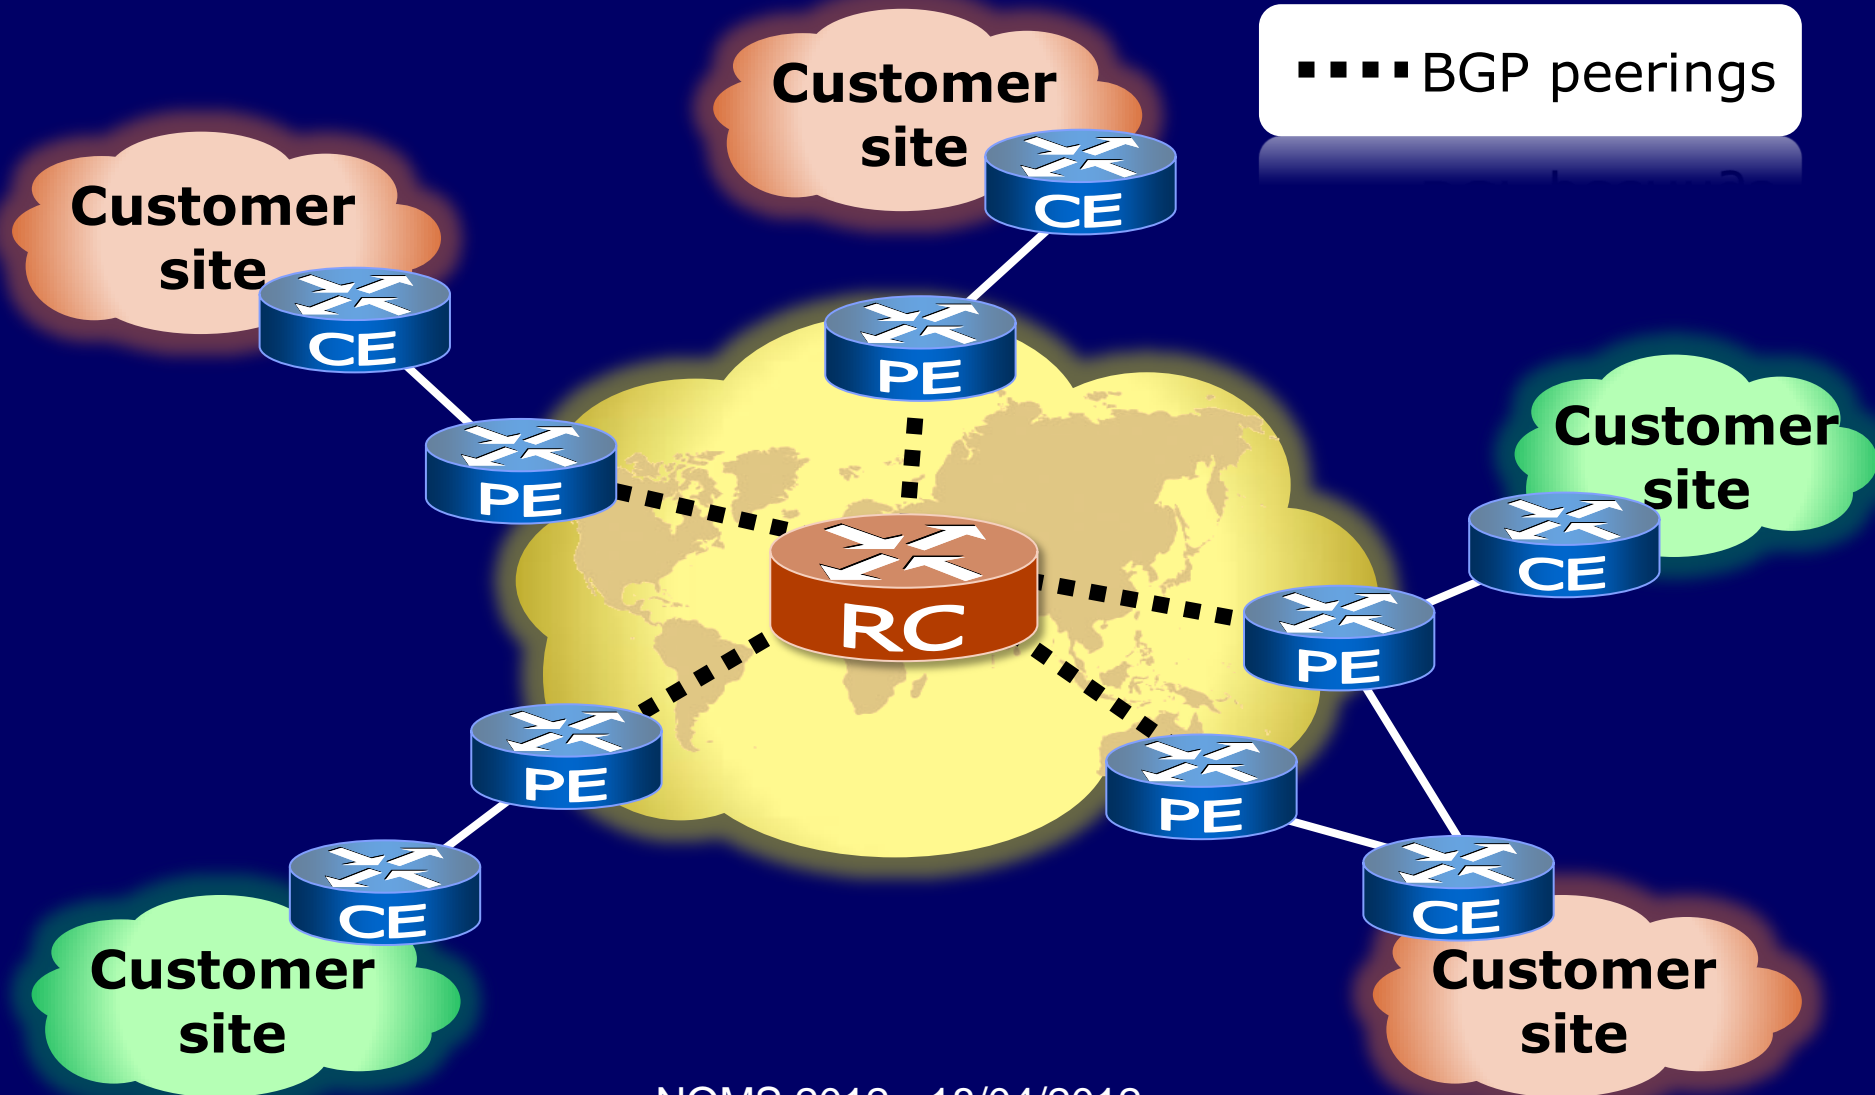
## 1. Collection





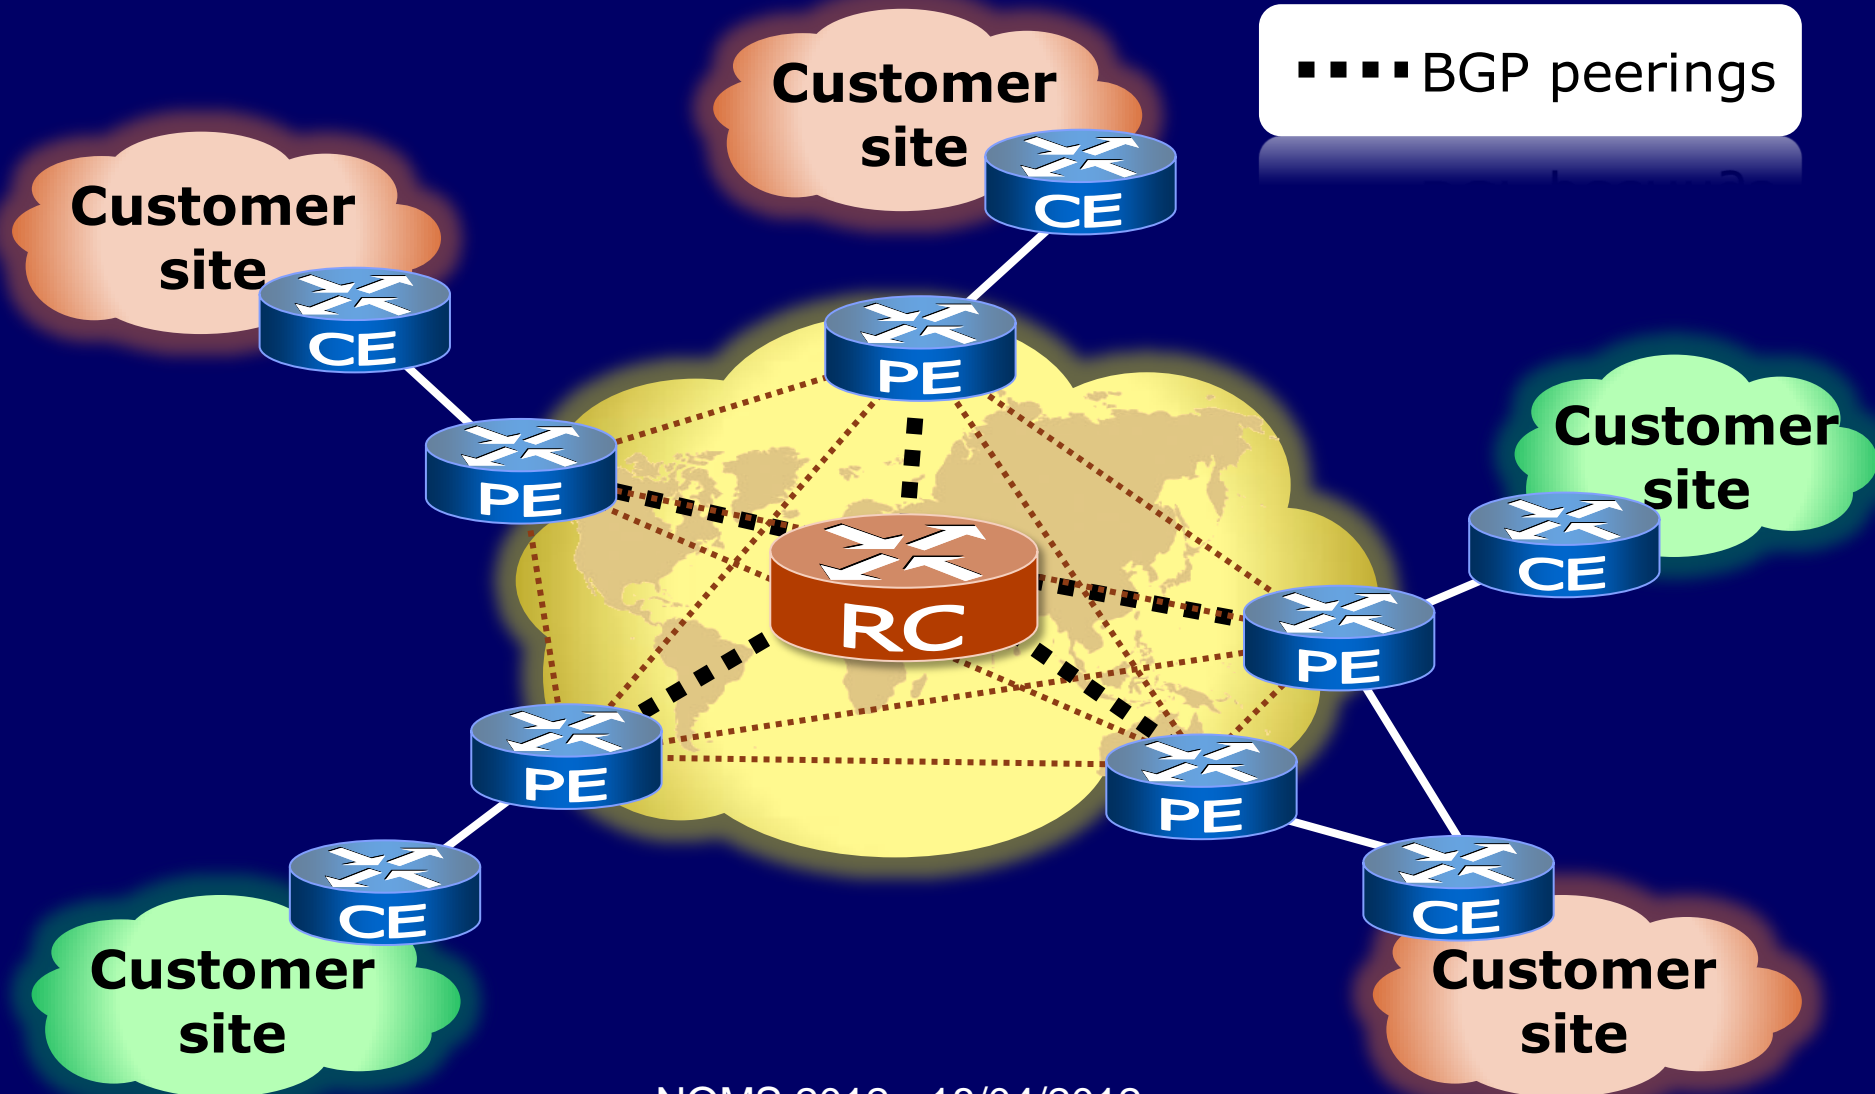
# Methodology

## 1. Collection



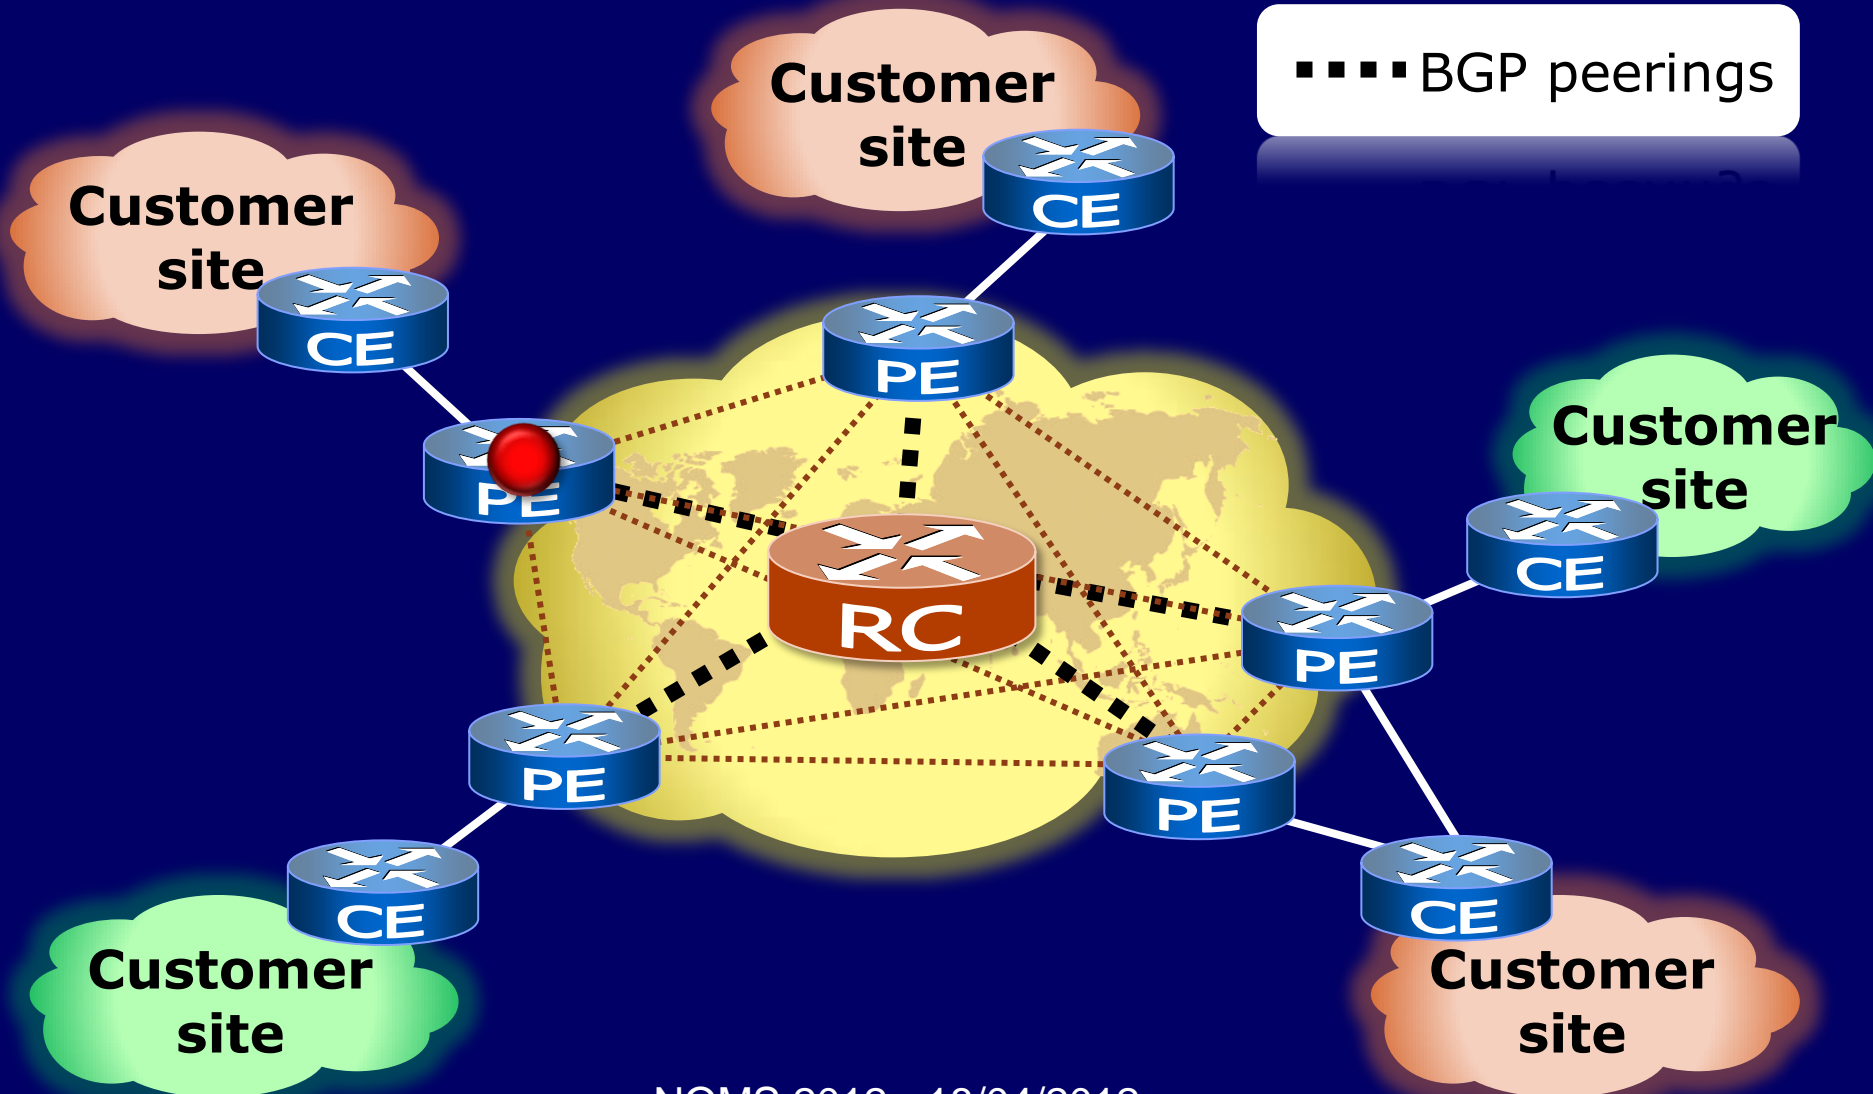
# Methodology

## 1. Collection



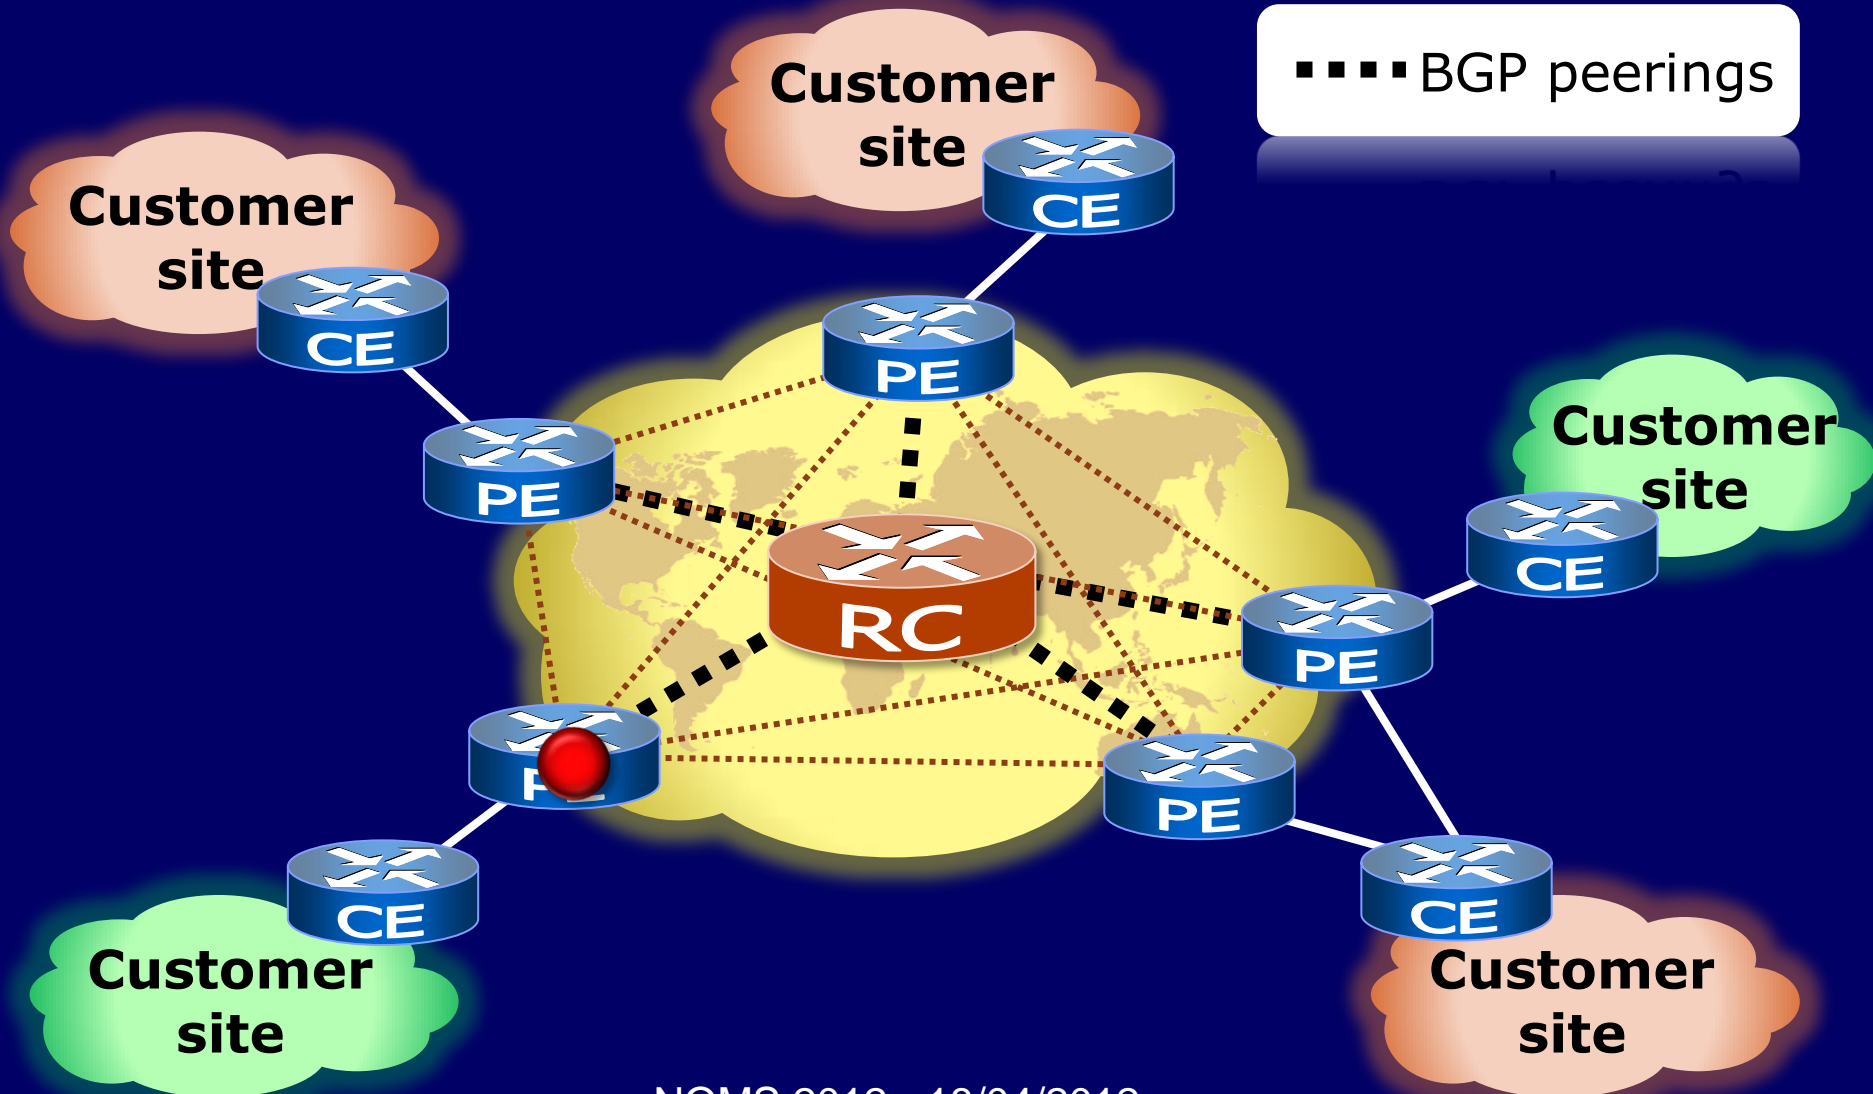
# Methodology

## 1. Collection



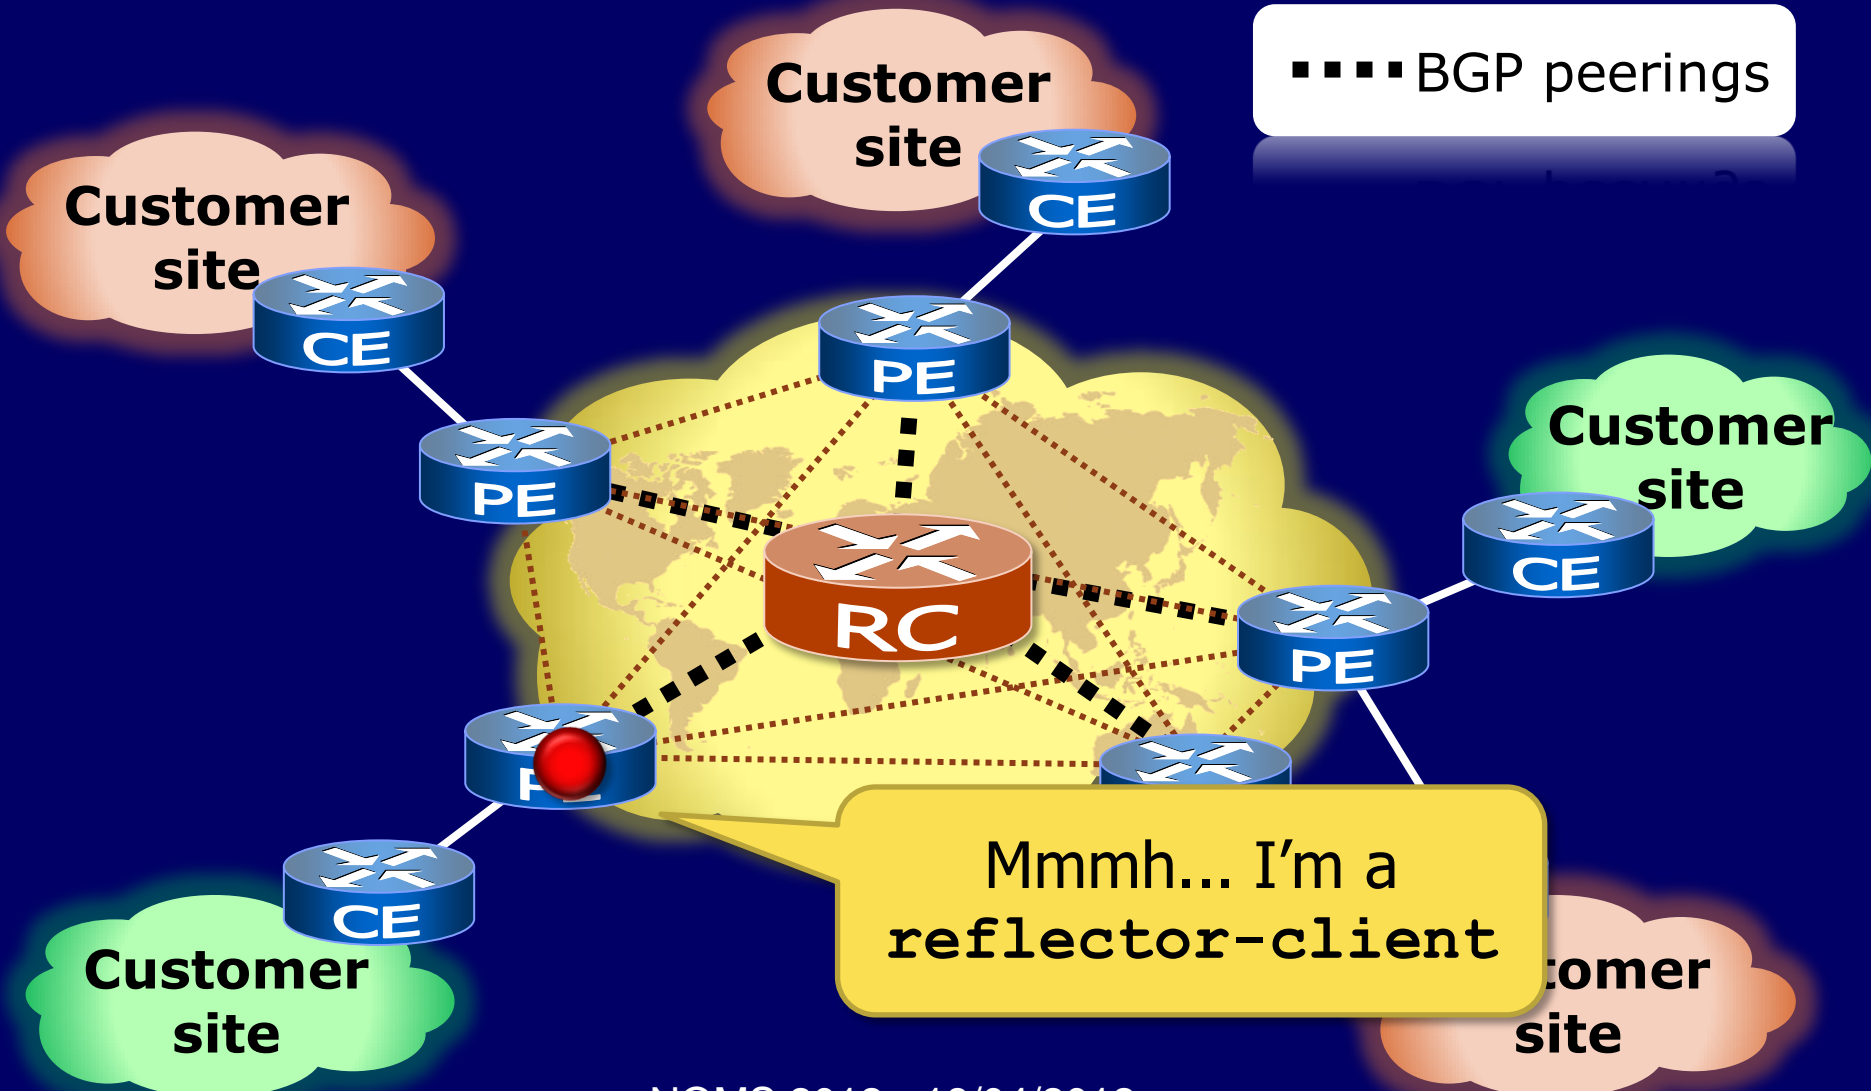
# Methodology

## 1. Collection



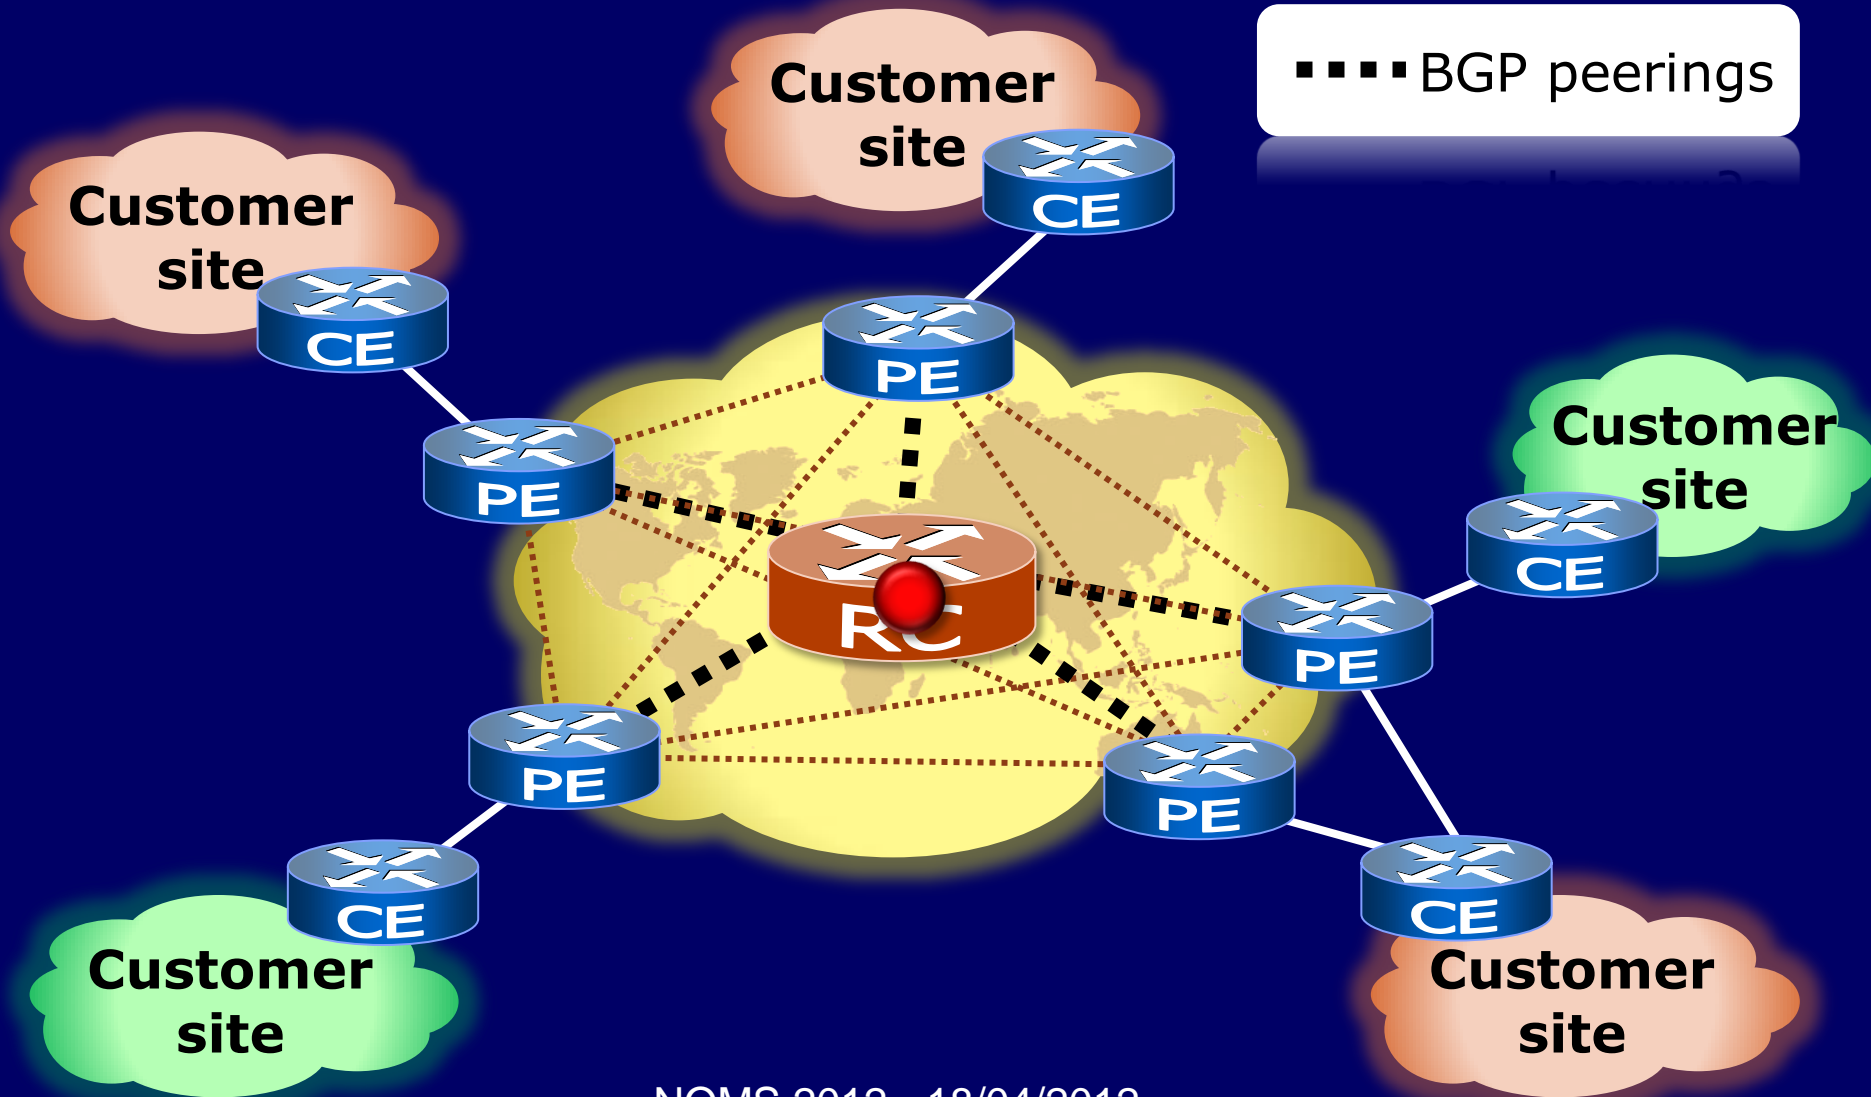
# Methodology

## 1. Collection



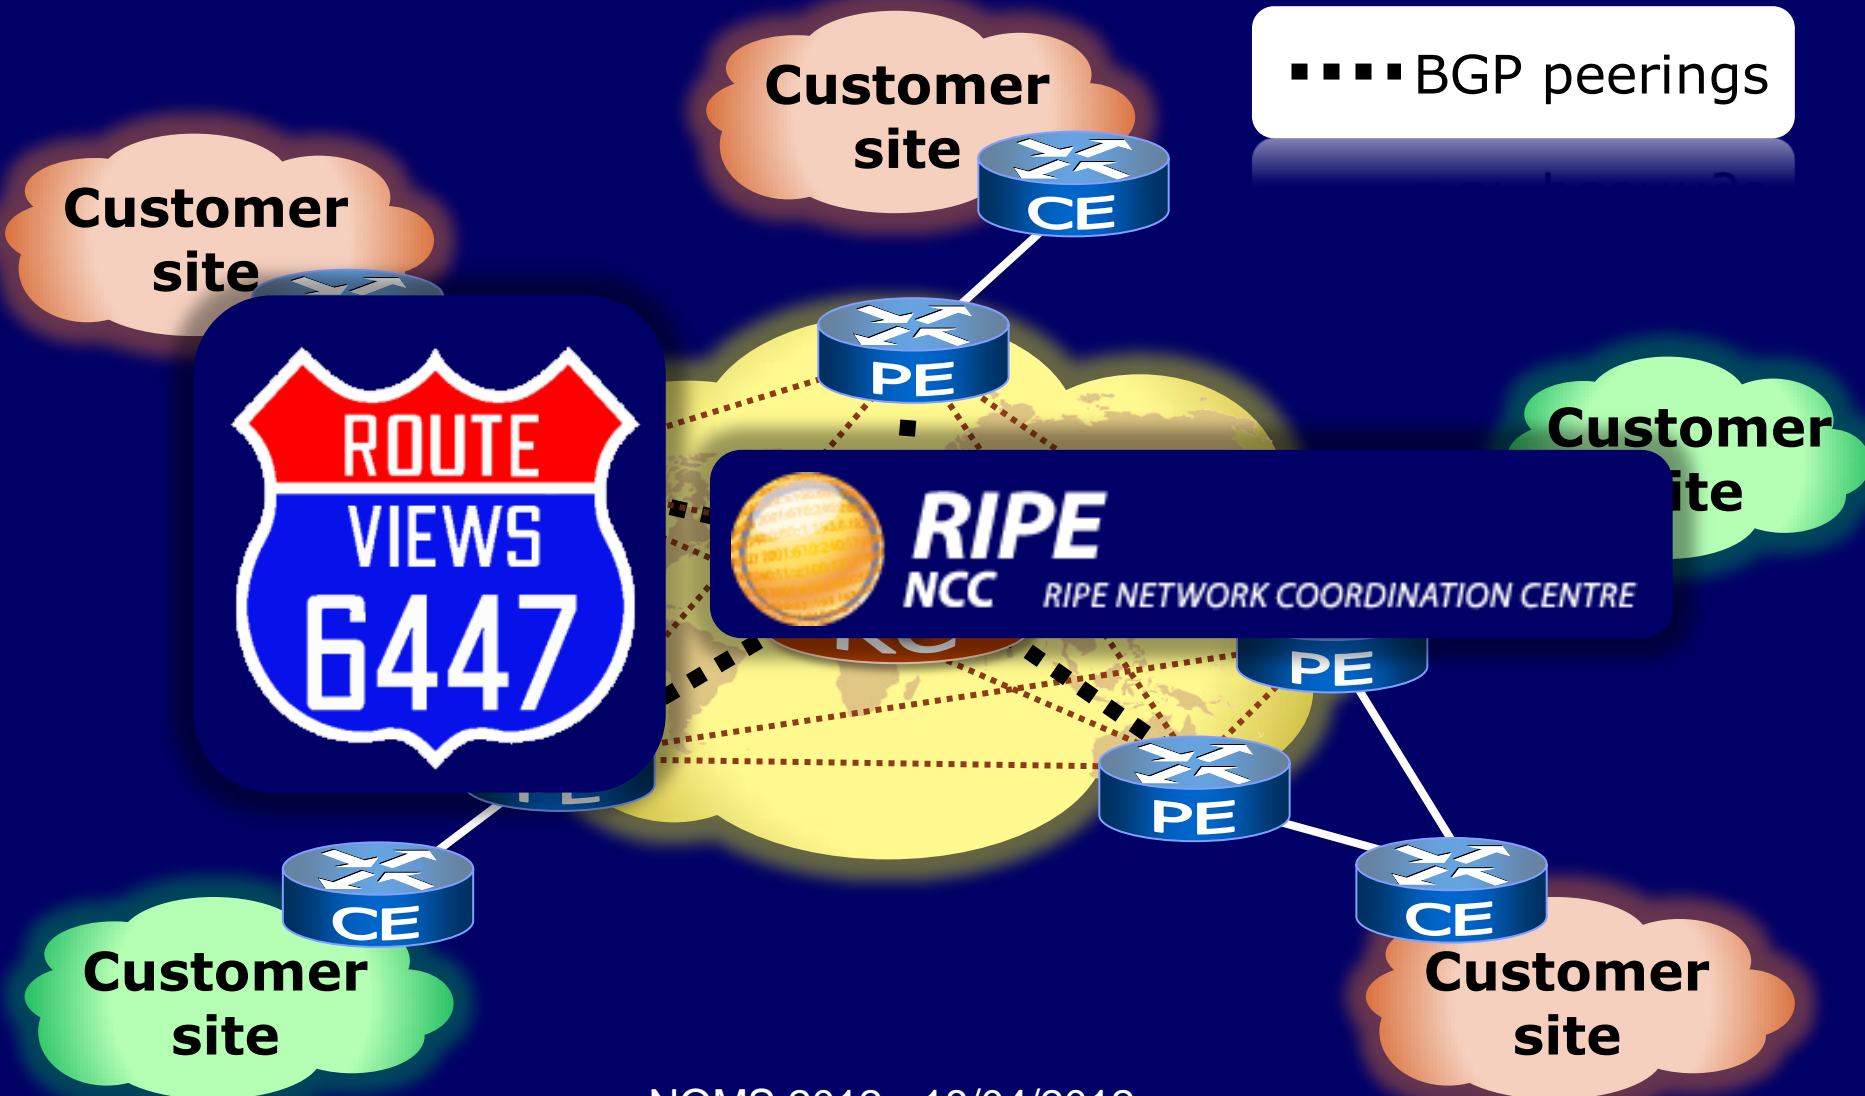
# Methodology

## 1. Collection



# Methodology

## 1. Collection



# Methodology

## 1. Collection





# Methodology

## 2. Reconstruction of VPN state

# Methodology

## 2. Reconstruction of VPN state



# Methodology

## 2. Reconstruction of VPN state



# Methodology

## 2. Reconstruction of VPN state

type (A/W)



# Methodology

## 2. Reconstruction of VPN state

type (A/W)



# Methodology

## 2. Reconstruction of VPN state

type (A/W)

NLRI

Extended  
communities

timestamp

BGP update



# Methodology

## 2. Reconstruction of VPN state

type (A/W)

NLRI

RD

+

prefix

Extended  
communities

timestamp

BGP update

# Methodology

## 2. Reconstruction of VPN state

type (A/W)

NLRI

RD

+

prefix

CE ID

Extended  
communities

timestamp

BGP update



# Methodology

## 2. Reconstruction of VPN state

type (A/W)

NLRI

RD

+

prefix

CE ID

Extended  
communities

RT

timestamp

BGP update

# Methodology

## 2. Reconstruction of VPN state

type (A/W)

NLRI

RD

+

prefix

CE ID

Extended  
communities

RT

timestamp

BGP update

- ✦ Exhaustive comparison of information from different BGP updates

# Methodology

## 2. Reconstruction of VPN state

### ◆ Example

# Methodology

## 2. Reconstruction of VPN state

### ◆ Example



# Methodology

## 2. Reconstruction of VPN state

### ✦ Example



RD<sub>1</sub>

+

px<sub>1</sub>

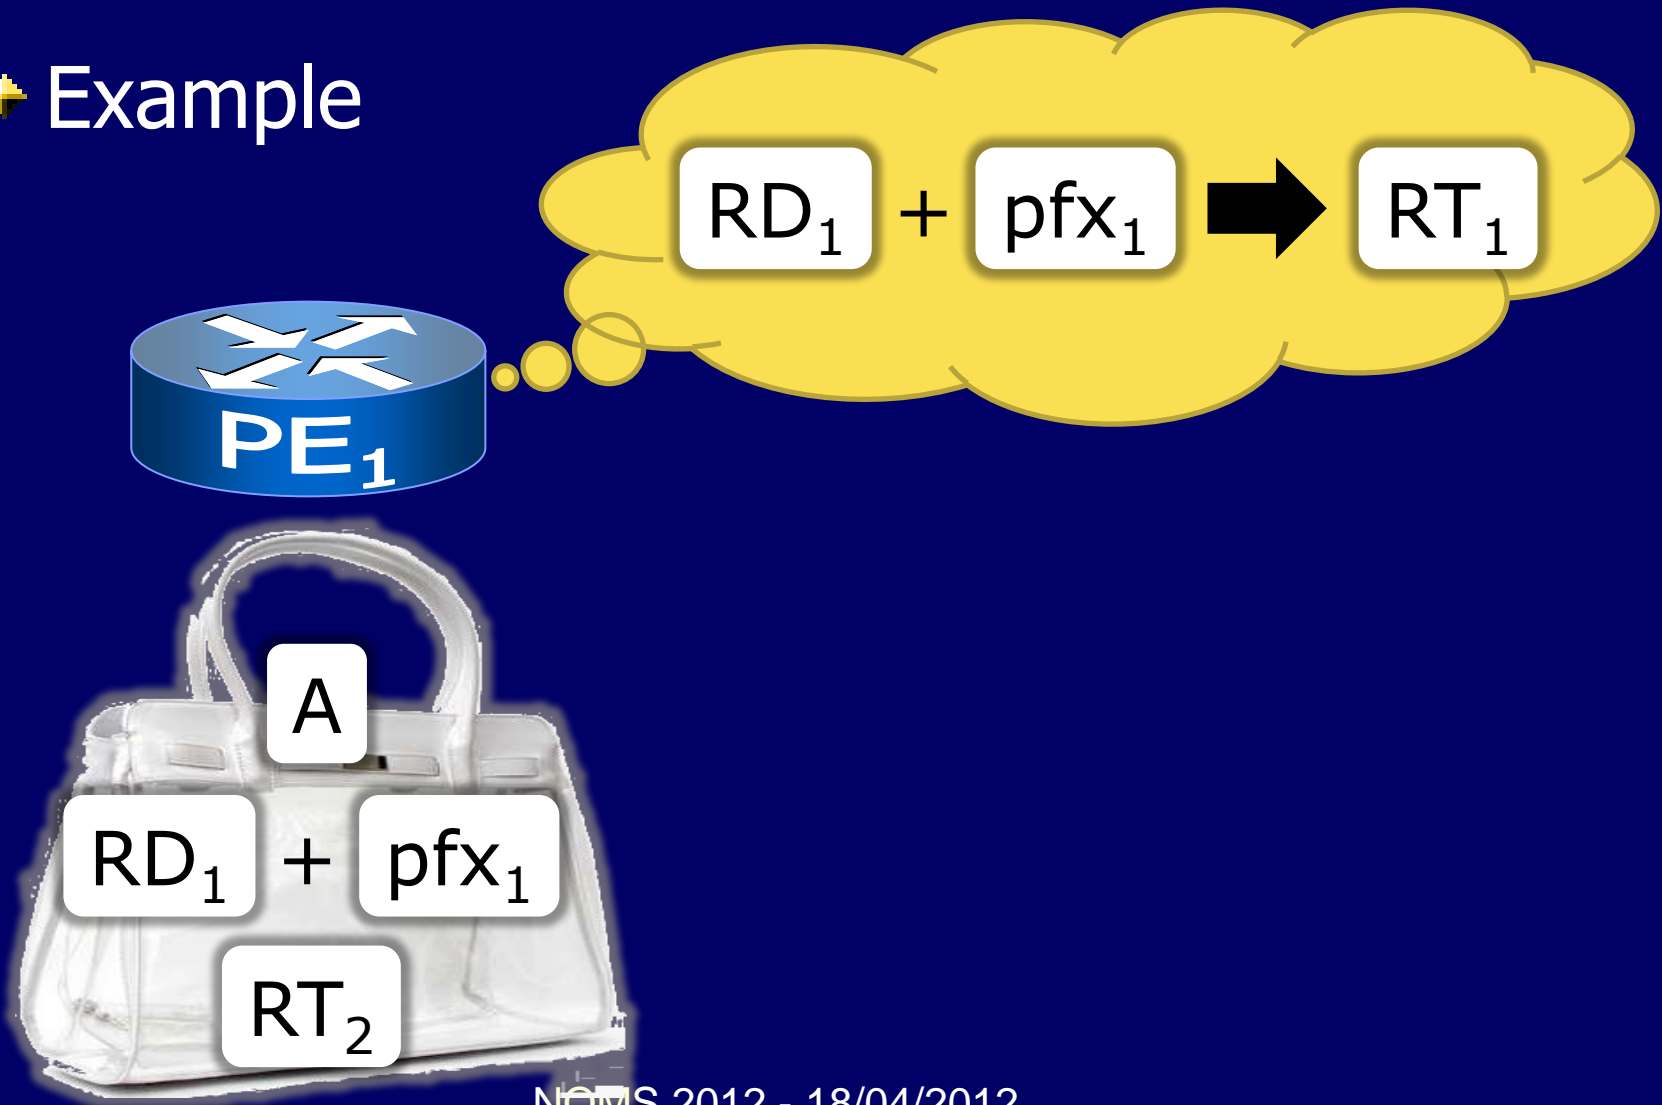


RT<sub>1</sub>

# Methodology

## 2. Reconstruction of VPN state

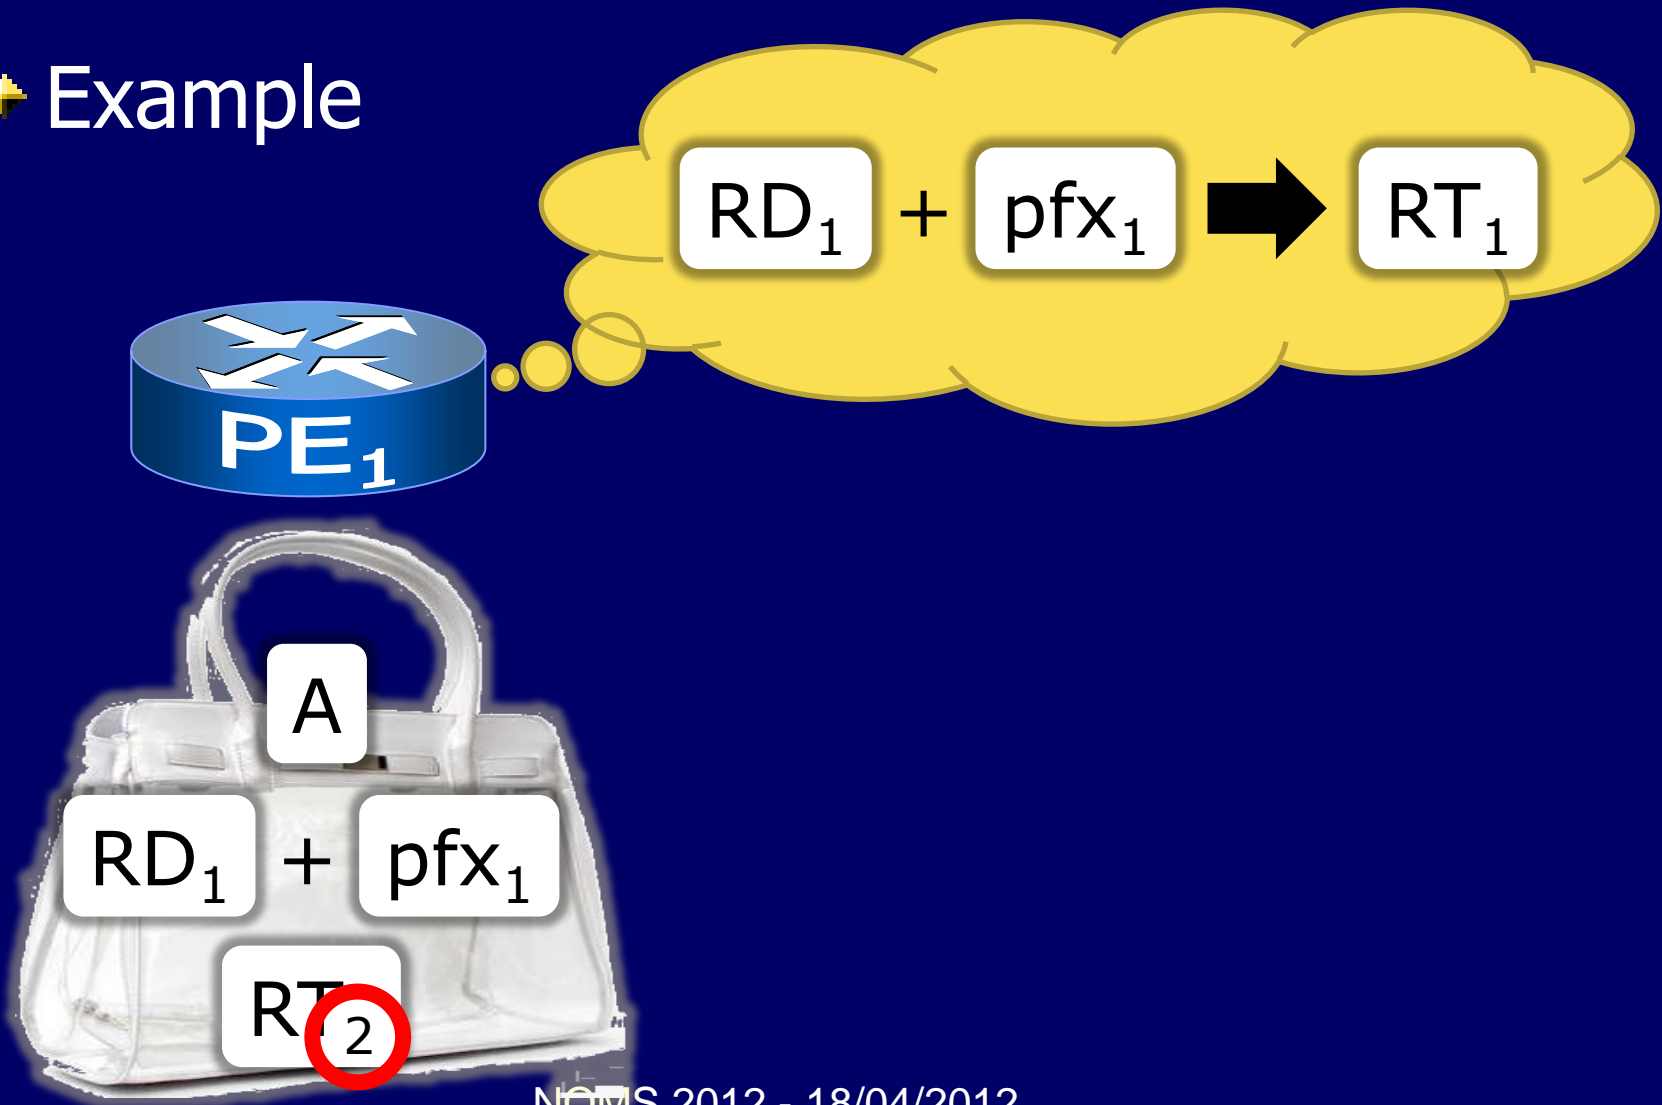
### ✦ Example



# Methodology

## 2. Reconstruction of VPN state

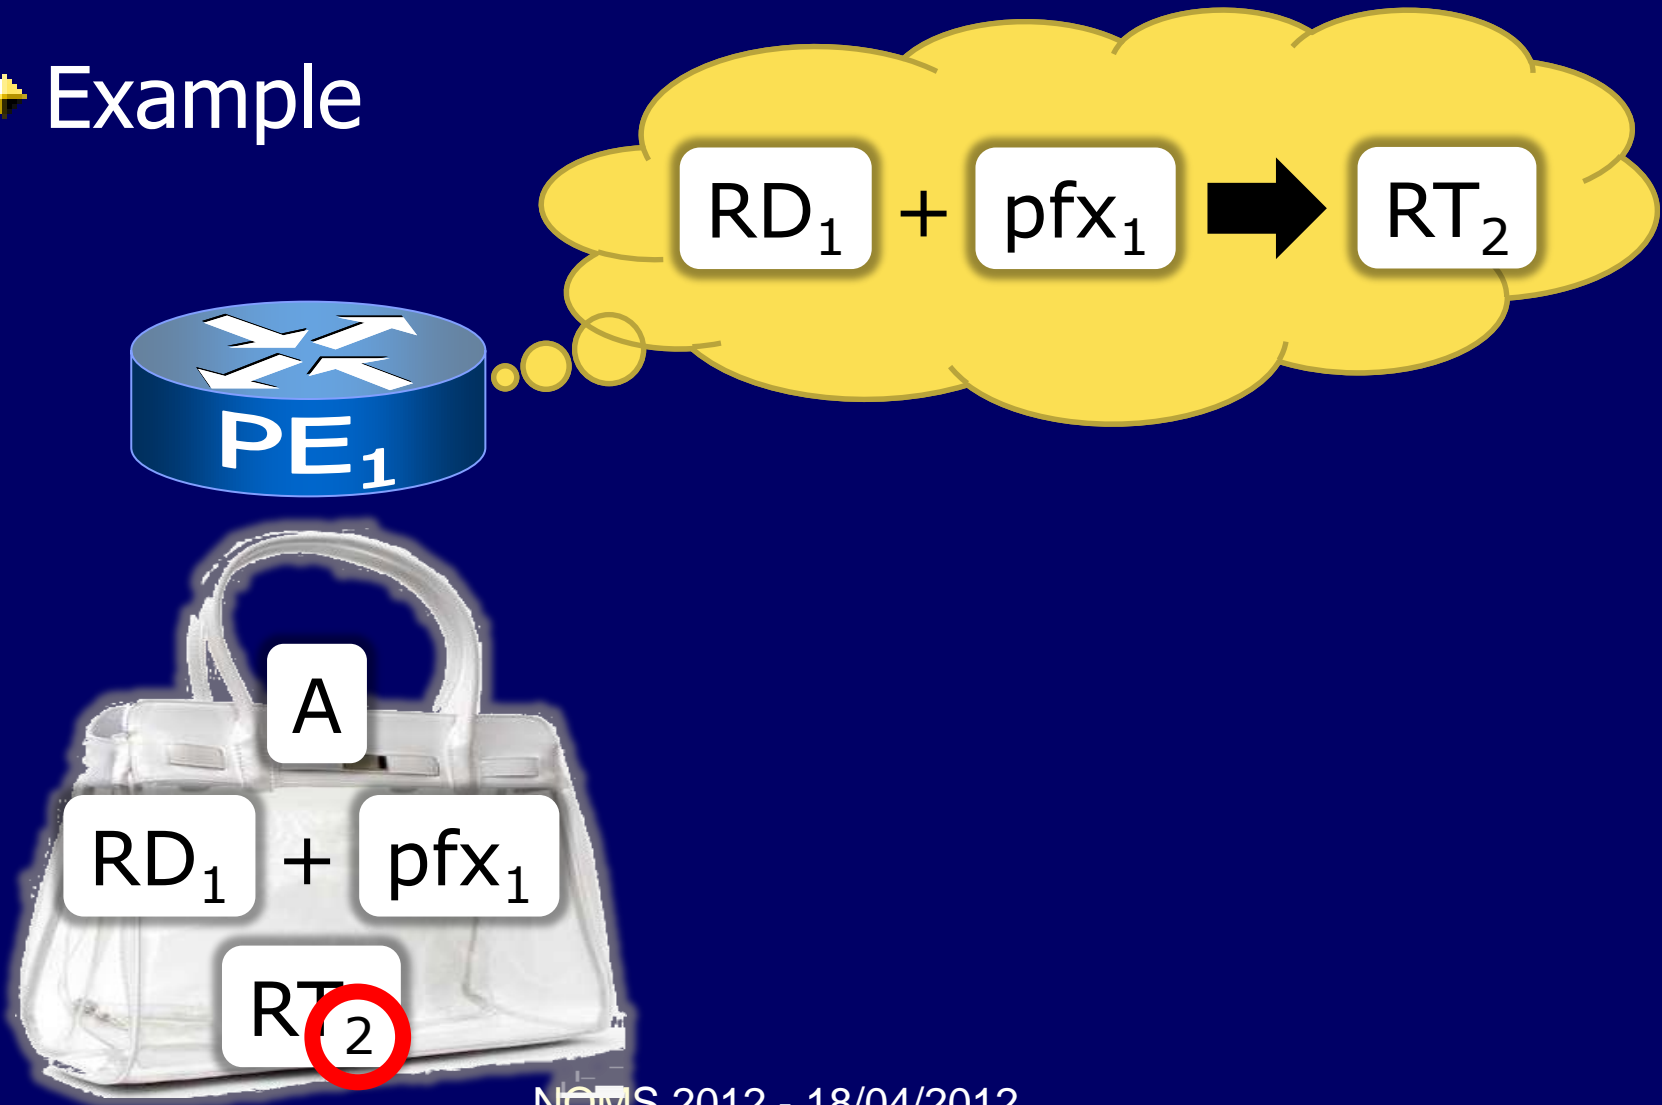
### ✦ Example



# Methodology

## 2. Reconstruction of VPN state

### ✦ Example

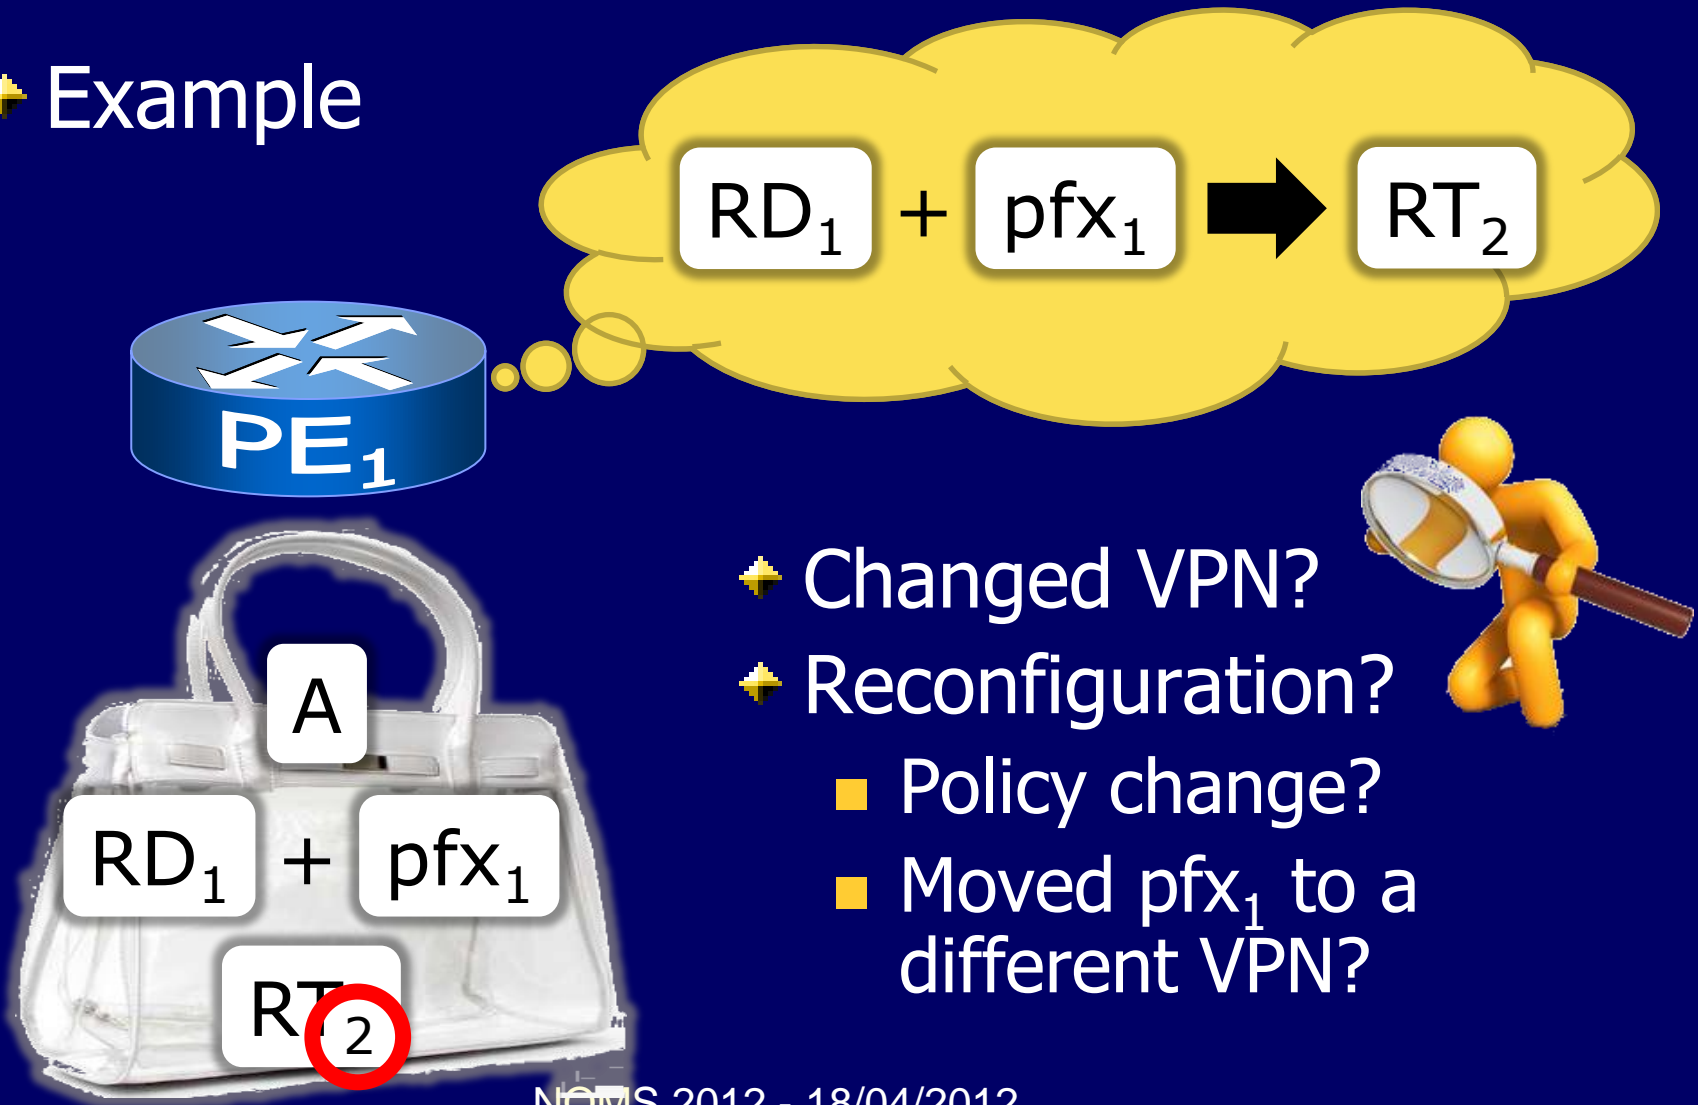




# Methodology

## 2. Reconstruction of VPN state

### ✦ Example



# Methodology

## 2. Reconstruction of VPN state

- ◆ Apply the method to a sequence of BGP updates



# Methodology

## 2. Reconstruction of VPN state

- Apply the method to a sequence of BGP updates




- Reconstruct history of VPN visibility at each PE



# Methodology

## 2. Reconstruction of VPN state


✦ A few difficulties:





# Methodology

## 2. Reconstruction of VPN state

- ✦ A few difficulties:
  - Investigation of the PE where the effect was first observed
  - Dealing with missing attributes in withdrawals
  - Inadmissible announcements [rfc4761]
  - Reannouncements
  - Synchronization with actual VPN states
  - Monitoring RC peering states





# Methodology

## 3. Visualization

# Methodology

## 3. Visualization


**Query:** visibility at each PE of

<b>RD</b>	12345:10011
<b>prefix</b>	172.16.110.0/30
<b>RT</b>	12345:111

# Methodology

## 3. Visualization

**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30  
**RT** 12345:111






# Methodology

## 3. Visualization

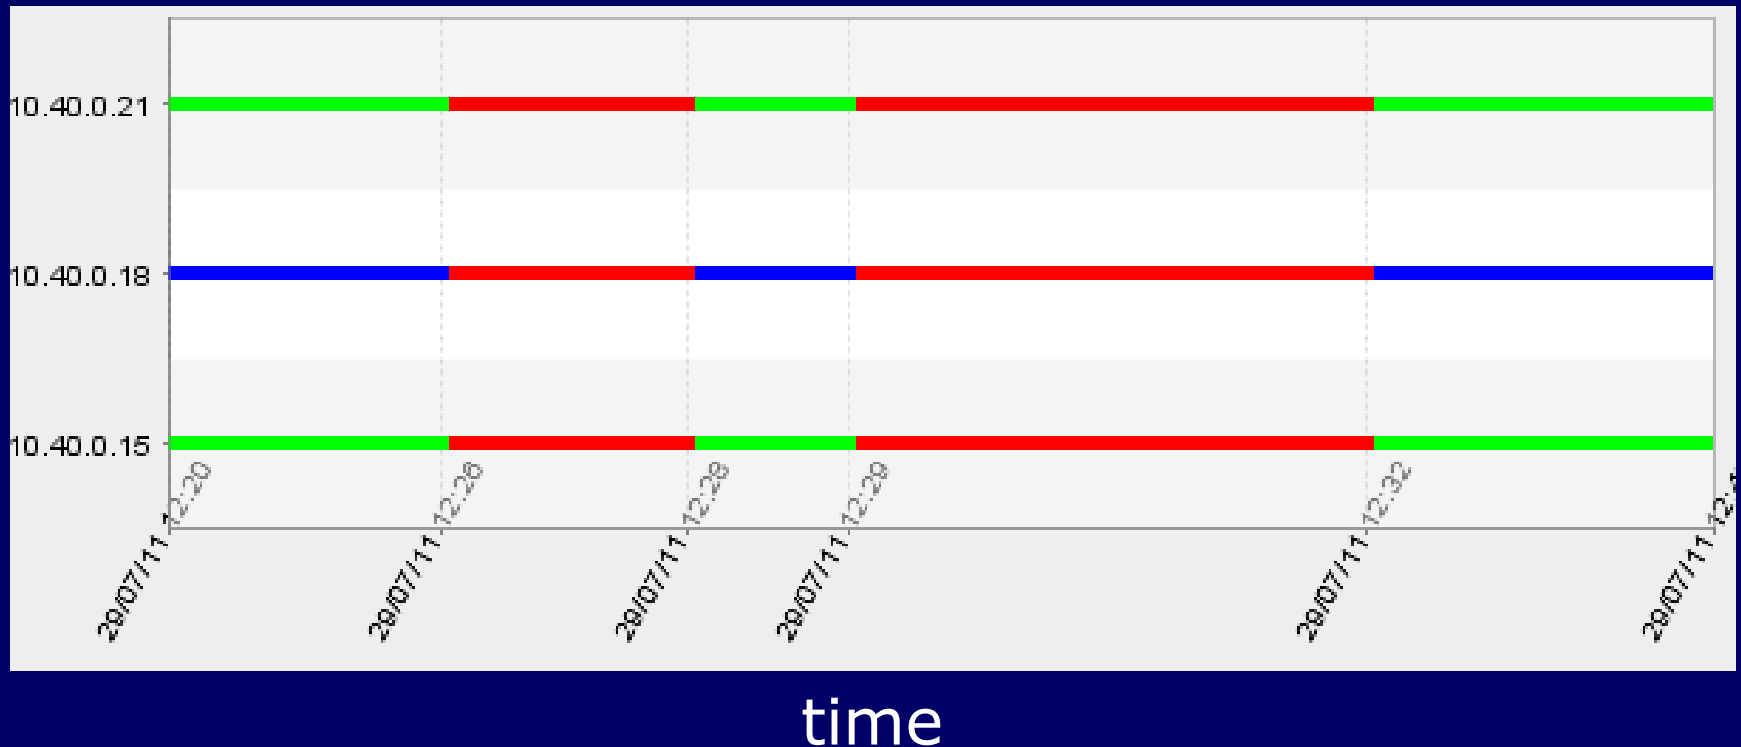
**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30  
**RT** 12345:111



# Methodology

## 3. Visualization

**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30  
**RT** 12345:111

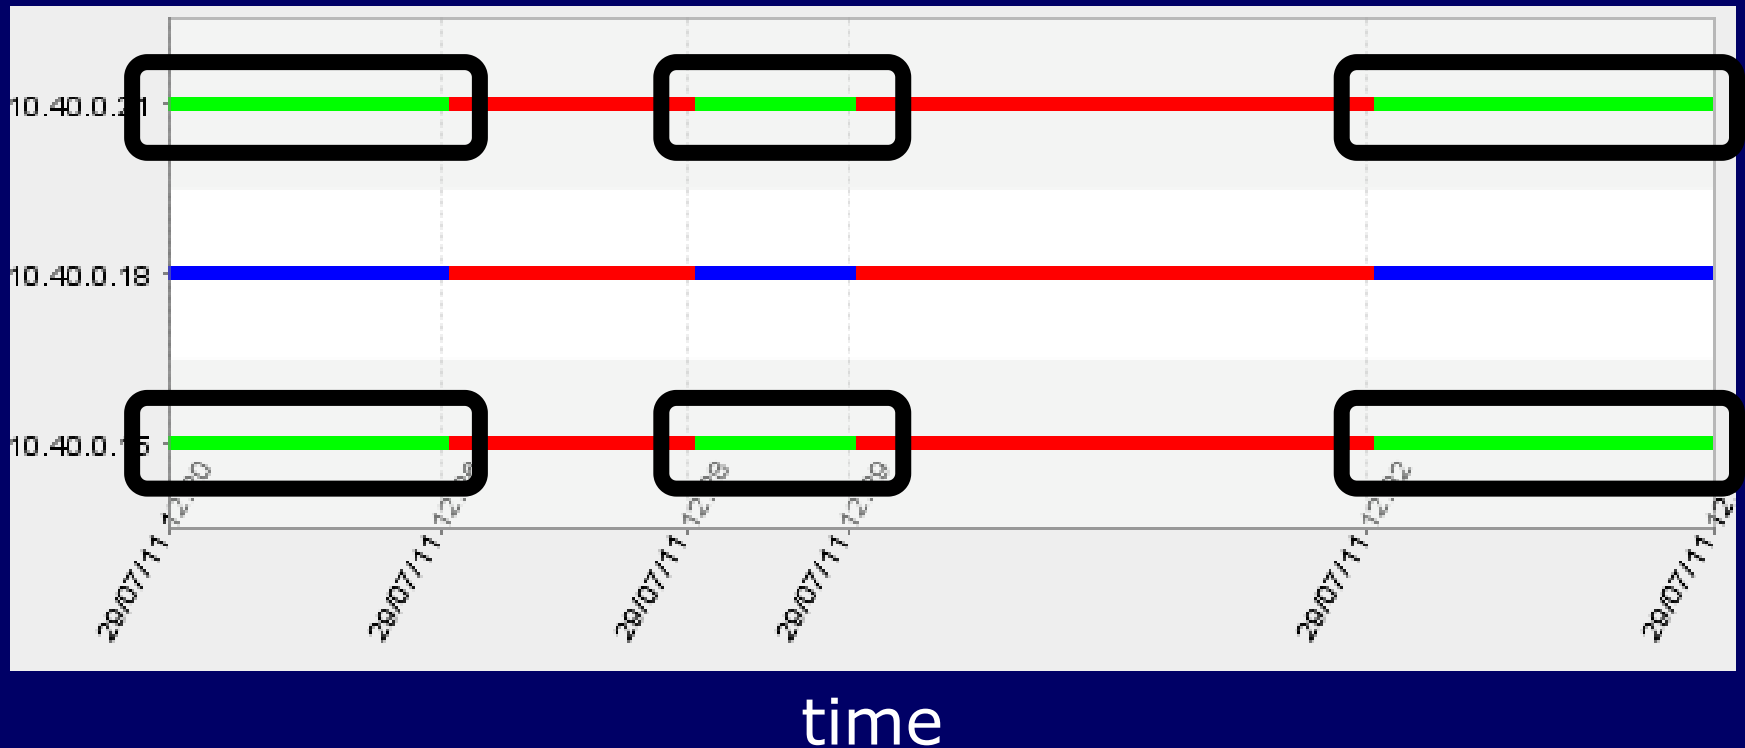


# Methodology

## 3. Visualization

**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30  
**RT** 12345:111

visible




# Methodology

## 3. Visualization

**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30  
**RT** 12345:111

visible

originated



# Methodology

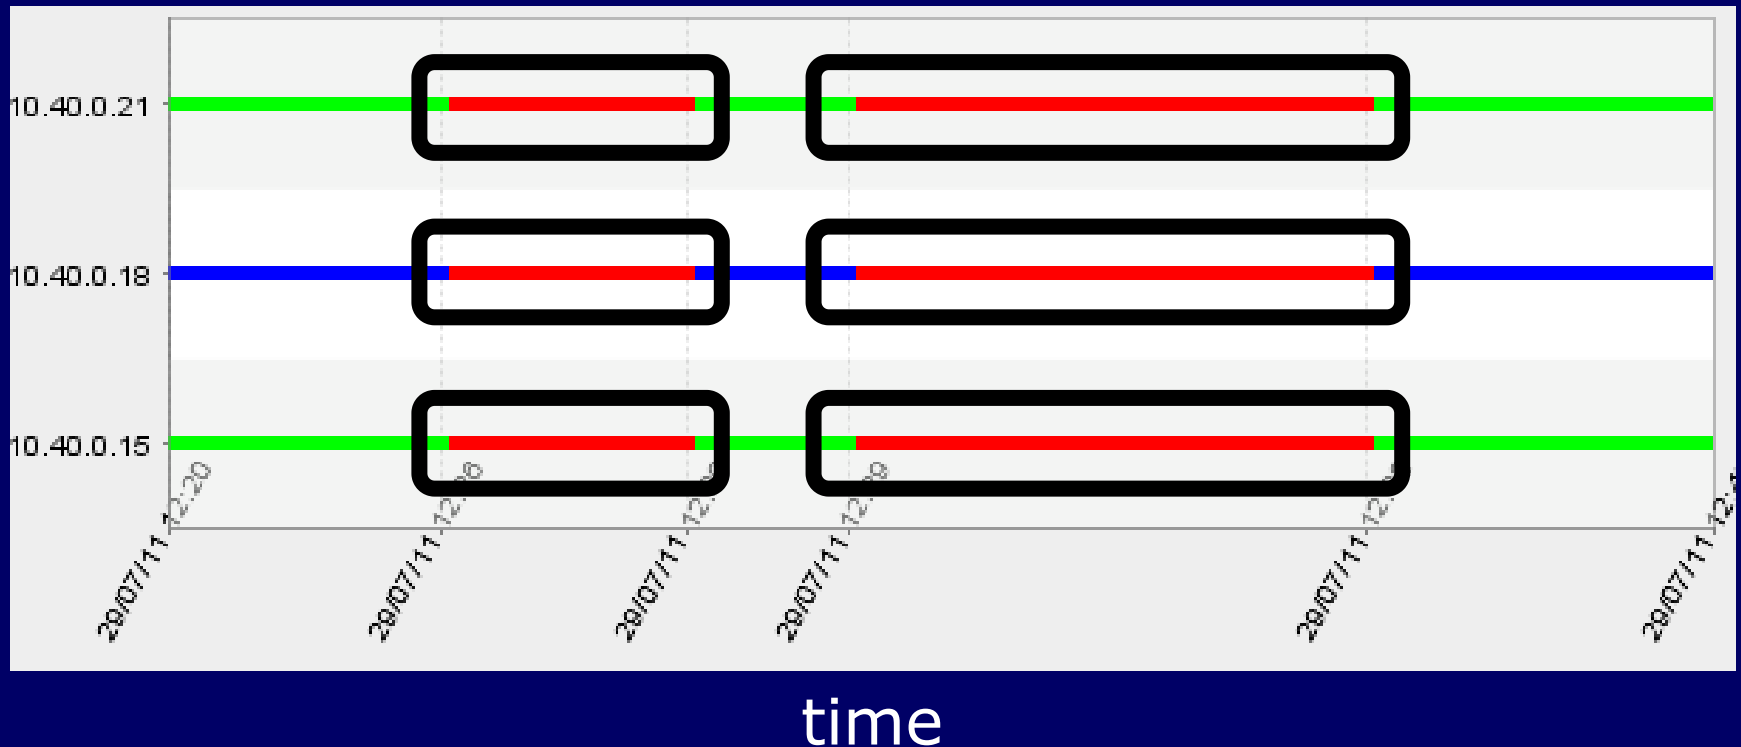
## 3. Visualization

**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30  
**RT** 12345:111

visible

originated

not visible






**Query:** visibility at each PE of

**RD** 12345:10011

**prefix** 172.16.110.0/30

**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111






**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111

...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110 0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111

...with **RT**  
12345:222






**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111


...with **RT**  
12345:222



**Query:** visibility at each PE of  
**RD** 12345:10011  
**prefix** 172.16.110.0/30

...with **RT**  
12345:111

...with **RT**  
12345:222



# Methodology

## 3. Visualization

✦ Queries



# Methodology

## 3. Visualization

### ◆ Queries

- Check information propagation
  - Input:  $RD + \{\text{prefix}, \text{CE ID}\}$ , RT
  - Output: Visibility from all PEs



# Methodology

## 3. Visualization

### Queries

#### ■ Check information propagation

- Input:  $RD+\{\text{prefix,CE ID}\}$ , RT
- Output: Visibility from all PEs

#### ■ Check a PE's visibility of a specific VPN

- Input: RT, PE
- Output: Visibility of all  $RD+\{\text{prefix,CE ID}\}$  with that RT at that PE



# Methodology

## 3. Visualization

### Queries

#### ■ Check information propagation

- Input:  $RD+\{\text{prefix}, \text{CE ID}\}$ , RT
- Output: Visibility from all PEs

#### ■ Check a PE's visibility of a specific VPN

- Input: RT, PE
- Output: Visibility of all  $RD+\{\text{prefix}, \text{CE ID}\}$  with that RT at that PE

#### ■ Highlight belonging of a prefix to a VPN

- Input:  $RD+\{\text{prefix}, \text{CE ID}\}$
- Output: Visibility of that  $RD+\{\text{prefix}, \text{CE ID}\}$  from all PEs, with each seen RT



# Methodology

## 3. Visualization

### Queries

#### ■ Check information propagation

- Input:  $RD+\{\text{prefix}, \text{CE ID}\}$ , RT
- Output: Visibility from all PEs

#### ■ Check a PE's visibility of a specific VPN

- Input: RT, PE
- Output: Visibility of all  $RD+\{\text{prefix}, \text{CE ID}\}$  with that RT at that PE

#### ■ Highlight belonging of a prefix to a VPN

- Input:  $RD+\{\text{prefix}, \text{CE ID}\}$
- Output: Visibility of that  $RD+\{\text{prefix}, \text{CE ID}\}$  from all PEs, with each seen RT

#### ■ Highlight participation of PEs in VPNs

- Input: RT
- Output: Visibility of that RT at each PE





# Methodology

## 3. Visualization

### Queries

- Check information propagation
  - Input:  $RD + \{\text{prefix}, \text{CE ID}\}$ , RT
  - Output: Visibility from all PEs
- Check a PE's visibility of a specific VPN\*
  - Input: RT, PE
  - Output: Visibility of all  $RD + \{\text{prefix}, \text{CE ID}\}$  with that RT at that PE
- Highlight belonging of a prefix to a VPN\*
  - Input:  $RD + \{\text{prefix}, \text{CE ID}\}$
  - Output: Visibility of that  $RD + \{\text{prefix}, \text{CE ID}\}$  from all PEs, with each seen RT
- Highlight participation of PEs in VPNs\*
  - Input: RT
  - Output: Visibility of that RT at each PE



# Methodology

## 3. Visualization

### Queries

#### ■ Check information propagation

- Input:  $RD + \{\text{prefix}, \text{CE ID}\}$ , RT
- Output: Visibility from all PEs

#### ■ Check a PE's visibility of a specific VPN\*

- Input: RT, PE
- Output: Visibility of all  $RD + \{\text{prefix}, \text{CE ID}\}$  with that RT at that PE

#### ■ Highlight belonging of a prefix to a VPN\*

- Input:  $RD + \{\text{prefix}, \text{CE ID}\}$
- Output: Visibility of that  $RD + \{\text{prefix}, \text{CE ID}\}$  from all PEs, with each seen RT

#### ■ Highlight participation of PEs in VPNs\*

- Input: RT
- Output: Visibility of that RT at each PE

\*  $VPN \equiv RT$



# Methodology

## 3. Visualization

### ◆ Queries

#### ■ Check information propagation

- Input:  $RD + \{\text{prefix}, \text{CE ID}\}$ , RT
- Output: Visibility<sup>#</sup> from all PEs

#### ■ Check a PE's visibility of a specific VPN\*

- Input: RT, PE
- Output: Visibility<sup>#</sup> of all  $RD + \{\text{prefix}, \text{CE ID}\}$  with that RT at that PE

#### ■ Highlight belonging of a prefix to a VPN\*

- Input:  $RD + \{\text{prefix}, \text{CE ID}\}$
- Output: Visibility<sup>#</sup> of that  $RD + \{\text{prefix}, \text{CE ID}\}$  from all PEs, with each seen RT

#### ■ Highlight participation of PEs in VPNs\*

- Input: RT
- Output: Visibility<sup>#</sup> of that RT at each PE

\*  $VPN \equiv RT$



# Methodology

## 3. Visualization

### Queries

#### ■ Check information propagation

- Input: RD+{prefix,CE ID}, RT
- Output: Visibility<sup>#</sup> from all PEs

#### ■ Check a PE's visibility of a specific VPN<sup>\*</sup>

- Input: RT, PE
- Output: Visibility<sup>#</sup> of all RD+{prefix,CE ID} with that RT at that PE

#### ■ Highlight belonging of a prefix to a VPN<sup>\*</sup>

- Input: RD+{prefix,CE ID}
- Output: Visibility<sup>#</sup> of that RD+{prefix,CE ID} from all PEs, with each seen RT

#### ■ Highlight participation of PEs in VPNs<sup>\*</sup>

- Input: RT
- Output: Visibility<sup>#</sup> of that RT at each PE

\* VPN $\equiv$ RT  
# over time



# Scalability

◆ Amount of routing updates



# Scalability



- ✦ Routing table size

- ✦ Amount of routing updates

# Scalability

- ✦ Routing table size

- ✦ Amount of routing updates

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

## ✦ Amount of routing updates



# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

😊 Only routing updates count

## ✦ Amount of routing updates

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

😊 Only routing updates count

😊 Same scalability of [ORV], [BGPlay], [iBGPlay]

## ✦ Amount of routing updates

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

😊 Only routing updates count

😊 Same scalability of [ORV], [BGPlay], [iBGPlay]

## ✦ Amount of routing updates

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

😊 Only routing updates count

😊 Same scalability of [ORV], [BGPlay], [iBGPlay]

## ✦ Amount of routing updates

☹ Lots of customers, prefixes, VPNs, etc.

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

😊 Only routing updates count

😊 Same scalability of [ORV], [BGPlay], [iBGPlay]

## ✦ Amount of routing updates

☹ Lots of customers, prefixes, VPNs, etc.

😊 Bursts (due to, e.g., configurations changes, faults) are unlikely

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

☺ Only routing updates count

☺ Same scalability of [ORV], [BGPlay], [iBGPlay]

## ✦ Amount of routing updates

☹ Lots of customers, prefixes, VPNs, etc.

☺ Bursts (due to, e.g., configurations changes, faults) are unlikely

☺ 2-3 orders of magnitude less than VPN routes  
[Ben-Houidi et al. 07]

# Scalability

## ✦ Routing table size

☹  $\gg$  #Internet prefixes:  $\sim k \times 10^5$   
[Ben-Houidi et al. 07]

😊 Only routing updates count

😊 Same scalability of [ORV], [BGPlay], [iBGPlay]

## ✦ Amount of routing updates


☹ Lots of customers, prefixes, VPNs, etc.

😊 Bursts (due to, e.g., configurations changes, faults) are unlikely

😊 2-3 orders of magnitude less than VPN routes  
[Ben-Houidi et al. 07]


😊 Our prototype works even for M/L ISPs

# Scalability vs Visibility







# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility



# Scalability vs Visibility




# Scalability vs Visibility

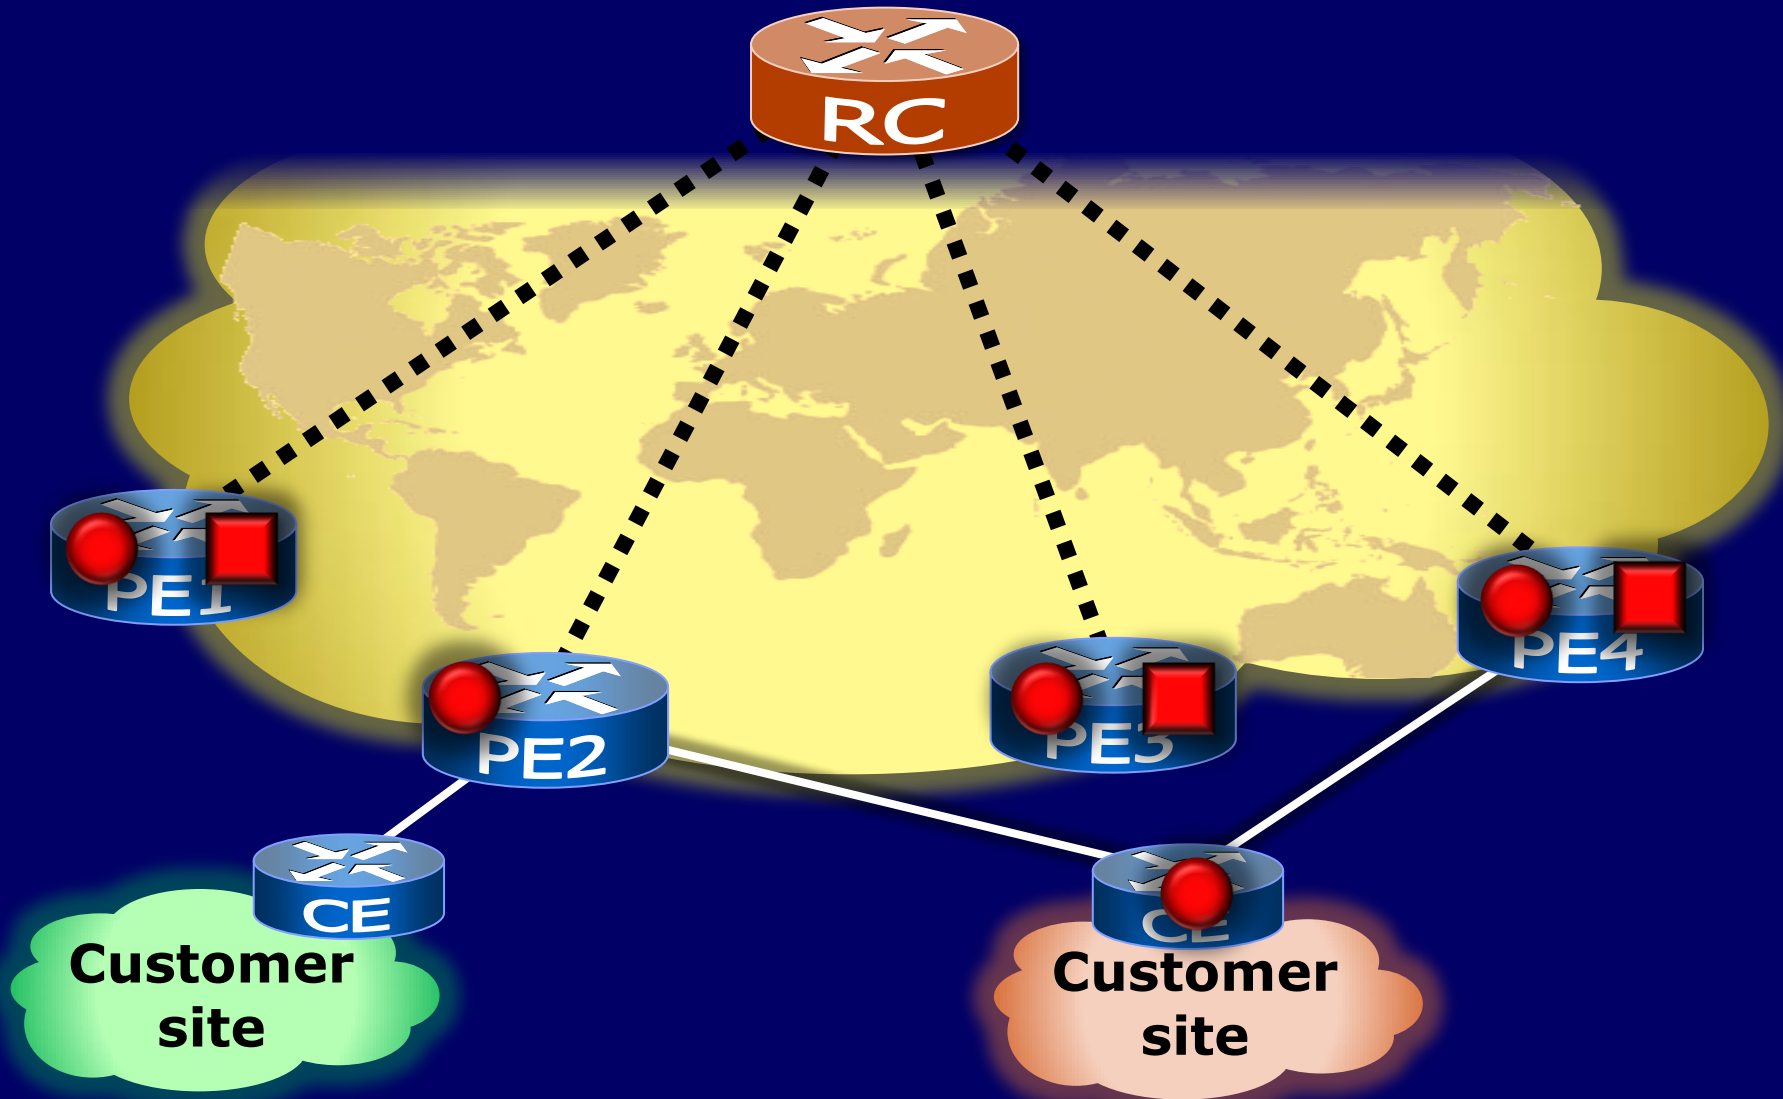





# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility



# Scalability vs Visibility




# Scalability vs Visibility

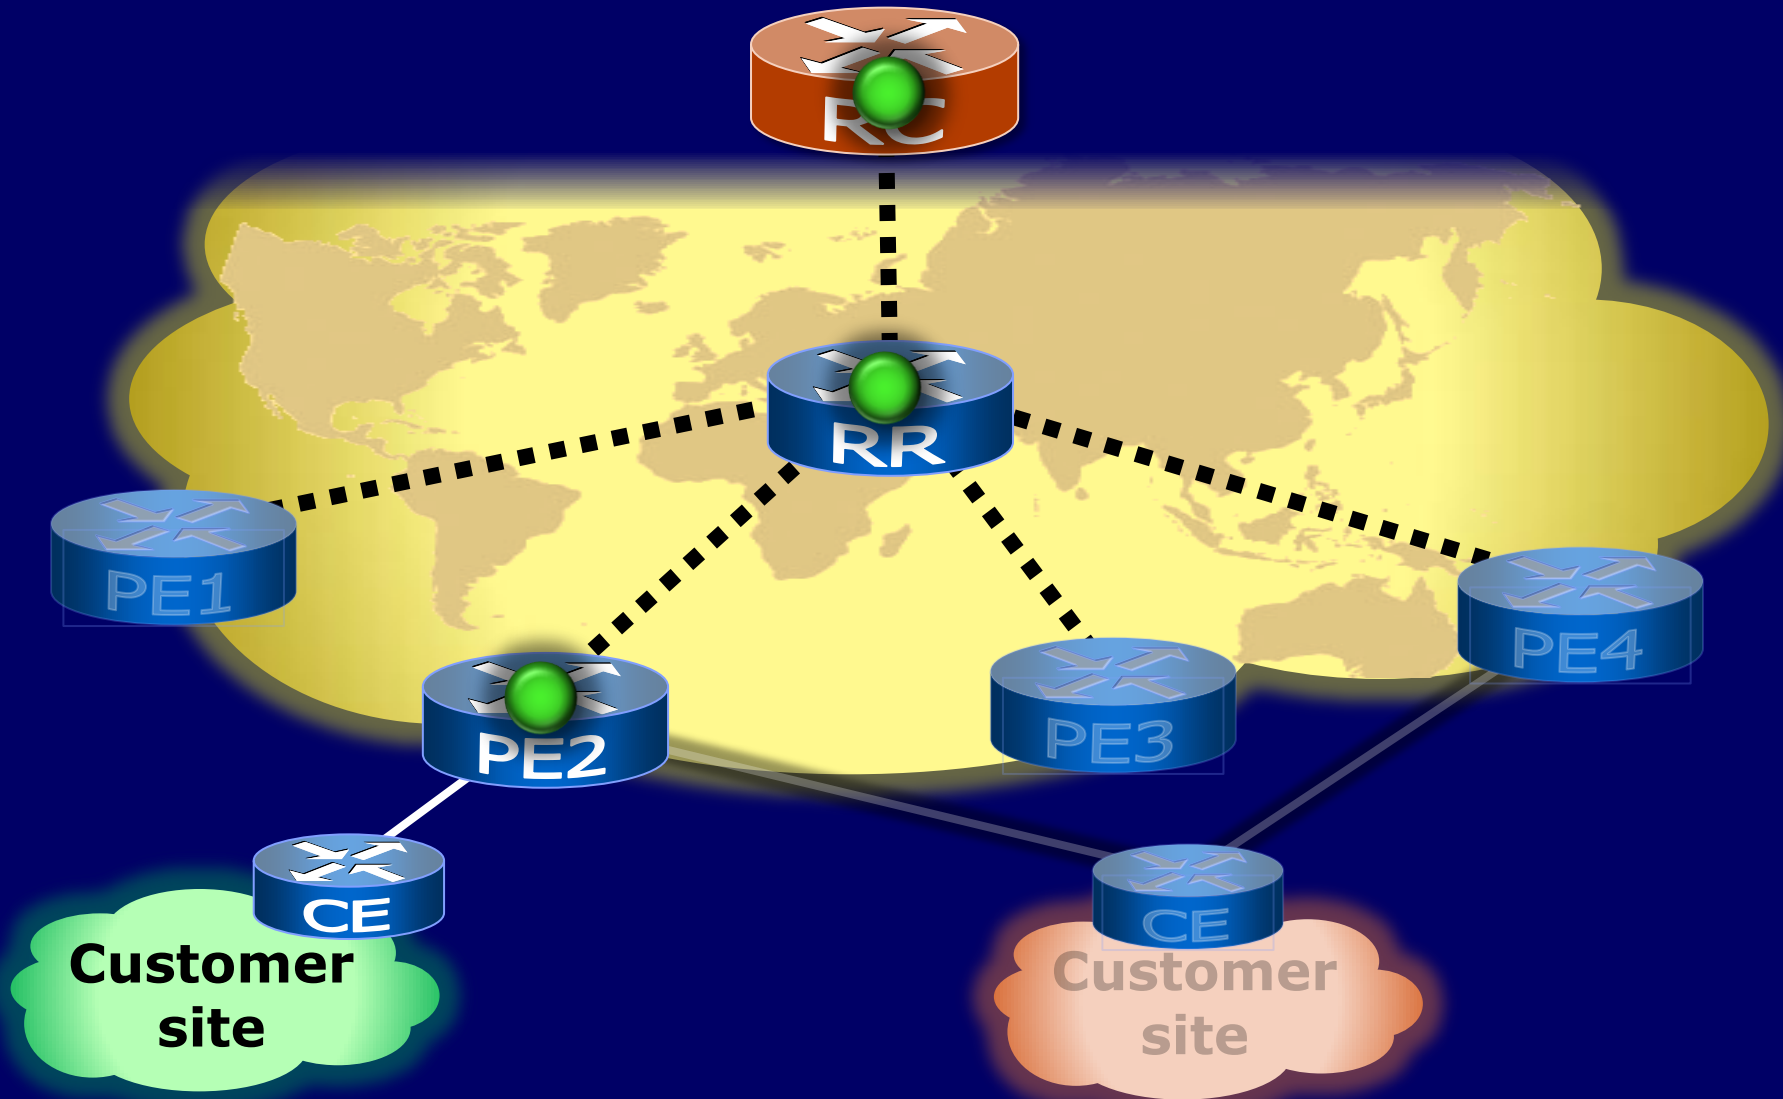





# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility




# Scalability vs Visibility



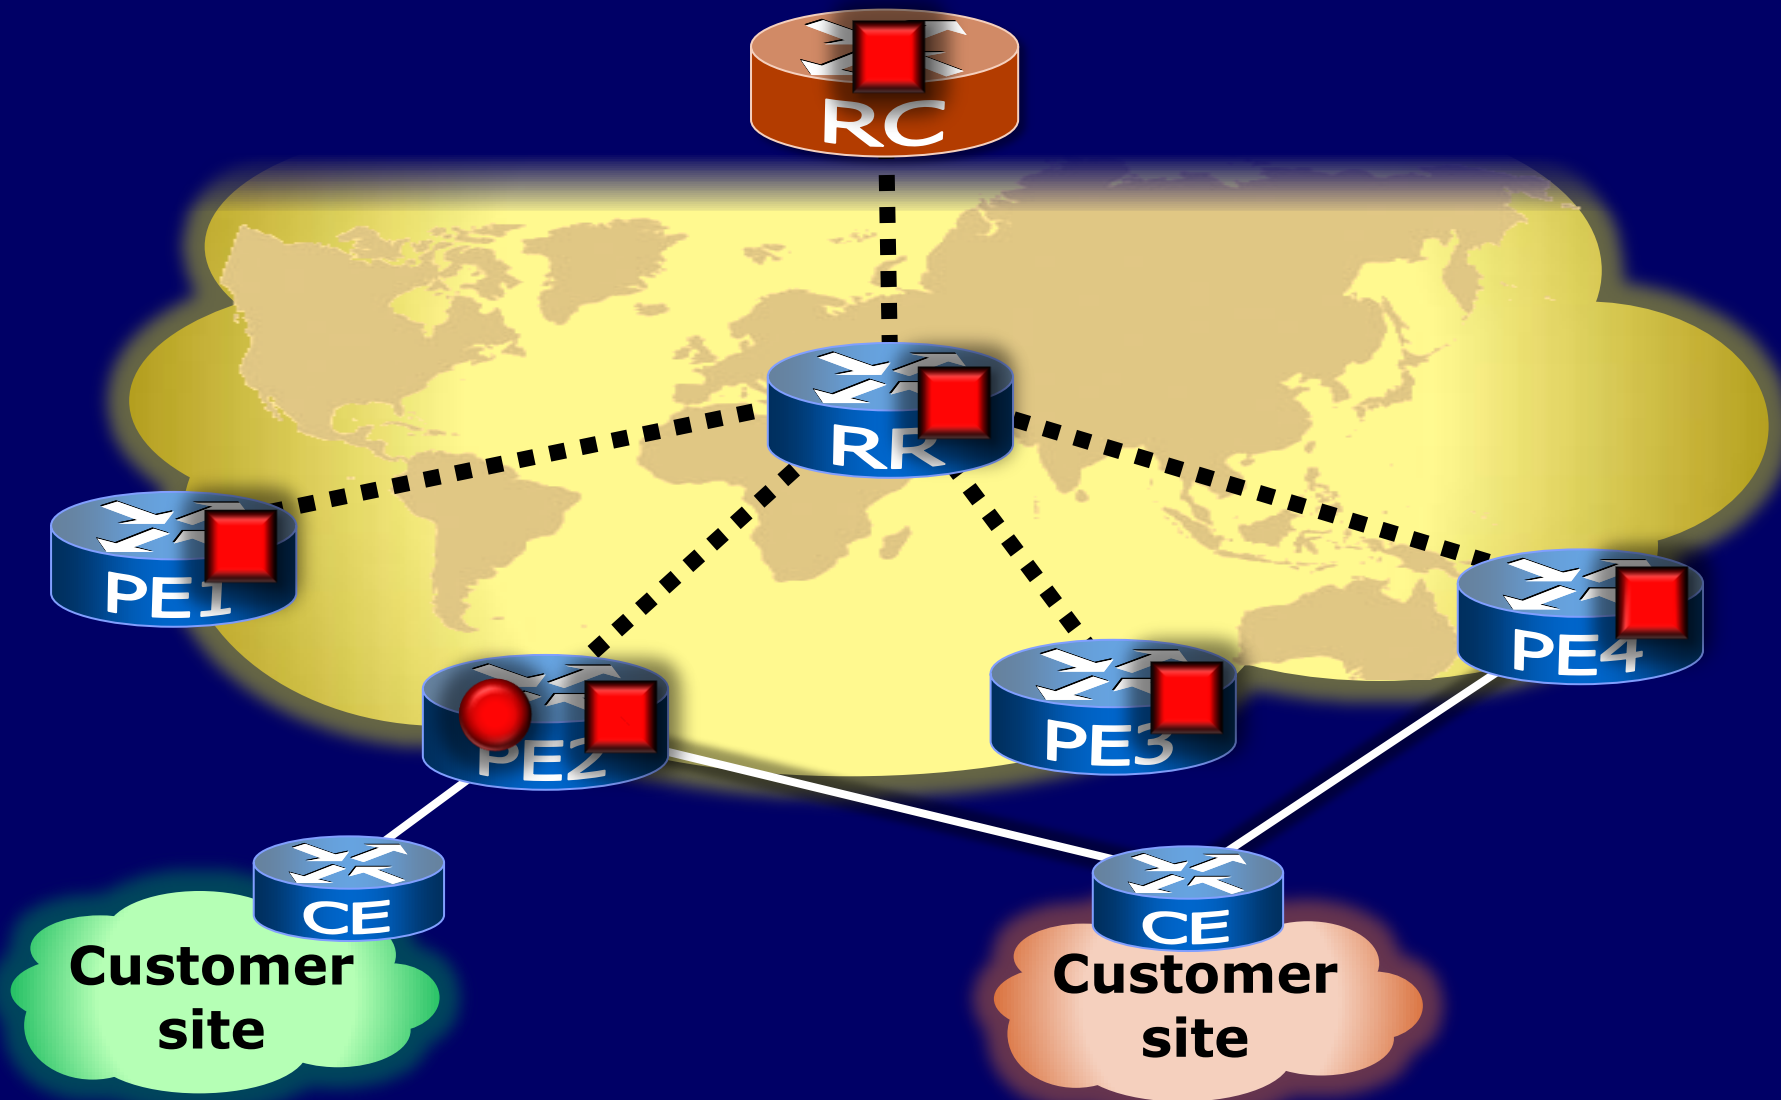
# Scalability vs Visibility



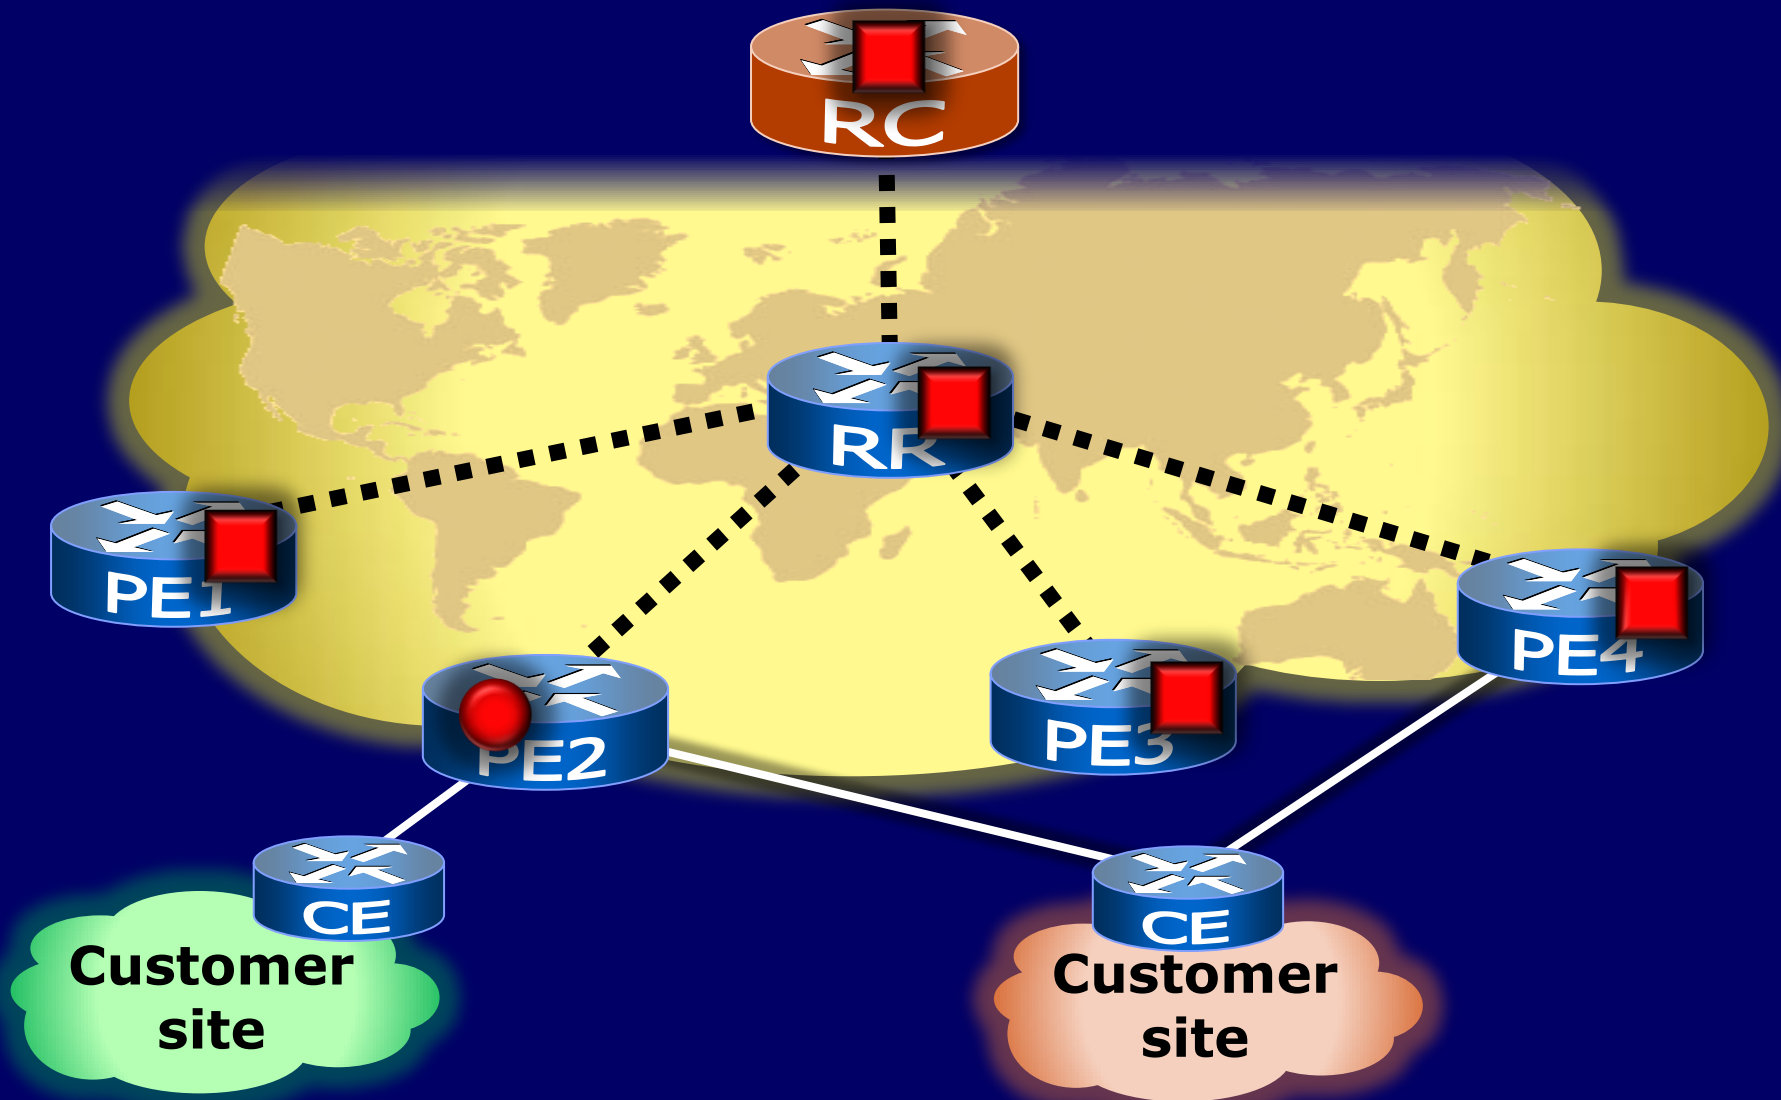
# Scalability vs Visibility



# Scalability vs Visibility

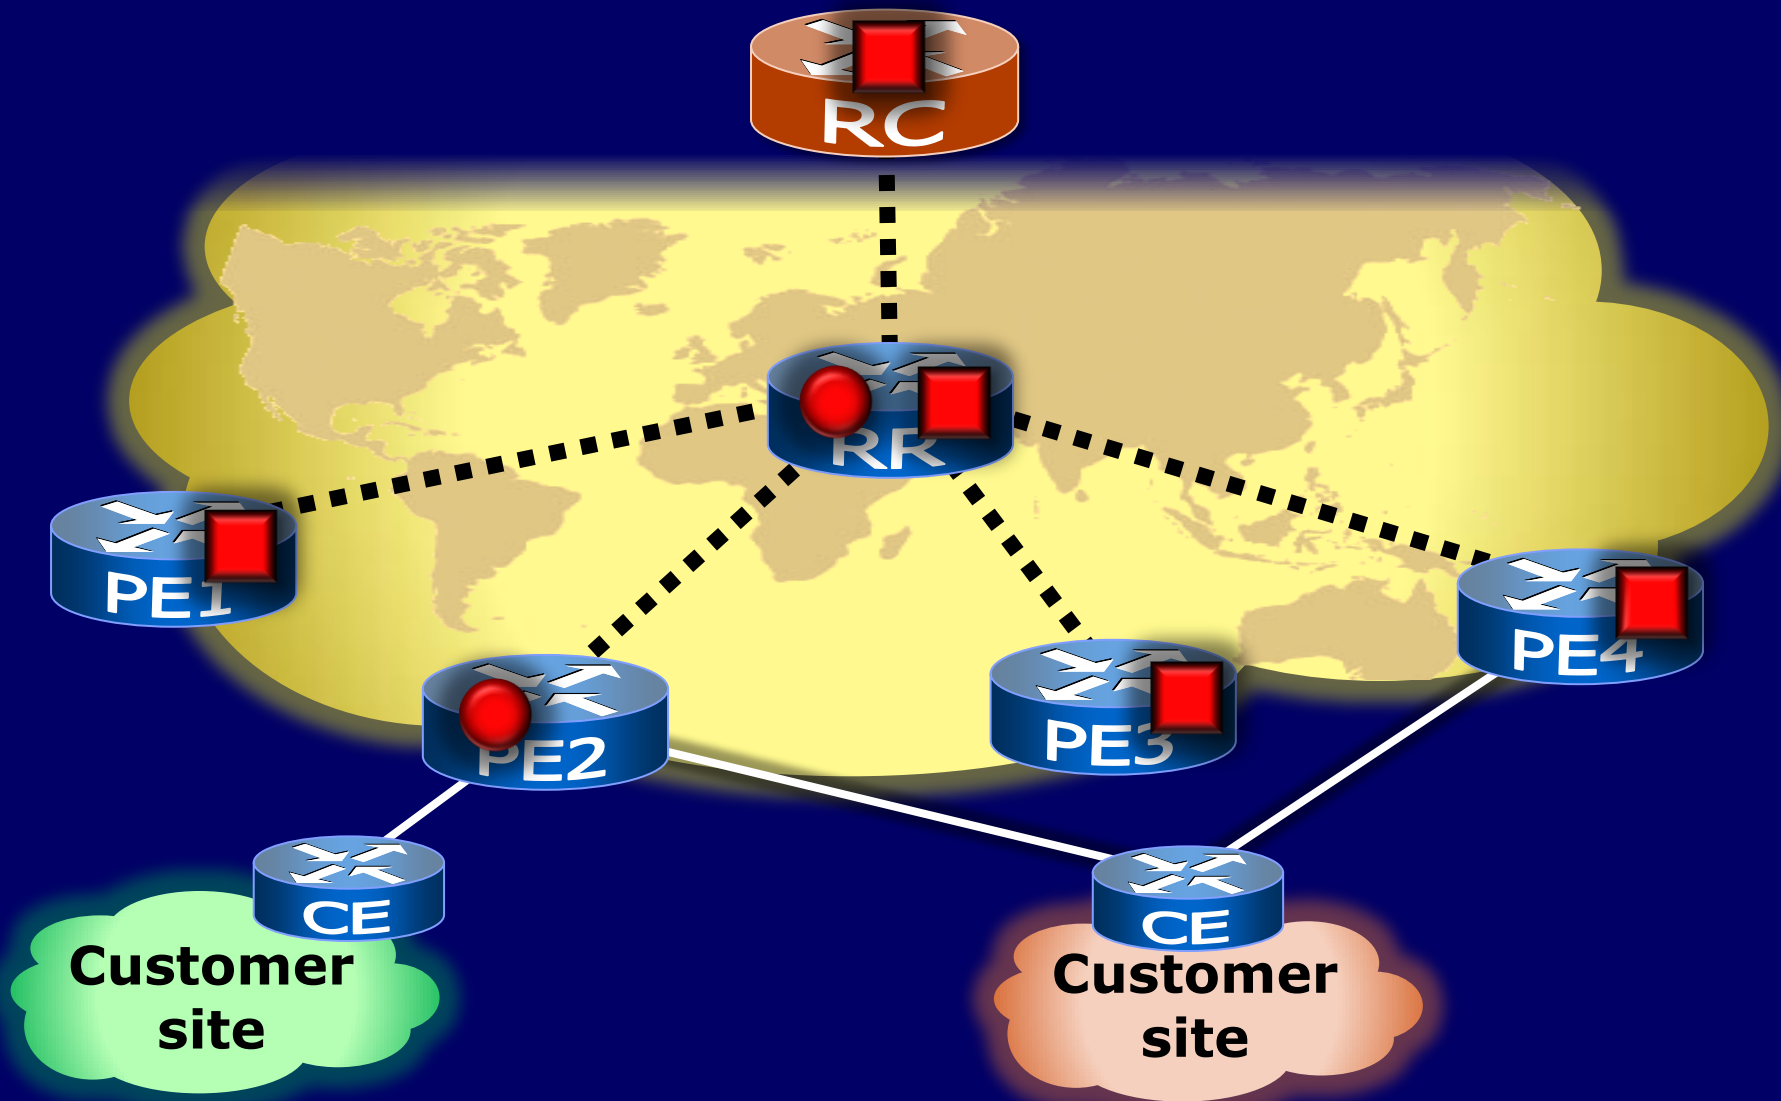


# Scalability vs Visibility

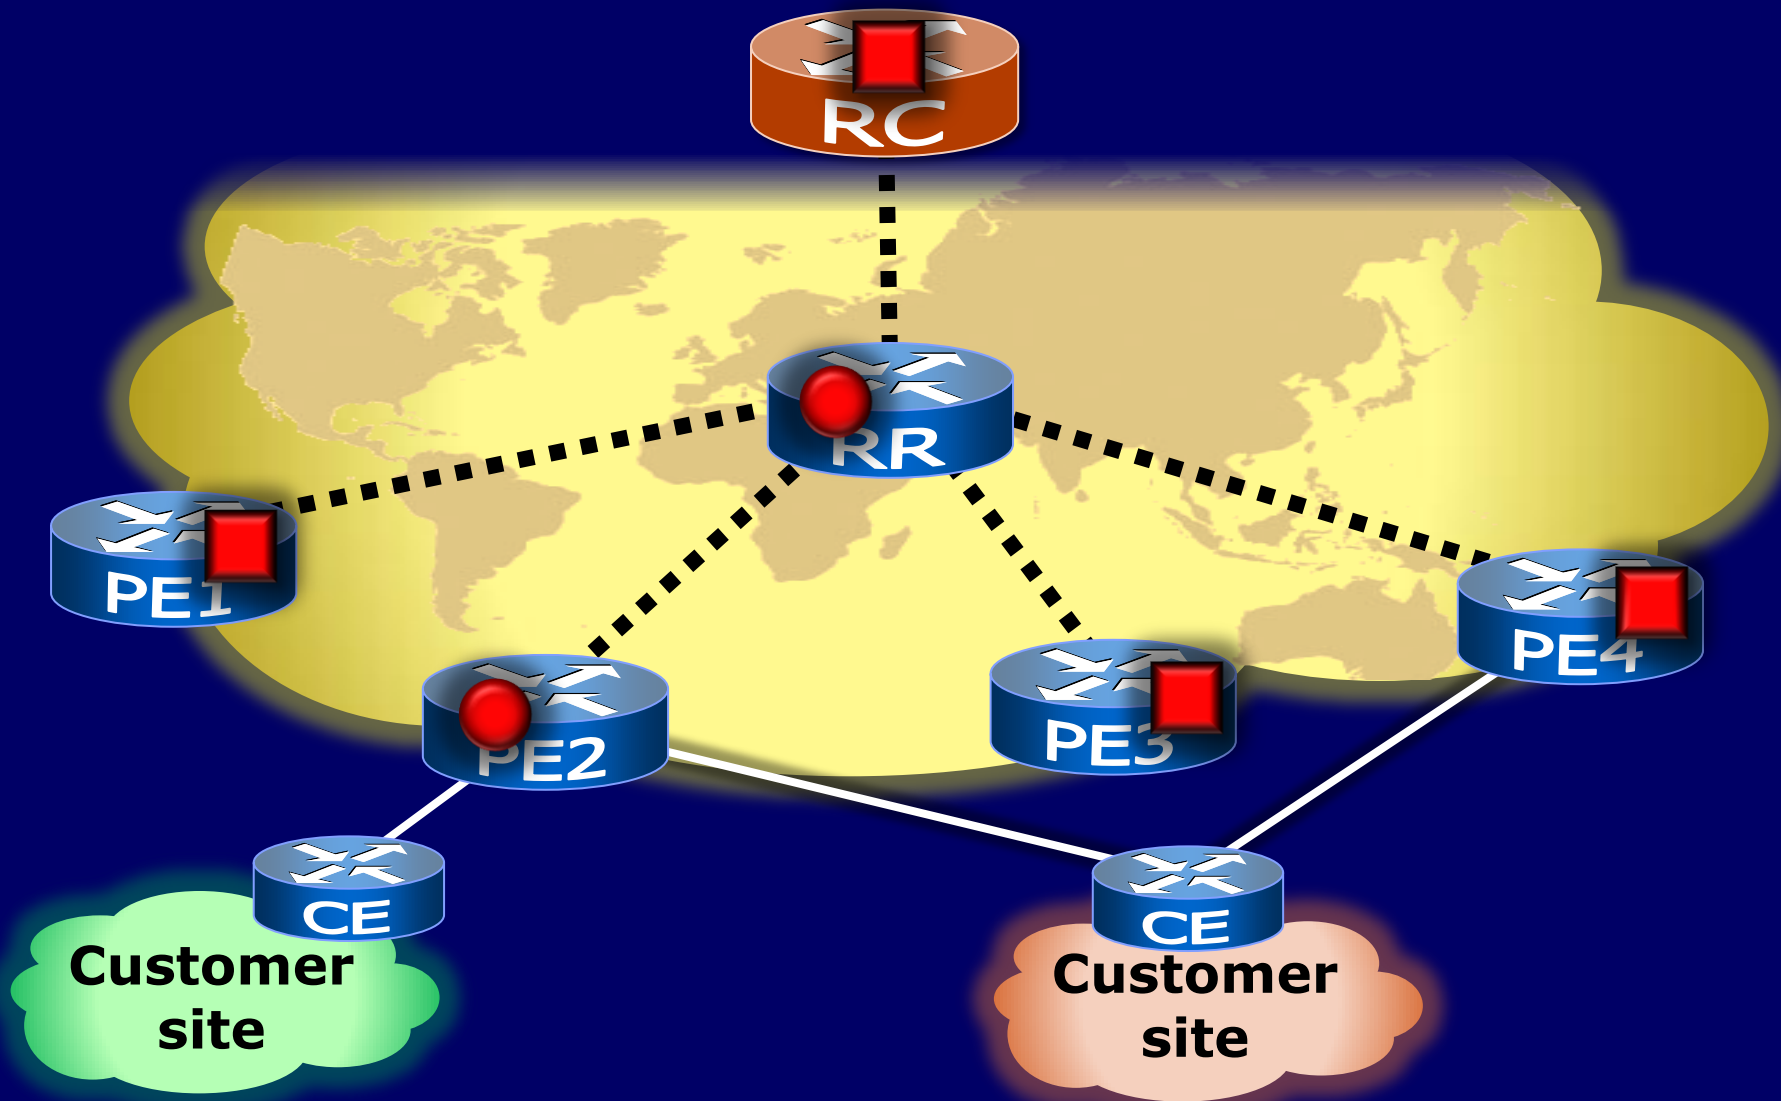




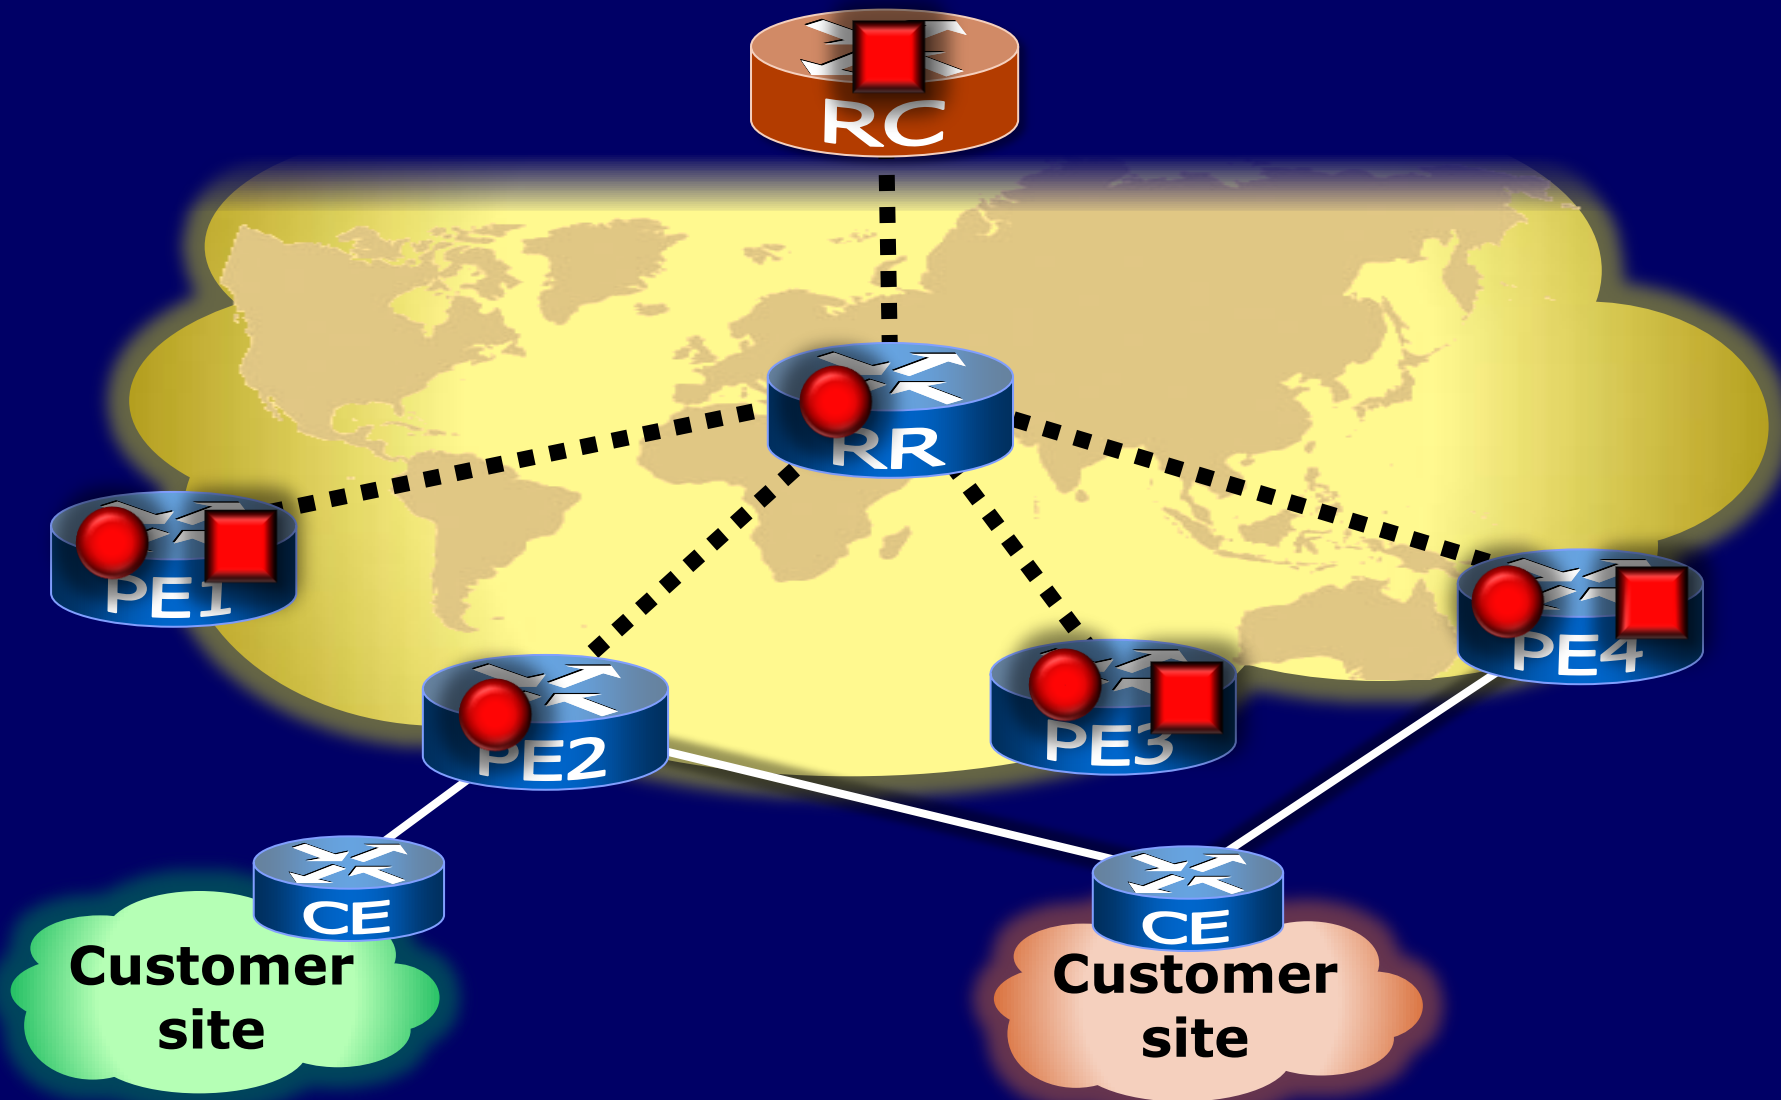
# Scalability vs Visibility



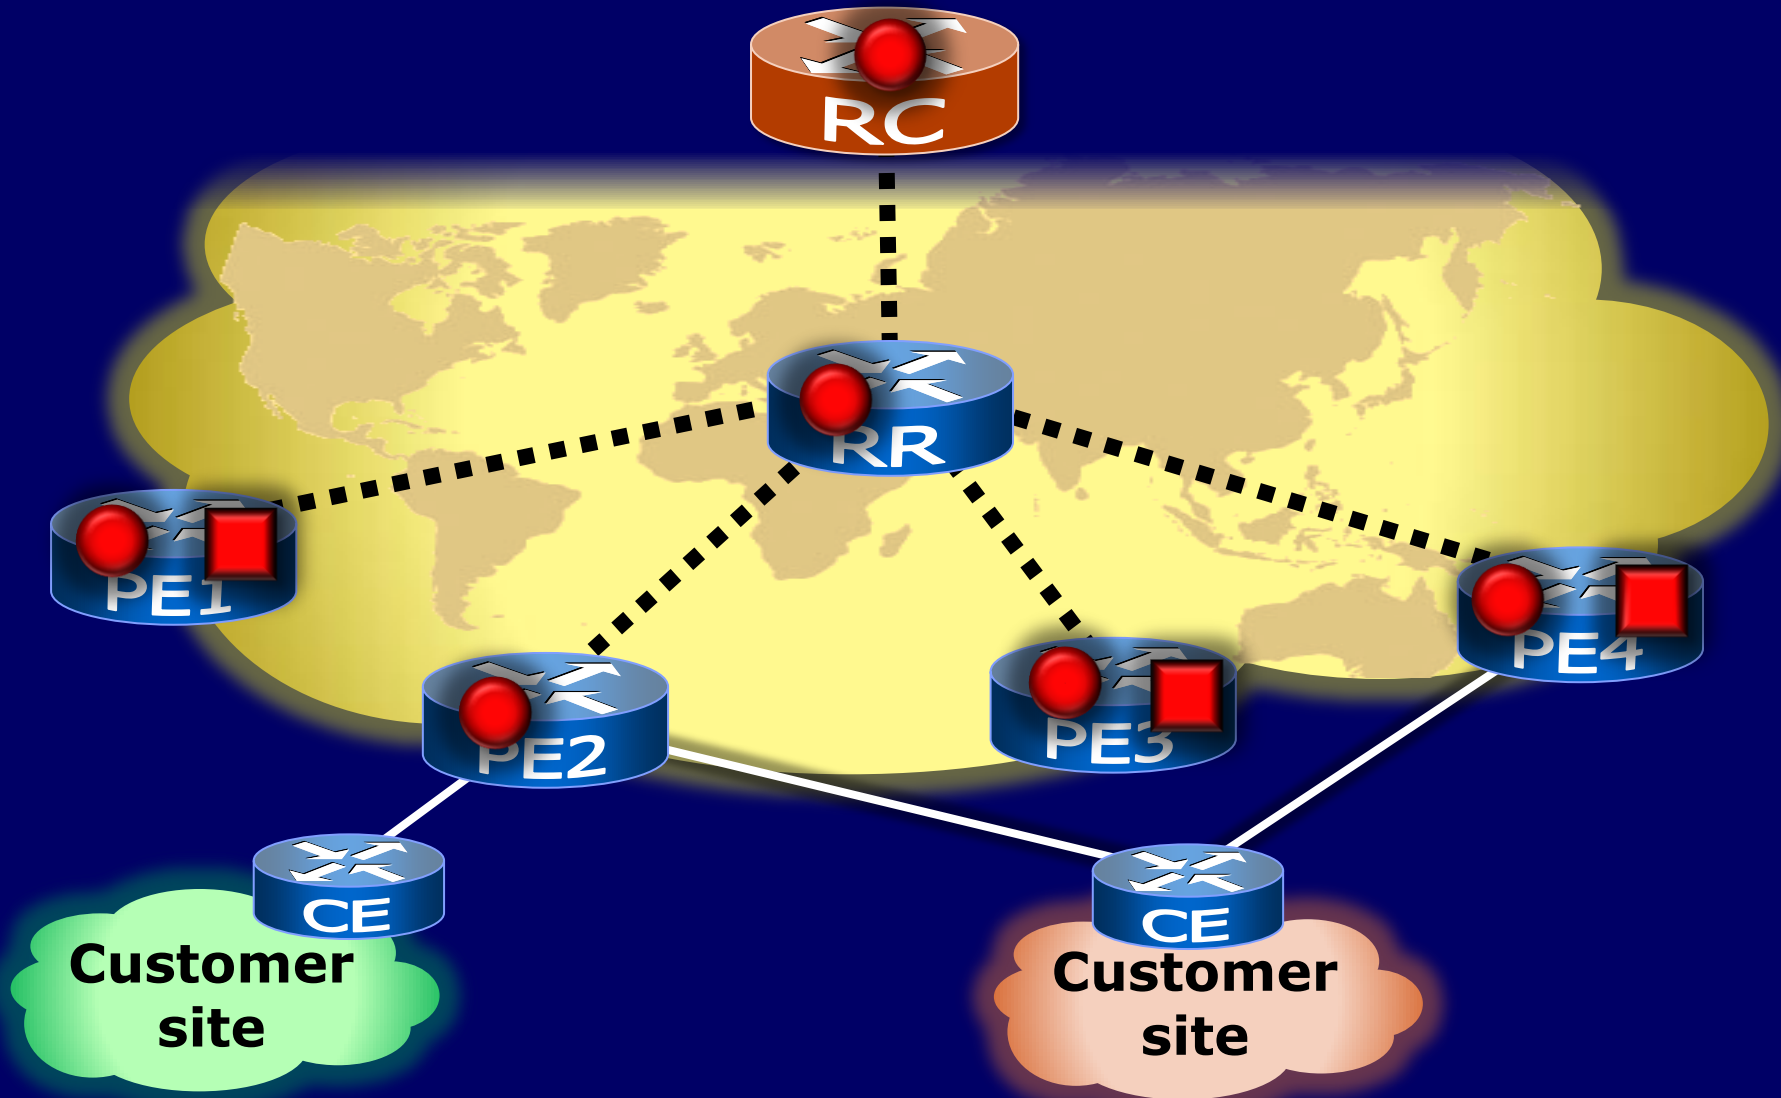
# Scalability vs Visibility



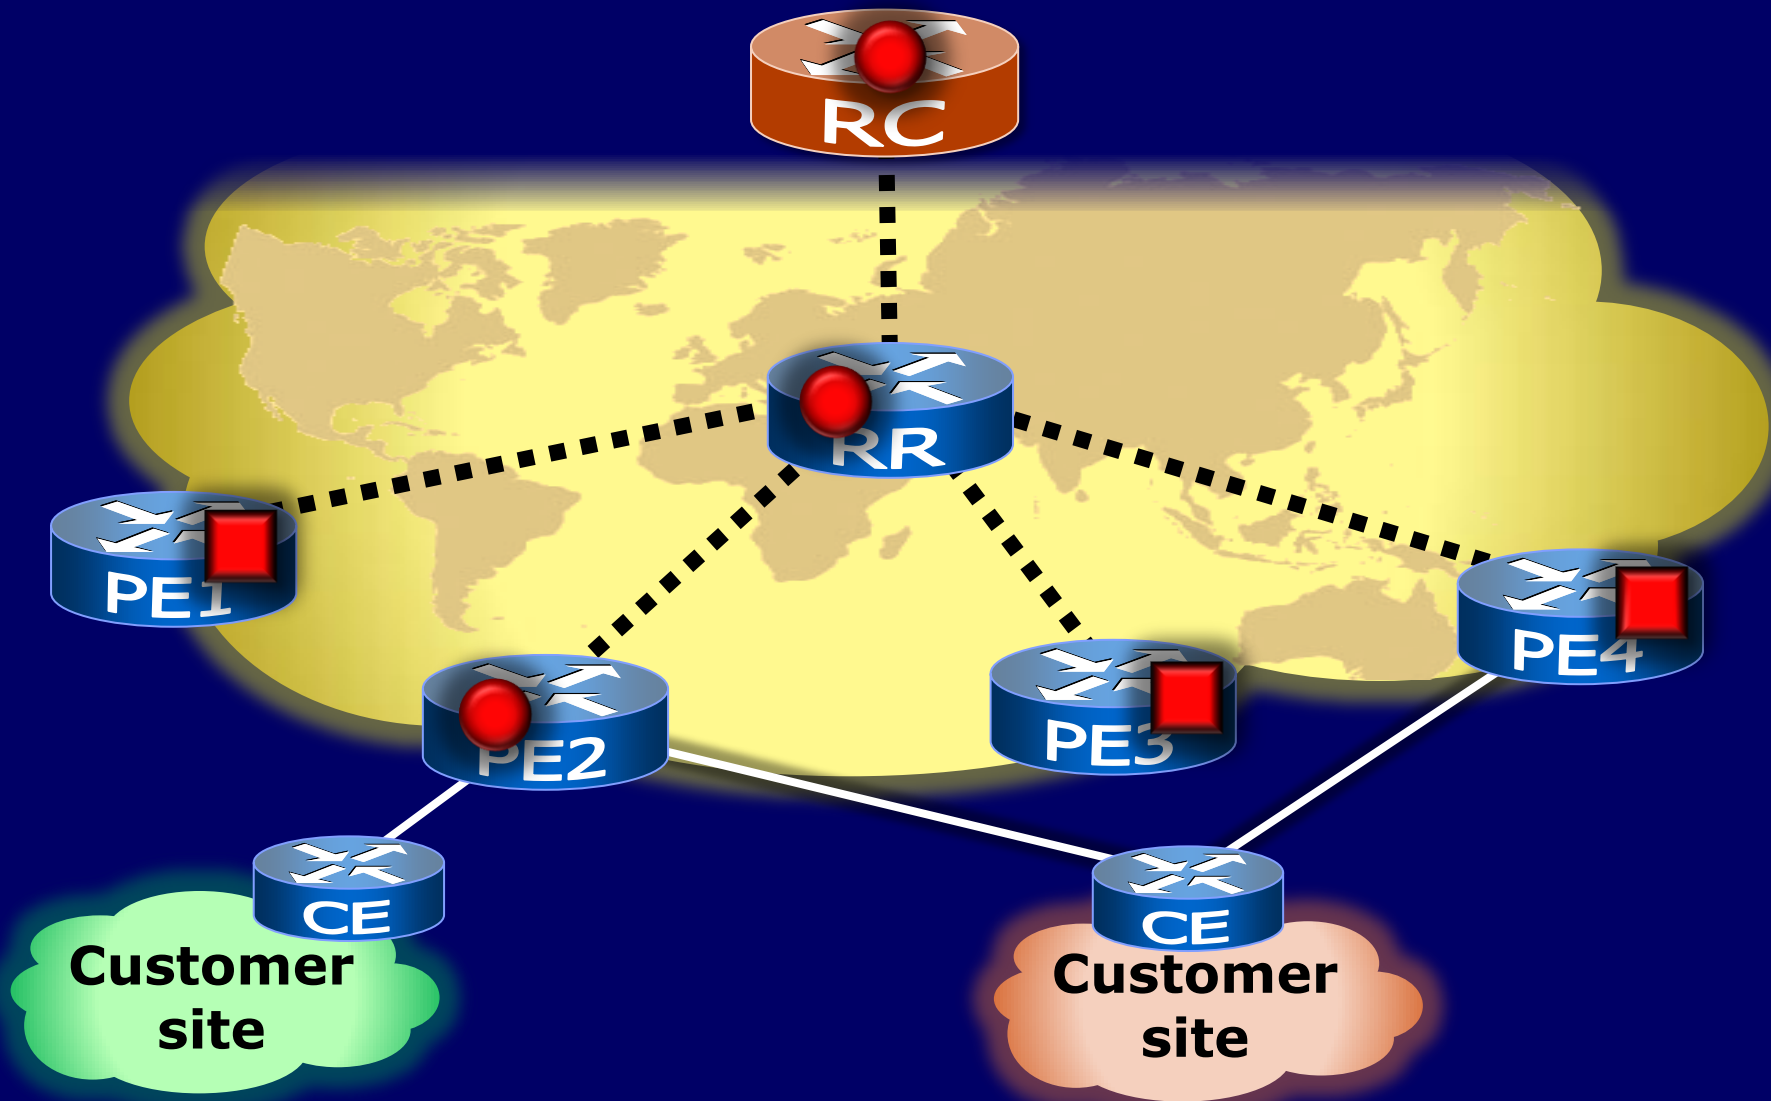
# Scalability vs Visibility



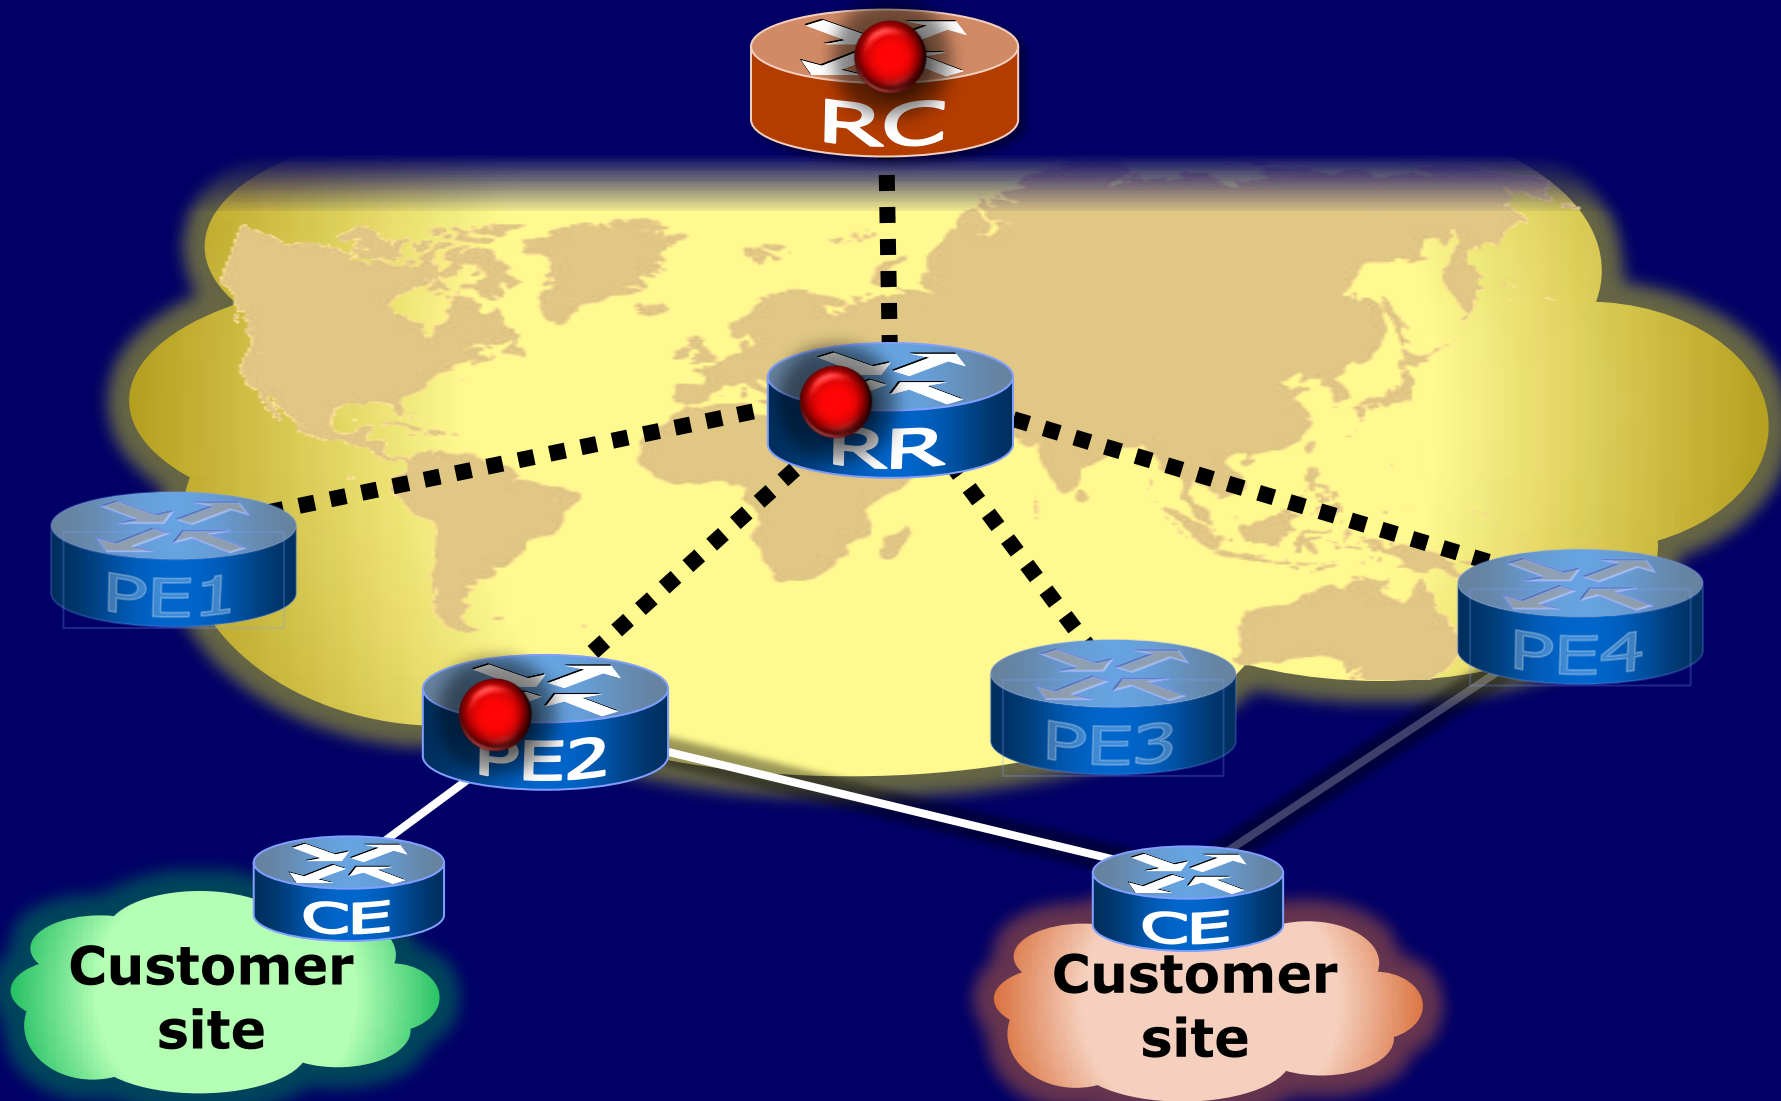
# Scalability vs Visibility




# Scalability vs Visibility



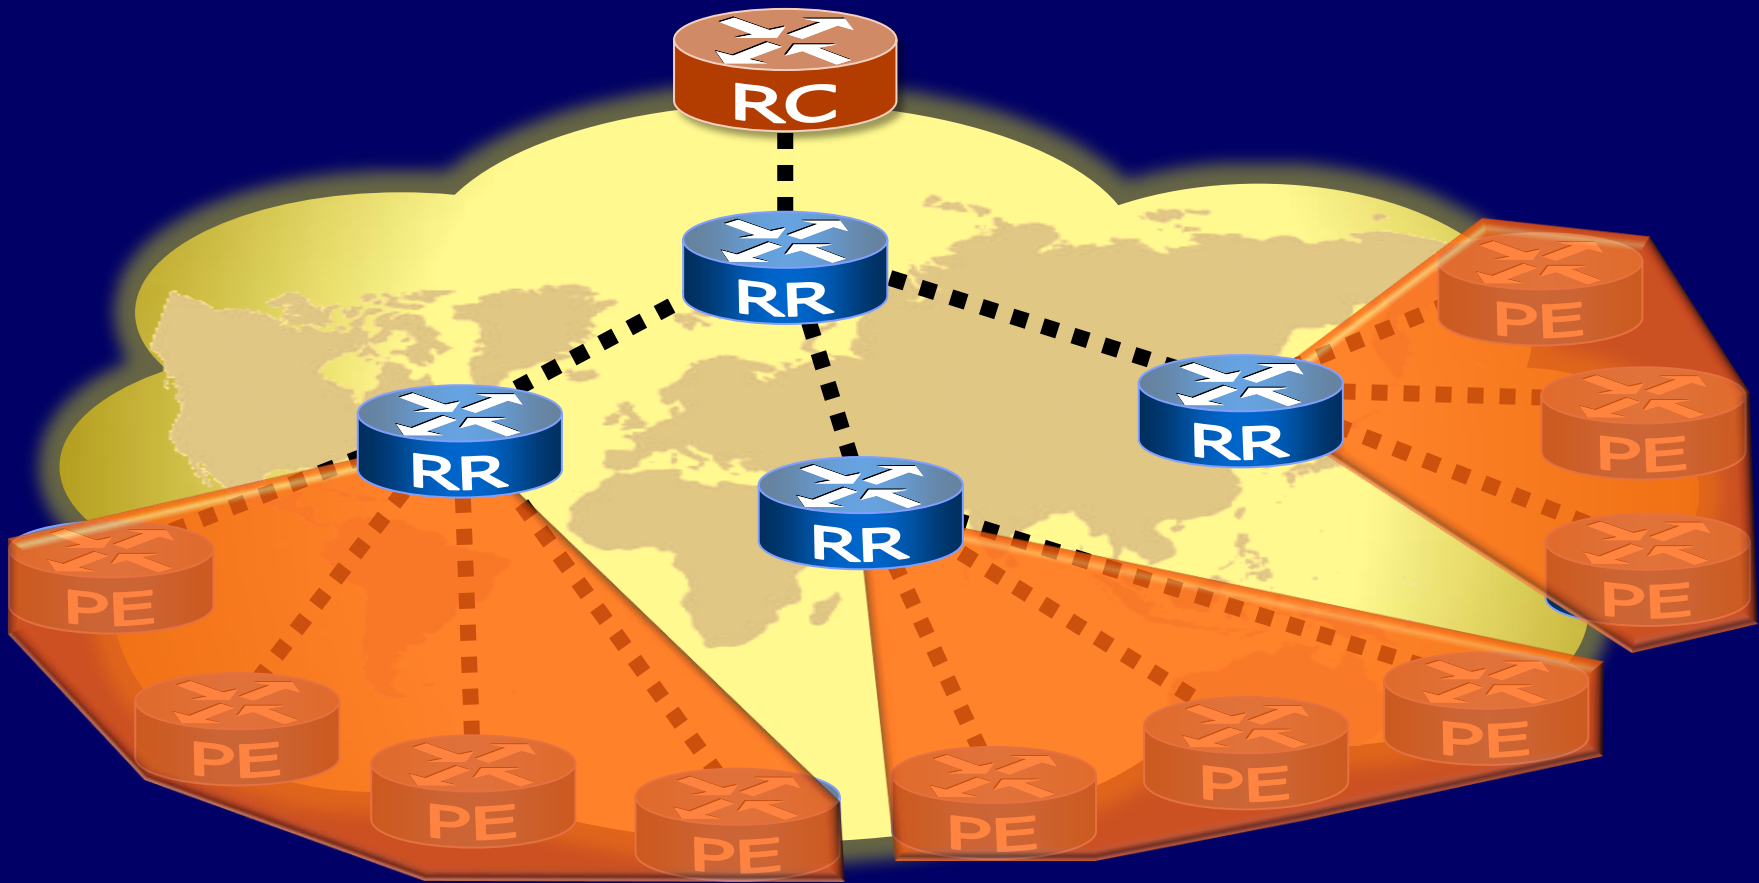
# Scalability vs Visibility



# Scalability vs Visibility

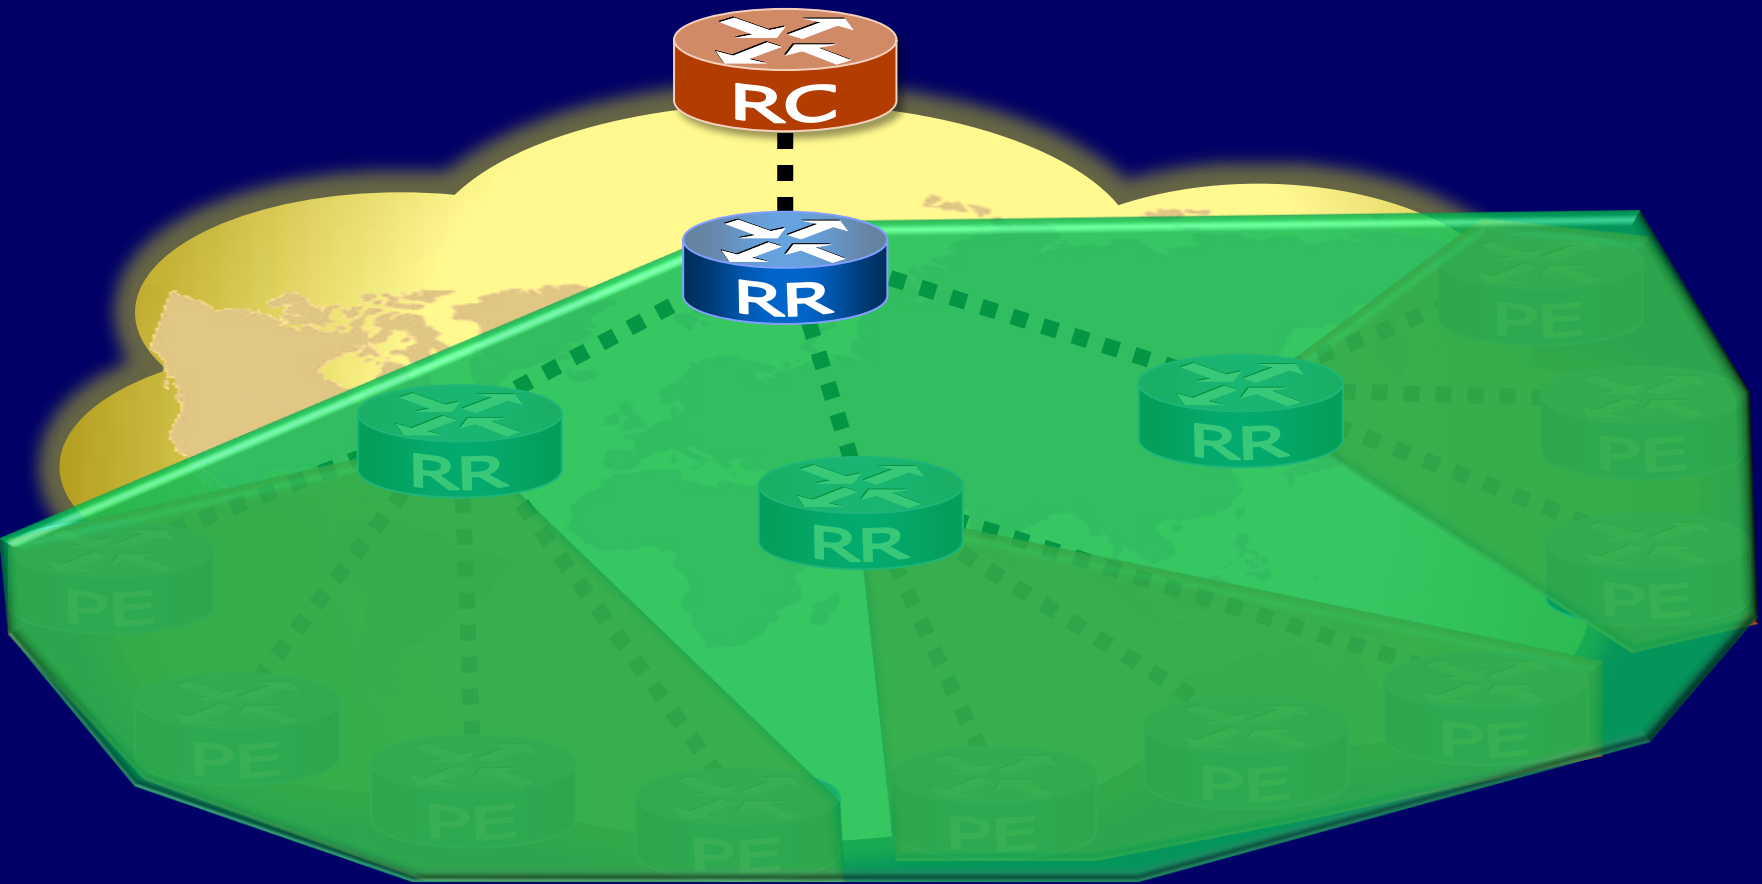


# Scalability vs Visibility






# Scalability vs Visibility



# Scalability vs Visibility



# Scalability vs Visibility



# Scalability vs Visibility



# Scalability vs Visibility



lower

**layer**

higher



lower

**scalability**

higher

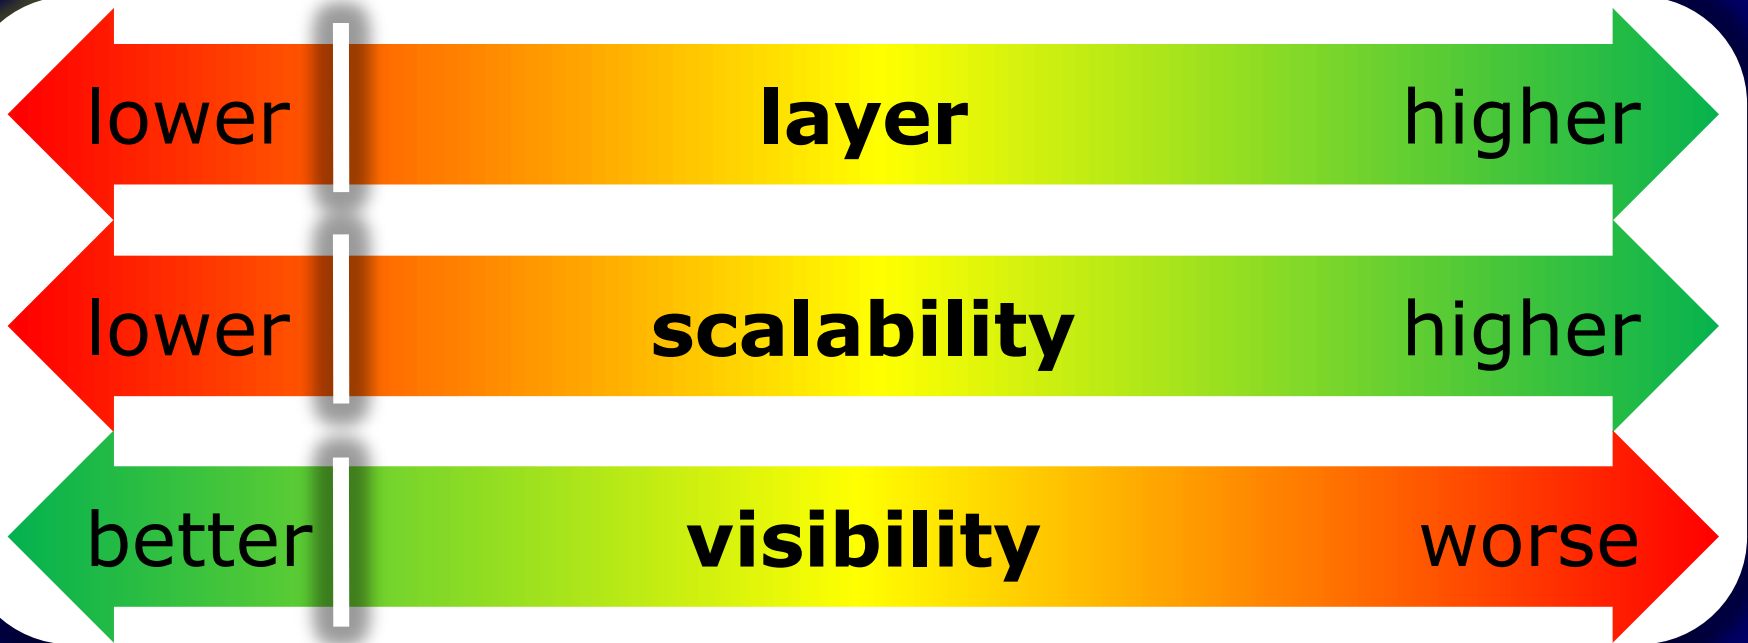


better


**visibility**

worse

# Scalability vs Visibility



# Scalability vs Visibility

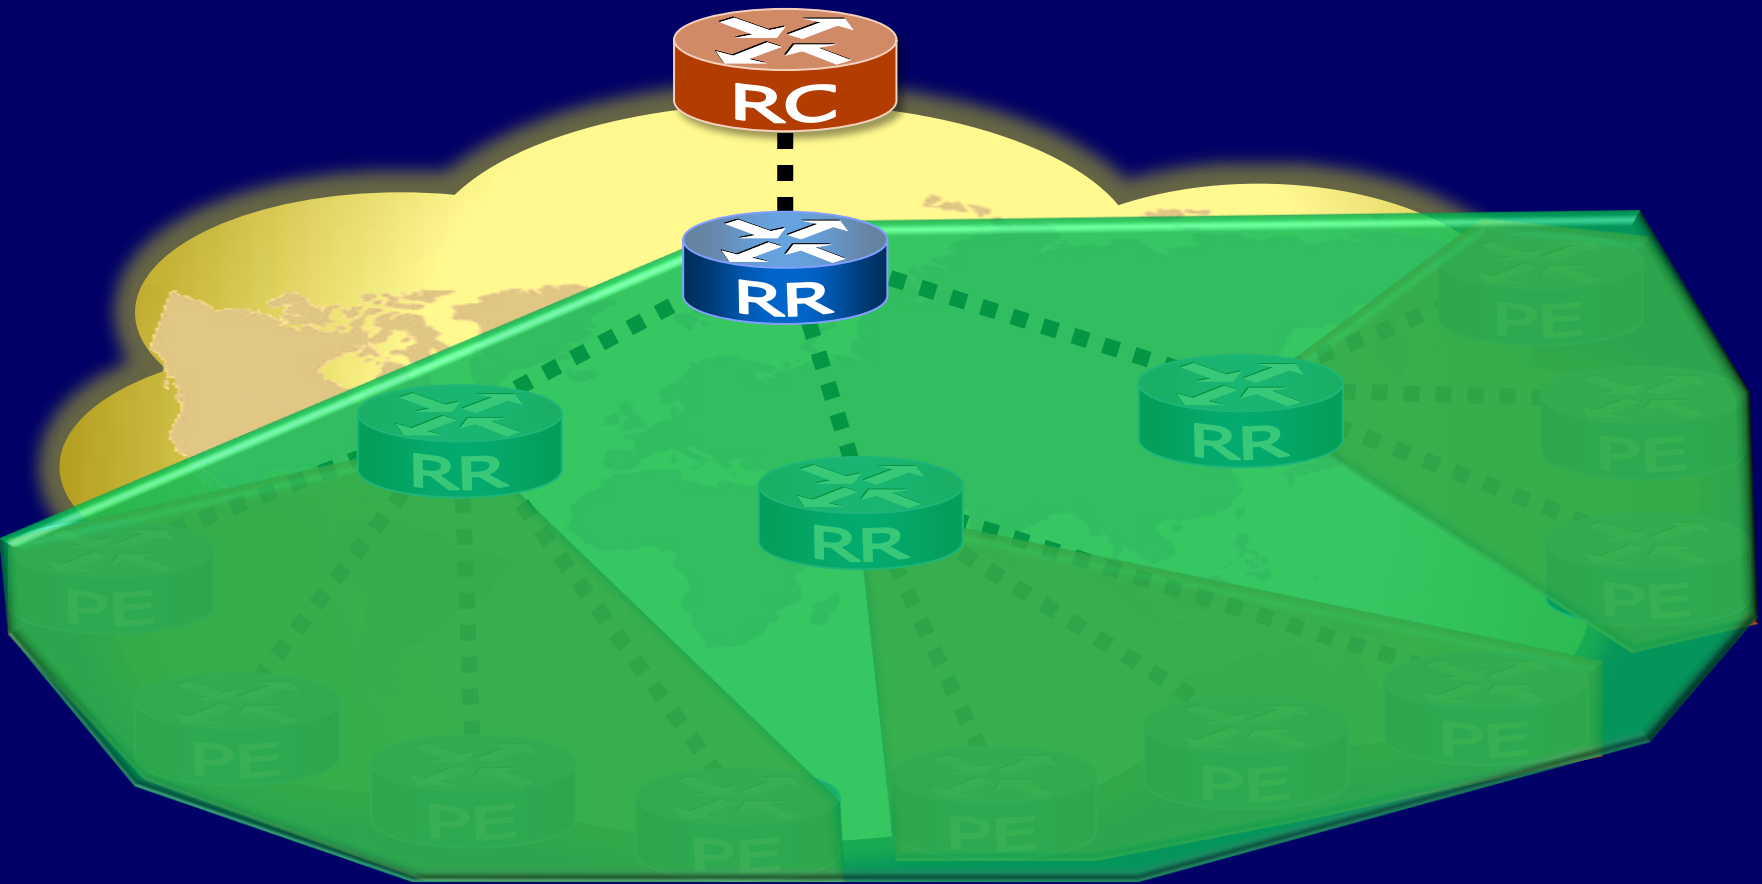


# Scalability vs Visibility







# Scalability vs Visibility



# Scalability vs Visibility



# Scalability vs Visibility



beware of matching updates

# Experimental Scenario



# Experimental Scenario



**ROUTE COLLECTOR**

# Experimental Scenario




**ROUTE COLLECTOR**

# Experimental Scenario




# Experimental Scenario

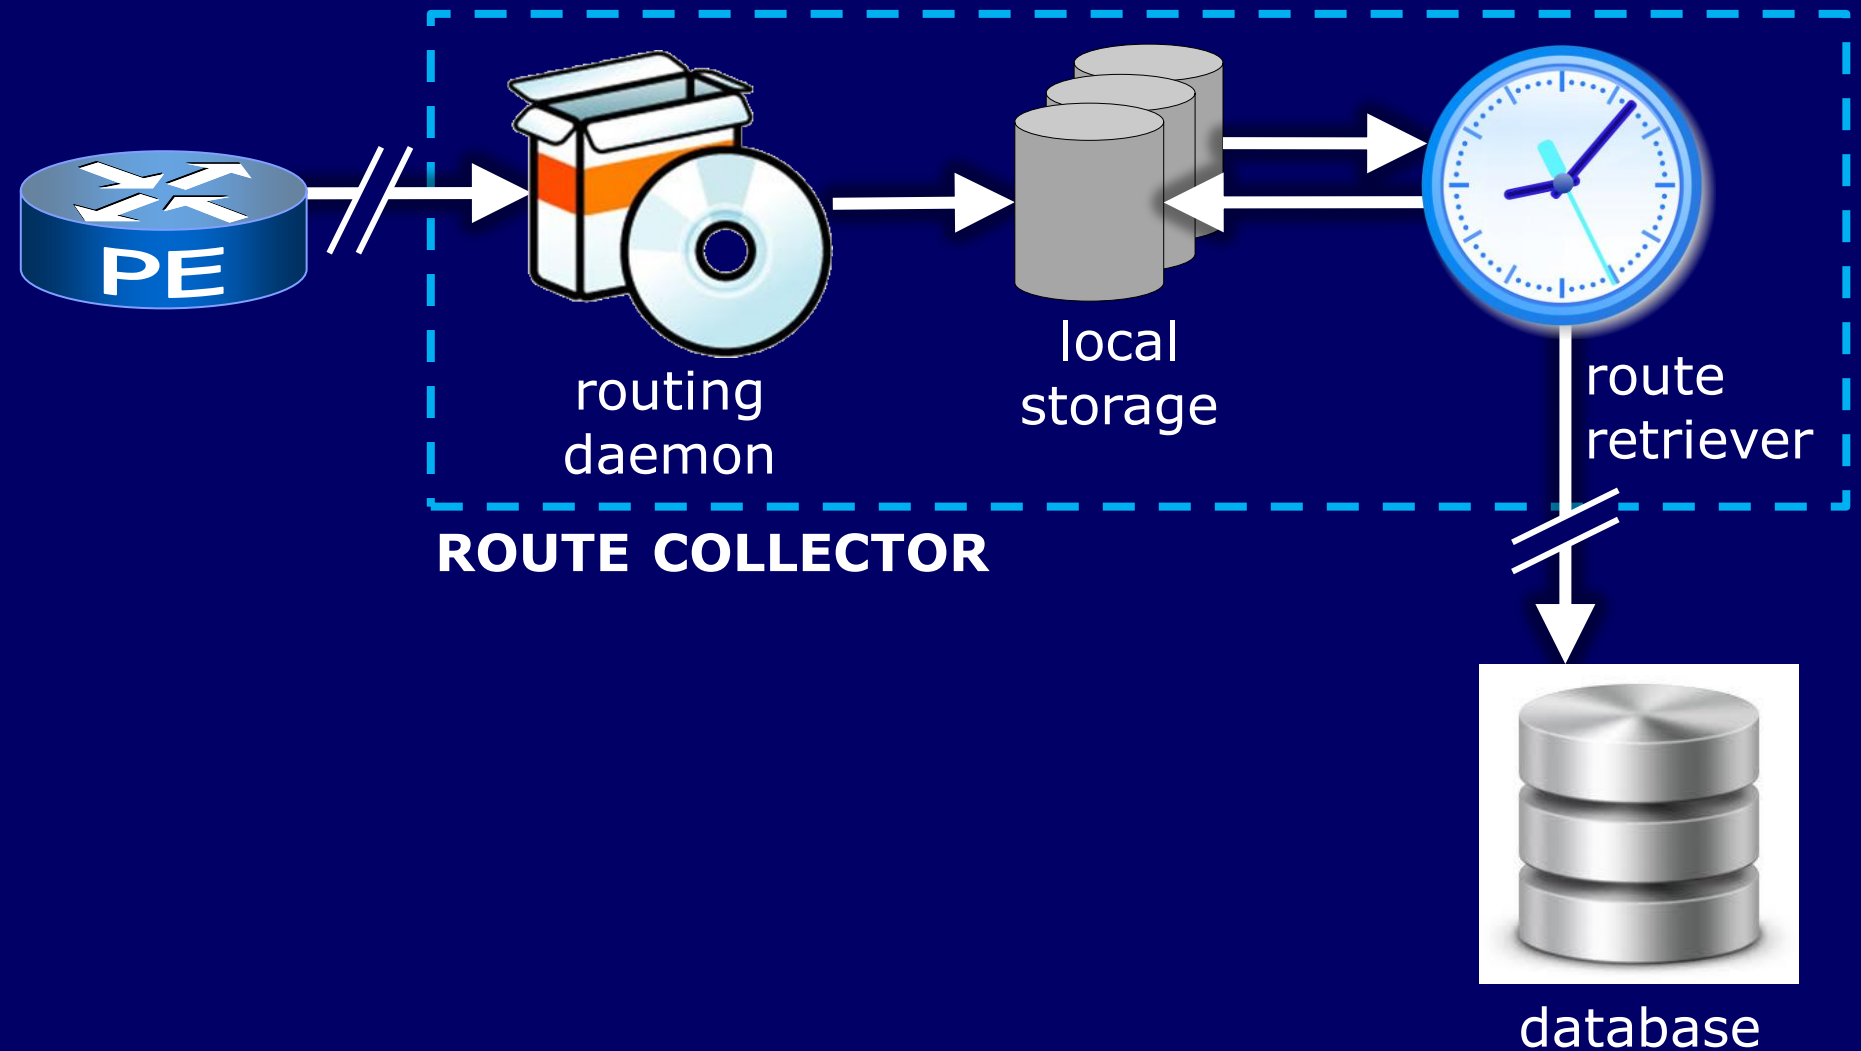




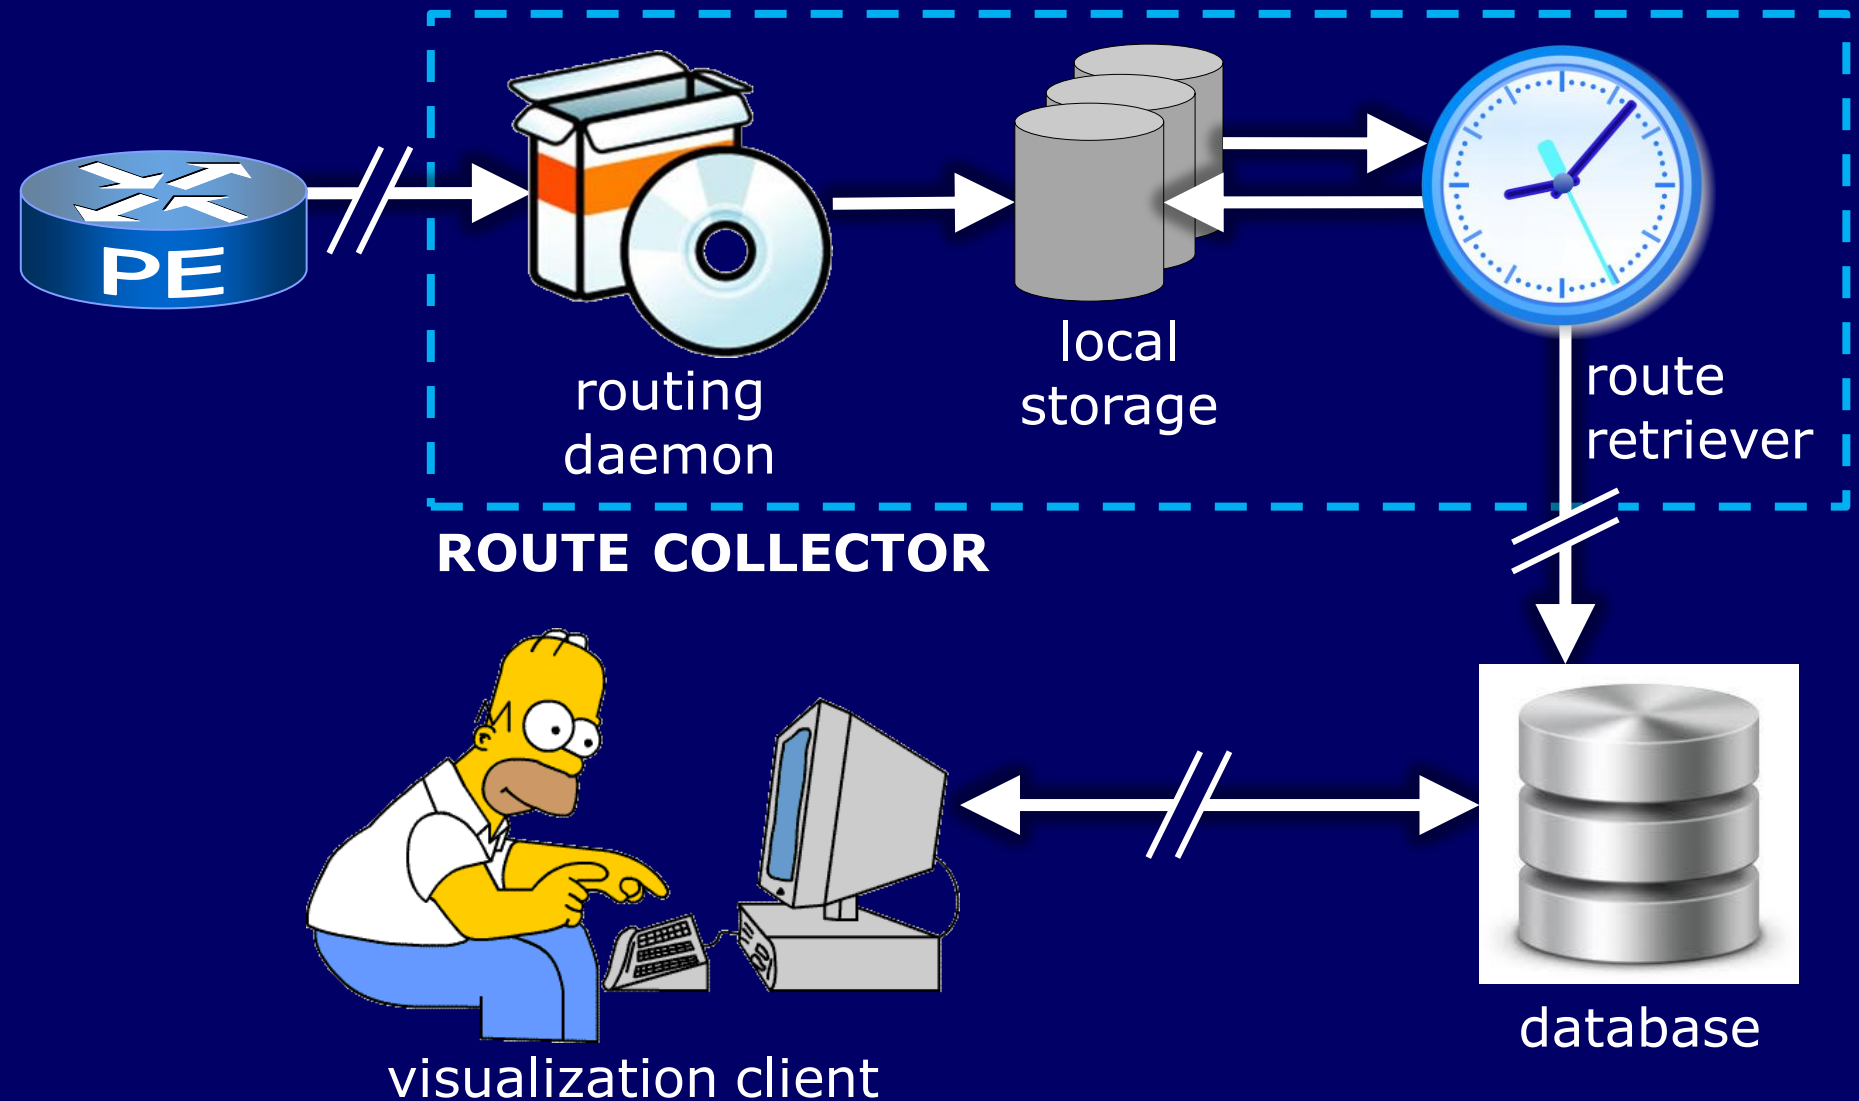
# Experimental Scenario



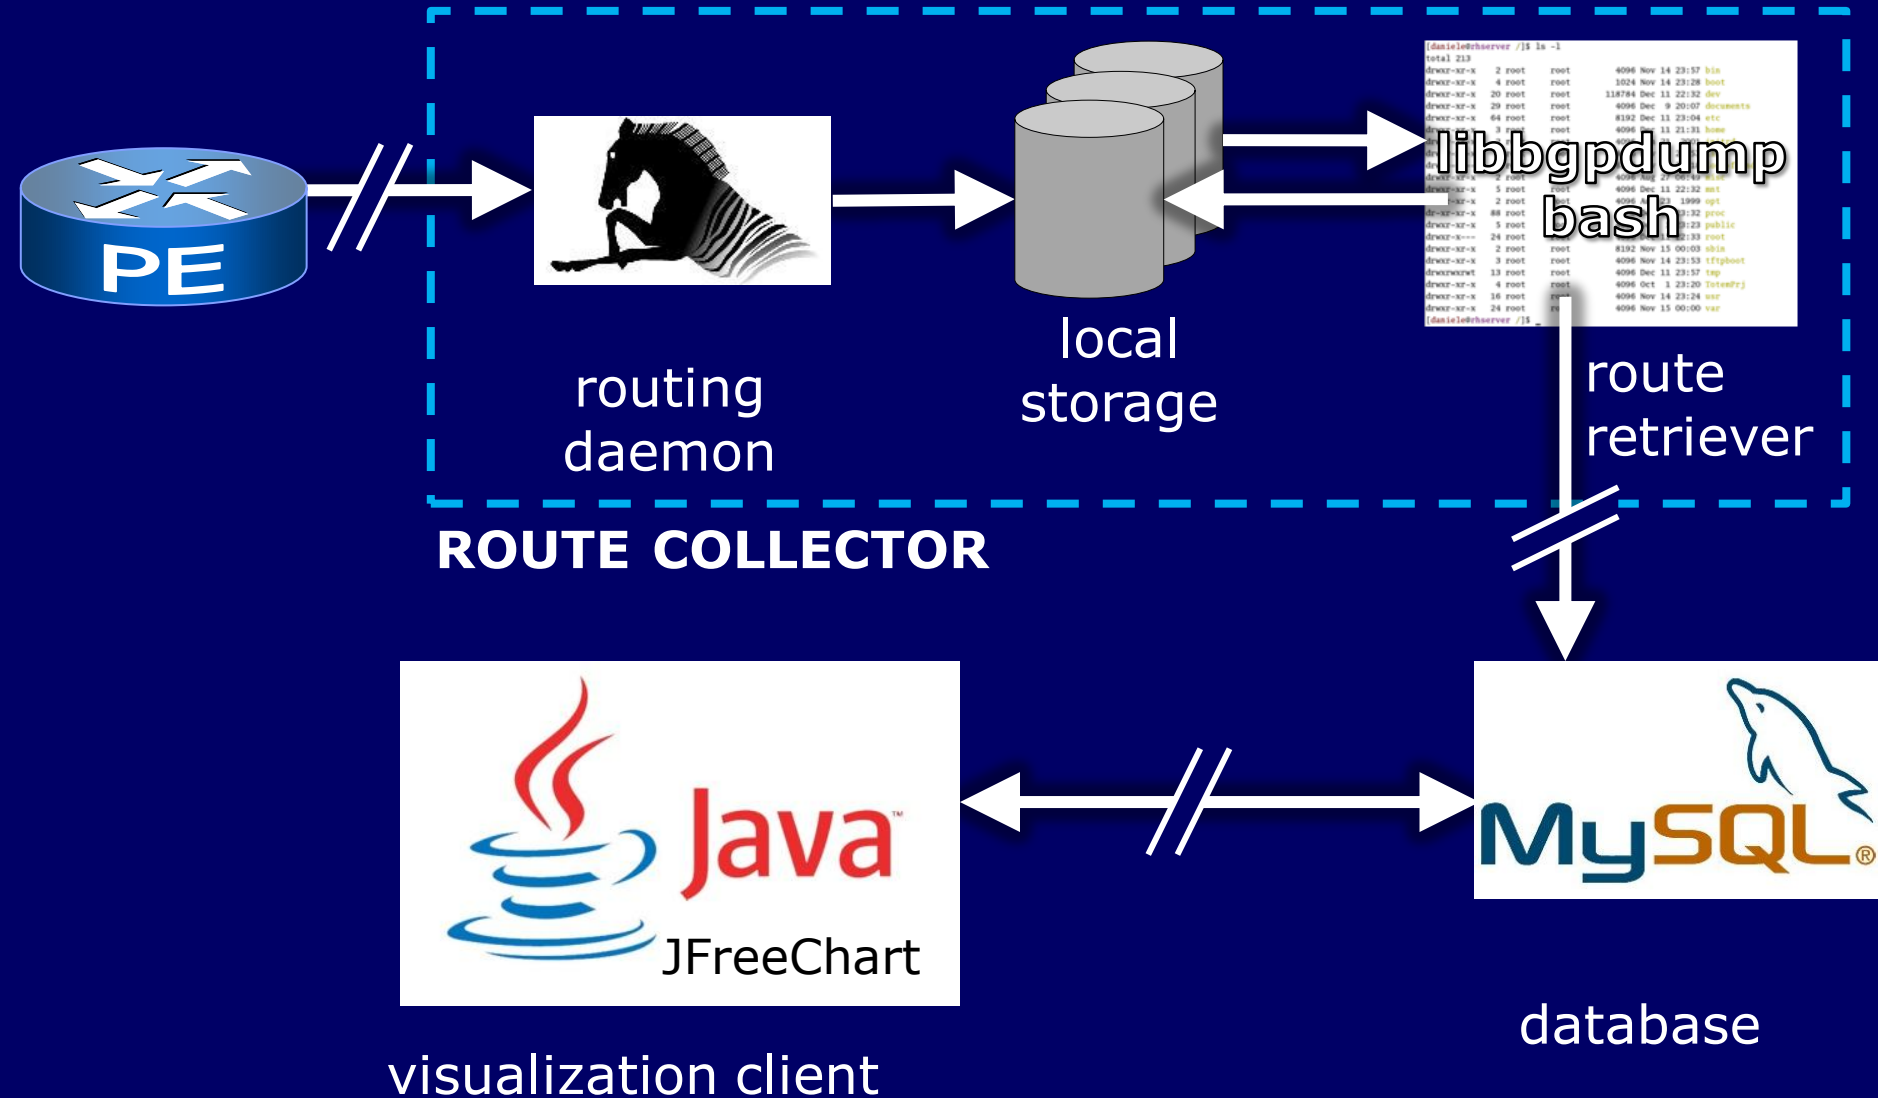
# Experimental Scenario



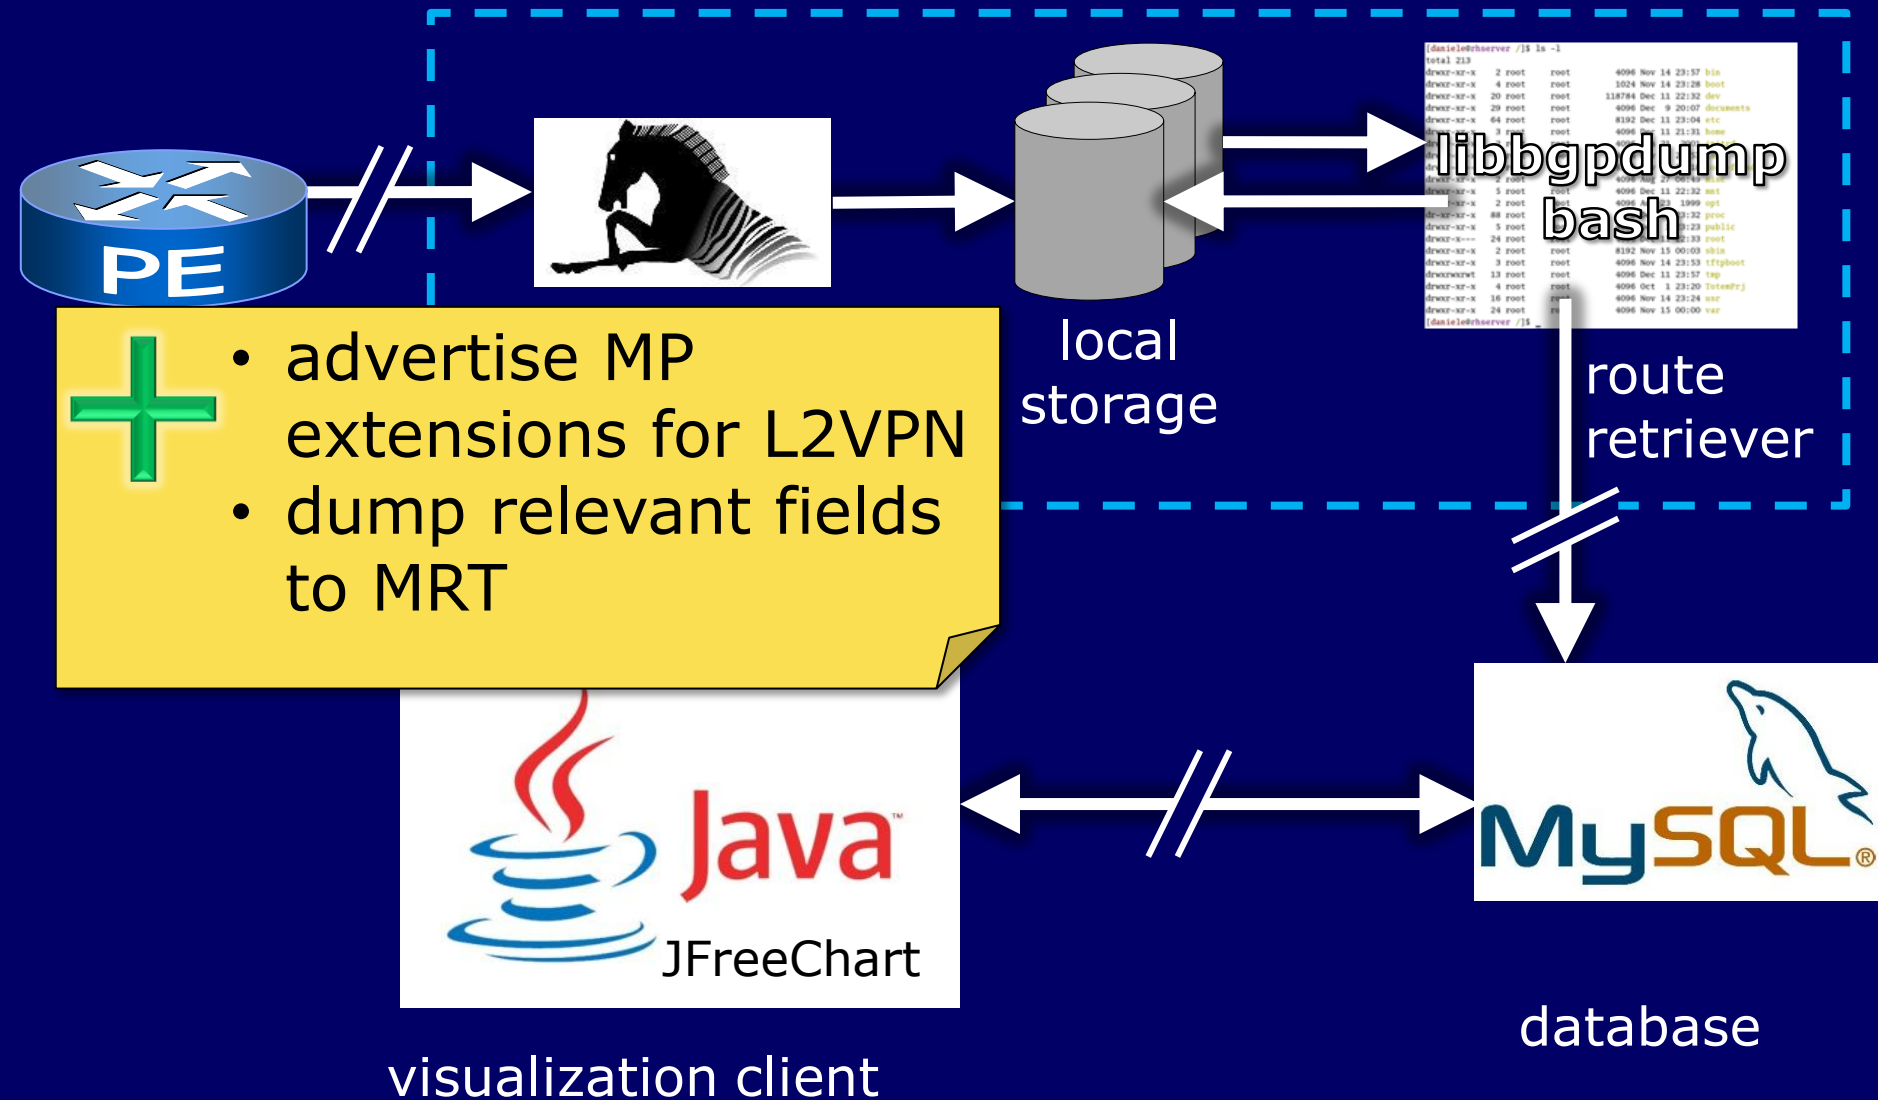
# Experimental Scenario



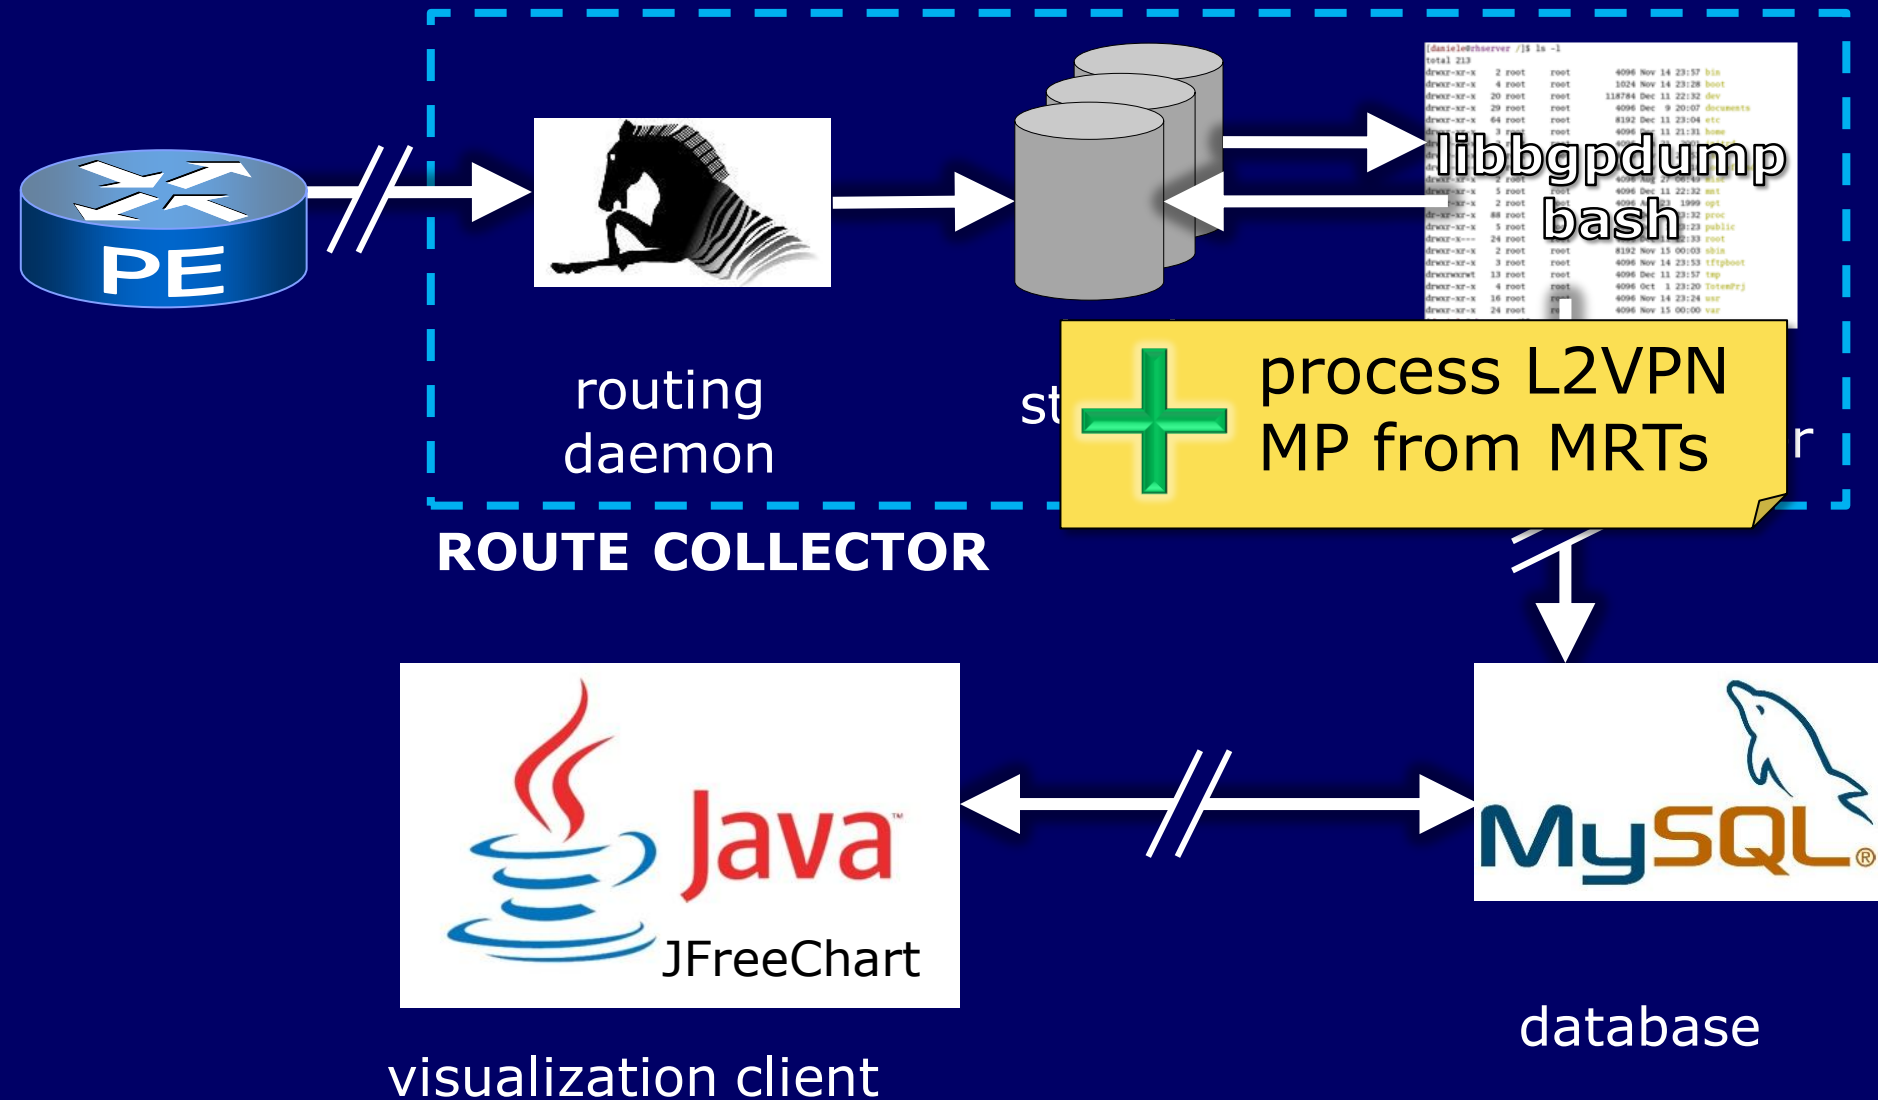
# Experimental Scenario



# Experimental Scenario



# Experimental Scenario



# Experimental Scenario



max lag: 3mins



visualization client



database



route  
retriever

# Experimental Scenario



max lag: 3mins




preliminary tests  
on Cisco routers









SEA




CHI




NYC




DEN




LAX




WAS



ATL




HOU







SEA



CHI




NYC




DEN




LAX




WAS

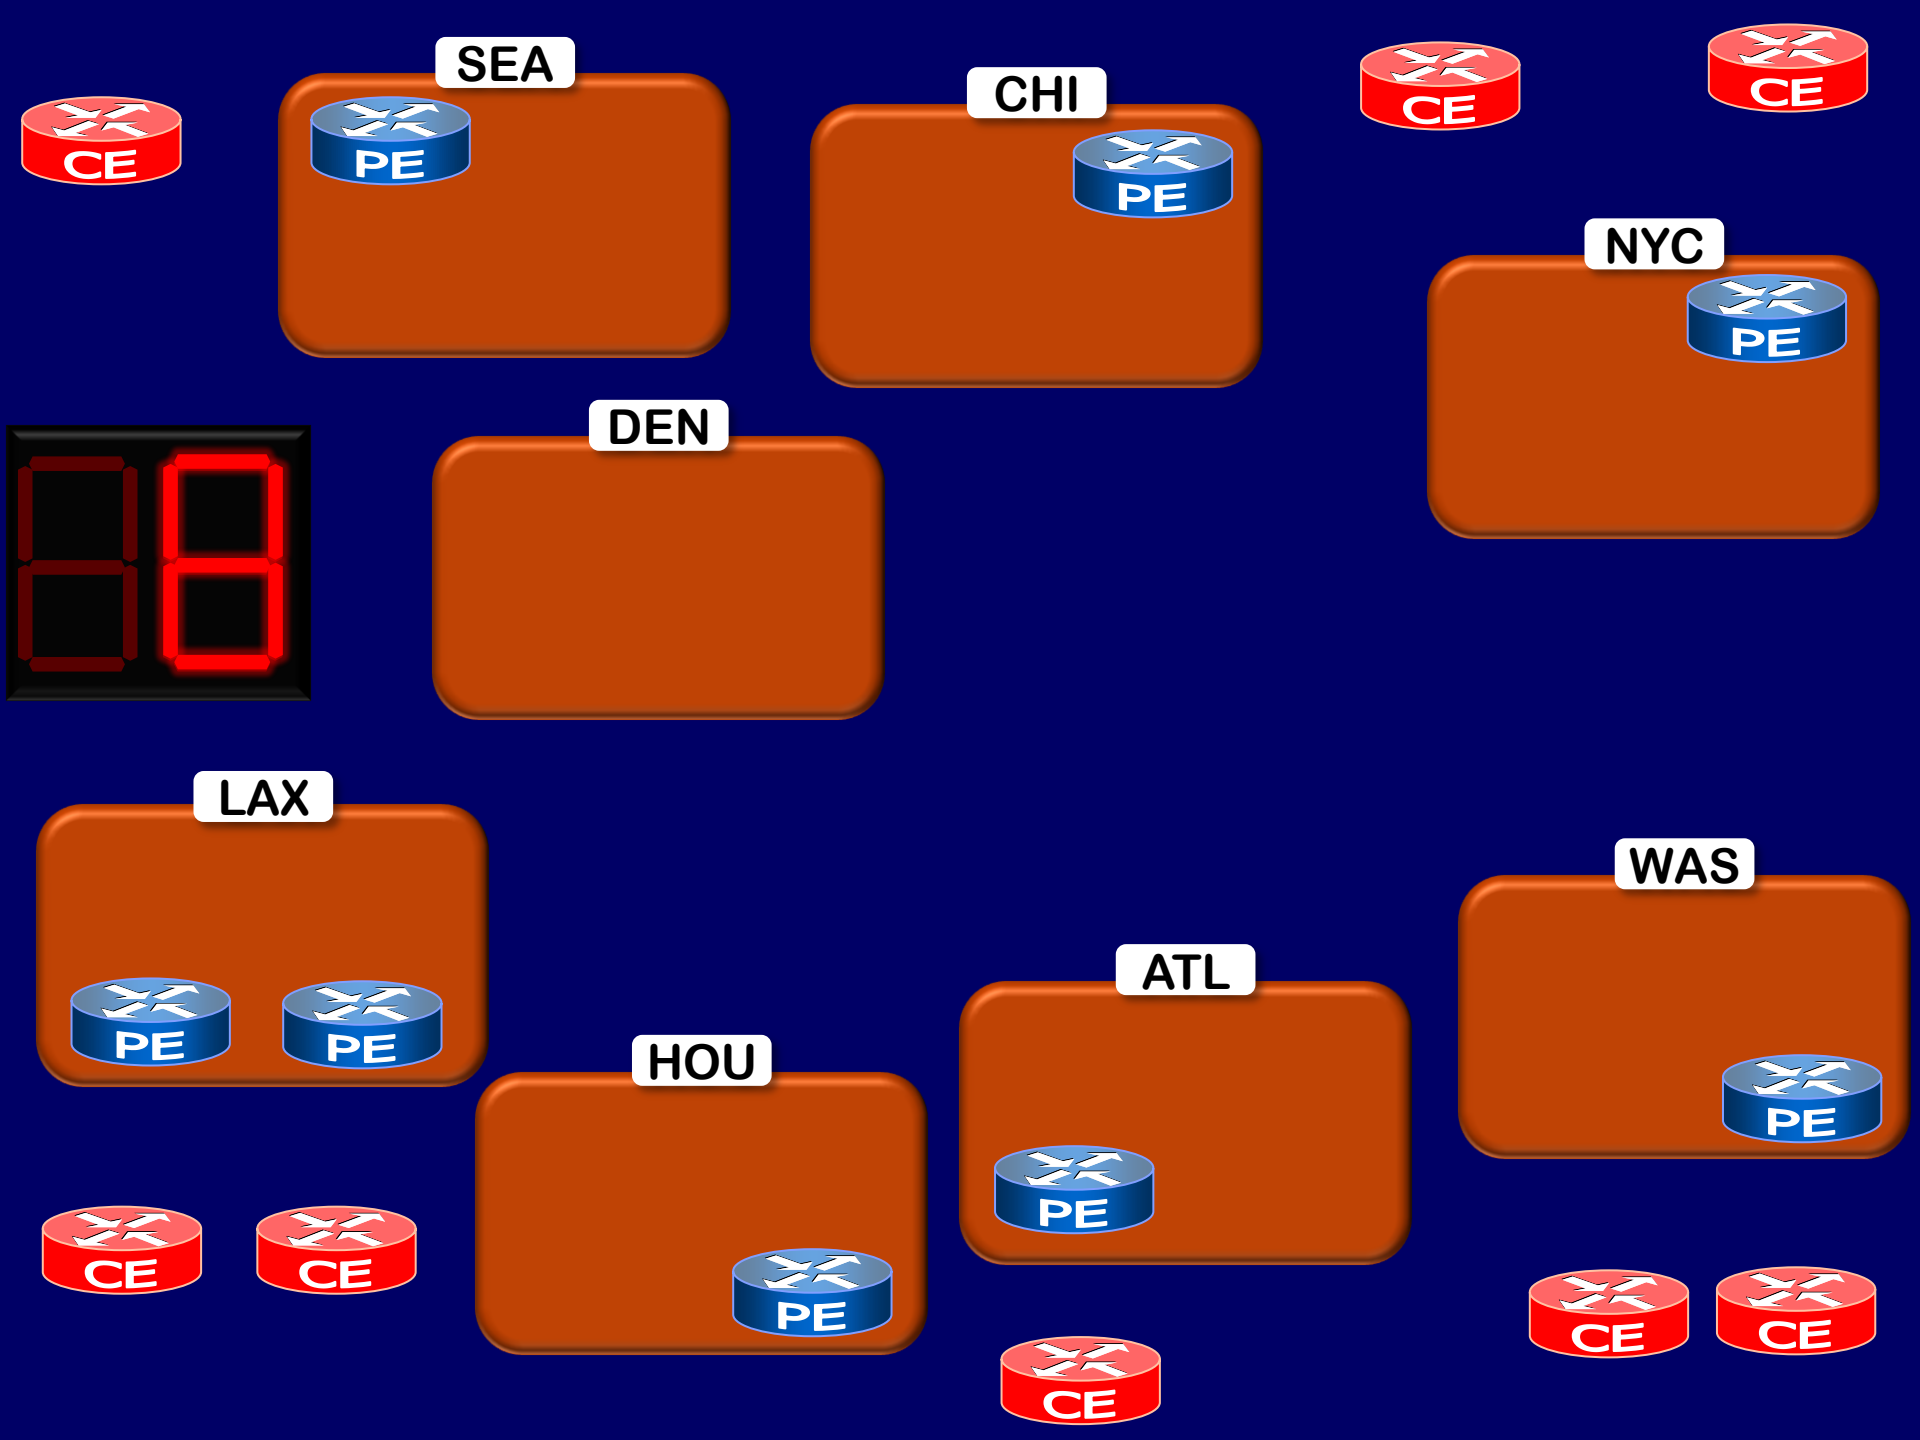


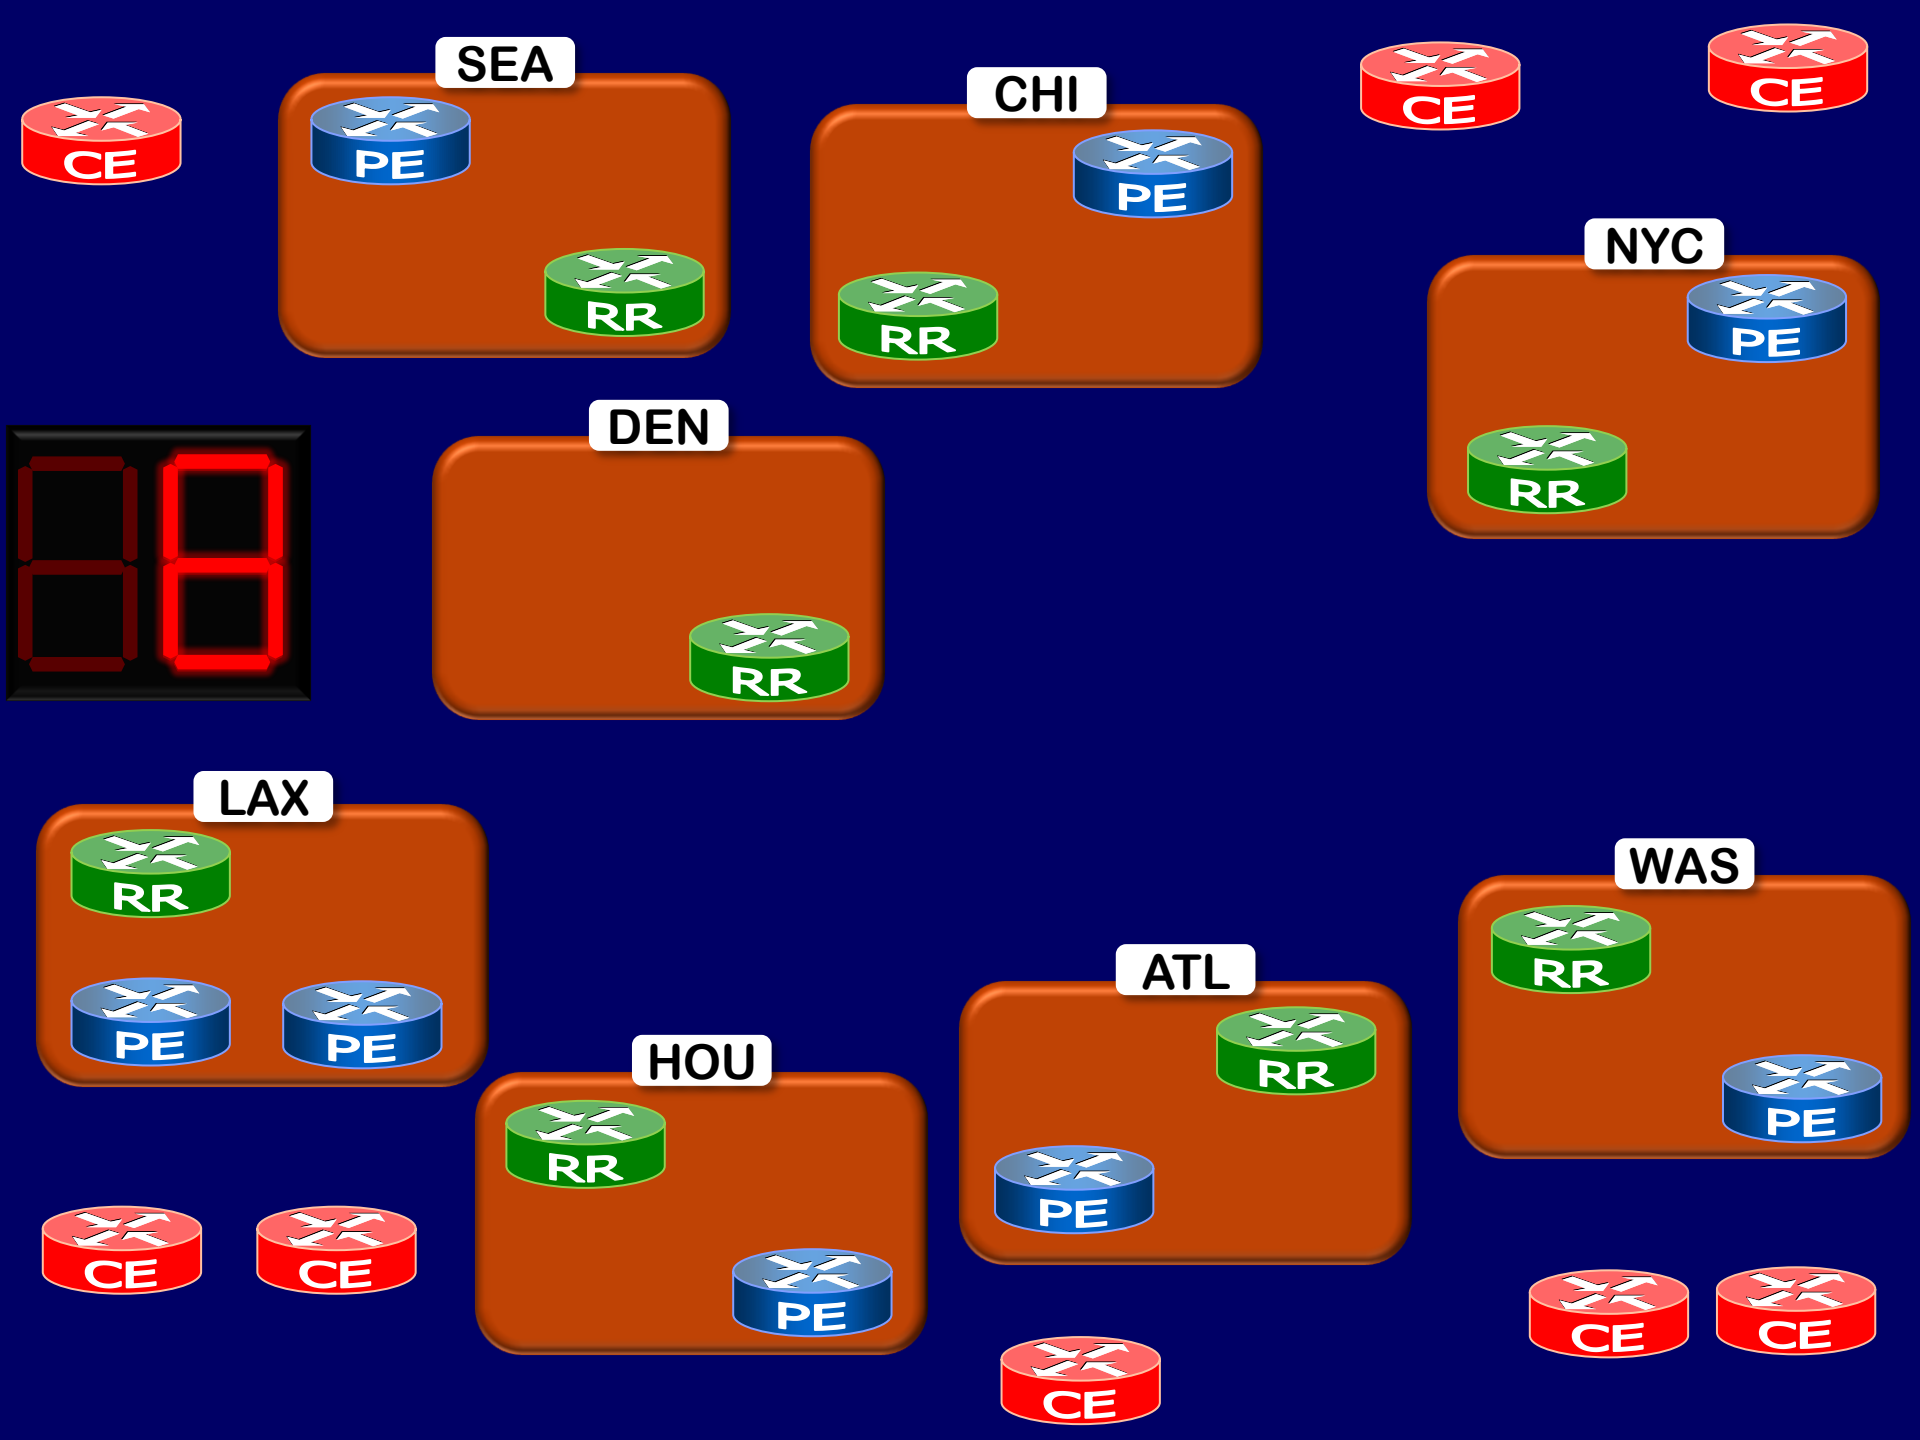
ATL

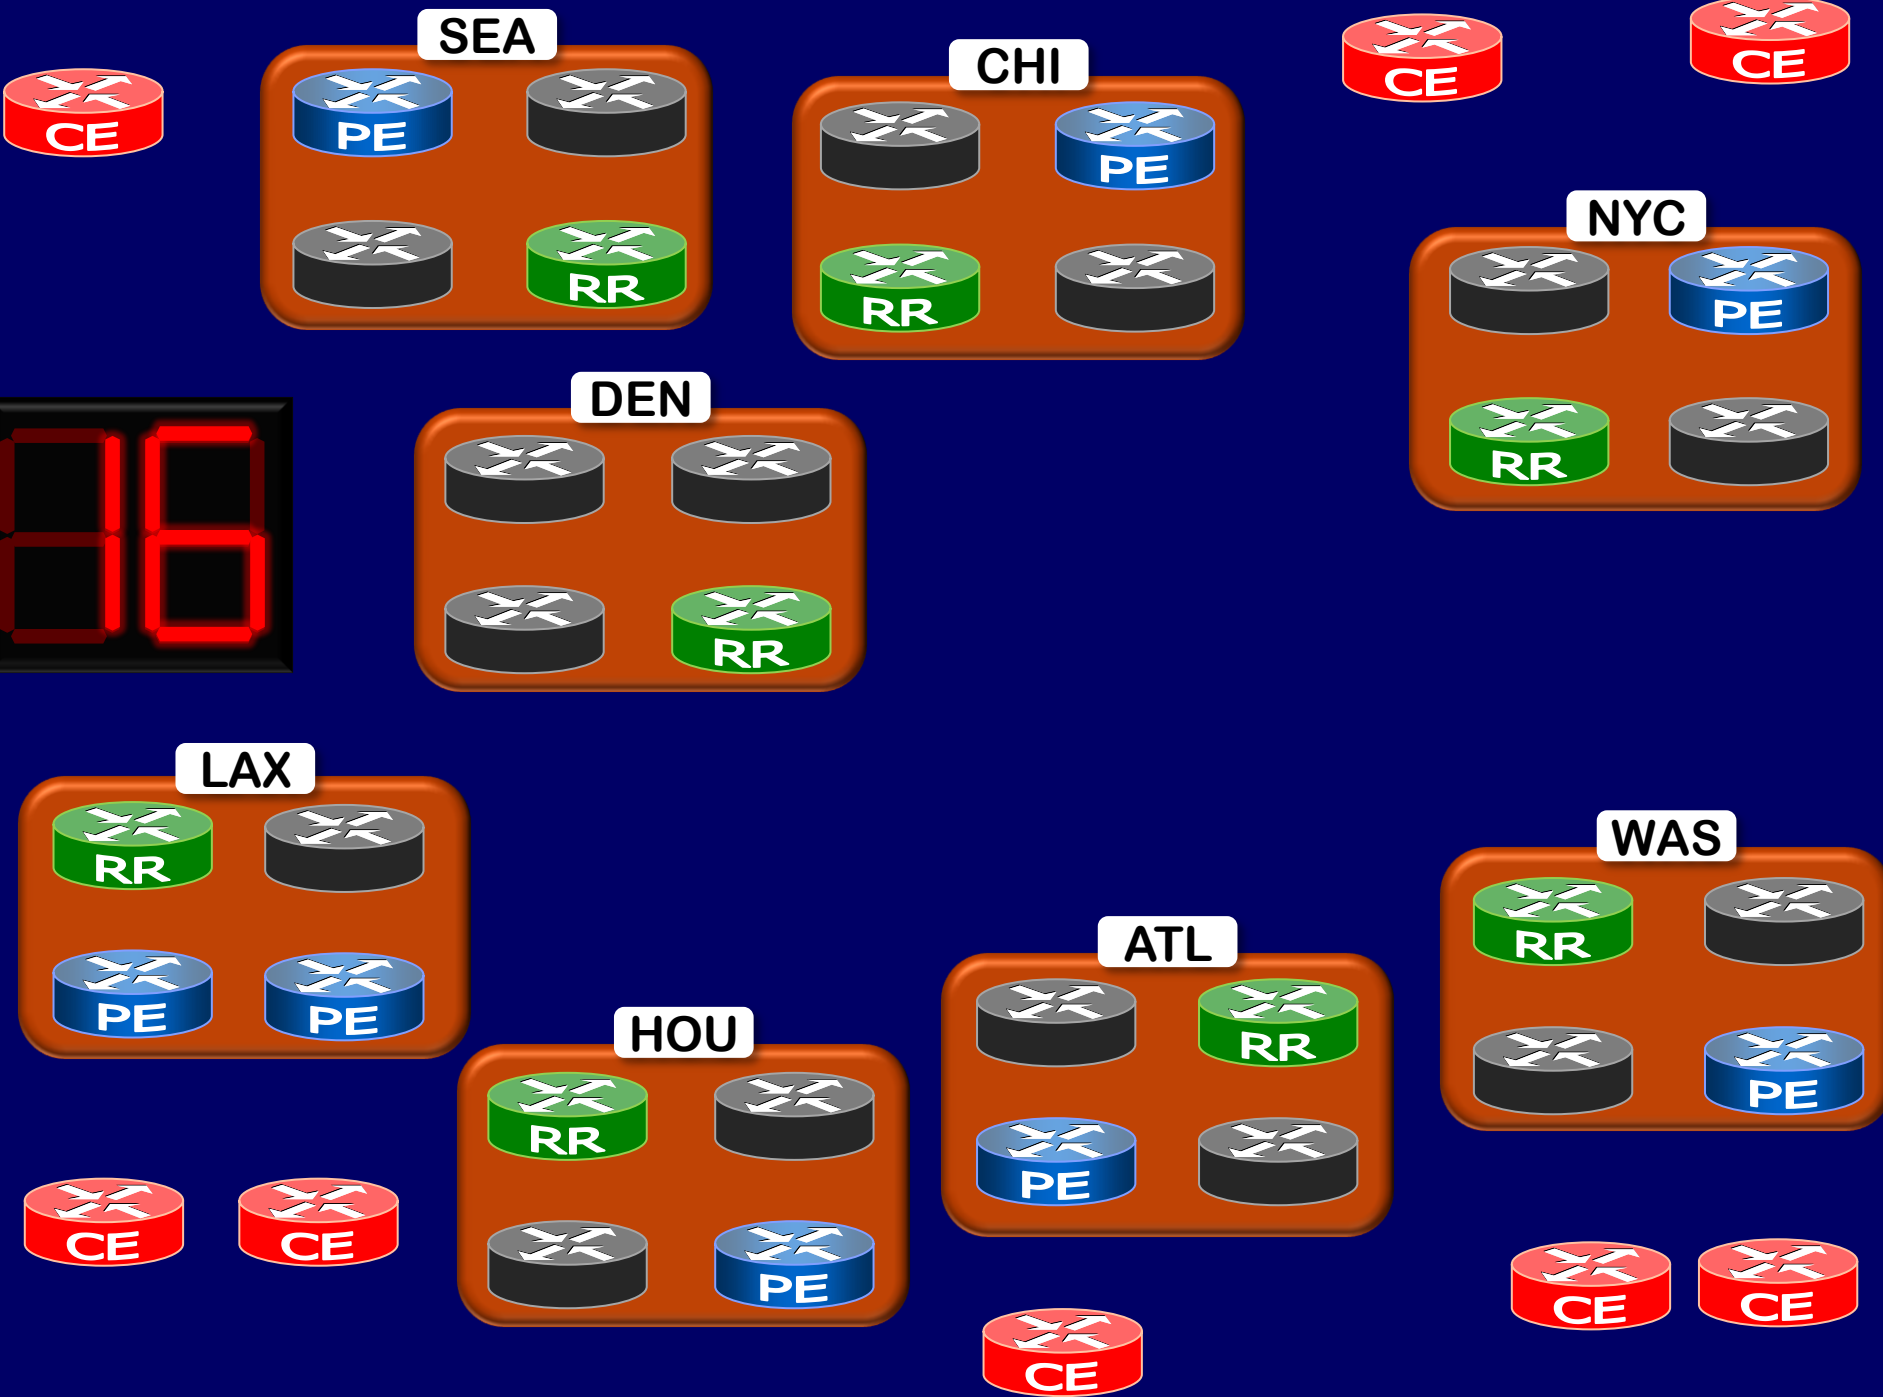



HOU






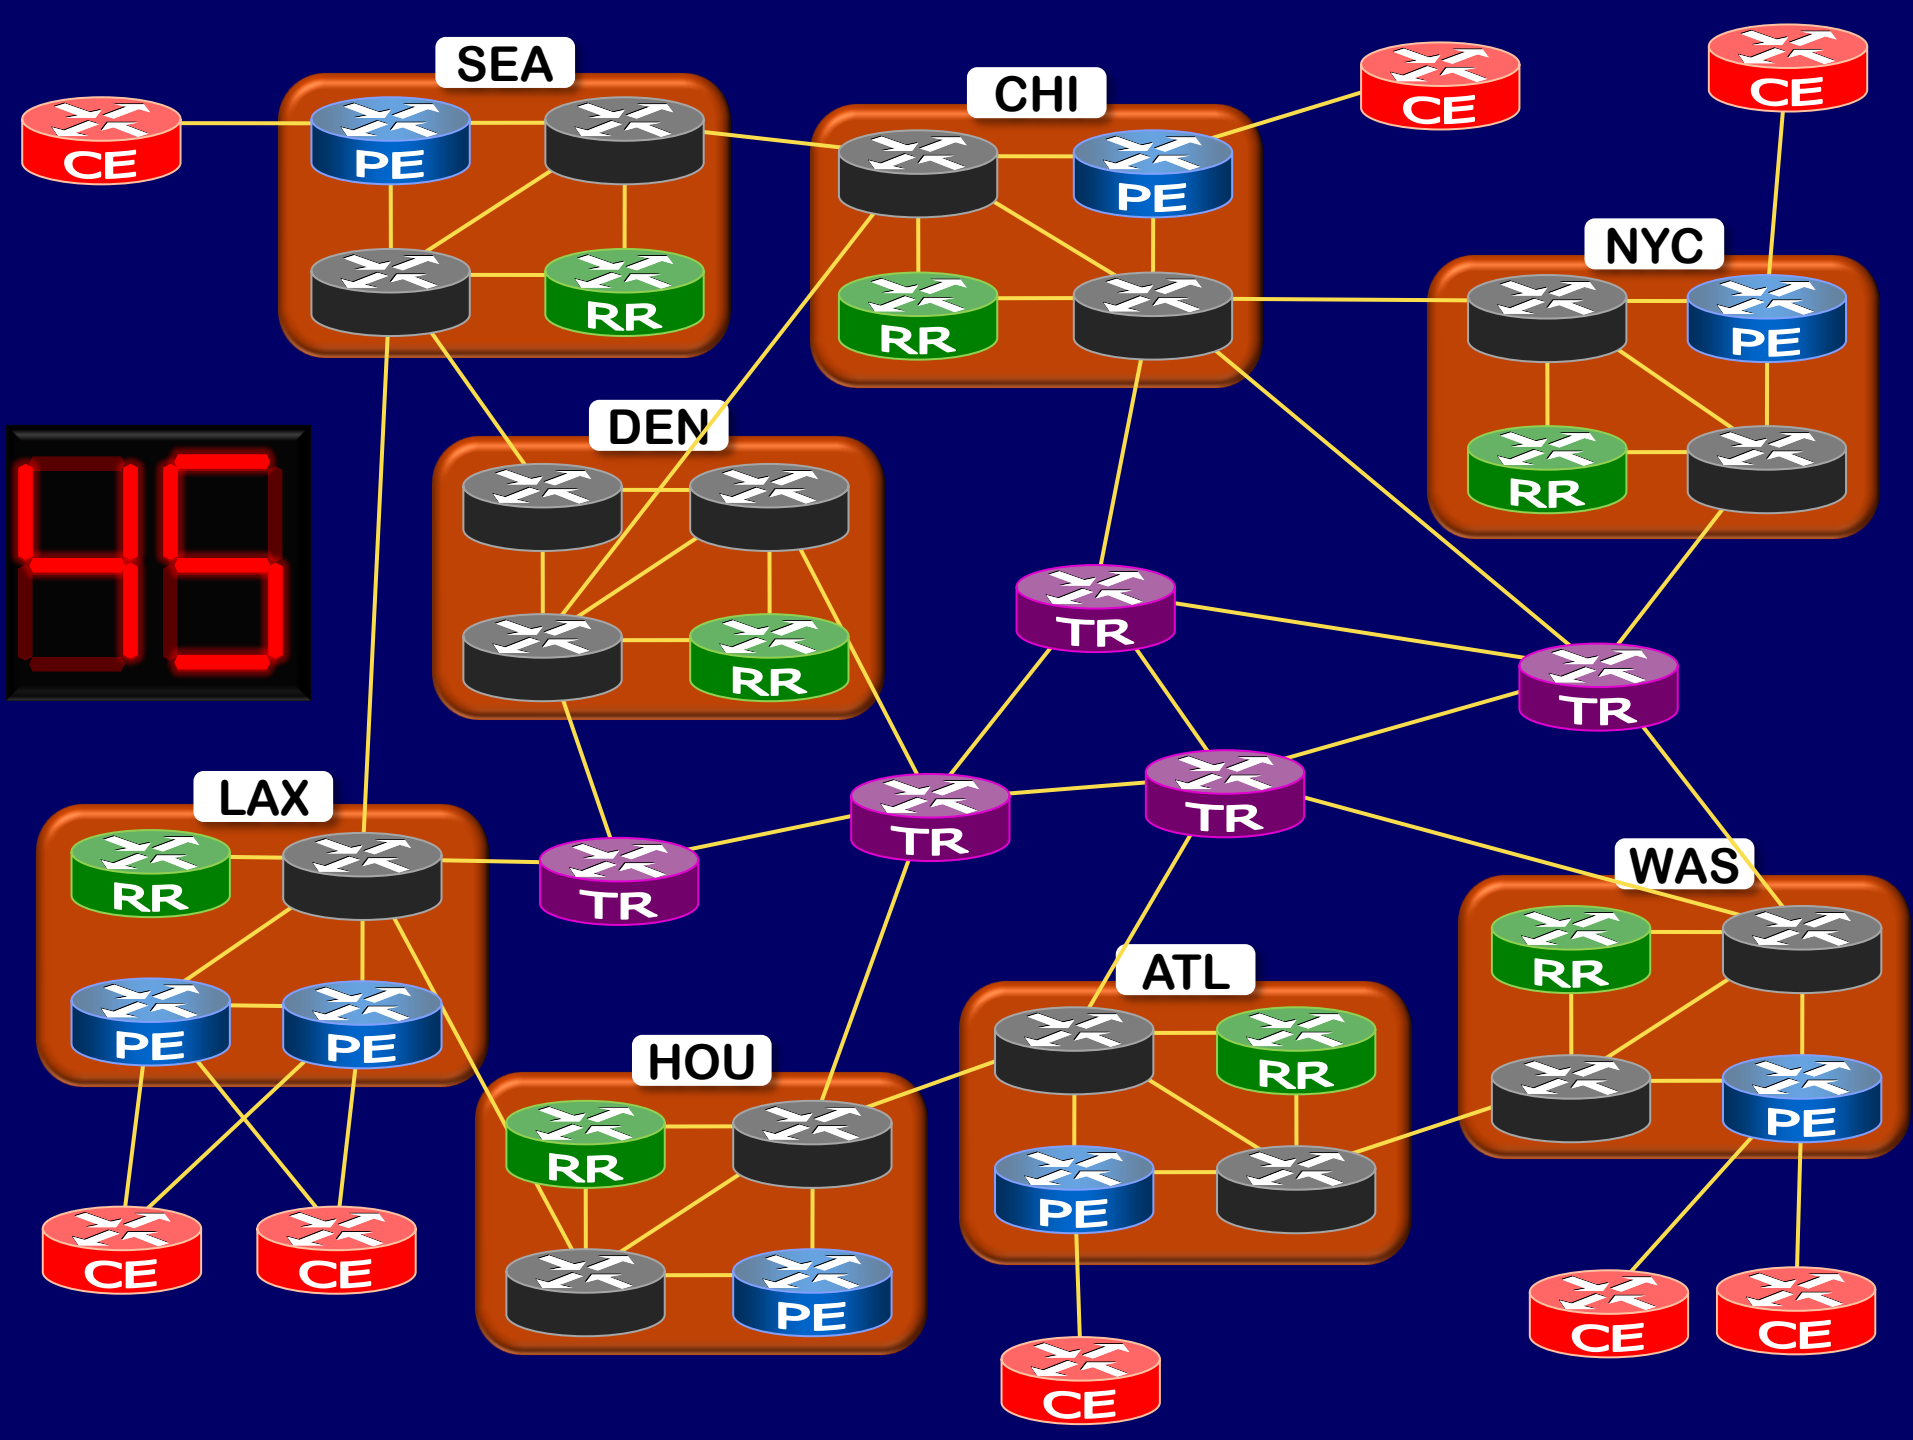


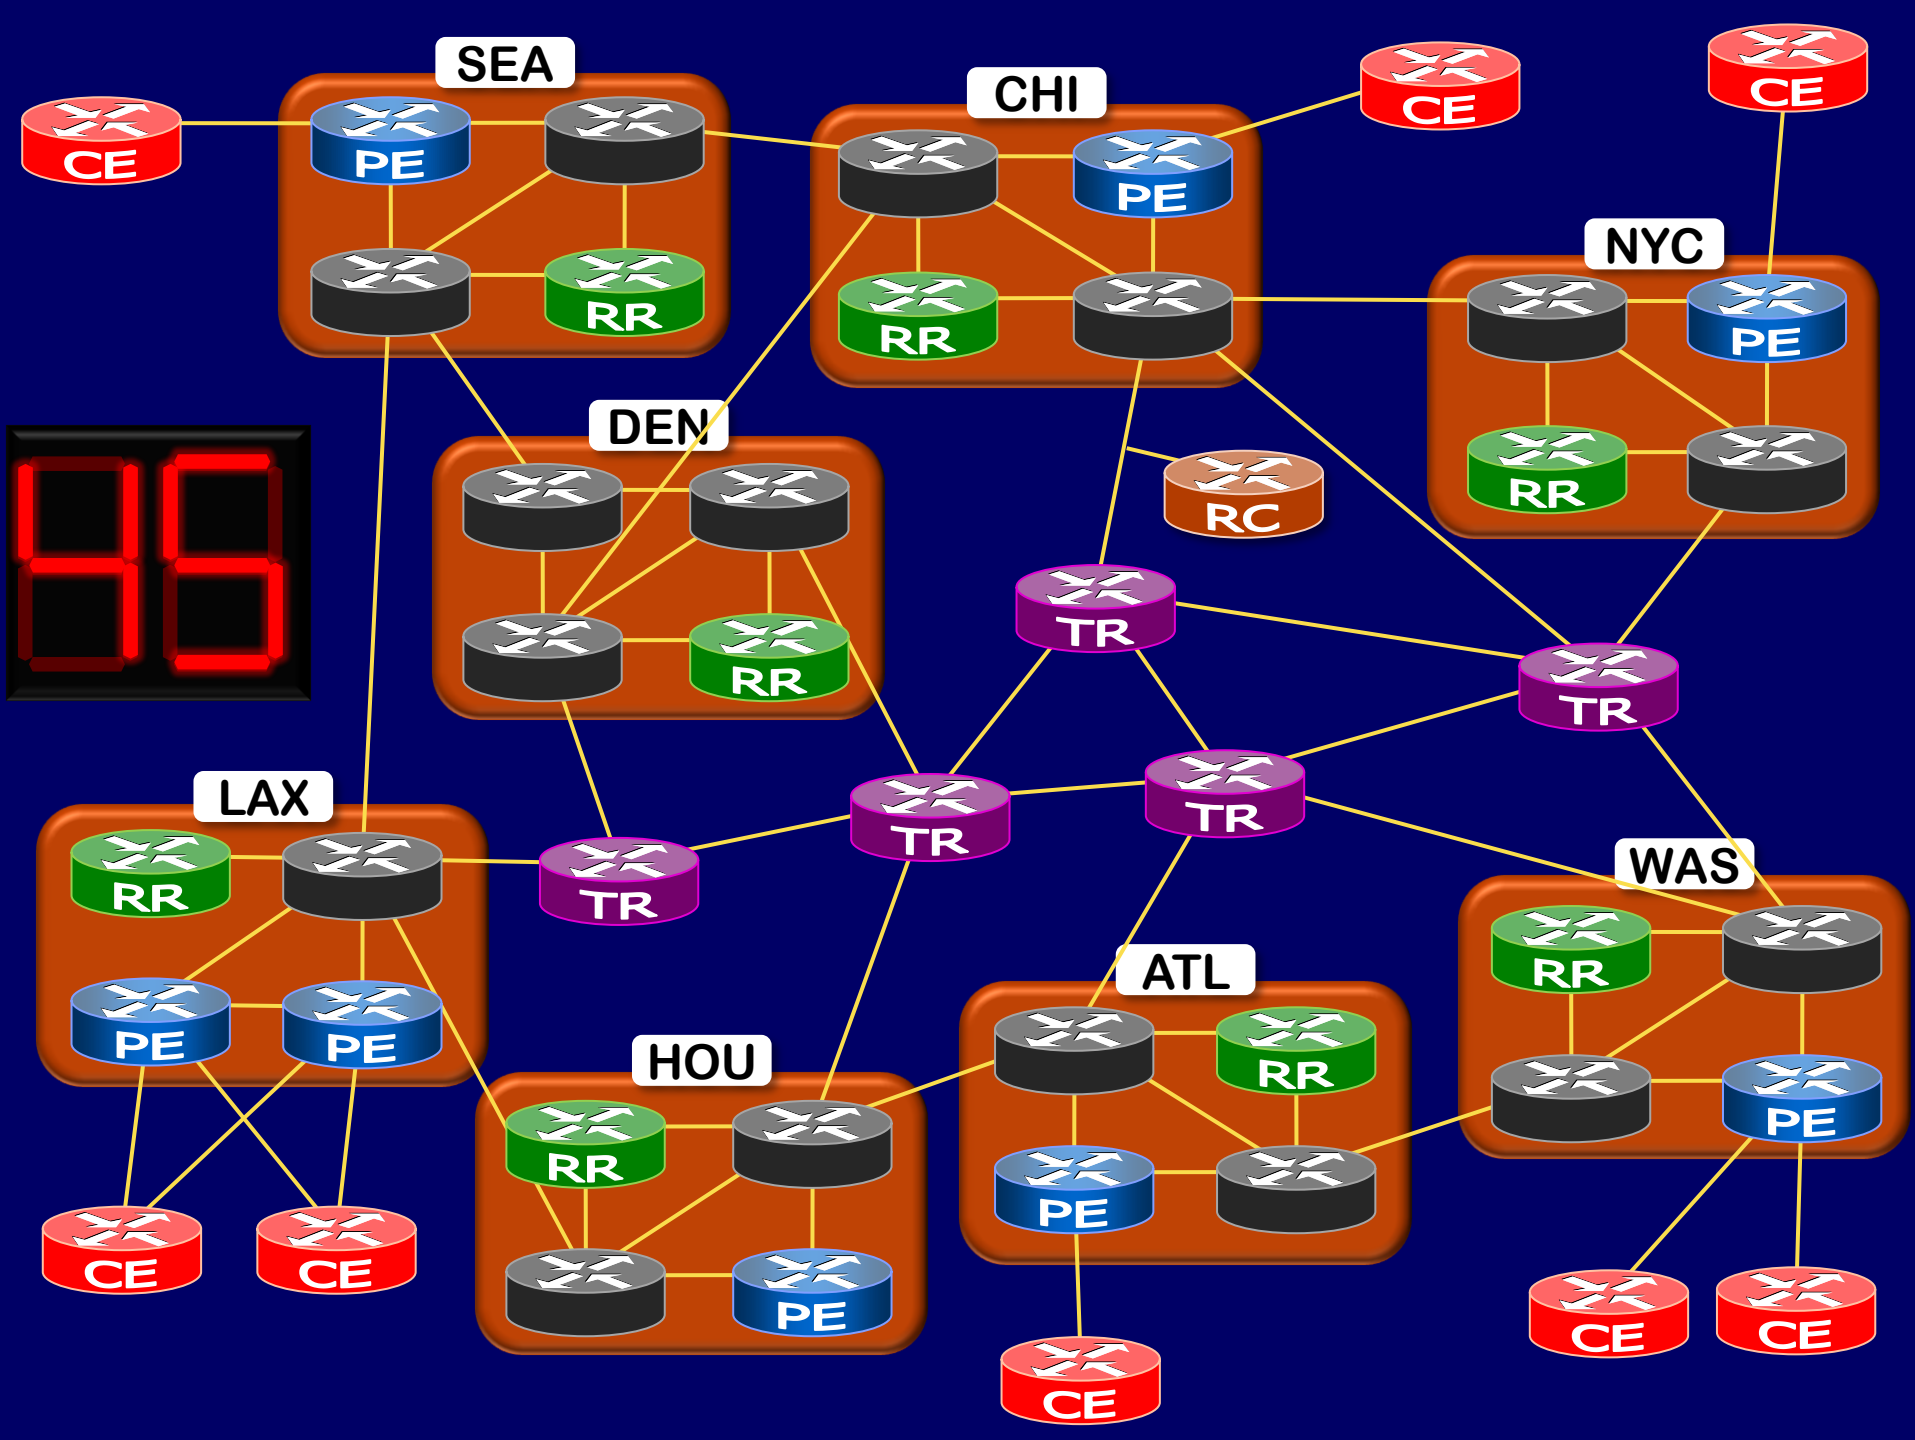








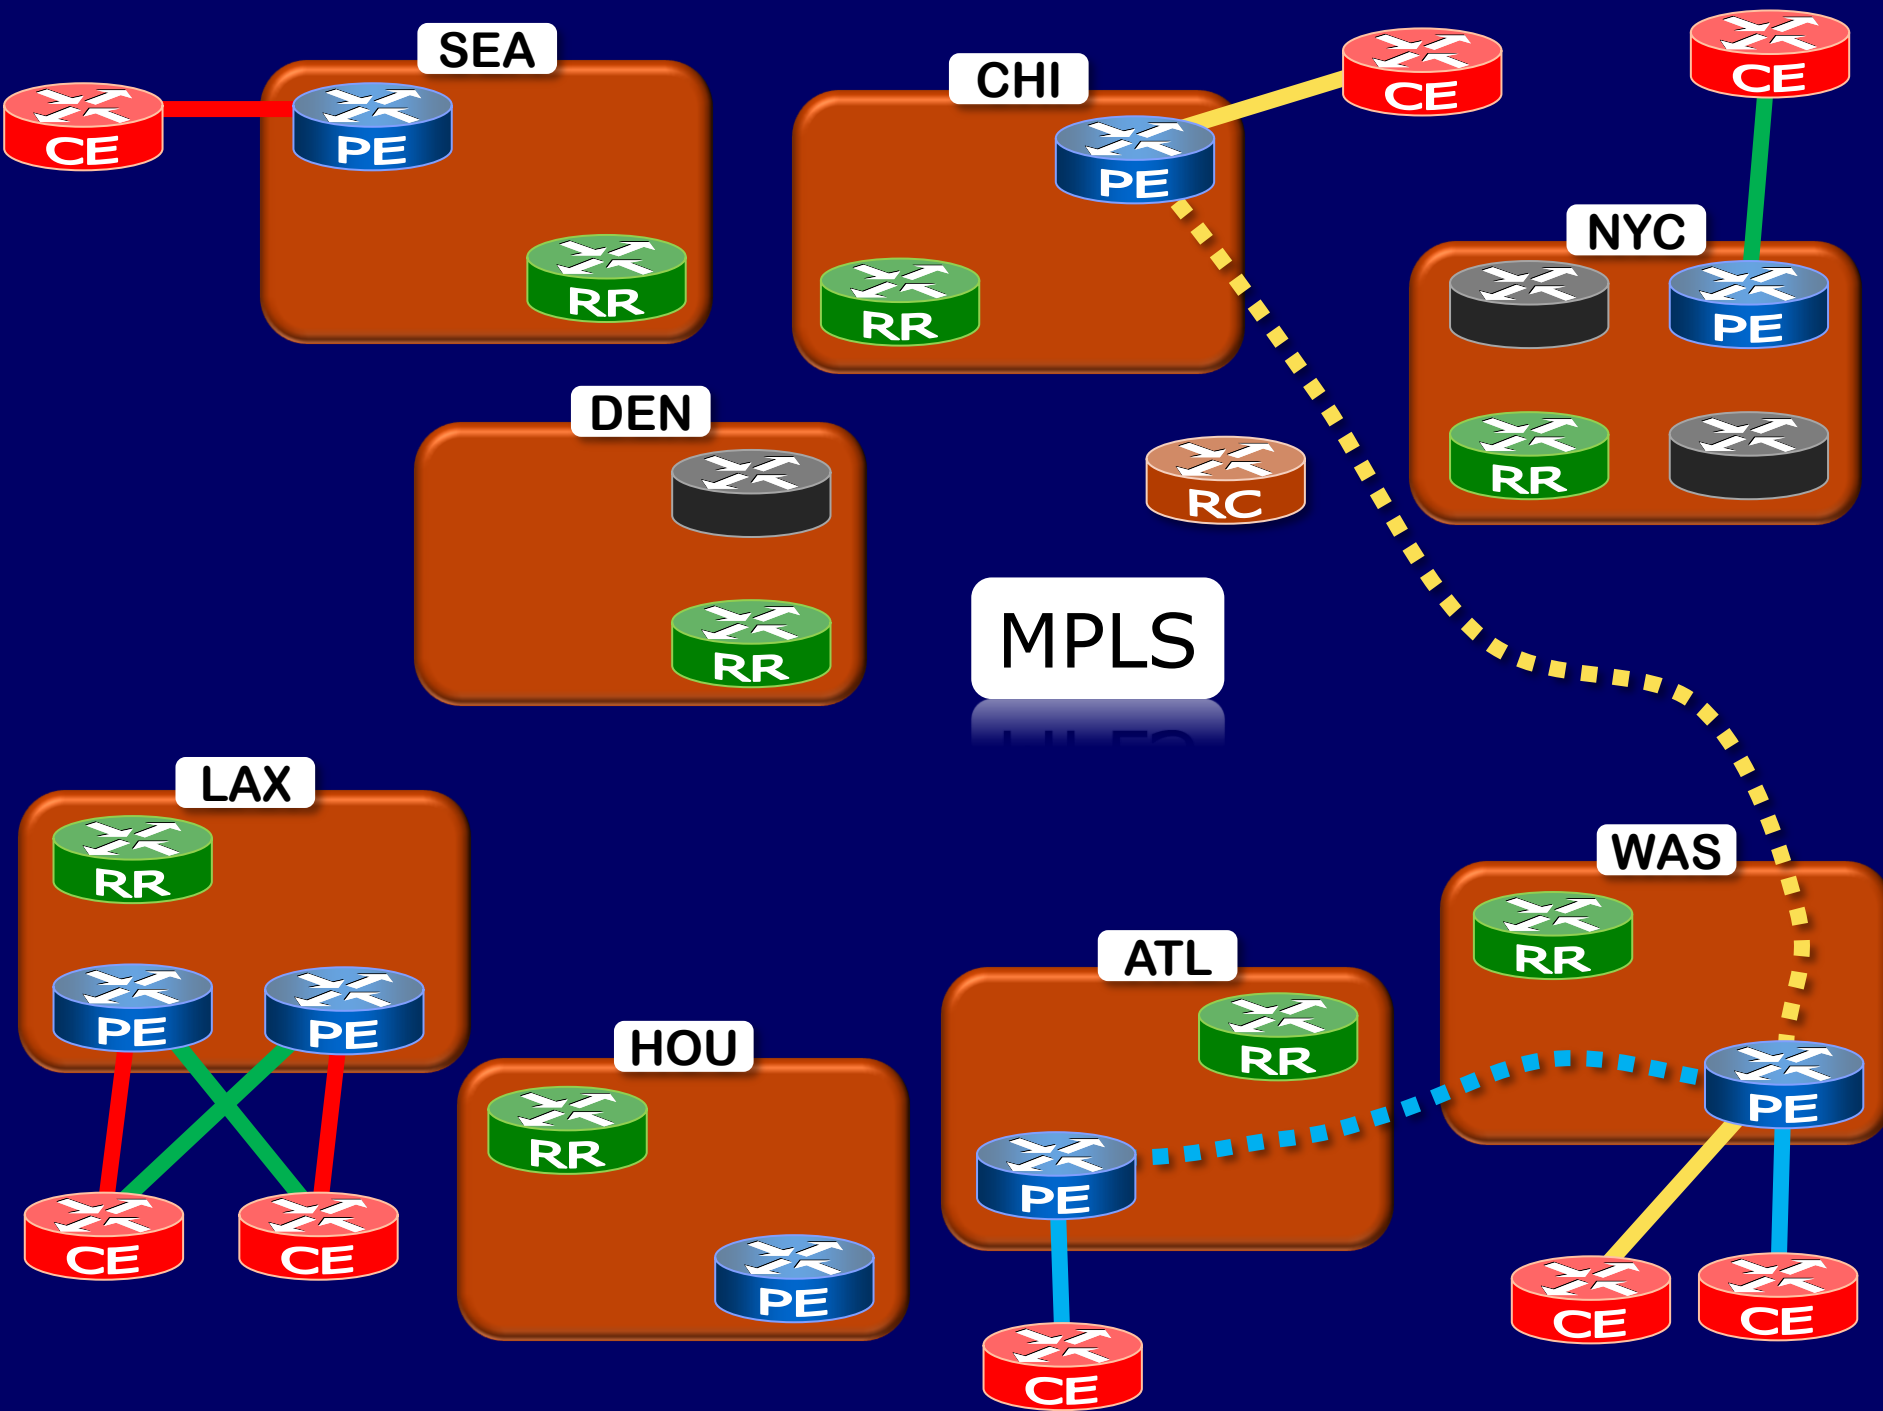





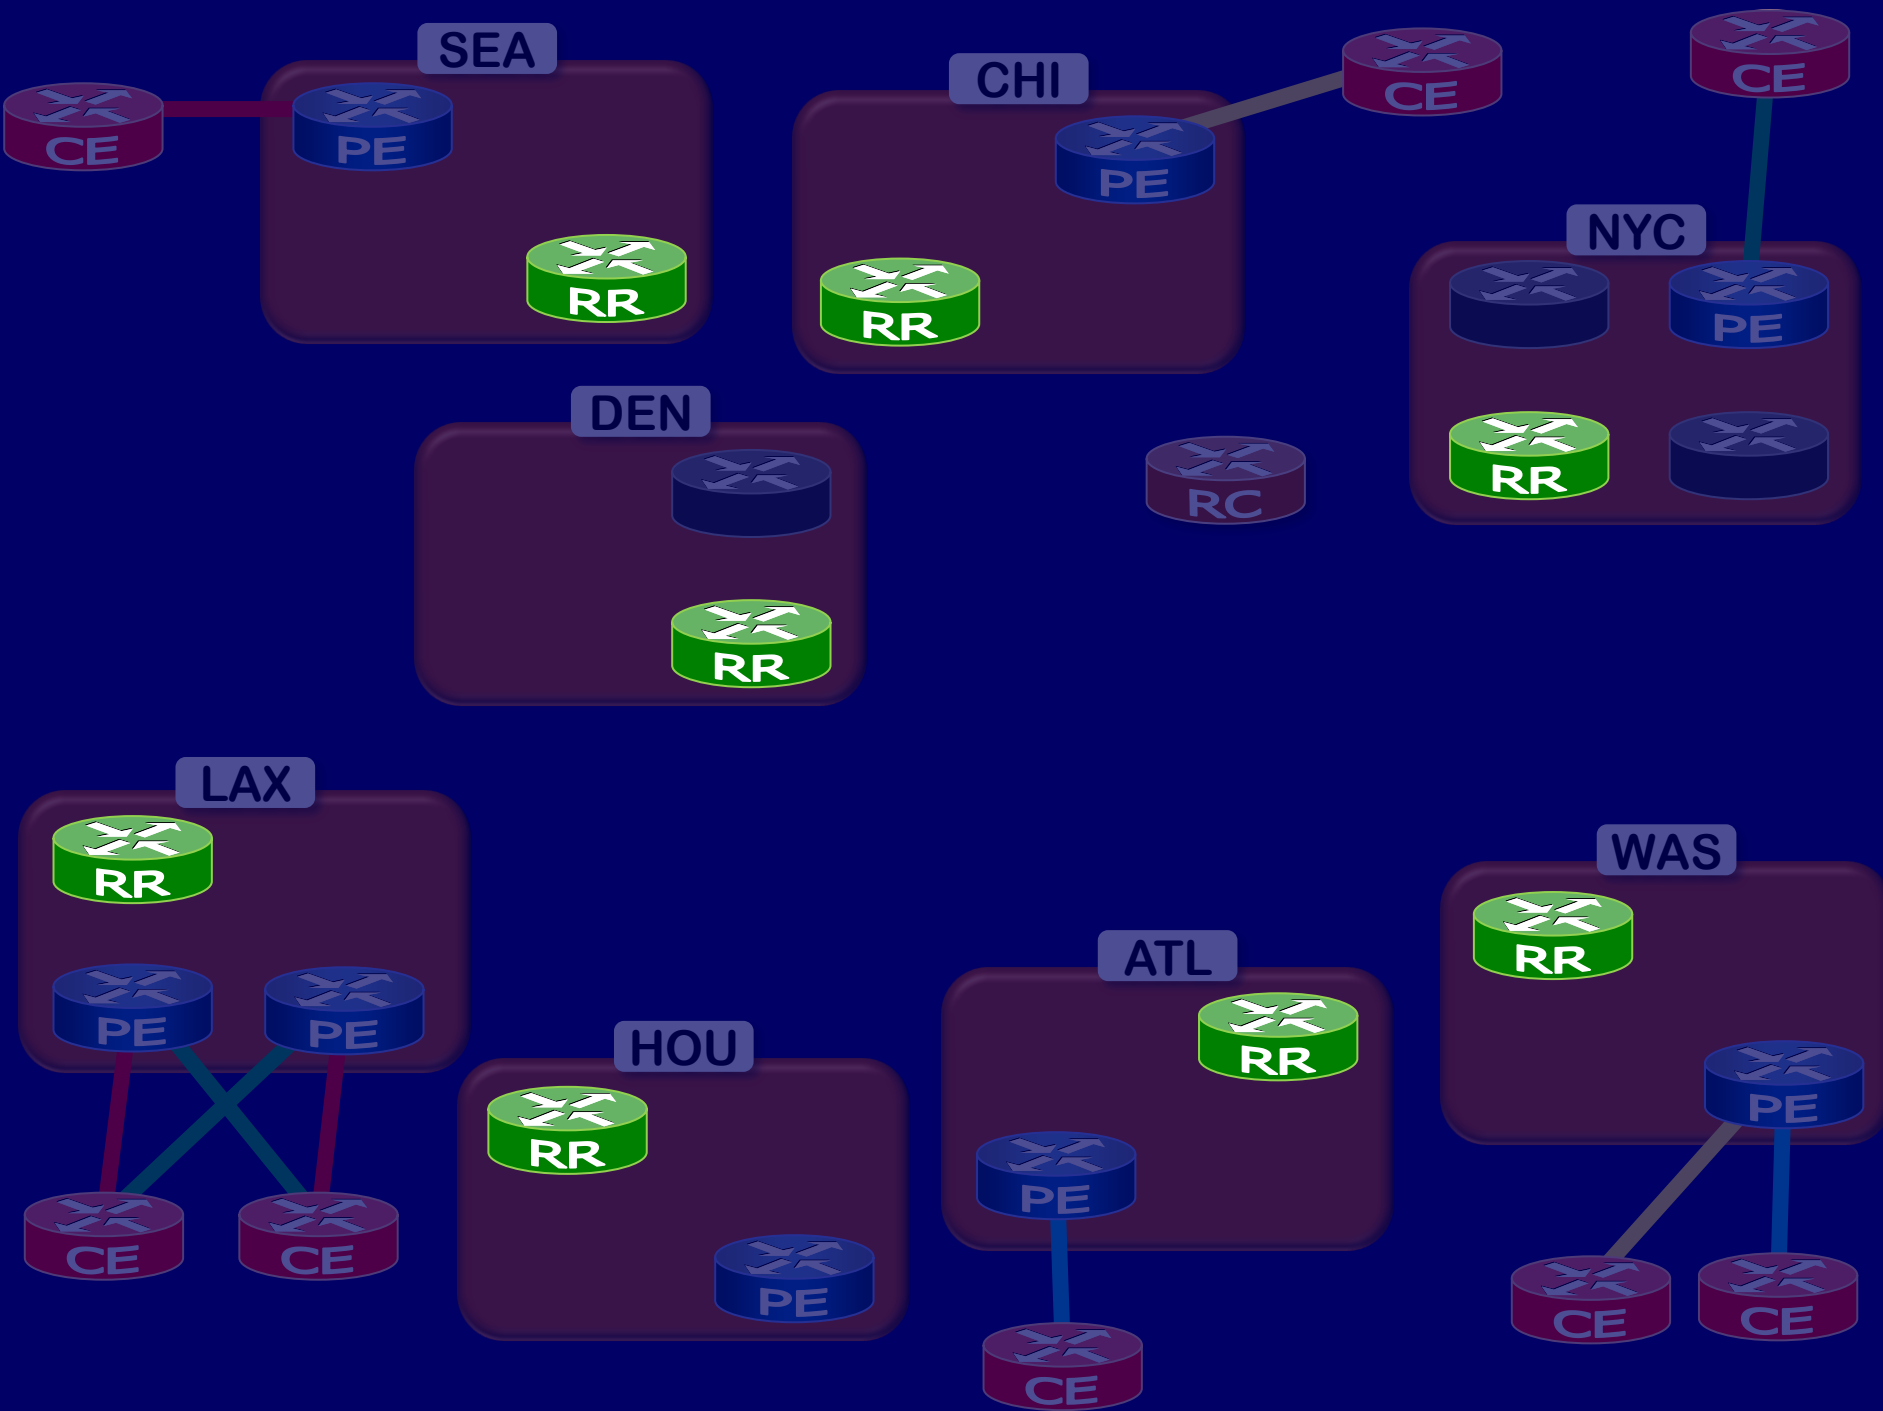







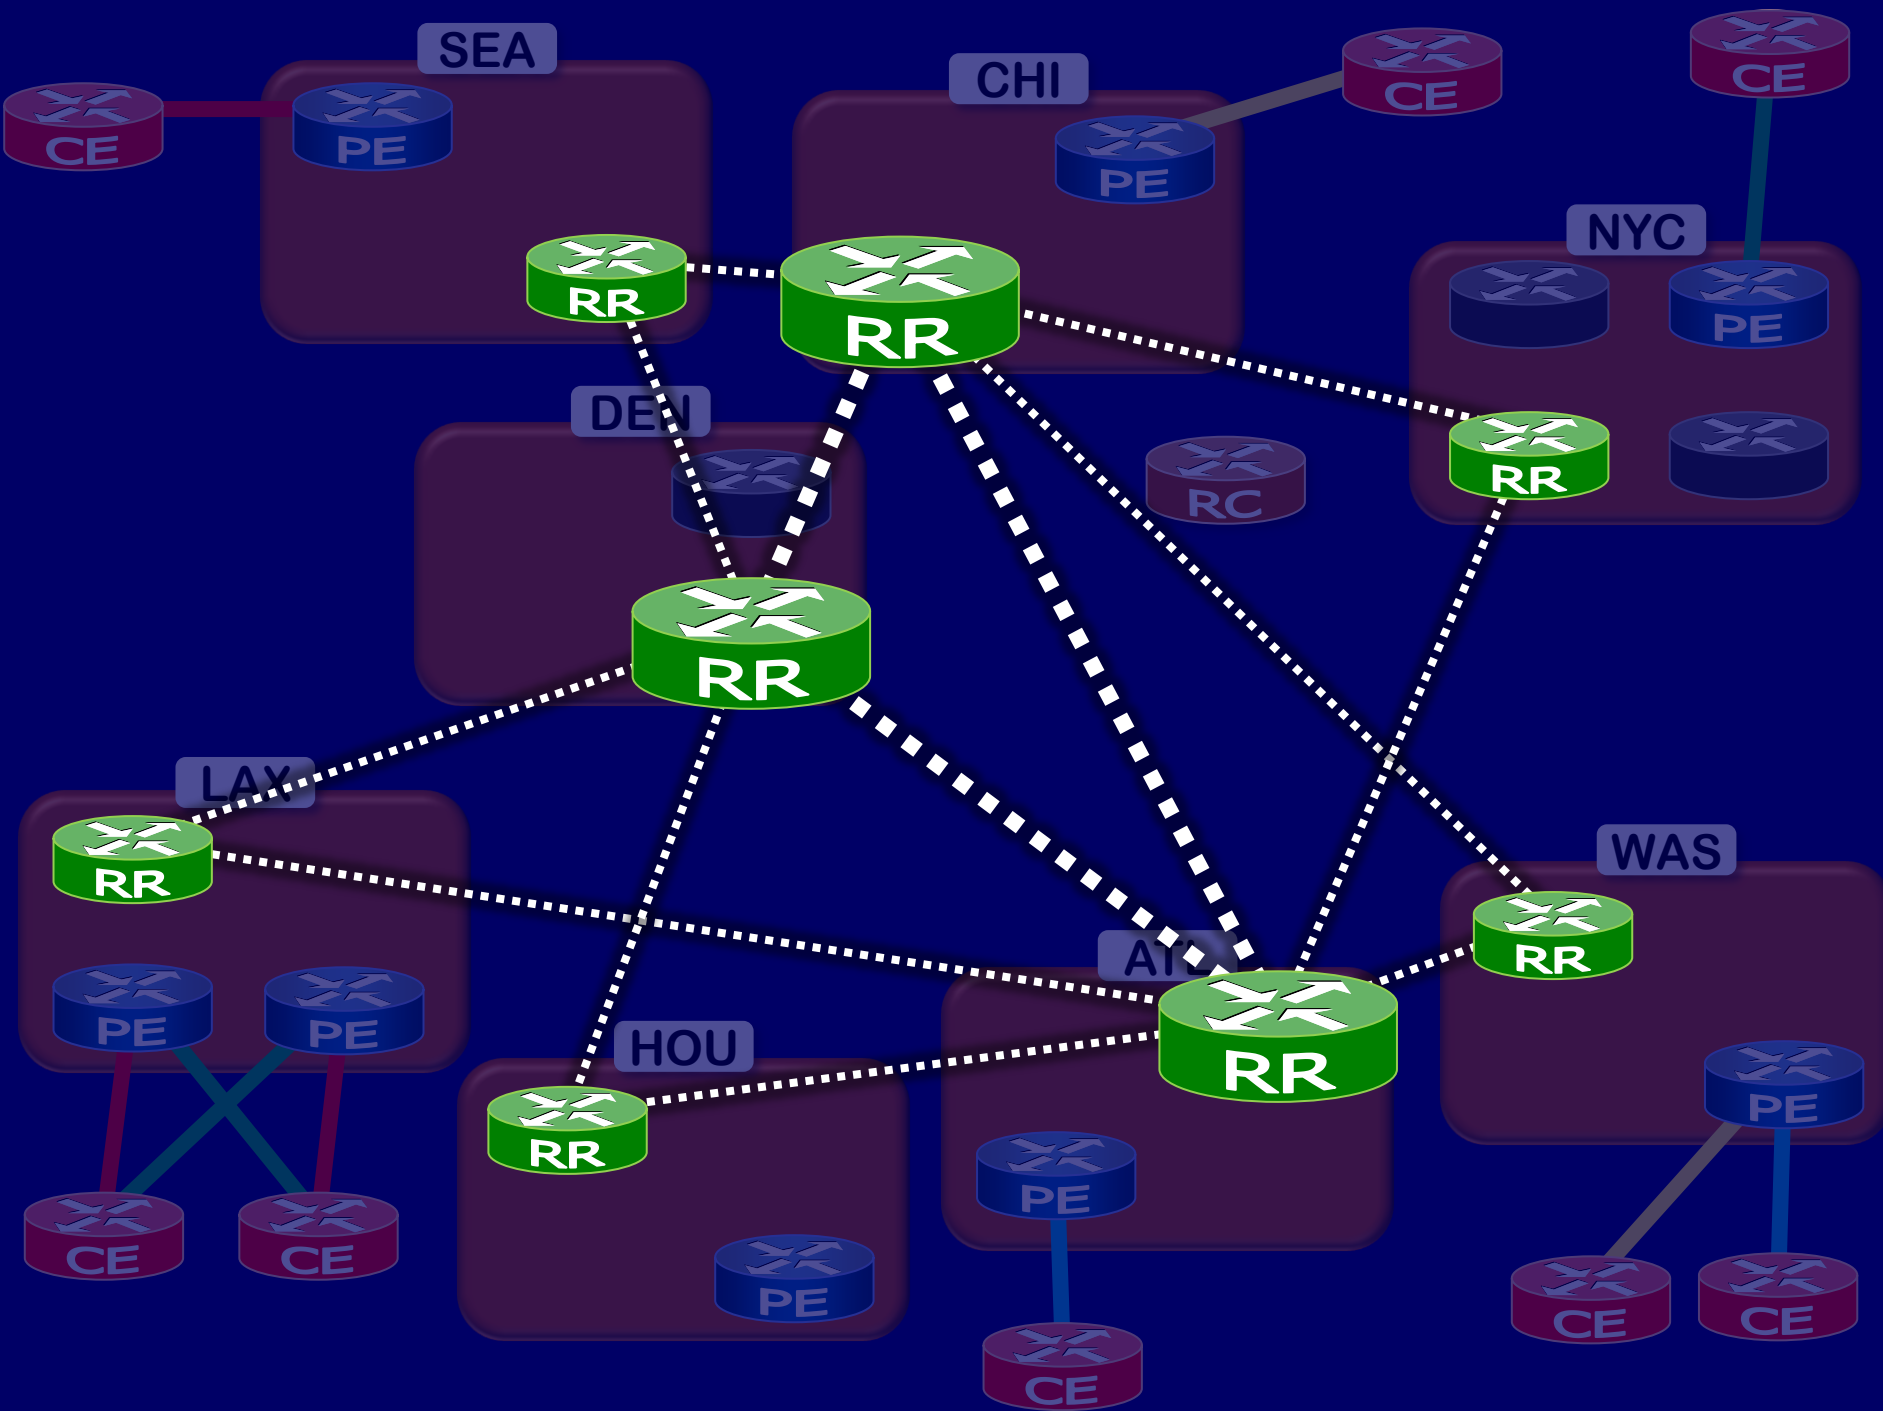





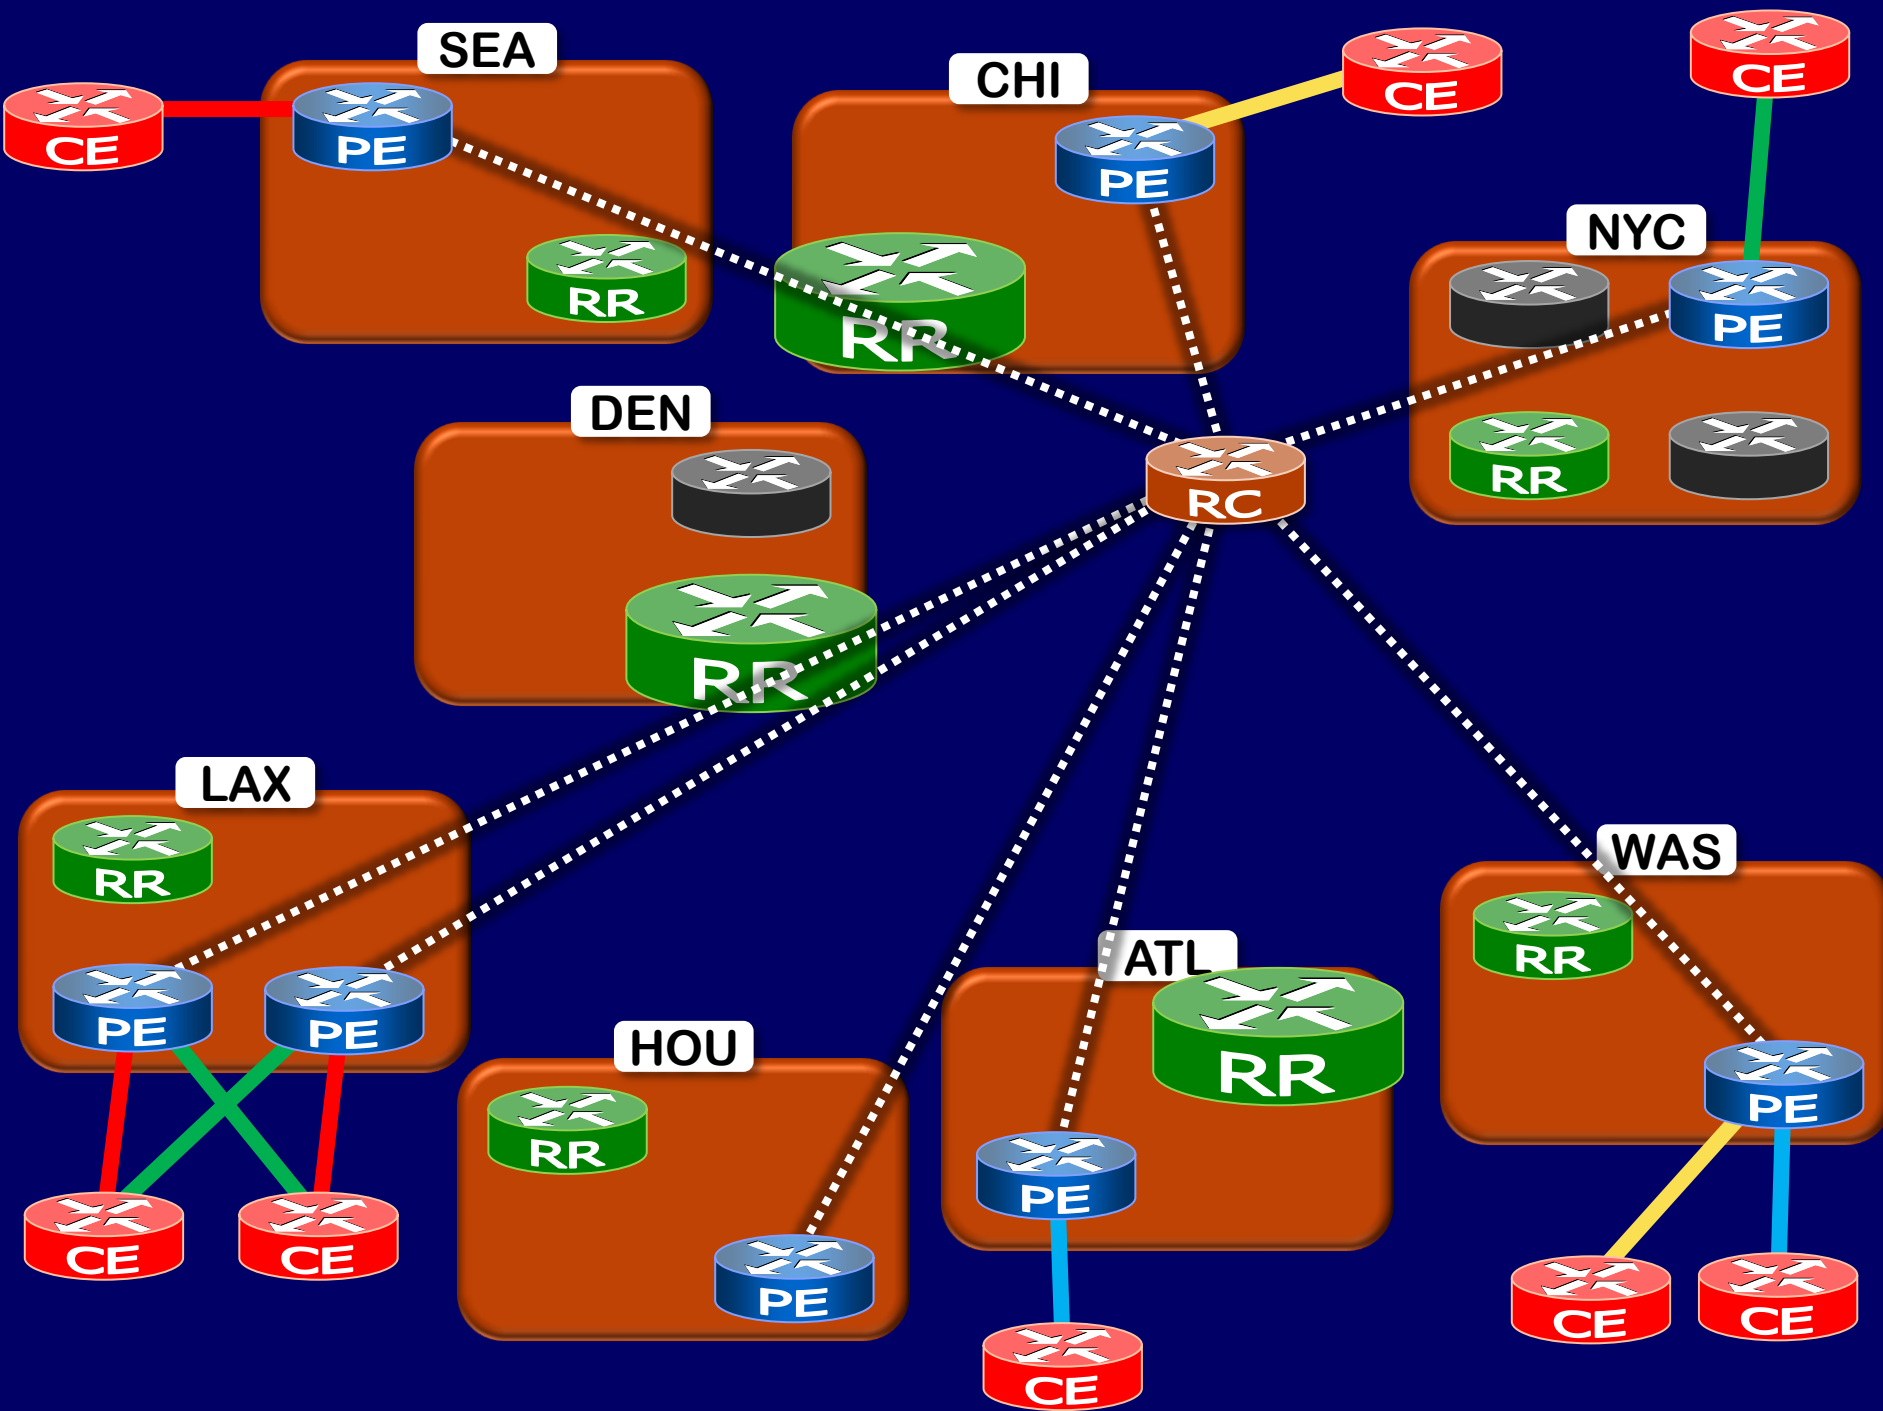















# Experiments

## ✦ Injected events:

- (De+re)activation of customer sites
- RT change
- (De+re)activation of multihoming
- Local preference change in a multihoming configuration



# Experiments

## ✦ Injected events:

- (De+re)activation of customer sites
- RT change
- (De+re)activation of multihoming
- Local preference change in a multihoming configuration

## ✦ Timing

- Random order
- Varying rate ( [1/hr...100/min] )



# Experiments

## ✦ Injected events:


- (De+re)activation of customer sites
- RT change
- (De+re)activation of multihoming
- Local preference change in a multihoming configuration


## ✦ Timing


- Random order
  - Varying rate ( [1/hr...100/min] )
- ✦ > 150,000 collected BGP updates
- Processing time: < 20s, without optimizations







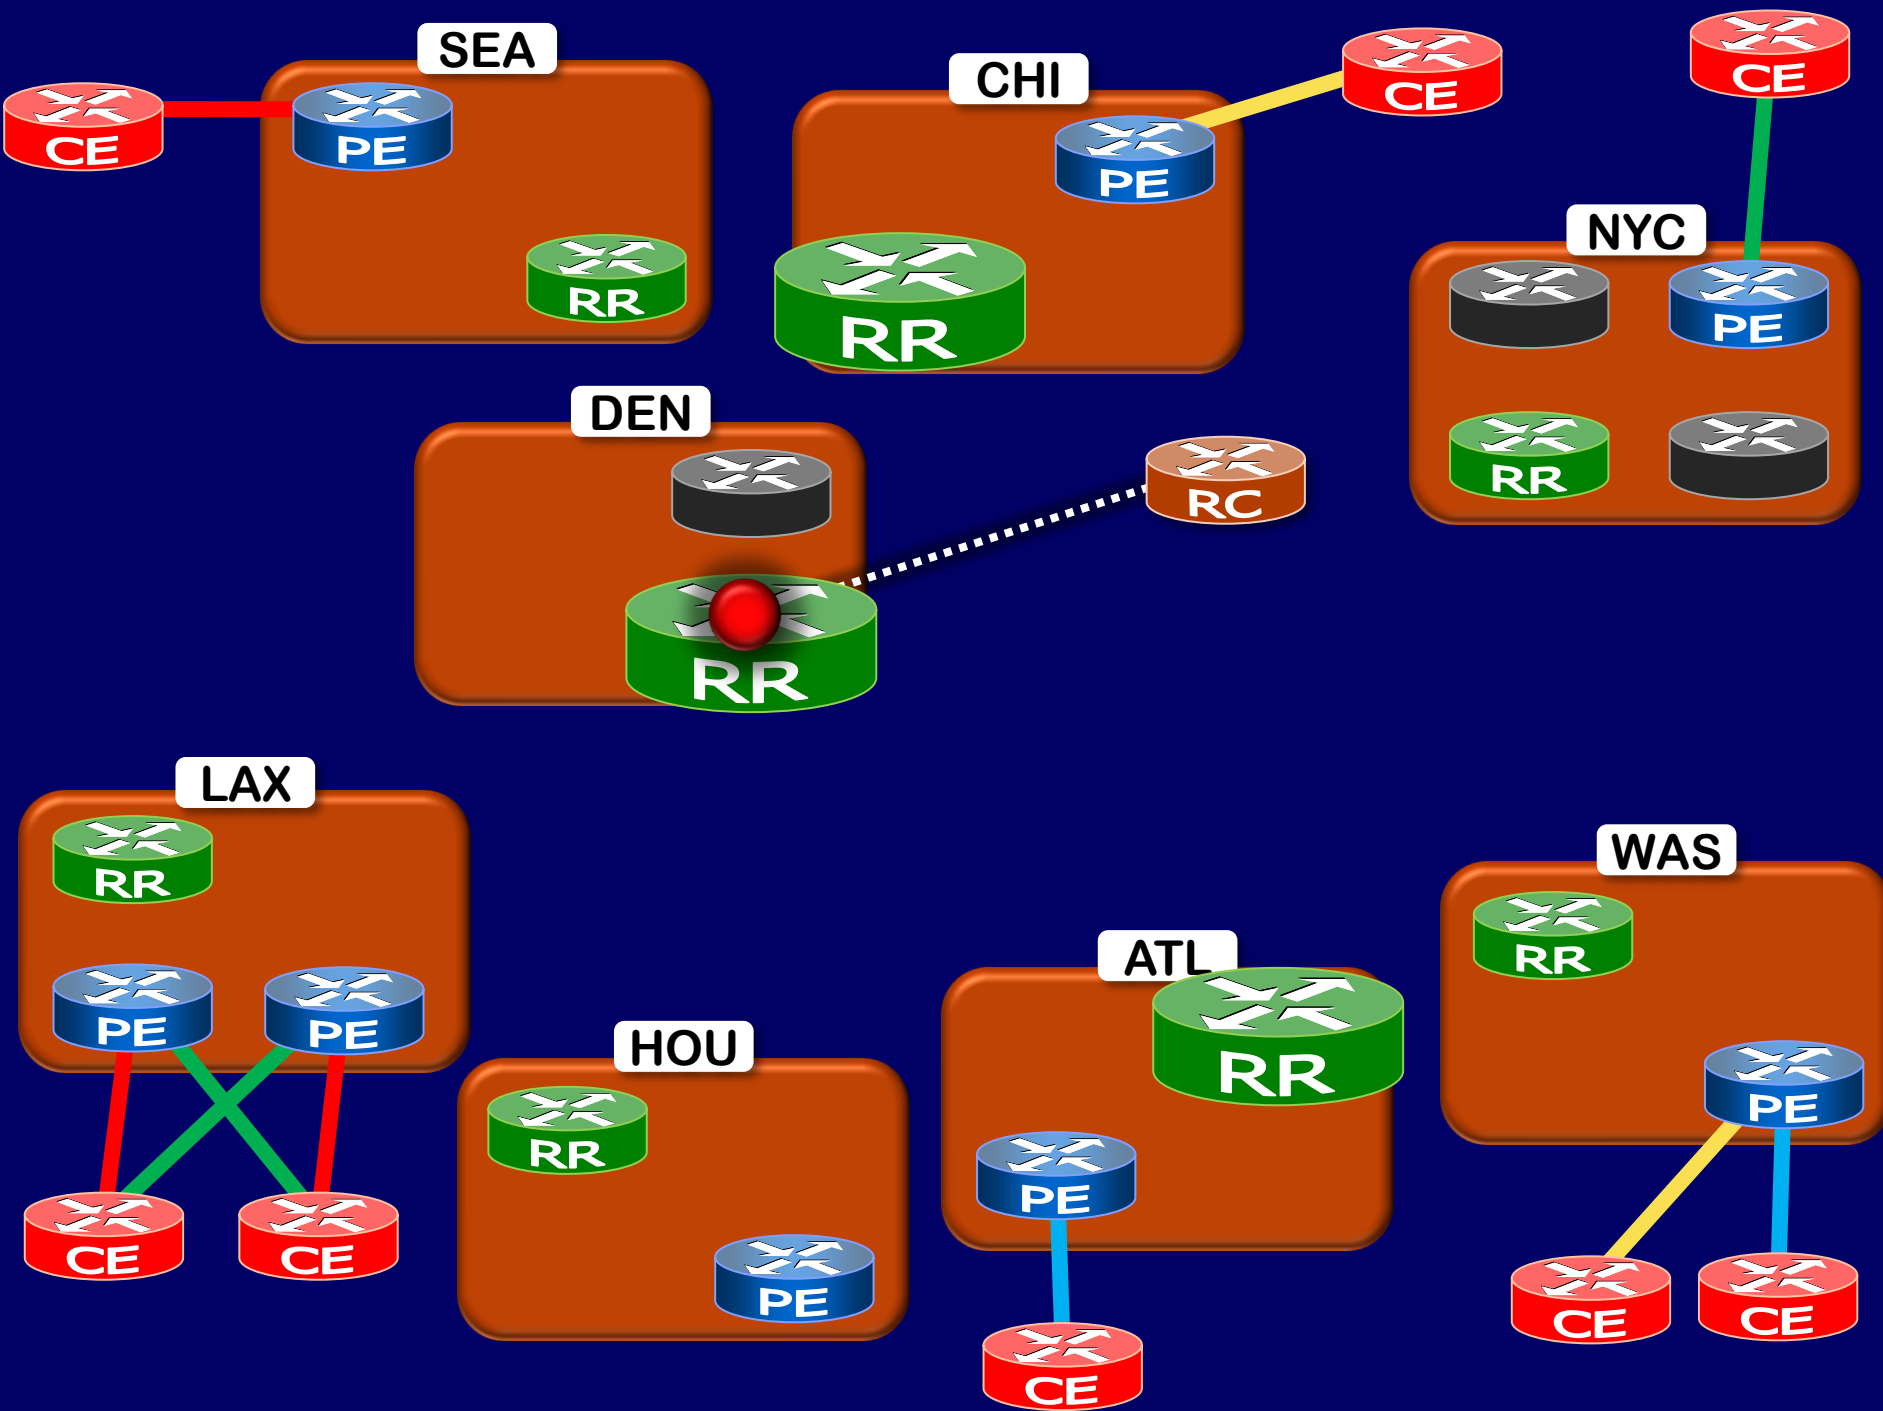






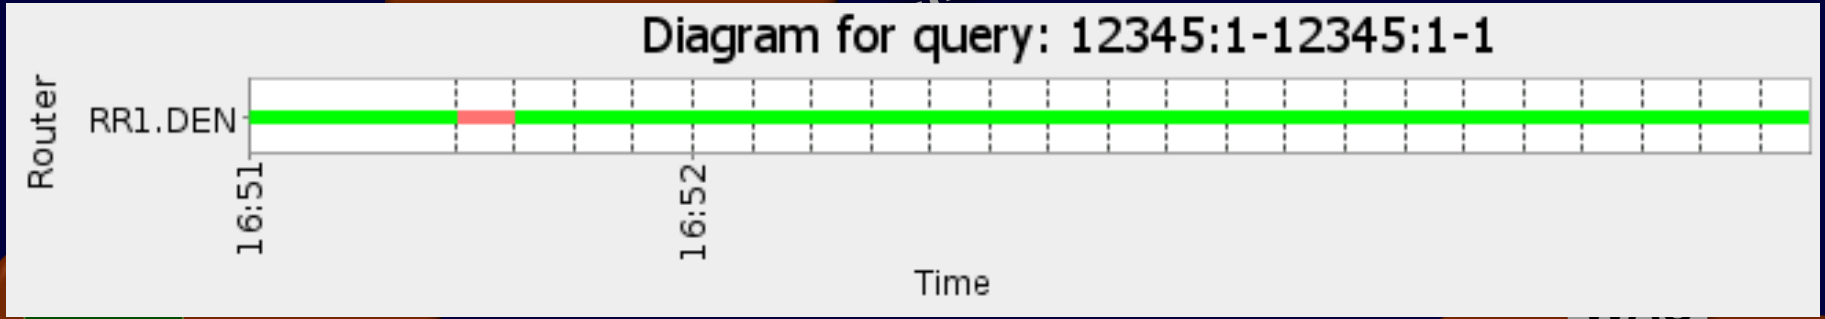








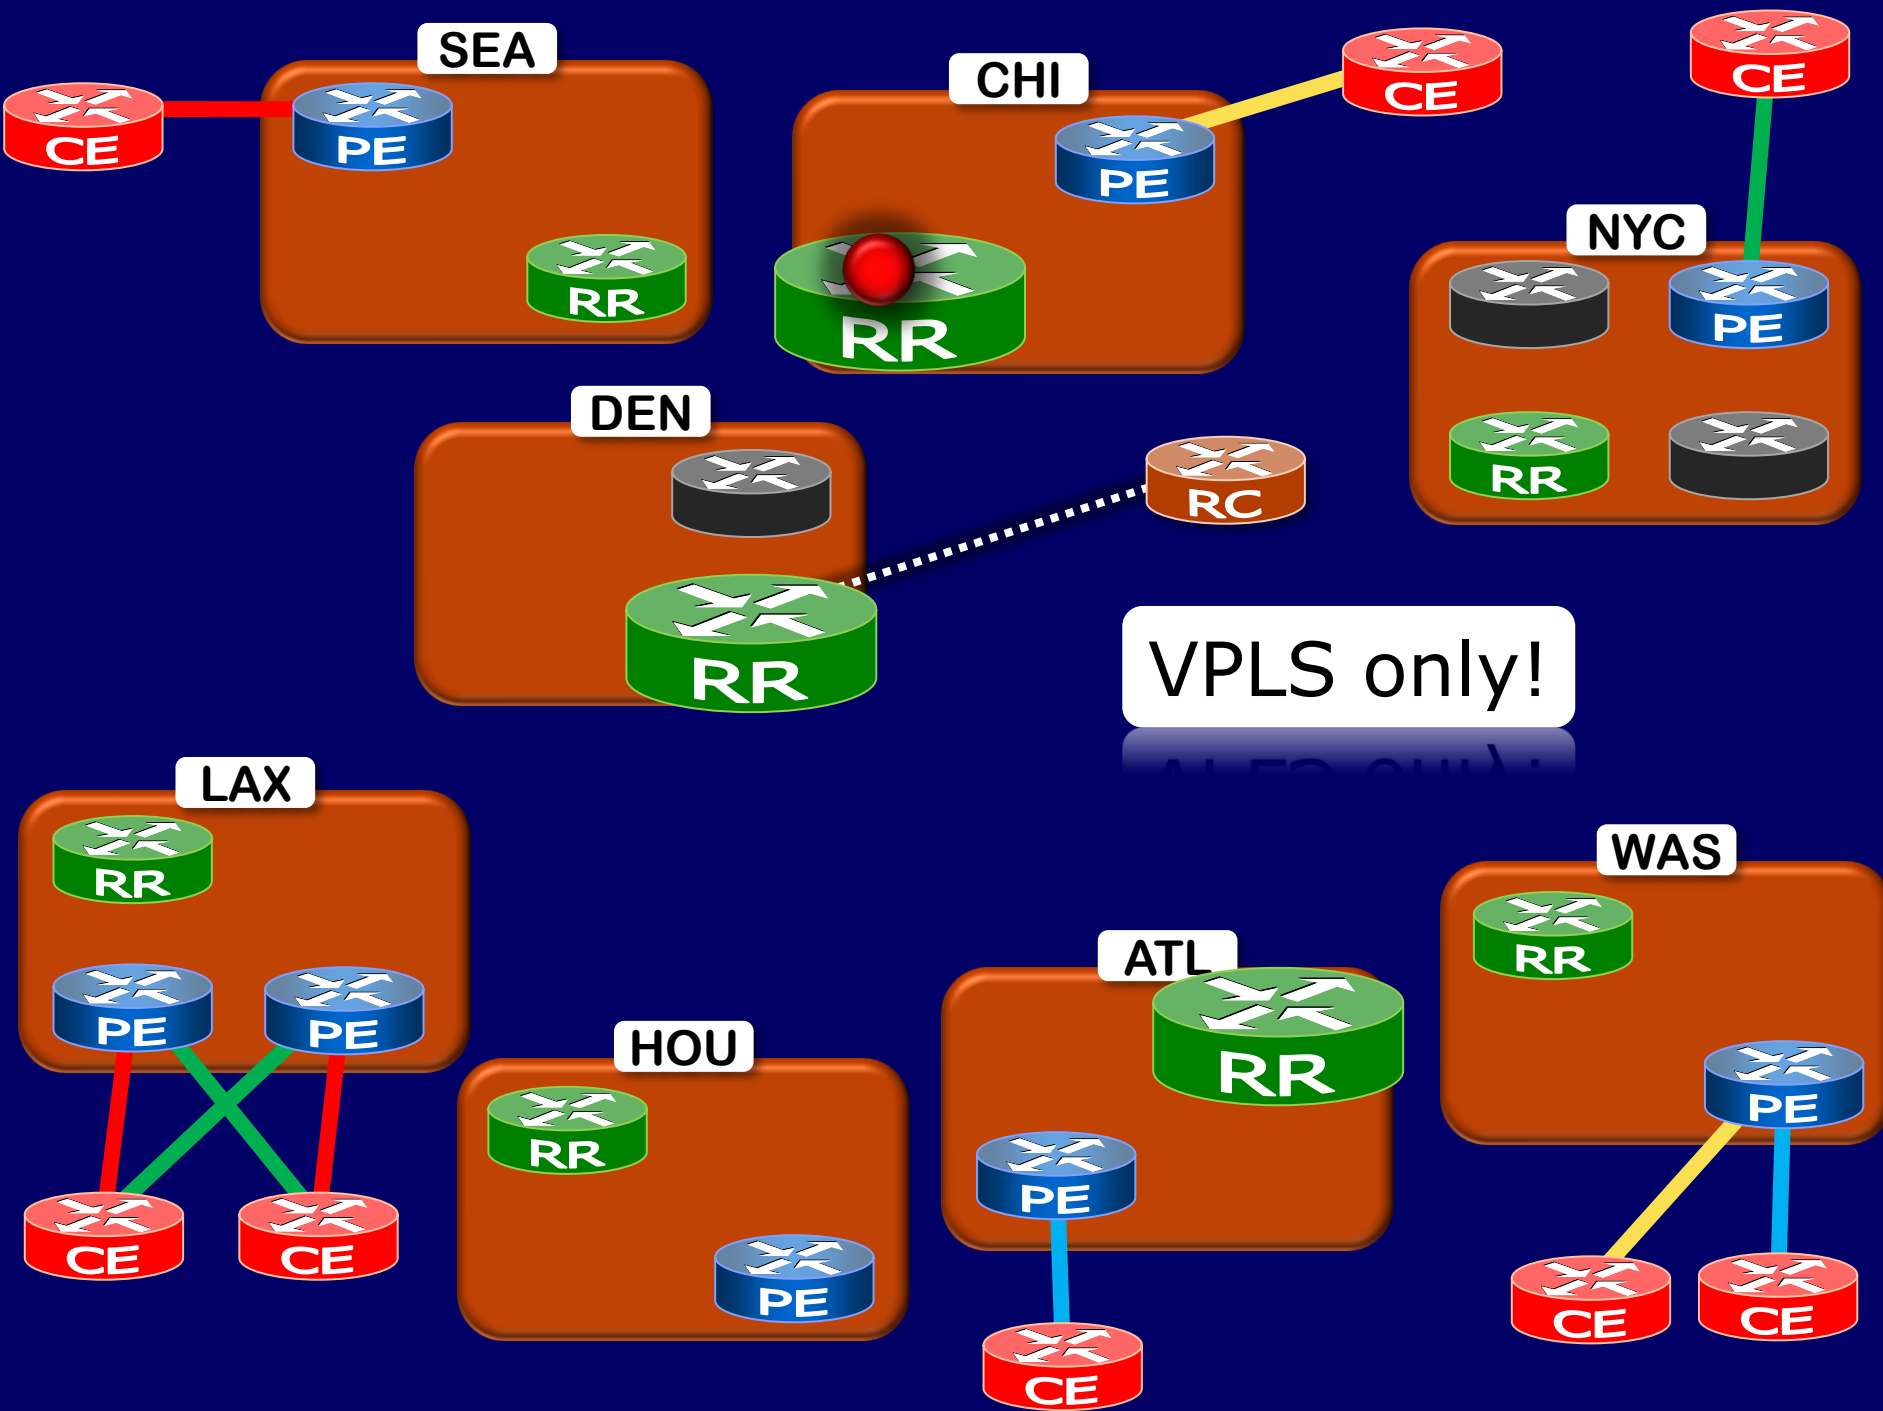












# The Oscillation Problem

- ◆ Did not affect forwarding



# The Oscillation Problem

- ◆ Did not affect forwarding
- ◆ Investigation with Juniper



# The Oscillation Problem



- ◆ Did not affect forwarding
- ◆ Investigation with Juniper
- ◆ Best route selection in VPLS only considered
  - VPLS control flags
  - site preference
  - PE router ID

# The Oscillation Problem



- ◆ Did not affect forwarding
- ◆ Investigation with Juniper
- ◆ Best route selection in VPLS only considered
  - VPLS control flags
  - site preference
  - PE router ID
  - ties were broken on most recent announcement (could carry updated labels)

# The Oscillation Problem




- ◆ Did not affect forwarding
- ◆ Investigation with Juniper
- ◆ Best route selection in VPLS only considered
  - VPLS control flags
  - site preference
  - PE router ID
  - ties were broken on most recent announcement (could carry updated labels)
- ◆ DISAGREE [Griffin et al. 02]

# The Oscillation Problem




- ◆ Did not affect forwarding
- ◆ Investigation with Juniper
- ◆ Best route selection in VPLS only considered
  - VPLS control flags
  - site preference
  - PE router ID
  - ties were broken on most recent announcement (could carry updated labels)
- ◆ DISAGREE [Griffin et al. 02]
- ◆ Fix (being) released






# Wrapping Up




# Wrapping Up



- ◆ A monitoring methodology


# Wrapping Up



- ◆ A monitoring methodology

Effects

# Wrapping Up




◆ A monitoring methodology

Effects

Signaling

# Wrapping Up




◆ A monitoring methodology

Effects

Signaling

MPLS+VPLS

# Wrapping Up



◆ A monitoring methodology


Effects

Signaling

MPLS+VPLS

Visualization

# Wrapping Up



◆ A monitoring methodology

Effects


Signaling

MPLS+VPLS

Visualization

Operation

# Wrapping Up



◆ A monitoring methodology

Effects

Signaling

MPLS+VPLS


Visualization

Operation

Reconfiguration



# Wrapping Up



◆ A monitoring methodology

Effects

Signaling

MPLS+VPLS


Visualization

Operation

Reconfiguration

Troubleshooting

# Wrapping Up



## ◆ A monitoring methodology

Effects

Signaling

MPLS+VPLS

Visualization


Operation

Reconfiguration

Troubleshooting

## ◆ Discussion on scalability vs visibility

# Wrapping Up



## ◆ A monitoring methodology

Effects

Signaling

MPLS+VPLS

Visualization

Operation


Reconfiguration

Troubleshooting

◆ Discussion on scalability vs visibility

◆ Architecture & prototype implementation

# Wrapping Up



## ◆ A monitoring methodology

Effects

Signaling

MPLS+VPLS

Visualization

Operation

Reconfiguration

Troubleshooting

- ◆ Discussion on scalability vs visibility
- ◆ Architecture & prototype implementation
- ◆ Experimentation revealing routing anomaly

# Future Work/Open Problems

- ◆ Monitor other protocols/kinds of information
- ◆ Collect non-best routes
- ◆ Improve the visualization
- ◆ Trigger alarms
- ◆ Improve inference of event causes



Acknowledgments to

JUNIPER<sup>®</sup>  
NETWORKS

Thank you

