https://t.me/learningnets

Deploying Cisco Service Provider Advanced Network Routing (SPADVROUTE)

v1.2

© 2014 Cisco and/or its affiliates. All rights reserved.

Learner Skills and Knowledge

- Students in this course should have attended the following classes or obtained an equivalent level of training:
  - *Building Cisco Service Provider Next-Generation Networks, Part 1* (SPNGN1) v1.2
  - *Building Cisco Service Provider Next-Generation Networks, Part 2* (SPNGN2) v1.2
  - *Deploying Cisco Service Provider Network Routing* (SPROUTE) v1.2

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Course Goal

Upon completing this course, you will be able to:

- Train service provider network professionals on the techniques to plan, implement, and monitor a scalable IP routing protocol

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Course Flow

	AM	PM
Day 1	Course Introduction Module 1: Service Provider Connectivity with BGP	Module 2: Scale Service Provider Networks
Day 2	Module 3: Secure and Optimize BGP	Module 3 (Cont.) Module 4: Multicast Overview
Day 3	Module 4 (Cont.)	Module 5: Intradomain and Interdomain Multicast Routing
Day 4	Module 5 (Cont.)	Module 5 (Cont.)
Day 5	Module 6: Service Provider IPv6 Transition Implementations	Module 6 (Cont.)

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Career Certifications

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Career Certifications

Expand Your Professional Options, Advance Your Career

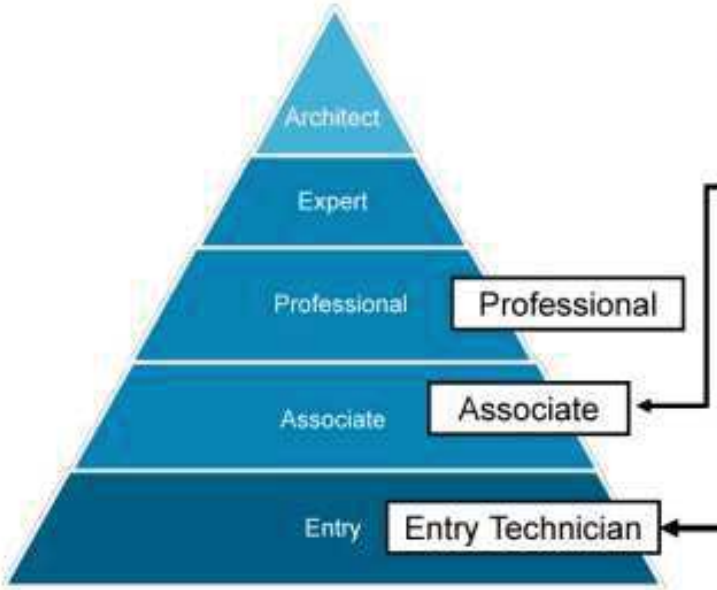

Cisco Certified Network Associate

Recommended Training Through Cisco Learning Partners

- Interconnecting Cisco Networking Devices Part 1
- Interconnecting Cisco Networking Devices Part 2

Cisco Certified Entry Network Technician

Recommended Training Through Cisco Learning Partners

- Interconnecting Cisco Networking Devices Part 1

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

CCNA Prep Center

https://t.me/learningnets

HOME
LEARNING AND EVENTS
CAREER CERTIFICATIONS AND PATHS
CERTIFICATION RESOURCES
CCNA Prep Center

Career Certifications & Paths

CCNA Prep Center

Get Ready for Your CCNA Certification

CCNA Prep Center provides certification candidates with resources including practice questions, labs, simulations, instructional videos, tips, advice, success stories and peer-discussion forums. The CCNA Prep Center also allows learners easy access to information and formal training from Cisco Learning Partners.

Search CCNA Prep Center

CCNA Prep Center

GO

Advanced Search

Print | Feedback | Help

Overview | CCNA Paths | Exam Study | CCNA TV | Discussions | Additional Information

Planning Your CCNA Preparation

CCNA Paths
CCNA Certification Paths, My Certification History, Pre-Assessment Exam and Exam Registration.

Exam Study
Study Tips, Practice Questions, Remote Labs, Simulations and more.

CCNA TV
Regular broadcasts with experts discussing CCNA topics and answering your questions.

Discussions
Engage in discussions with Cisco experts on technical questions or program issues.

Additional Information
Articles and information about certifications and the job market.

Featured Content

Cisco Subnet Game Beta
Sharpen your subnetting skills as you set up subnetworks in the mysterious Area 51 Military Air Base.
Play the Game
About the Game

Related Tools

Certifications Online Support
Certifications Tracking System Tool
Global Learning Partner Locator
Cisco Learning Connection
Products & Services Tool Index

Related Links

CCNP Prep Center
Certifications Community
Networking Academy

Learning and Events
Career Certification and Paths
About Learning Partners

About Cisco
Cisco Press

Your Opinion Counts

Please take a moment to tell us your thoughts about the CCNA Prep Center.

© 2014 Cisco and/or its affiliates. All rights reserved.

Learner Introductions

- Your name
- Your company
- Job responsibilities
- Skills and knowledge
- Brief history
- Objective

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Icons and Symbols

Router

Wireless Router

Secure Router

Firewall

Home Office

Workgroup Switch

Access Point

Cisco IOS XR Router

Cisco IOS XE Router

Wireless Connectivity

Line: Serial

Line: Ethernet

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

https://t.me/learningnets

https://t.me/learningnets

Defining Customer-to-Provider Connectivity Requirements

Service Provider Connectivity with BGP

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco IP NGN Architecture

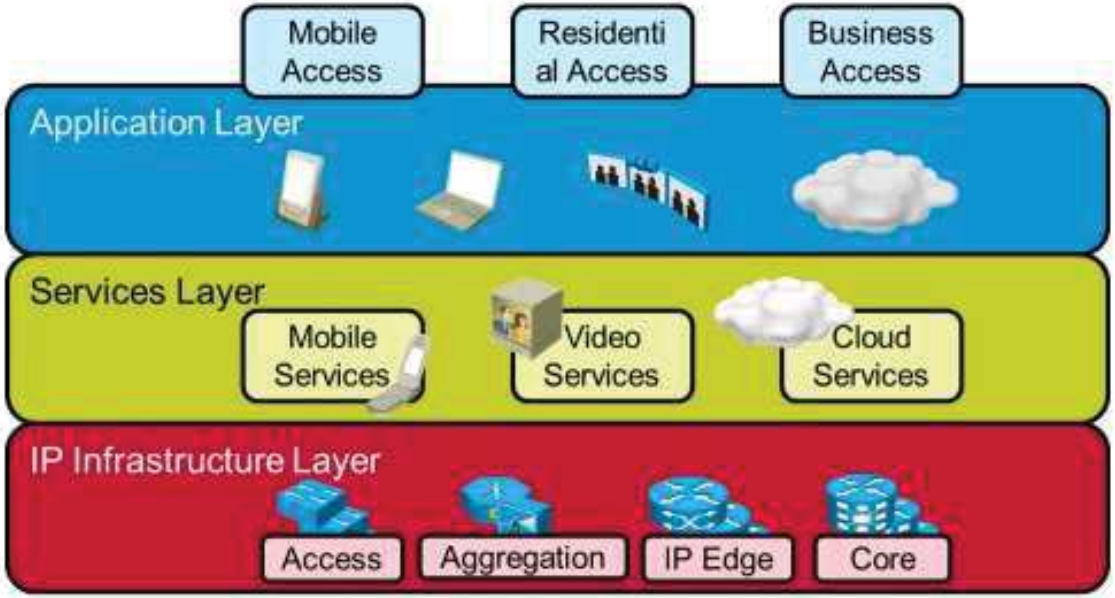

- The Cisco IP NGN is a next-generation service provider infrastructure for video, mobile, and cloud or managed services.
- The Cisco IP NGN provides an all-IP network for services and applications, regardless of access type.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco IP NGN Architecture (Cont.)

Cisco IP NGN Infrastructure Layer

Customer-to-provider connectivity focuses on the:

- IP infrastructure layer of the Cisco IP NGN.
- Service provider edge and customer devices.

https://t.me/learningnets

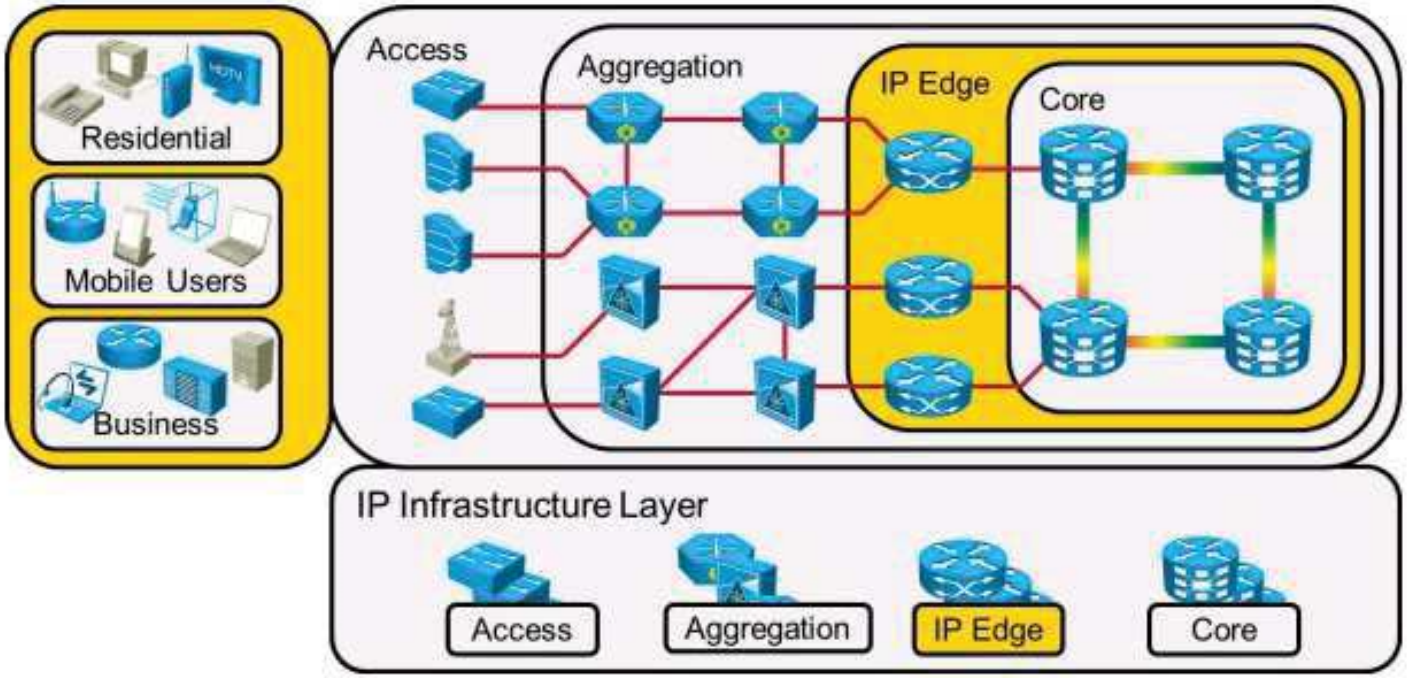

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-to-Service Provider Connectivity Types

Connectivity types:

- **Single-homed customers:** A customer is connected to a single service provider using one link.
- **Dual-attached customers:** A customer is connected to a single service provider using two links.
- **Multihomed customers:** A customer is connected to two service providers, using one link for each service provider.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer Connectivity

- The simplest setup is a single link between the customer network and the service provider.
- There is no redundancy for link, equipment, or service provider failure.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customer Connectivity

Dual-attached customer connectivity characteristics:

- Customers wanting increased redundancy install several physical links to the service provider.
- Redundant links are used as follows:
  - Primary and backup.
  - For load sharing.
- Redundancy is for link or equipment failure.
- There is no redundancy for service provider failure.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Multihomed Customer Connectivity

In a multihomed connection:

- Customers with maximum redundancy requirements install physical links to multiple service providers.
- Redundant links are used as follows:
  - Primary and backup.
  - Primary and backup with direct traffic.
  - For load sharing.
- There is redundancy for link, equipment, or service provider failure.
- Multihomed customers should connect to independent service providers.

https://t.me/learningnets

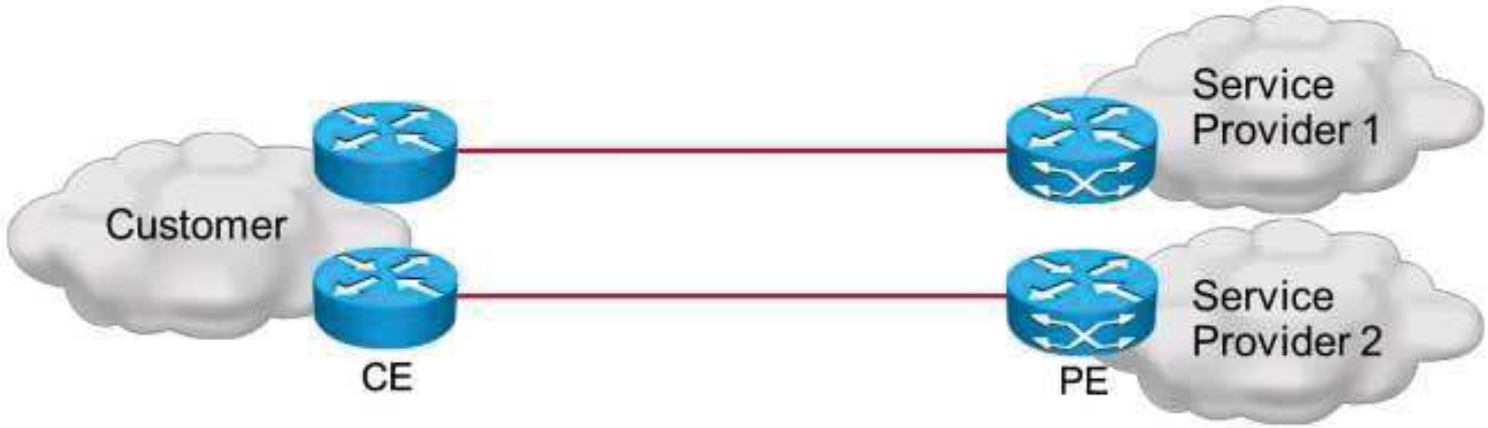

© 2014 Cisco and/or its affiliates. All rights reserved.

Routing Schemes

Customer to service provider routing scheme:

- Static or dynamic routing can be used between a customer and a service provider.
- BGP is the only acceptable dynamic routing protocol.
- Because of its lower complexity, static routing is preferred where possible.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer Routing Schemes

- Single-homed customers typically do not require BGP:
  - The static route for the provider-assigned address space of the customer is on the PE router.
  - The static default route is on the CE router.
- BGP can be used to detect link failures and trigger dial backup:
  - The service provider originates only the default route.
  - The customer originates its address space.

https://t.me/learningnets

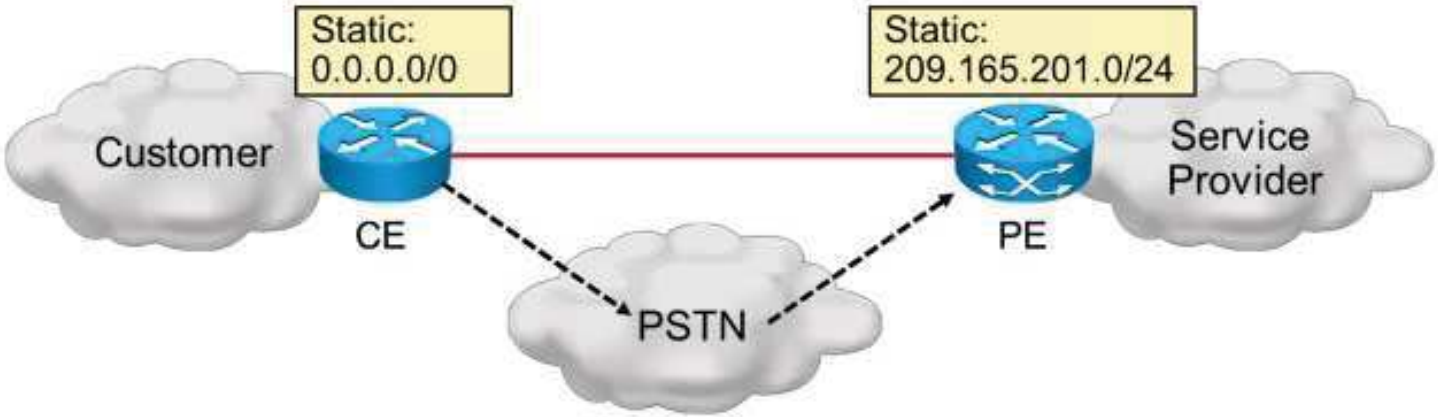

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customer Routing Schemes

Information about dual-attached customer routing schemes:

- Static routing can be used if link and service provider device failure can be detected.
- BGP between the customer and service provider is usually used in this setup:
  - The service provider originates the default route (in the primary or secondary scenario) or specific routes (in the load-balancing scenario).
  - The customer originates its address space.

https://t.me/learningnets

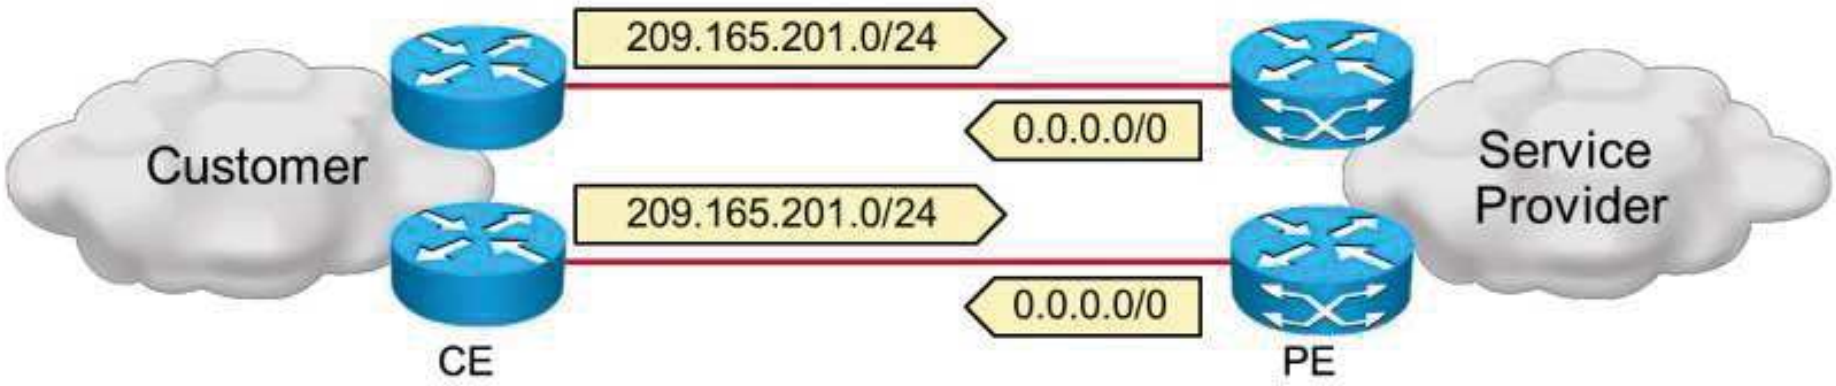

© 2014 Cisco and/or its affiliates. All rights reserved.

Multihomed Customer Routing Schemes

Multihomed routing schemes characteristics:

- Static routing is not possible.
- BGP must be used in this setup:
  - The service provider originates the default route (in the primary or secondary scenario), the default route and service provider-owned routes (in the primary or secondary with direct traffic scenario), or all routes (in the load-balancing scenario).
  - The customer originates its address space.

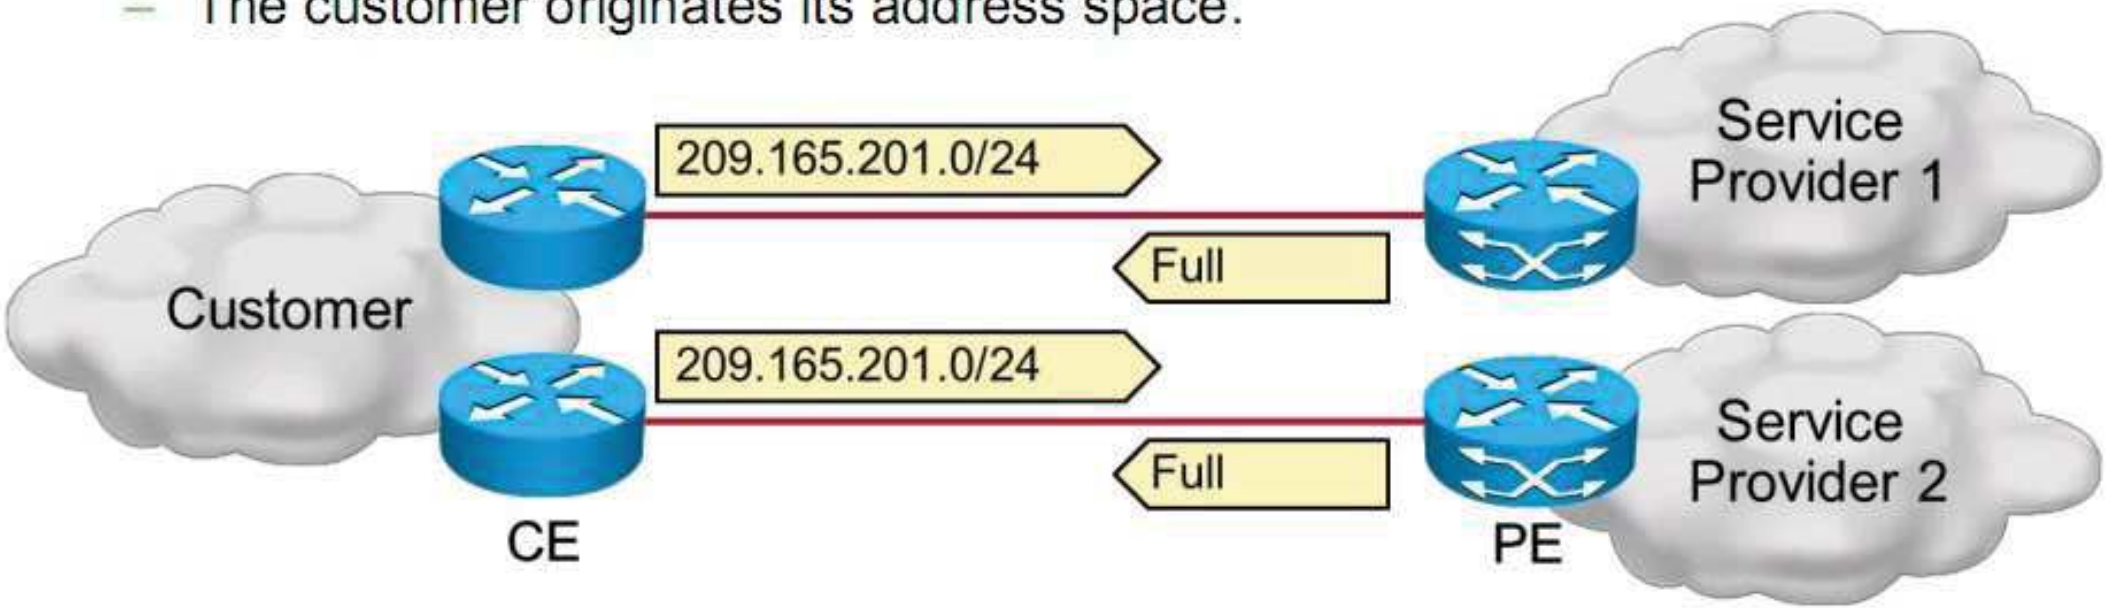

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Addressing and AS Number Allocation

Single-homed customers:

- PA address space.
- Single (link) IP address or multiple (subnet) IP addresses.
- If service provider changes, readdressing is needed.
- If BGP is desired, private AS number can be used.

Multihomed customers:

- PI address space.
- Customer has to advertise address space to both service providers—BGP.
- If service providers change, readdressing is not needed.
- Use of public AS number is recommended.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer IP Addressing Schemes

Single-homed customer public IP addressing scheme characteristics:

- In IPv4, the customer uses PAT for outbound sessions.
- In IPv4, there is limited support for inbound sessions (static PAT).
- A single IP address requires public addressing on the access link.
- If static routing is used, the AS number is not needed.

https://t.me/learningnets

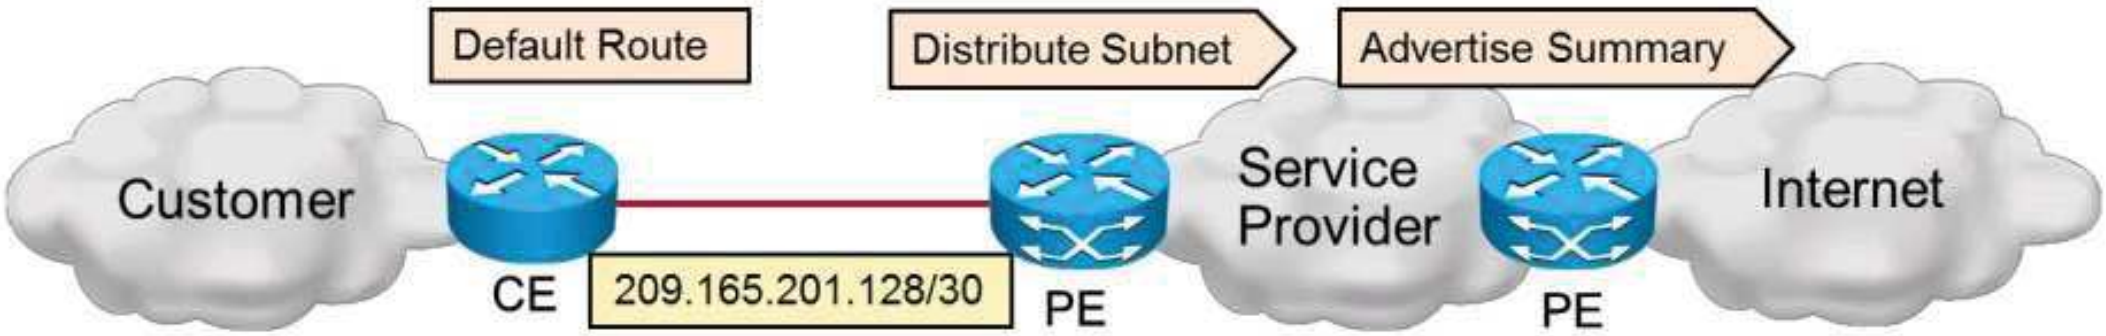

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer IP Addressing Schemes (Cont.)

Multiple IP addresses schemes:

- The public subnet is assigned to the customer from the service provider pool.
- The access link uses public addresses.
- Multiple IP addresses support customer services that require inbound sessions.
- If static routing is used, the AS number is not needed.

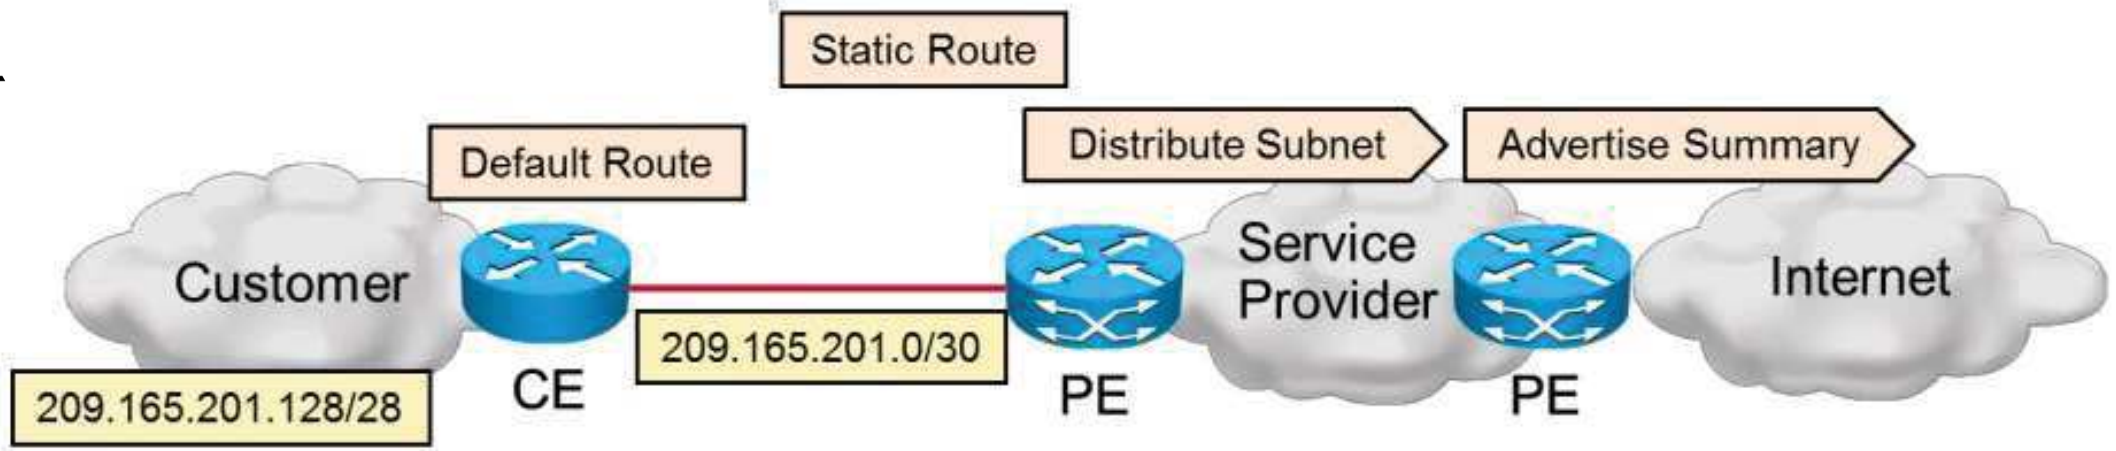

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customer IP Addressing and AS Number Schemes

Dual-attached customer public IP addressing AS number schemes characteristics:

- The public subnet is assigned to the customer from the service provider pool.
- The access link uses public addresses.
- The customer needs the AS number if BGP is used.
- The private AS number (64512-65535) can be used.

https://t.me/learningnets

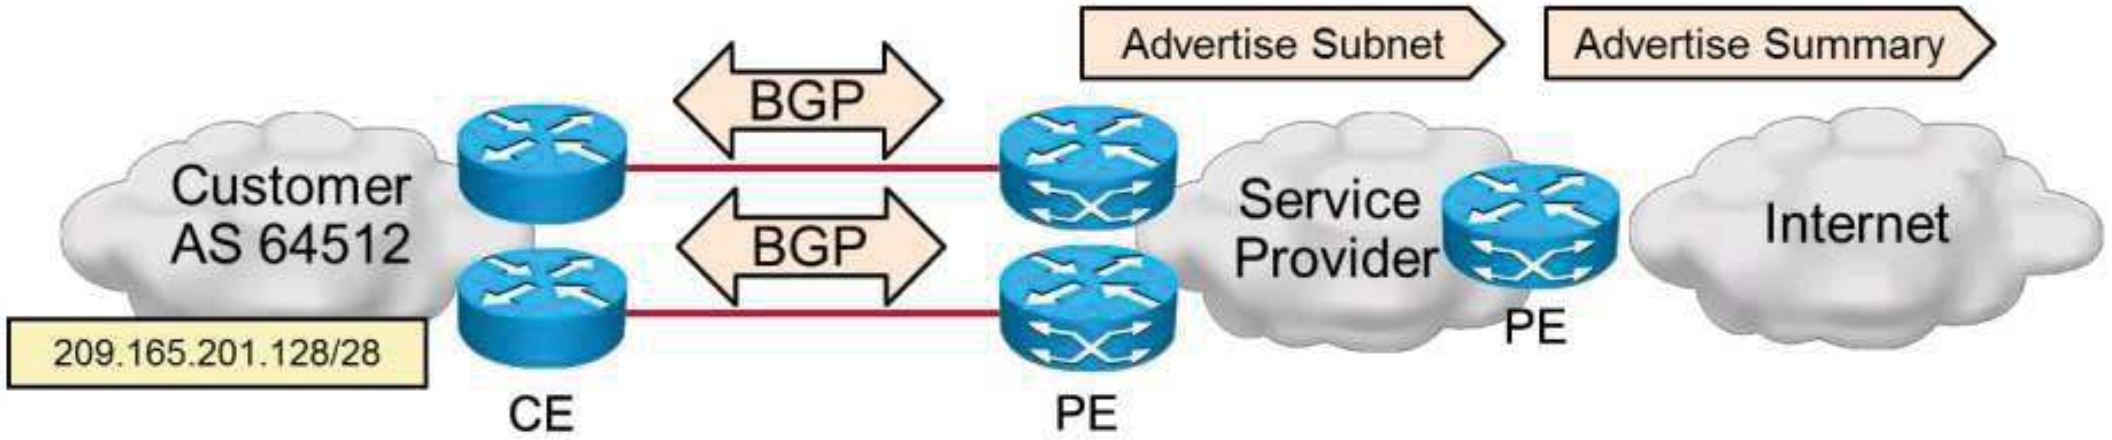

© 2014 Cisco and/or its affiliates. All rights reserved.

Multihomed Customer IP Addressing and AS Number Schemes

Multihomed Customers:

- The public subnet is assigned to the customer by a LIR.
- BGP must be used.
- The public AS number, assigned by the LIR, has to be used.
- The access link uses public addresses.

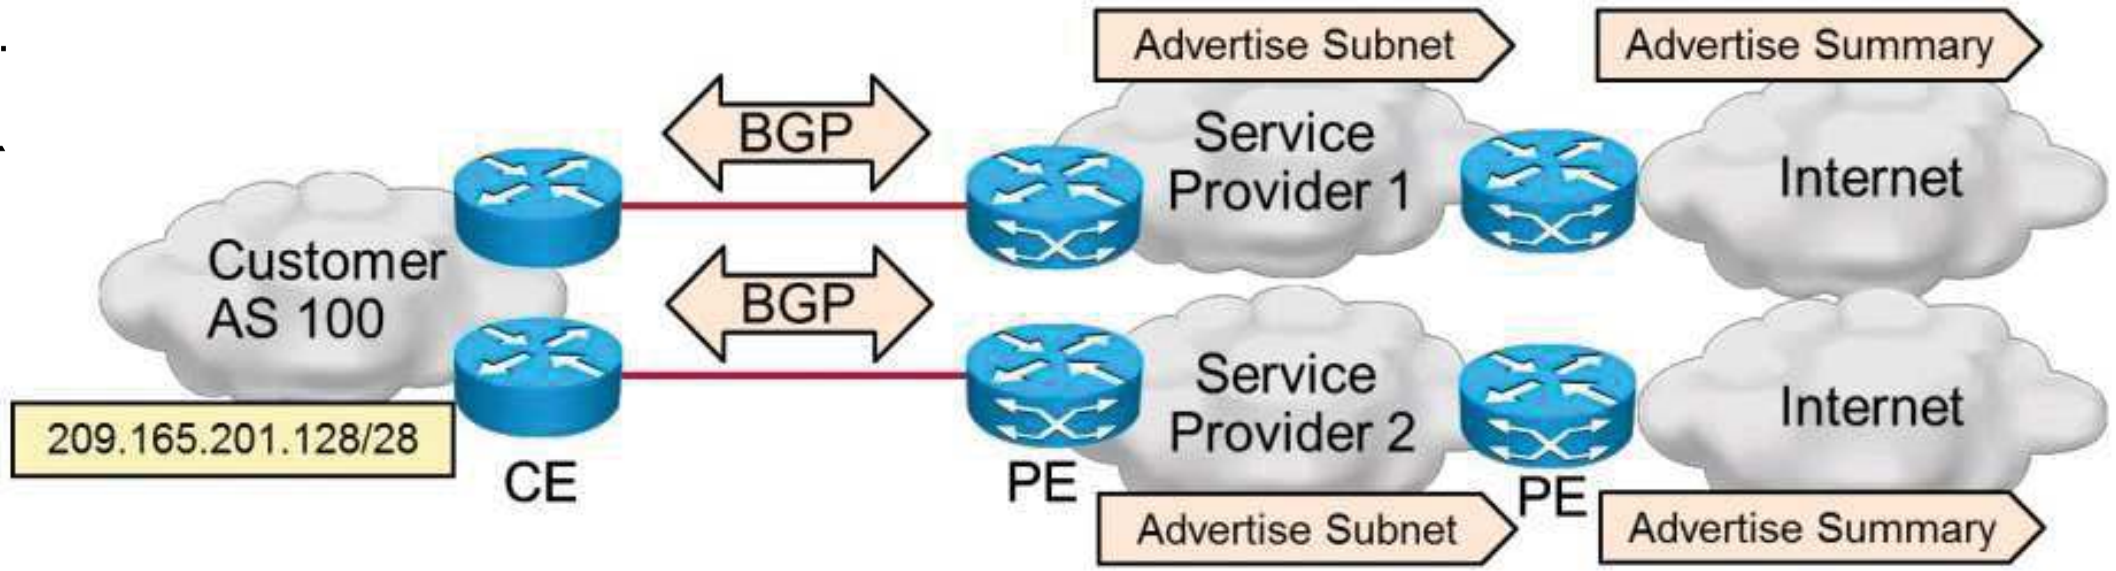

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- Customer requirements dictate the use of different connectivity types; these are:
  - Single-homed and dual-attached connection type.
  - Multihomed connection type.
- Routing that is used between a customer and service provider depends on the selected connectivity type and differs for:
  - Single-homed connection type.
  - Dual-attached connection type.
  - Multihomed connection type.
- IP addressing and AS numbering depend on the selected connectivity type:
  - Single-homed connection type requires IP address scheme.
  - Dual-attached connection type requires IP address scheme and AS number scheme.
  - Multihomed connection type requires IP address scheme and AS number scheme.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

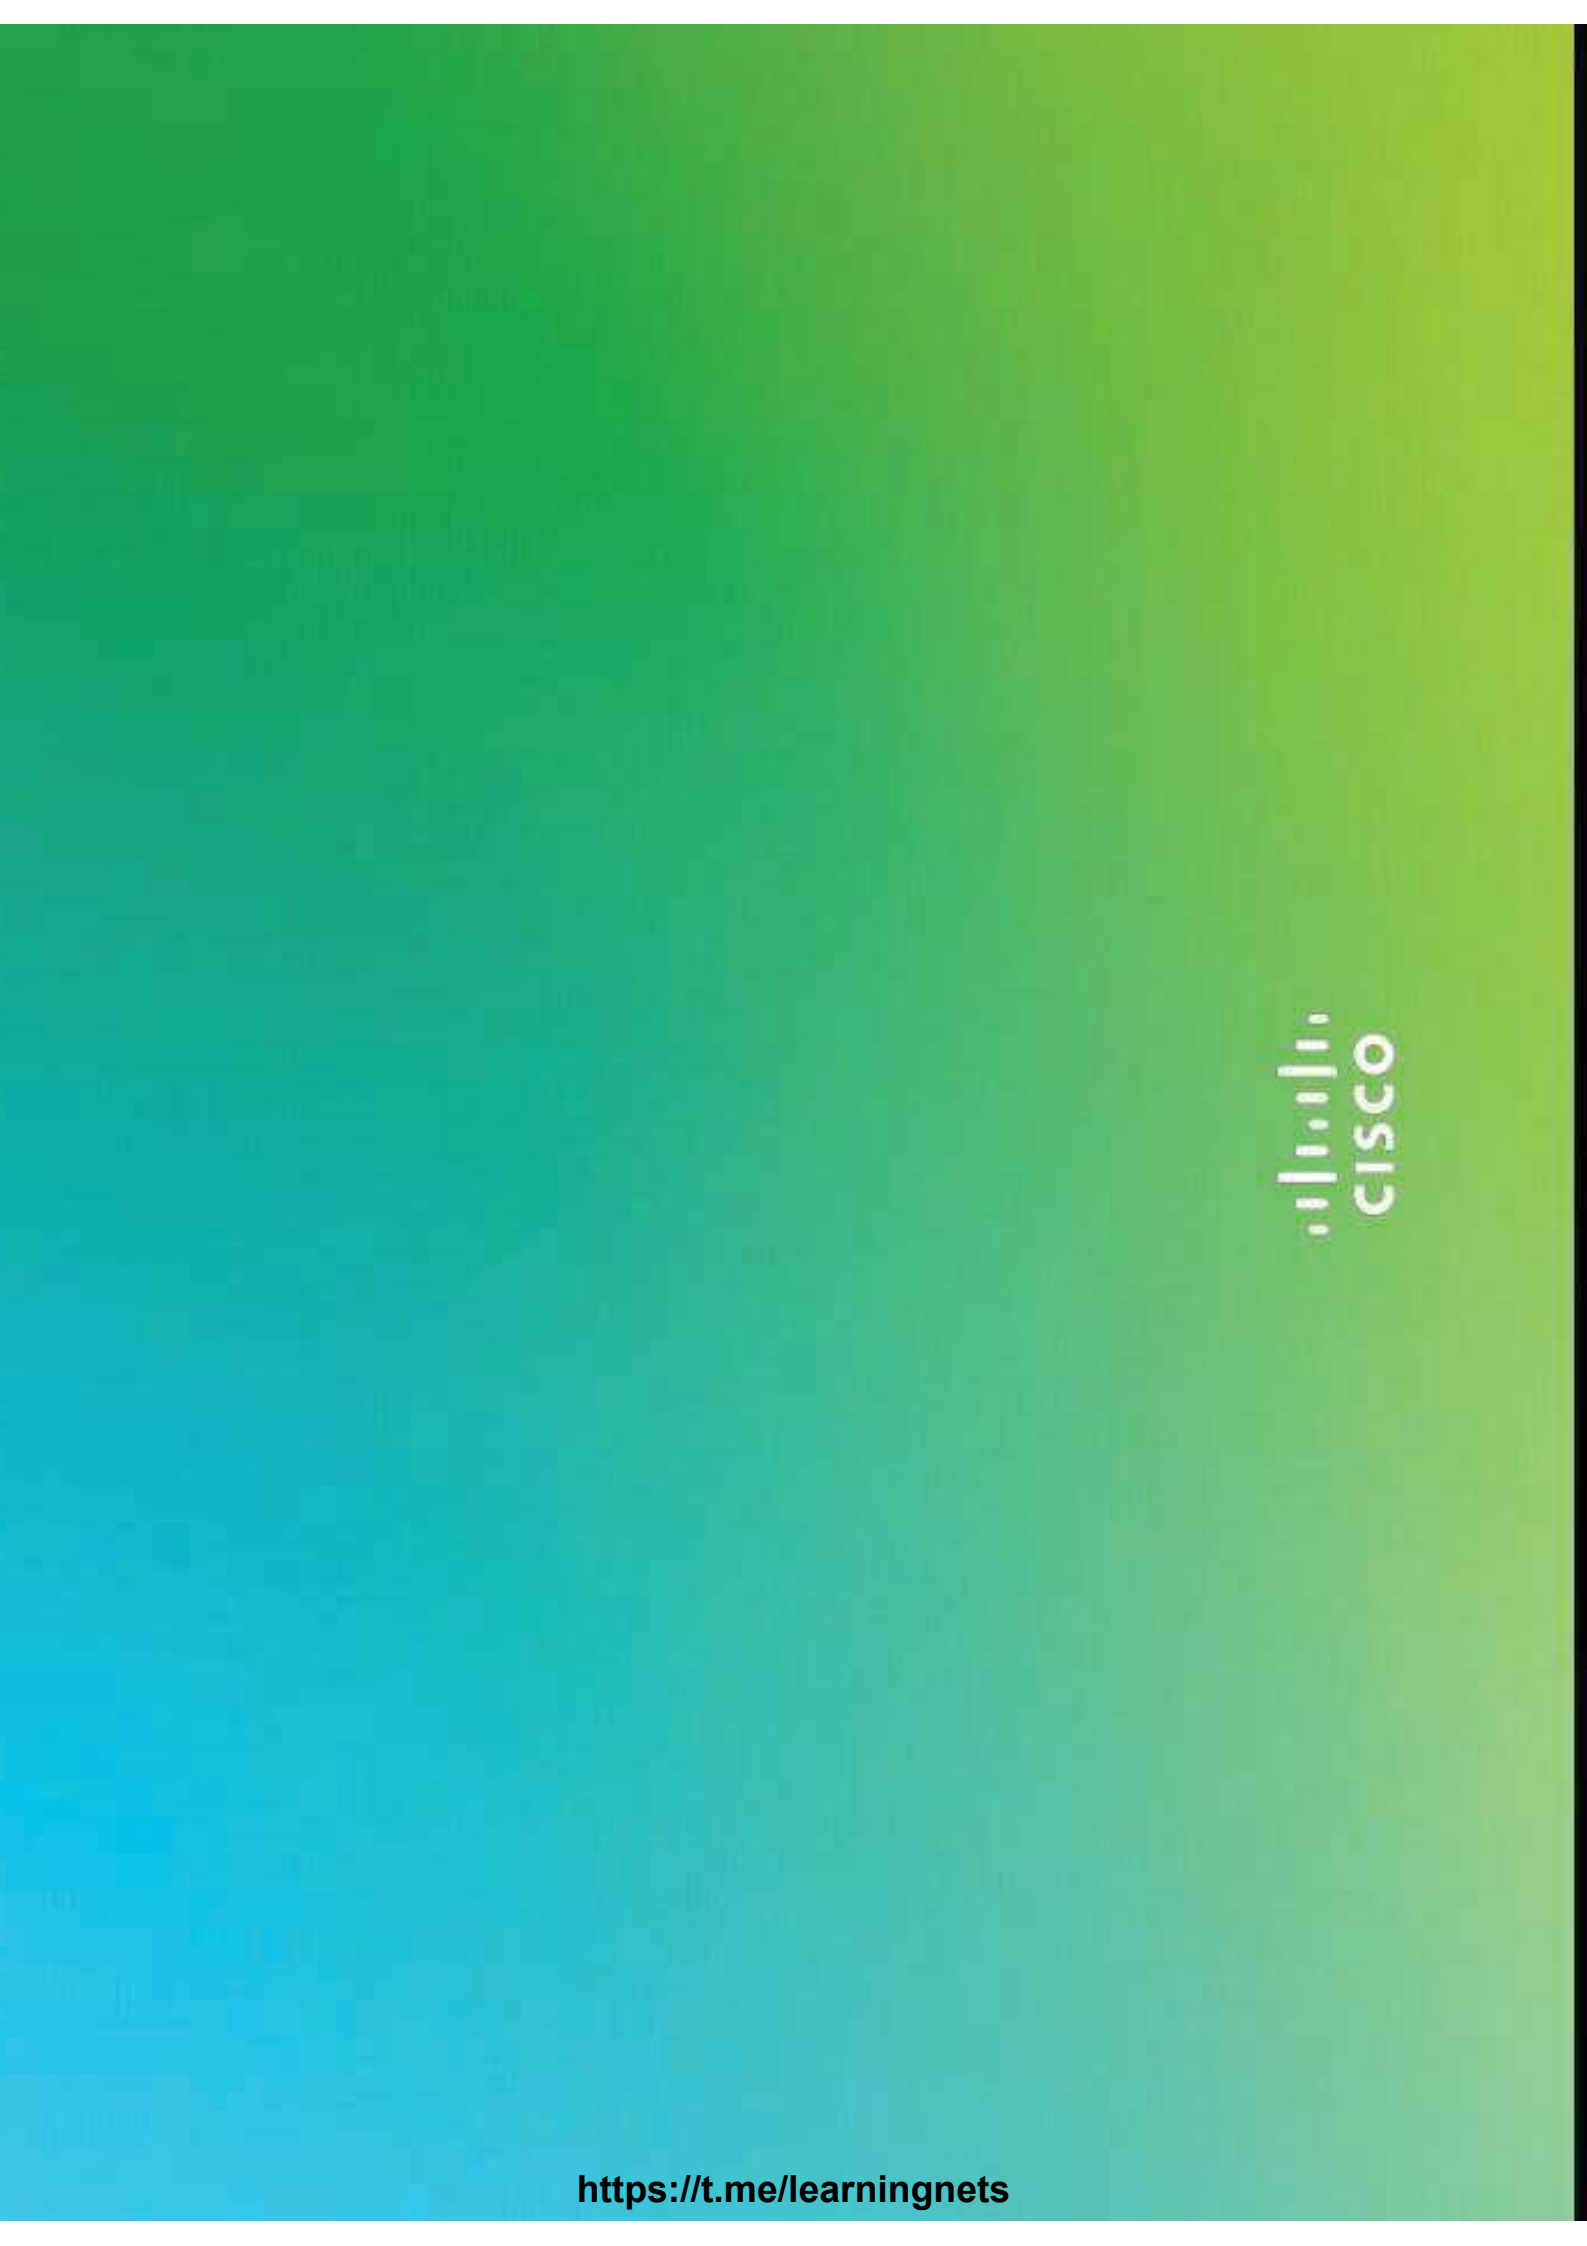

https://t.me/learningnets

https://t.me/learningnets

Connecting a Customer to a Service Provider

Service Provider Connectivity with BGP

© 2014 Cisco and/or its affiliates. All rights reserved.

Implementing Customer Connectivity Using Static Routing

Static Routing:

- Static routing should be used only in the following examples:
  - Single-homed customers.
  - Dual-attached customers, where link and equipment failure can be detected.
- With multihomed customers, dynamic routing with BGP should be used.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer Using Static Routing and a Single IP Address

Main characteristics about single-homed customer using static routing and a single IP address:

- Default route on the CE router and redistribution into customer IGP.
- NAT with PAT on the CE for outbound connectivity.
- Redistribution of connected routes into service provider BGP:
  - Match only routes that should be redistributed.
- Specific customer routes should not be advertised to upstream service providers:
  - Summarization of routes on the PE facing upstream service providers.
  - Customer routes with no-export community when redistributed into BGP.

https://t.me/learningnets

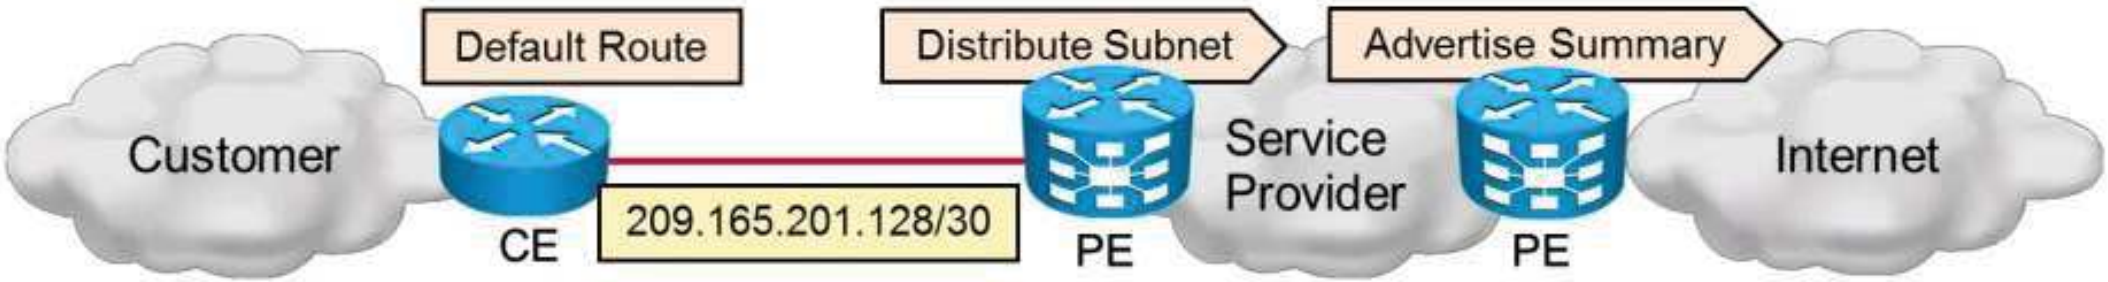

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer Using Static Routing and a Single IP Address (Cont.)

The figure shows the configuration.

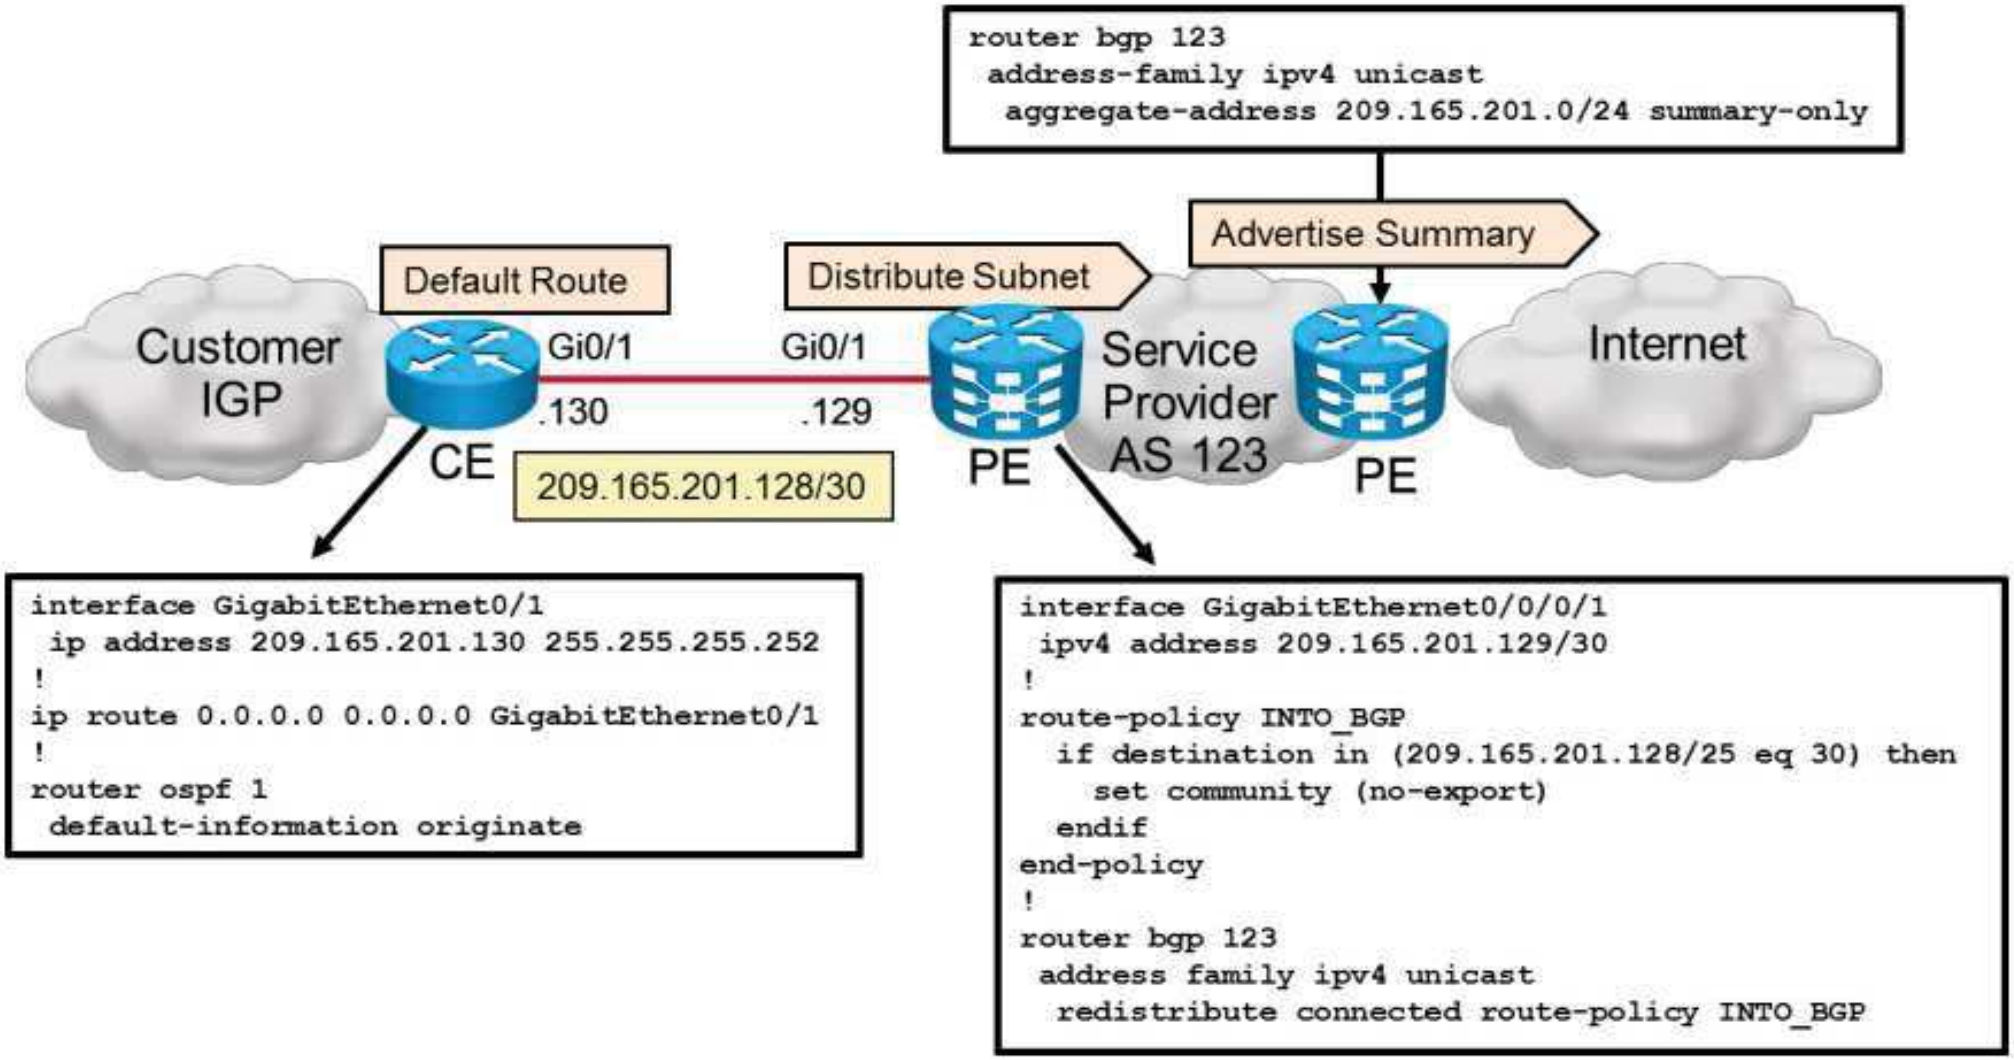

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer Using Static Routing and Multiple IP Addresses

Using static routing and multiple IP addresses characteristics:

- Default route on the CE router and redistribution into customer IGP.
- Static route on the PE router and redistribution into service provider BGP.
- Specific customer routes should not be advertised to upstream service providers:
  - Summarization of routes on the PE facing upstream service providers.
  - Customer routes with no-export community when redistributed into BGP

https://t.me/learningnets

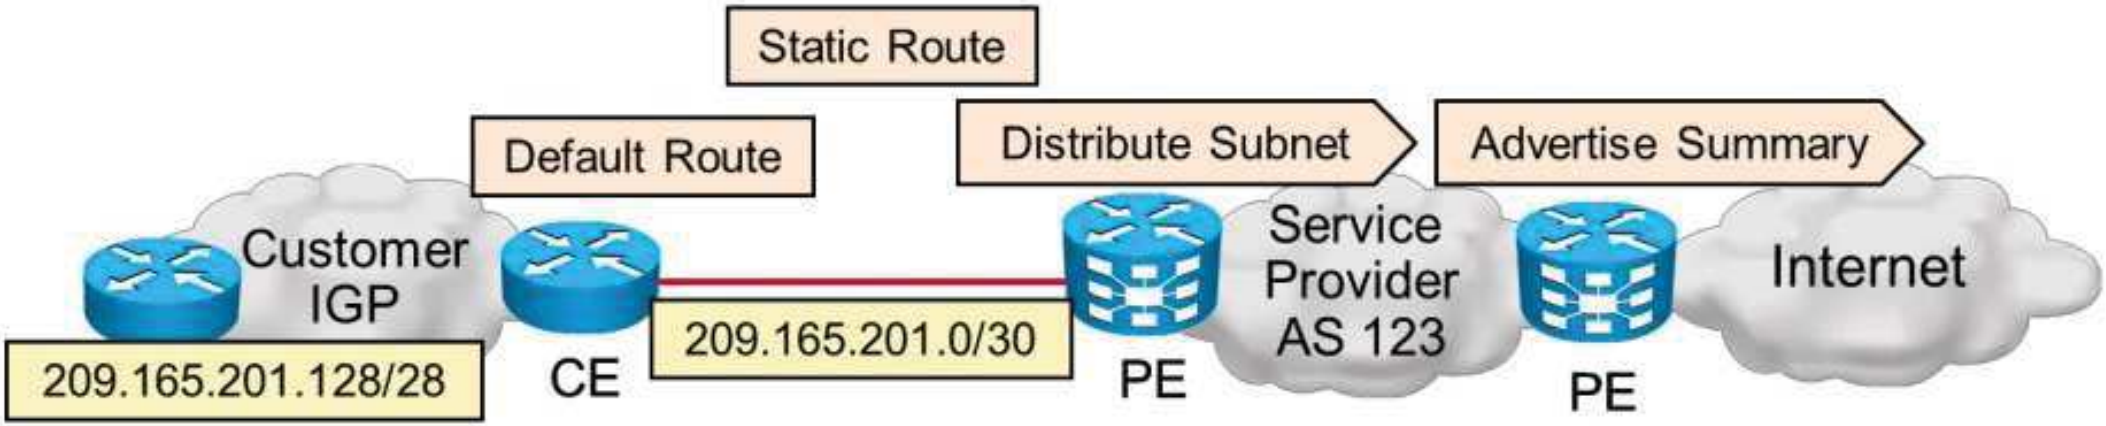

© 2014 Cisco and/or its affiliates. All rights reserved.

Single-Homed Customer Using Static Routing and Multiple IP Addresses (Cont.)

Configuration:

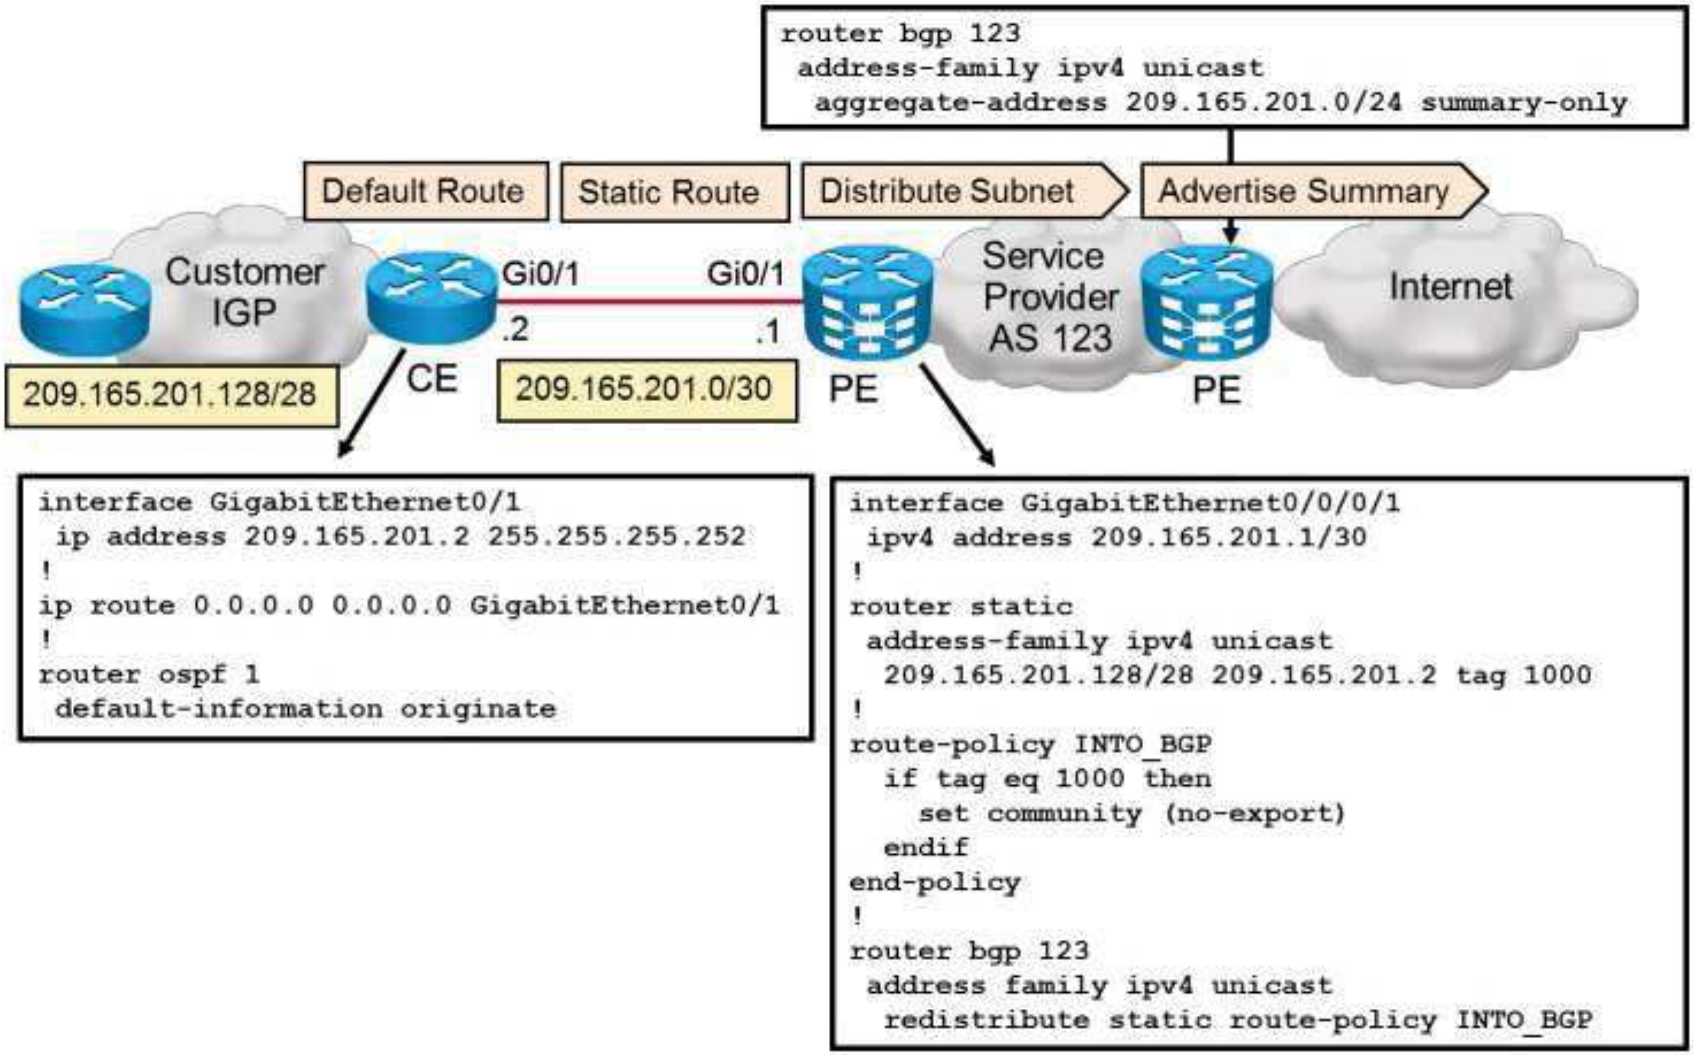

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Primary and Backup Scenario

Default routes on the CE router and redistribution into customer IGP:

- Floating static on the backup CE router.

Static route on the PE router and redistribution into service provider BGP:

- Floating static on the backup PE router.
- Separately tag primary and backup static routes and match them for redistribution.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Primary and Backup Scenario (Cont.)

Specific customer routes should not be advertised to upstream service providers:

- Summarization of routes on the PE facing upstream service providers.
- Customer routes with no-export community when redistributed into BGP.

https://t.me/learningnets

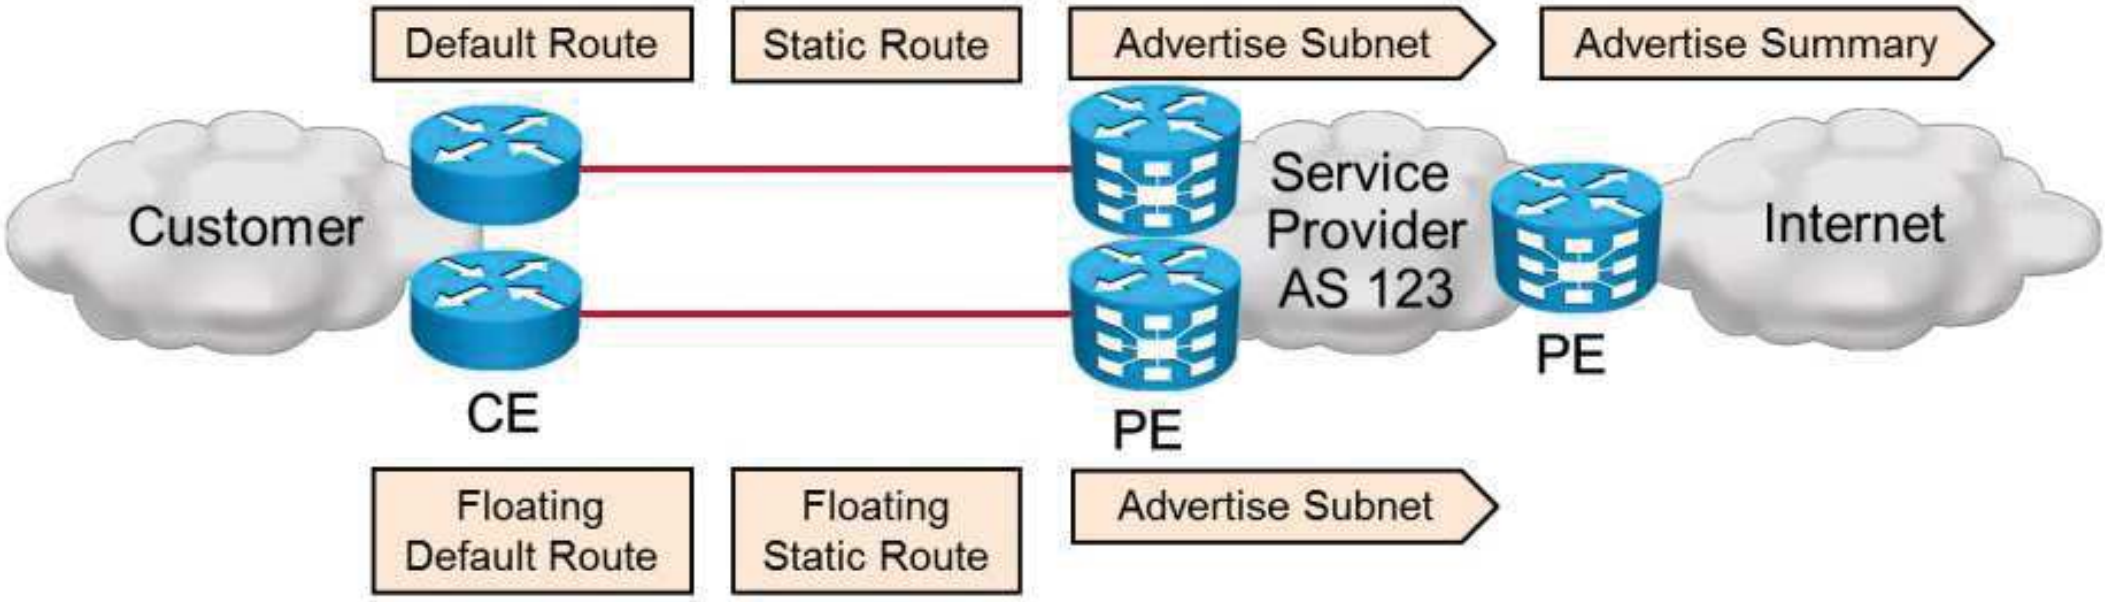

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Primary and Backup Scenario (Cont.)

The backup should not be considered unless the primary path is not available:

- Lower local preference than primary.
- Lower or same weight than primary (and locally originated route has weight of 32768).

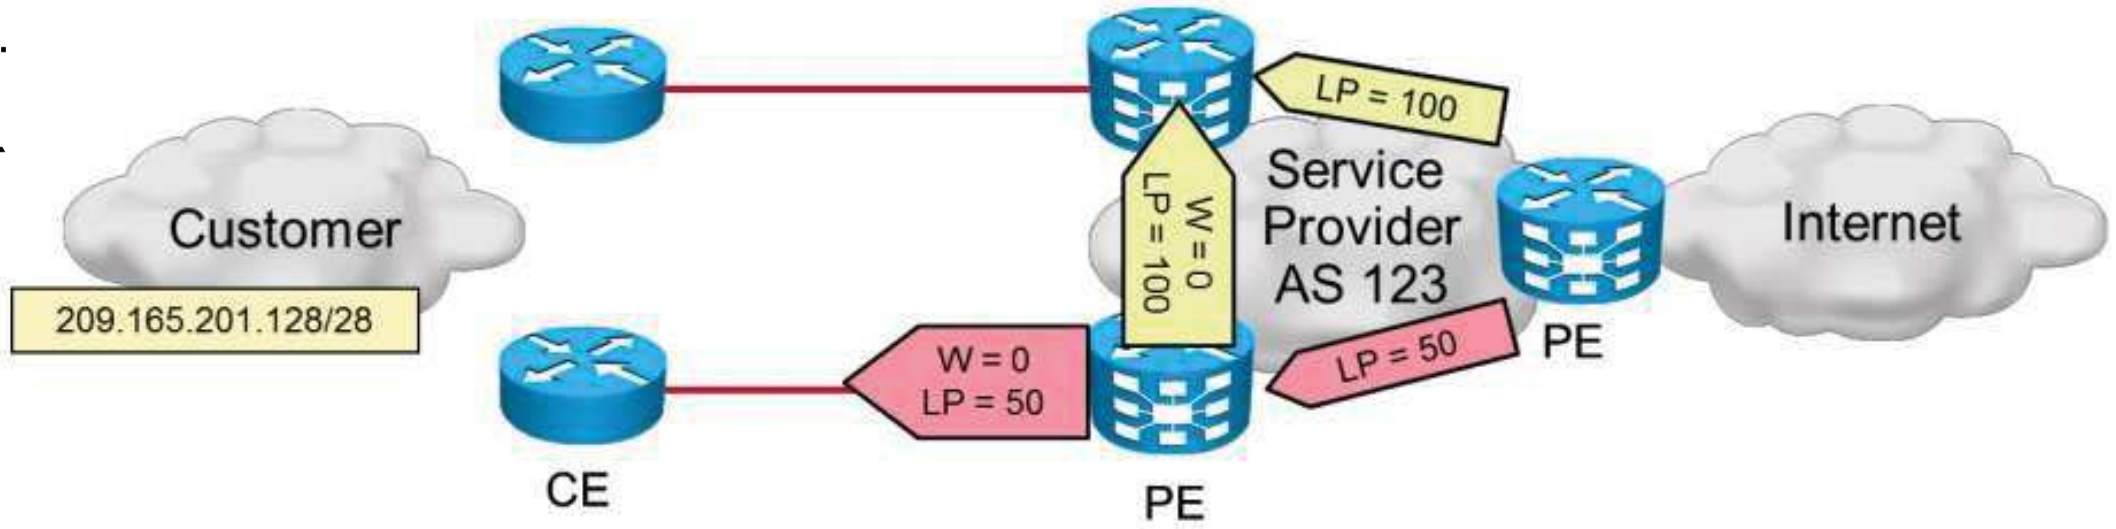

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Primary and Backup Scenario (Cont.)

Configuration of CE:

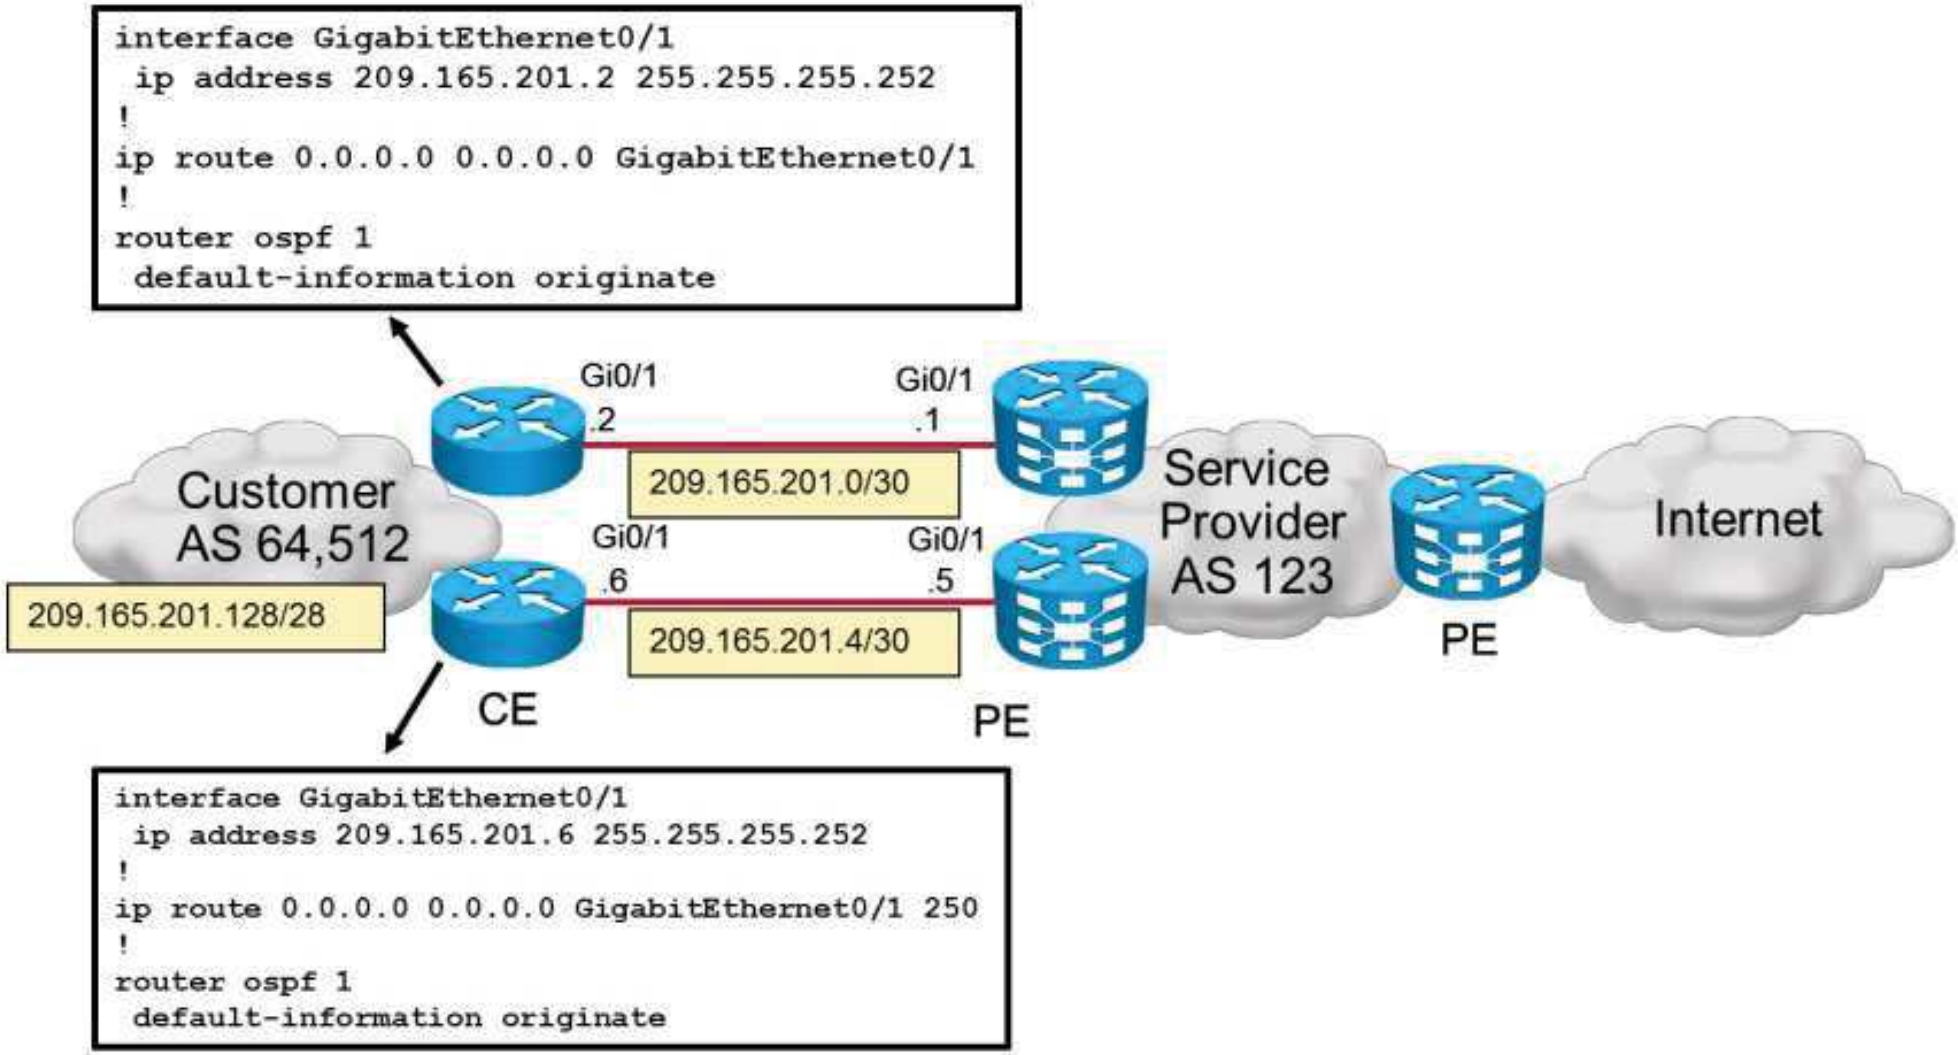

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Primary and Backup Scenario (Cont.)

Configuration of PE:

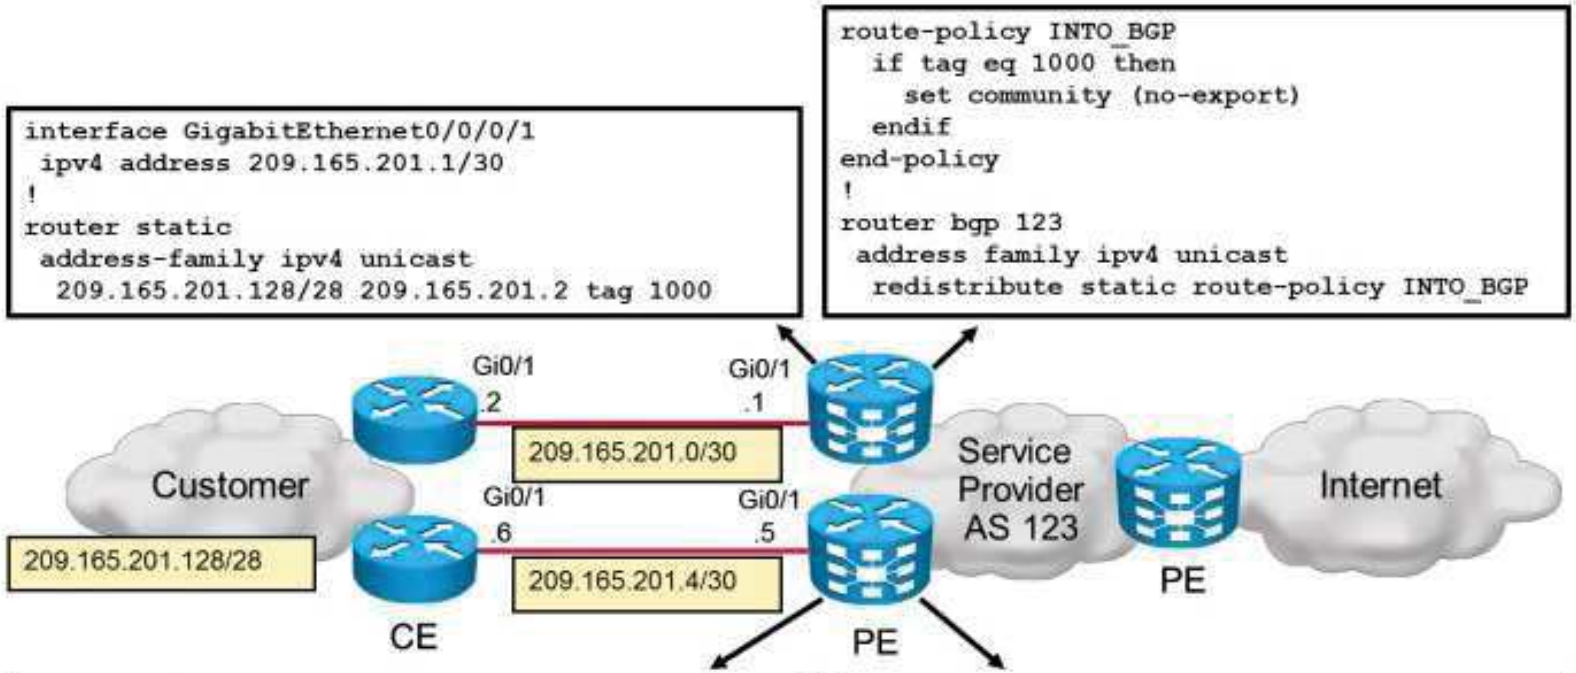

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Load-Balancing Scenario

- Outgoing traffic load balancing:
  - On each CE router, the default route is redistributed into IGP. The exit point is determined by IGP.
- Incoming traffic load balancing:
  - Each PE router advertises only part of the customer address space.
  - Each PE router advertises the whole address space for backup purposes.

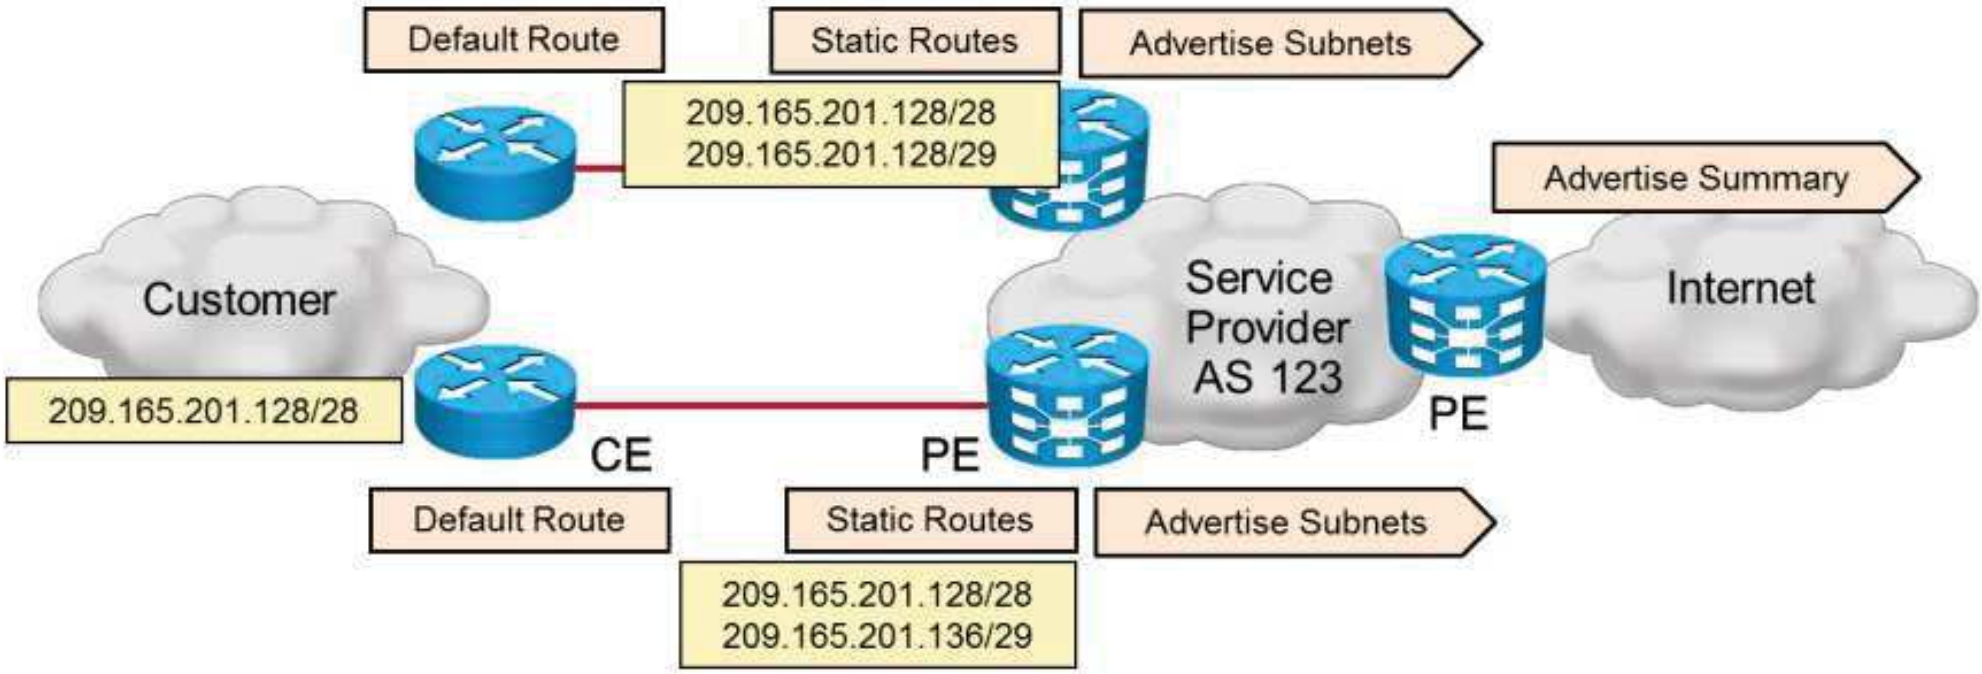

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using Static Routing in a Load-Balancing Scenario (Cont.)

Configuration:

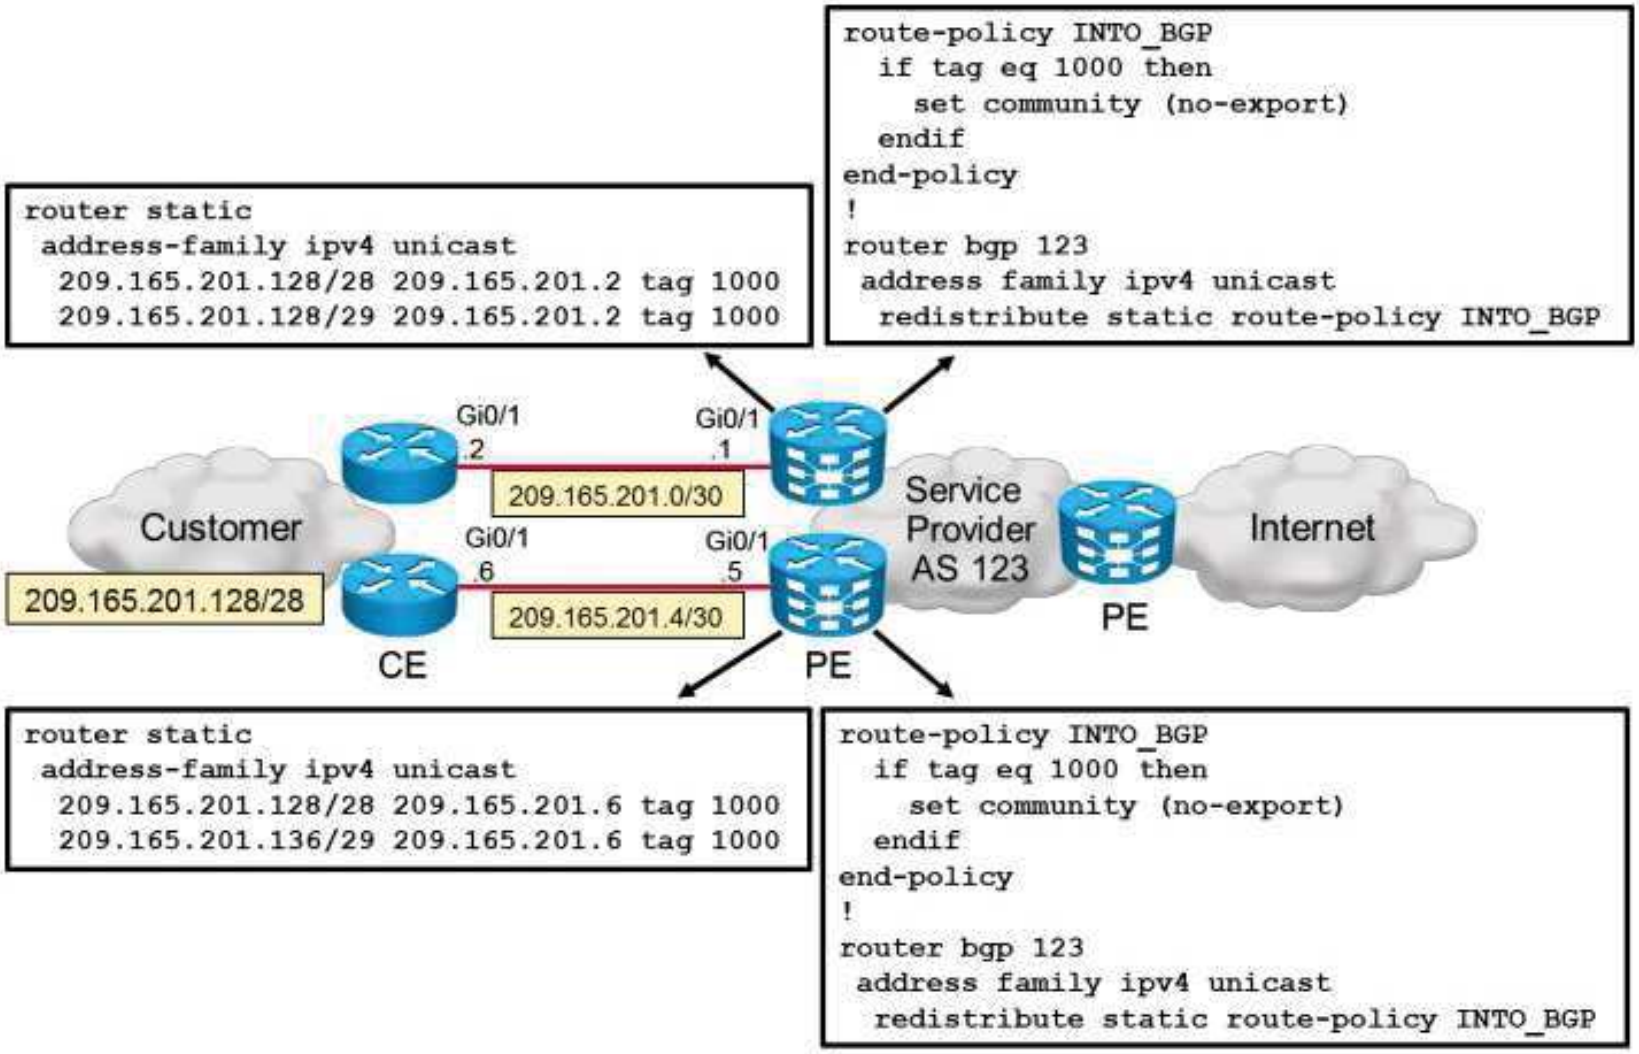

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Connecting a Dual-Attached Customer to a Single Service Provider

Connecting a dual-attached customer characteristics:

- BGP is run between the customer and service provider.
- The customer advertises allocated address space into BGP. The customer can use private AS numbers.
- The service provider advertises the default route to the customer. CE routers advertise the default route to the customer network using IGP.
- The service provider has to deploy inbound BGP filters.

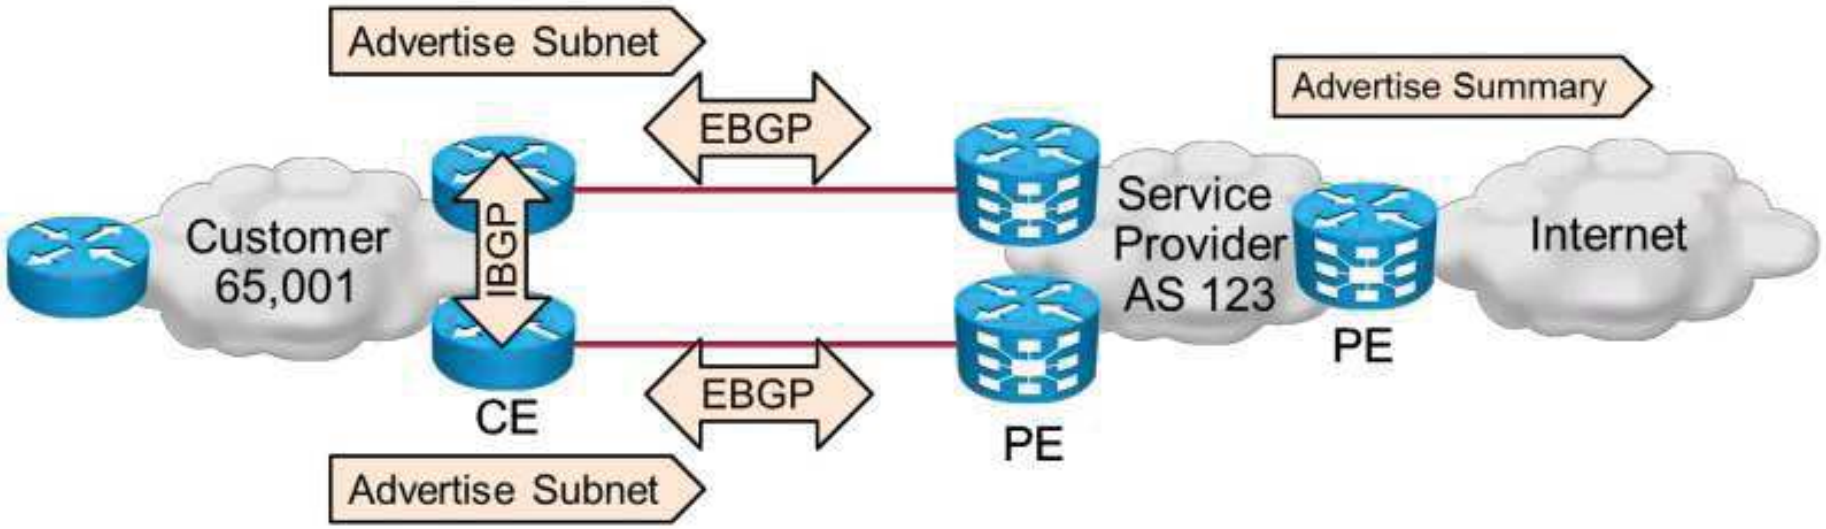

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Conditional BGP Advertising

Conditional BGP advertising characteristics:

- The CE router should advertise the whole customer space into BGP.
- The CE router should advertise customer space only if the customer network is reachable from the CE router. This is called conditional advertising.
- The CE router should stop advertising customer address space if the customer core network is not available.
- Reachability of the customer core network is tested using static route. The route should point to a core network next hop that is learned via IGP.

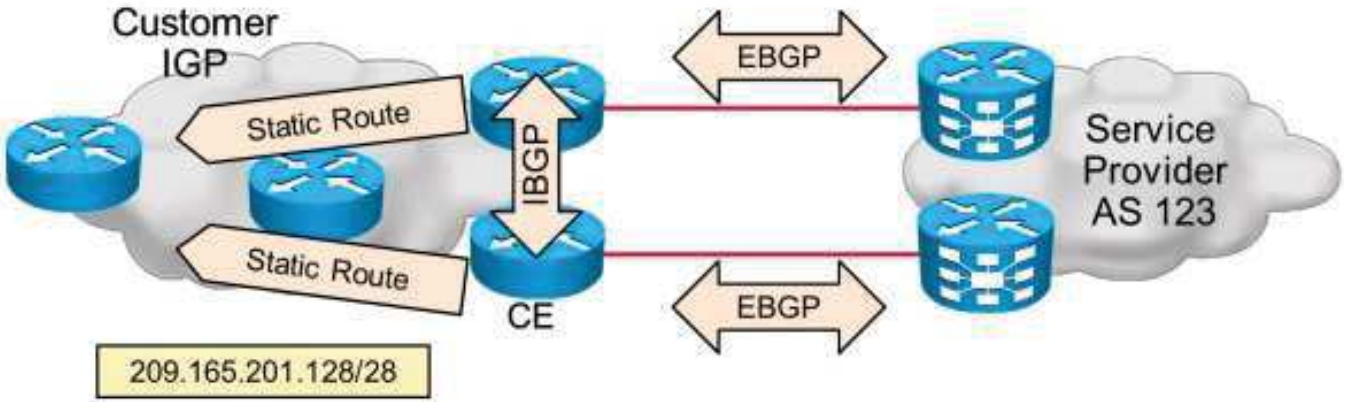

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configurations on the Service Provider PE Router

The service provider must:

- (Conditionally) Advertise a default route to the customer through BGP.
- Filter incoming BGP updates with a prefix-list to verify that the customer announces only the assigned address space.
- Filter incoming BGP updates with an AS-path filter list to verify that the customer uses only its own AS number.
- Optionally, the no-export community should be set on customer routes.

https://t.me/learningnets

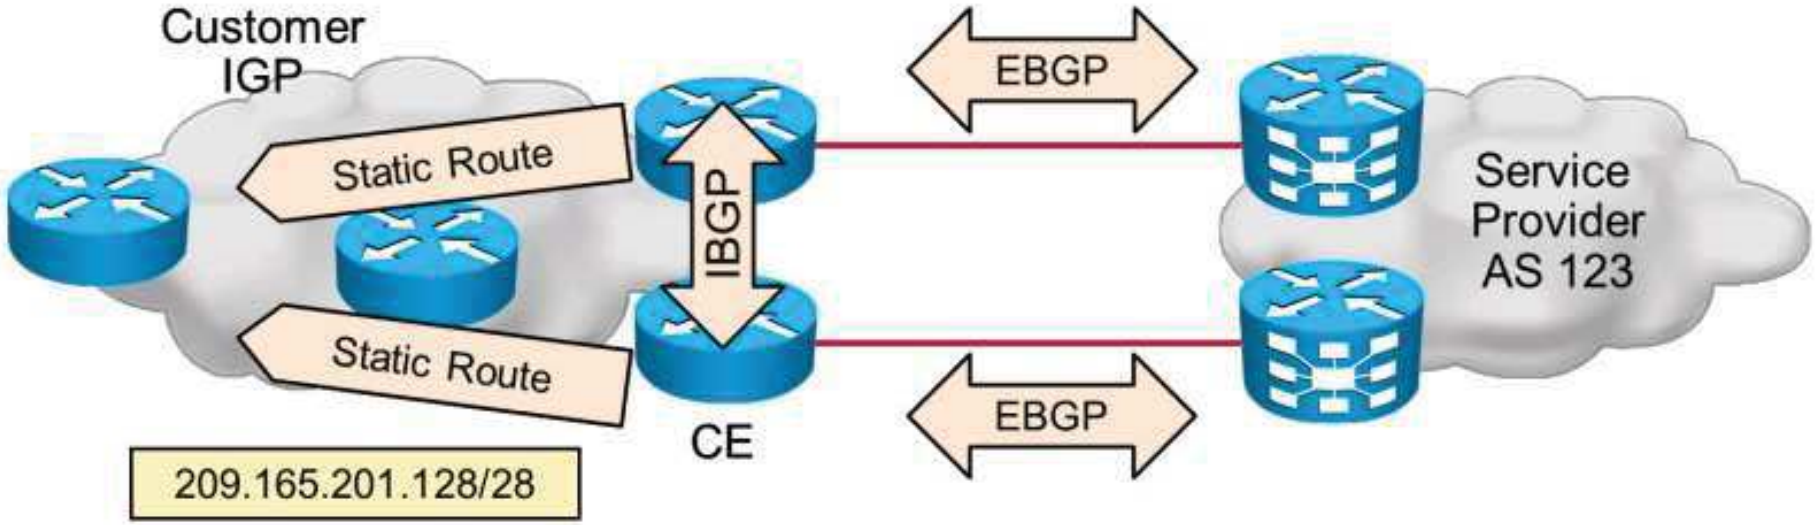

© 2014 Cisco and/or its affiliates. All rights reserved.

Removing a Private AS Number

Key Features and benefits for removing a private AS number:

- A dual-attached customer can use private AS numbers.
- The private AS numbers must be removed from the AS path before the customer routes are advertised to other service providers.
- Private AS numbers can be removed from the tail of the AS path using the remove-private-AS keyword.

https://t.me/learningnets

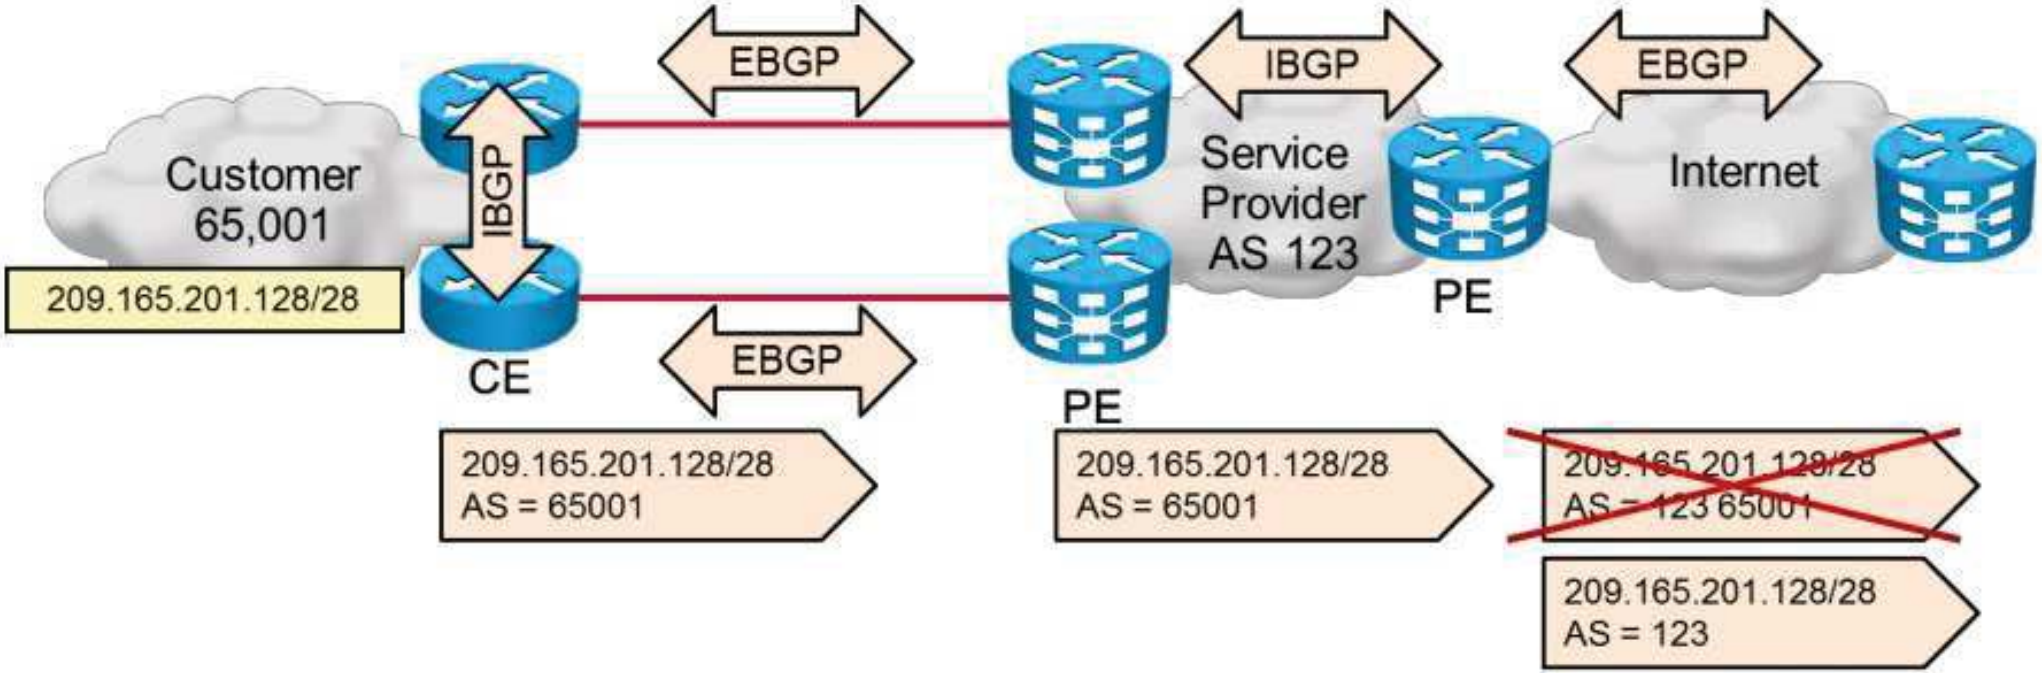

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customer Using BGP

Configuration of CE:

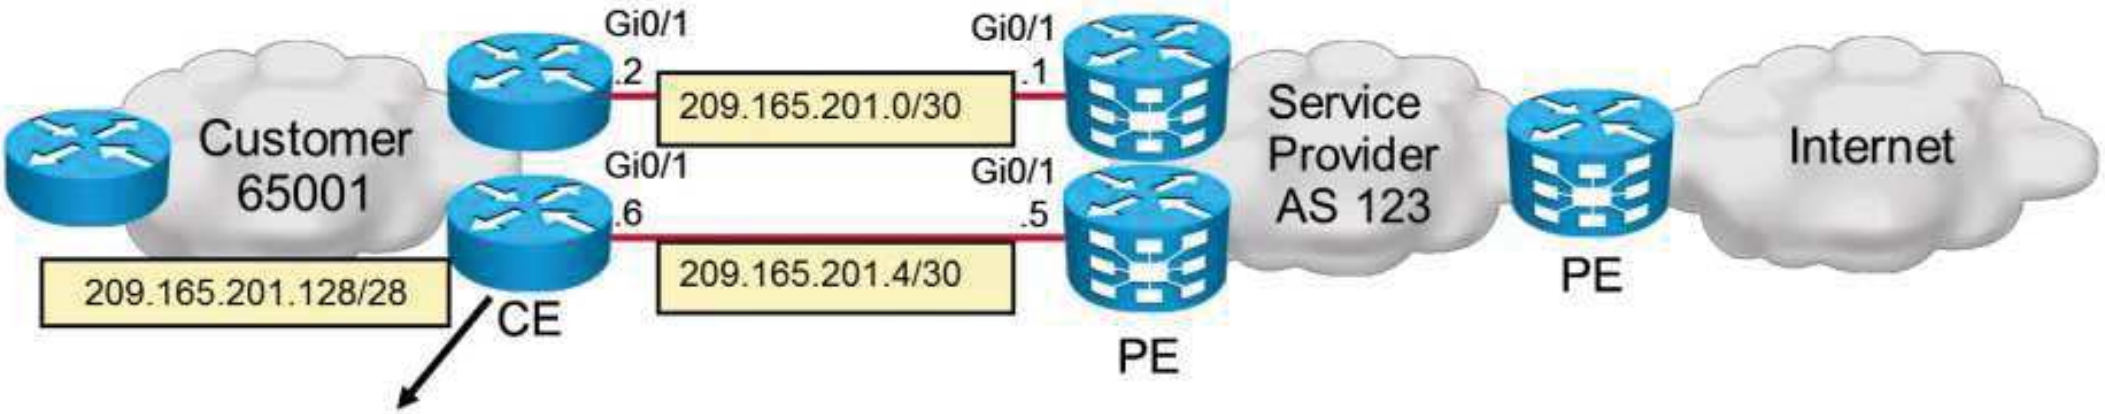


```
ip route 209.165.201.128 255.255.255.240 209.165.201.135
!
router bgp 65001
 neighbor 209.165.201.5 remote-as 123
 neighbor 209.165.201.129 remote-as 65001
 network 209.165.201.128 mask 255.255.255.240
!
router ospf 1
 default-information originate
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customer Using BGP (Cont.)

Configuration of PE:

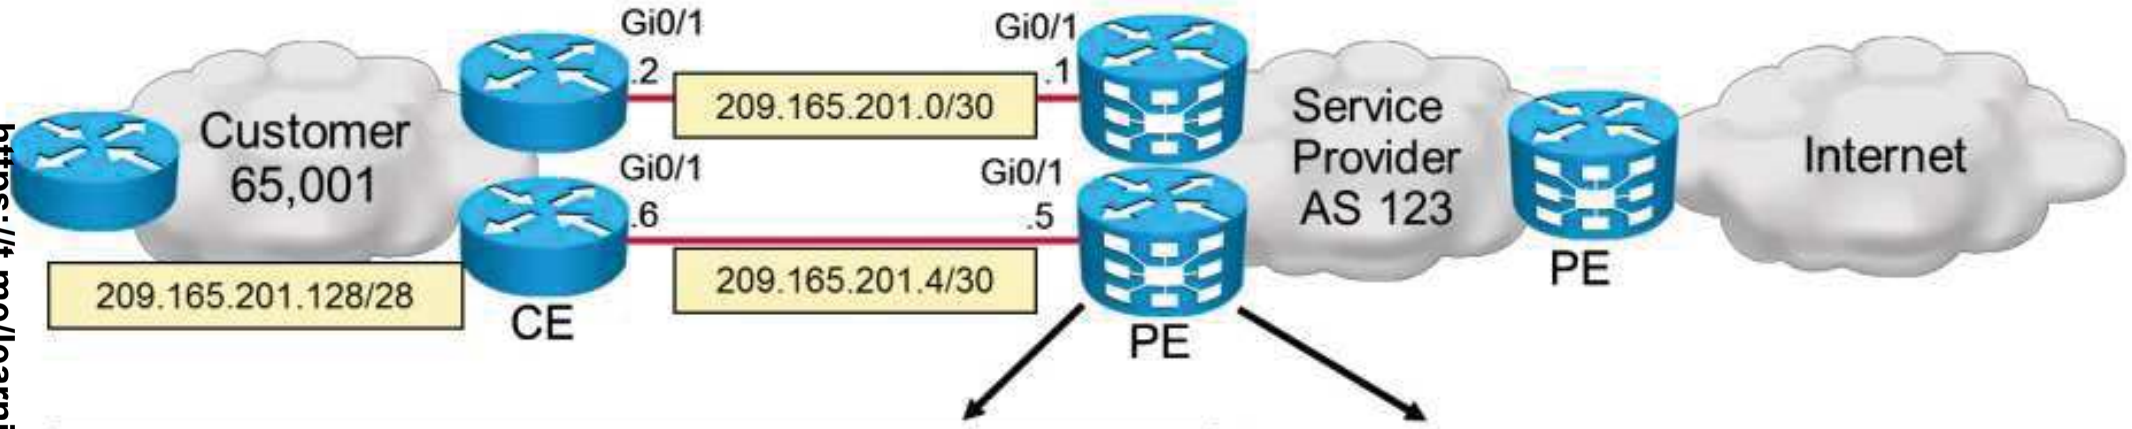


```
route-policy DEFAULT
 if destination in (0.0.0.0/0) then
  pass
 endif
end-policy
!
route-policy CUSTOMER
 if destination in (209.165.201.128/28) and
 as-path in (ios-regex '^65001(_65001)*$') then
  set community (no-export)
 endif
end-policy
```

```
router bgp 123
 neighbor 209.165.201.6
  remote-as 65001
  address-family ipv4 unicast
   default-originate
   route-policy DEFAULT out
   route-policy CUSTOMER in
   remove-private-AS
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider Migrations Using Local AS

Advantage for service provider migrations using local AS:

- If a service provider migrates to a new AS number, customers have to change the BGP configuration.
- A merged service provider can present itself to the customer using the old AS number, so the customer is not required to change the BGP configuration.

https://t.me/learningnets

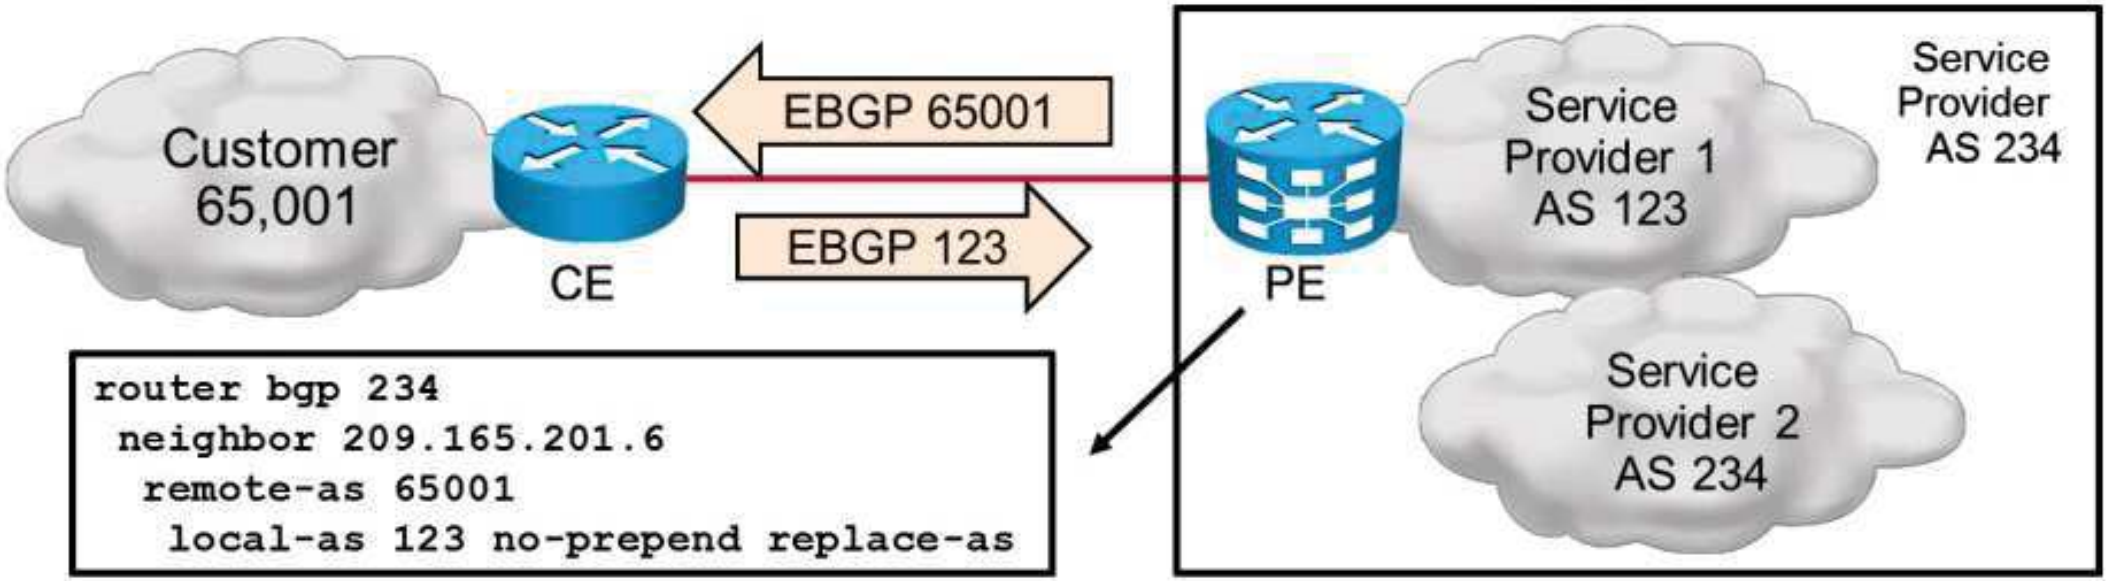

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using BGP in a Primary and Backup Scenario

The route selection is controlled by the CE routers.

- Outgoing traffic:
  - Local preference is used to select the primary or backup link.
- Incoming traffic:
  - MED is used to advertise the primary or backup link to the service provider. This is not reliable because the service provider can change the local preference.

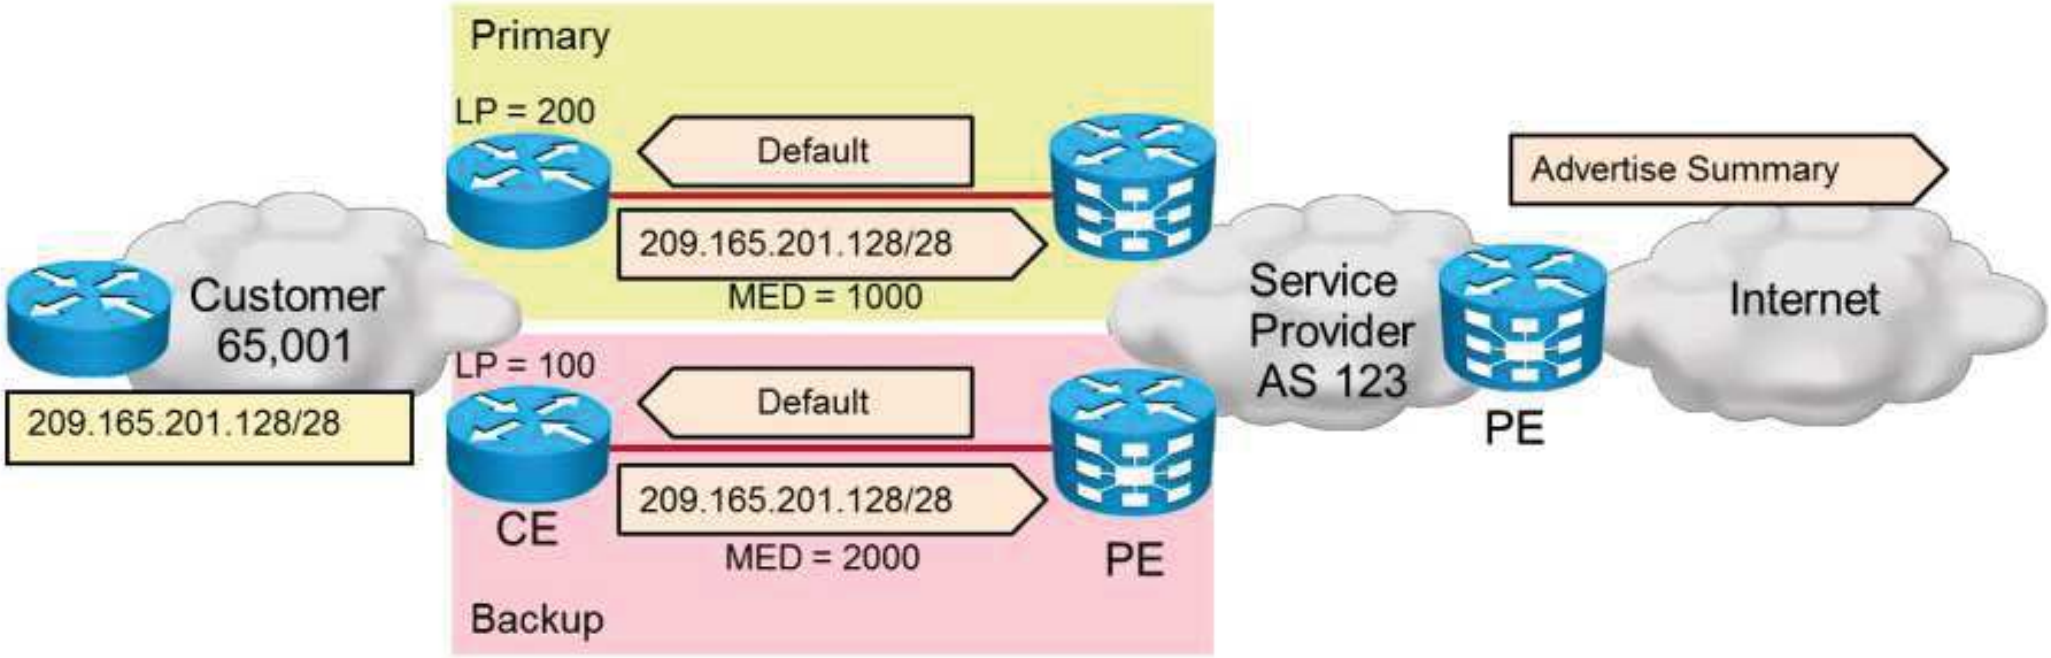

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using BGP in a Primary and Backup Scenario (Cont.)

Configuration:

```
router bgp 65001
 neighbor 209.165.201.1 remote-as 123
 neighbor 209.165.201.1 route-map HIGH_LOC_PREF in
 neighbor 209.165.201.130 remote-as 65001
 network 209.165.201.128 mask 255.255.255.240
!
route-map HIGH_LOC_PREF permit 10
 set local-preference 200
```

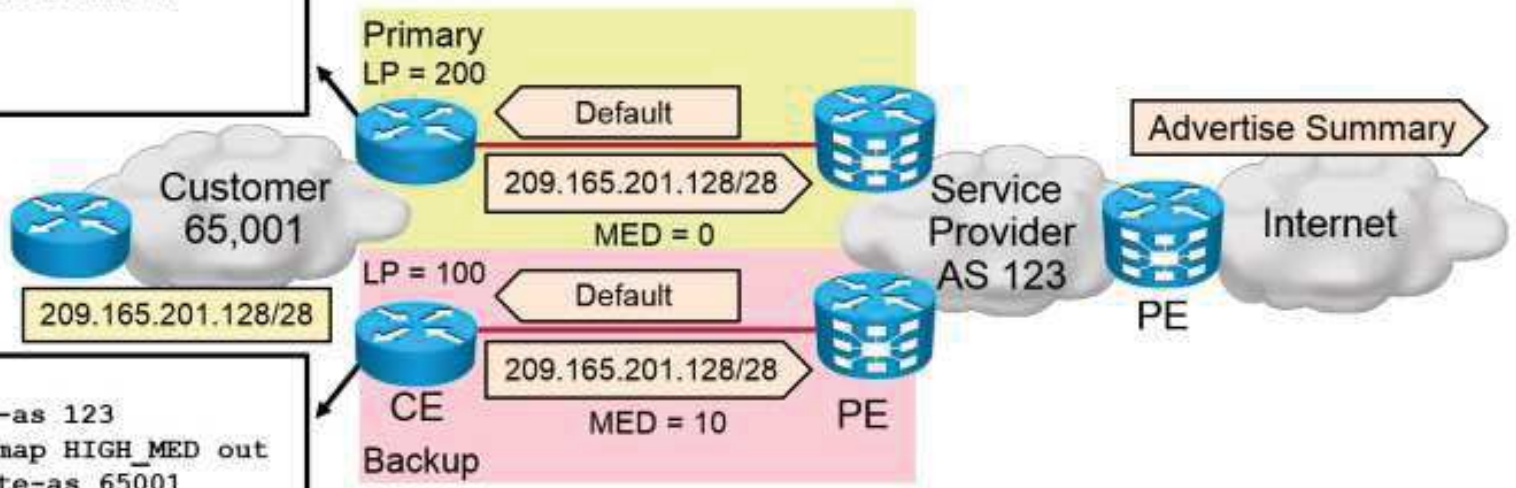


```
router bgp 65001
 neighbor 209.165.201.5 remote-as 123
 neighbor 209.165.201.5 route-map HIGH_MED out
 neighbor 209.165.201.129 remote-as 65001
 network 209.165.201.128 mask 255.255.255.240
!
route-map HIGH_MED permit 10
 set metric 10
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using BGP in a Load-Balancing Scenario

- Outgoing traffic:
  - Load sharing is identical to the static routing scenario. It is determined by IGP.
- Incoming traffic:
  - Announce portions of the customer address space to each upstream router.
  - Configure BGP multipath support in the provider network.
  - Use EBGP multihop in environments where parallel links run between a pair of routers.

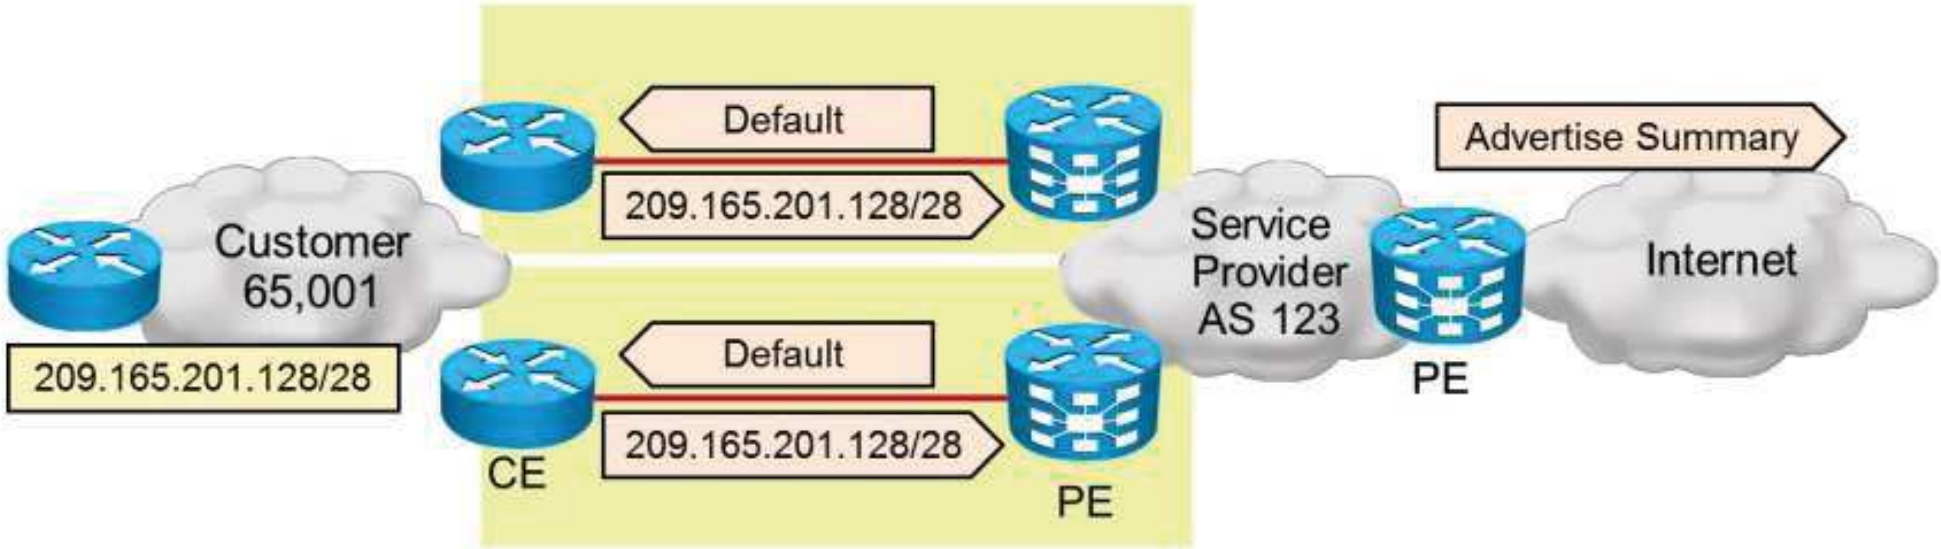

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using BGP in a Load-Balancing Scenario (Cont.)

BGP Multipath:

- By default, BGP selects a single path as the best path and installs it in the IP routing table.
- If configured, a BGP router can select up to eight identical BGP routes as the best routes and install them in the IP routing table for load-sharing purposes.

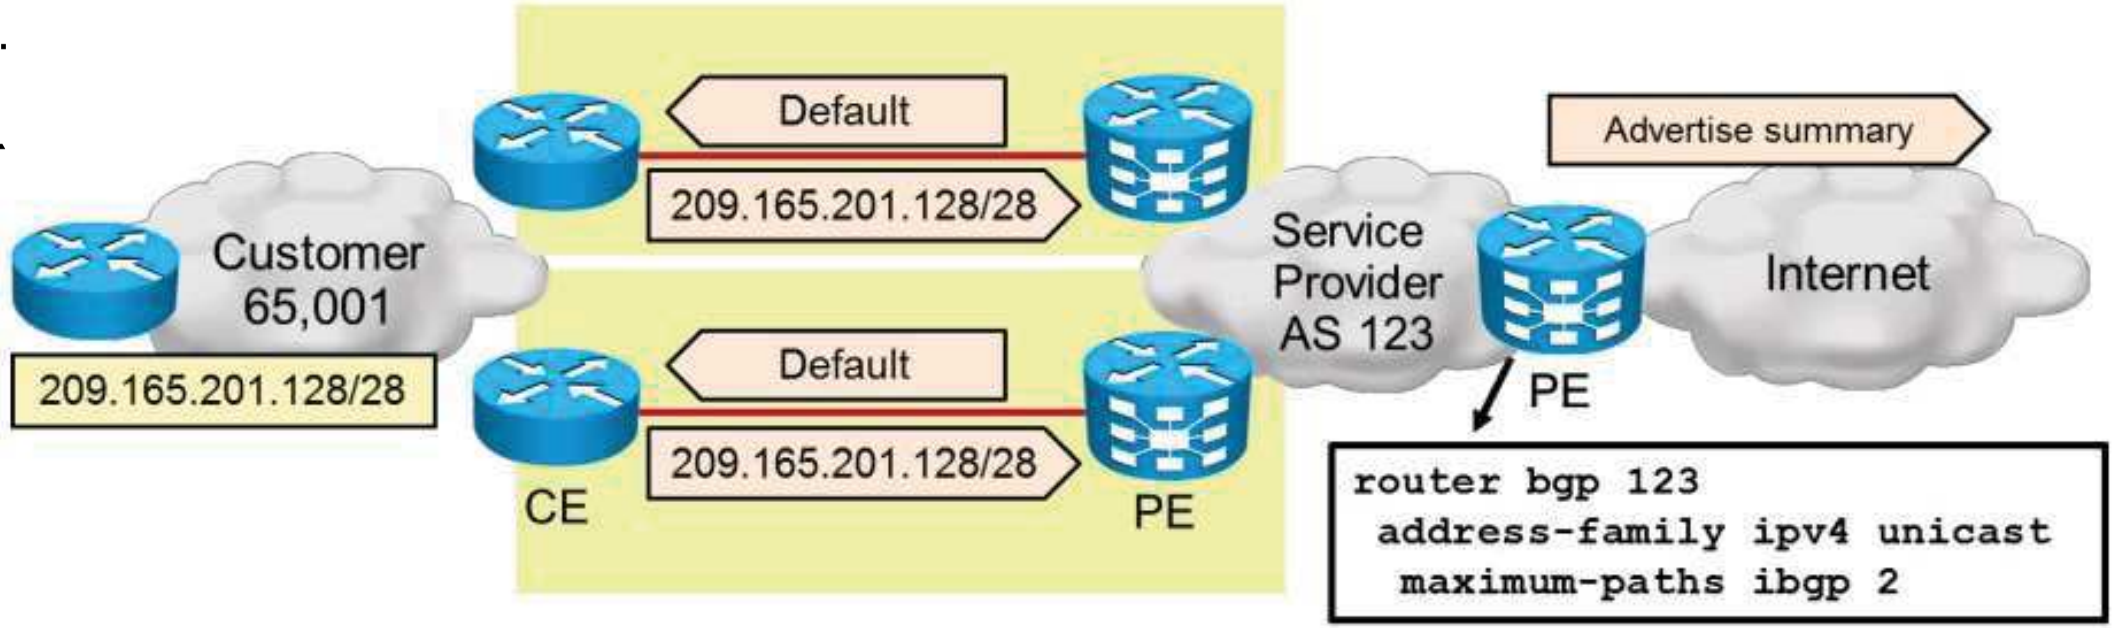

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-Attached Customers Using BGP in a Load-Balancing Scenario (Cont.)

EBGP Multihop:

- Used when parallel links between routers exist. In this case, loopback interfaces have to be used for BGP peering.
- Because of recursive lookup, load sharing toward a BGP destination always occurs if there are several equal-cost paths to the BGP next hop.
- By default, EBGP neighbors must be directly connected.
- The **ebgp-multihop** command declares an EBGP neighbor to be distant.

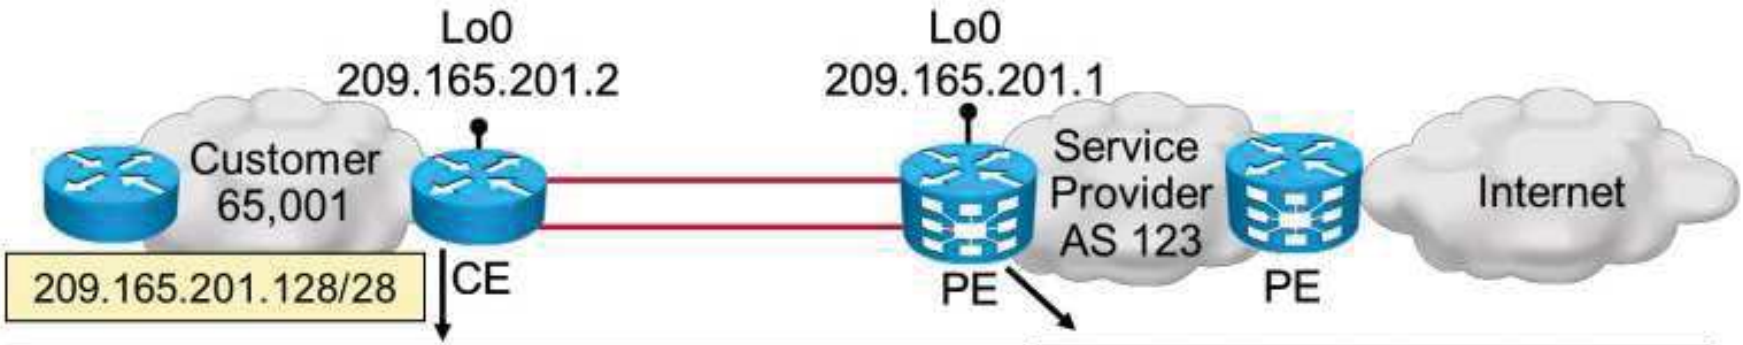


```
router bgp 65001
 neighbor 209.165.201.1 remote-as 123
 neighbor 209.165.201.1 ebgp-multihop 2
 neighbor 209.165.201.1 update-source Loopback0
```

```
router bgp 123
 neighbor 209.165.201.2
  remote-as 65001
  ebgp-multihop 2
  update-source Loopback0
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Connecting a Multihomed Customer to Multiple Service Providers

Multihomed Customer Using BGP:

- BGP is run between the customer and service provider.
- The customer advertises allocated address space into BGP. It is recommended to use public AS numbers.
- The service provider has to deploy inbound BGP filters.
- The service provider advertises the following:
  - Default route, default route + local routes, full Internet routing table

https://t.me/learningnets

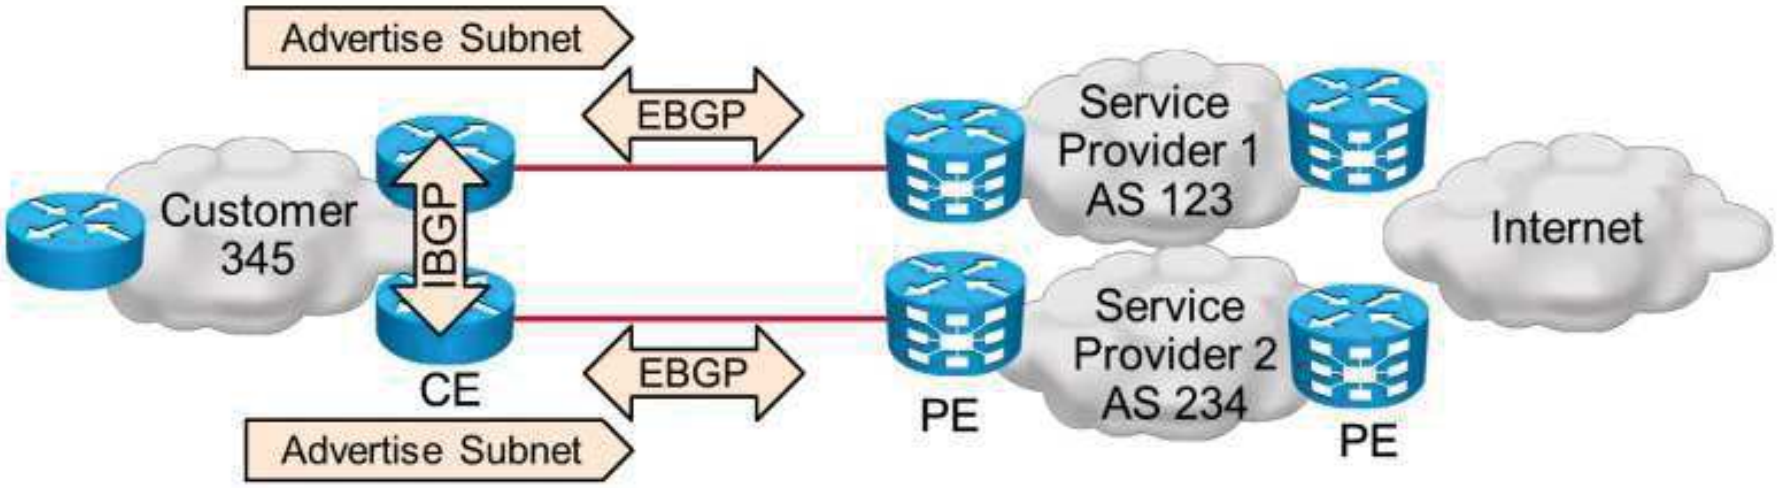

© 2014 Cisco and/or its affiliates. All rights reserved.

Connecting a Multihomed Customer to Multiple Service Providers (Cont.)

Multihomed Customer Routing Policies:

- Implemented by the customer:
  - Local preference usage for outbound path.
  - AS path prepending for return path.
- Optionally aided by the service provider through signaling:
  - Translating BGP communities to local preference and/or AS-path prepending.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-Implemented BGP Routing Policies in a Primary and Backup Scenario

- Outbound path:
  - A higher local preference for the default route comes from the primary service provider.
- Return path:
  - AS-path prepending of the customer route is sent to the backup service provider.
- Only the default route is required from both service providers.

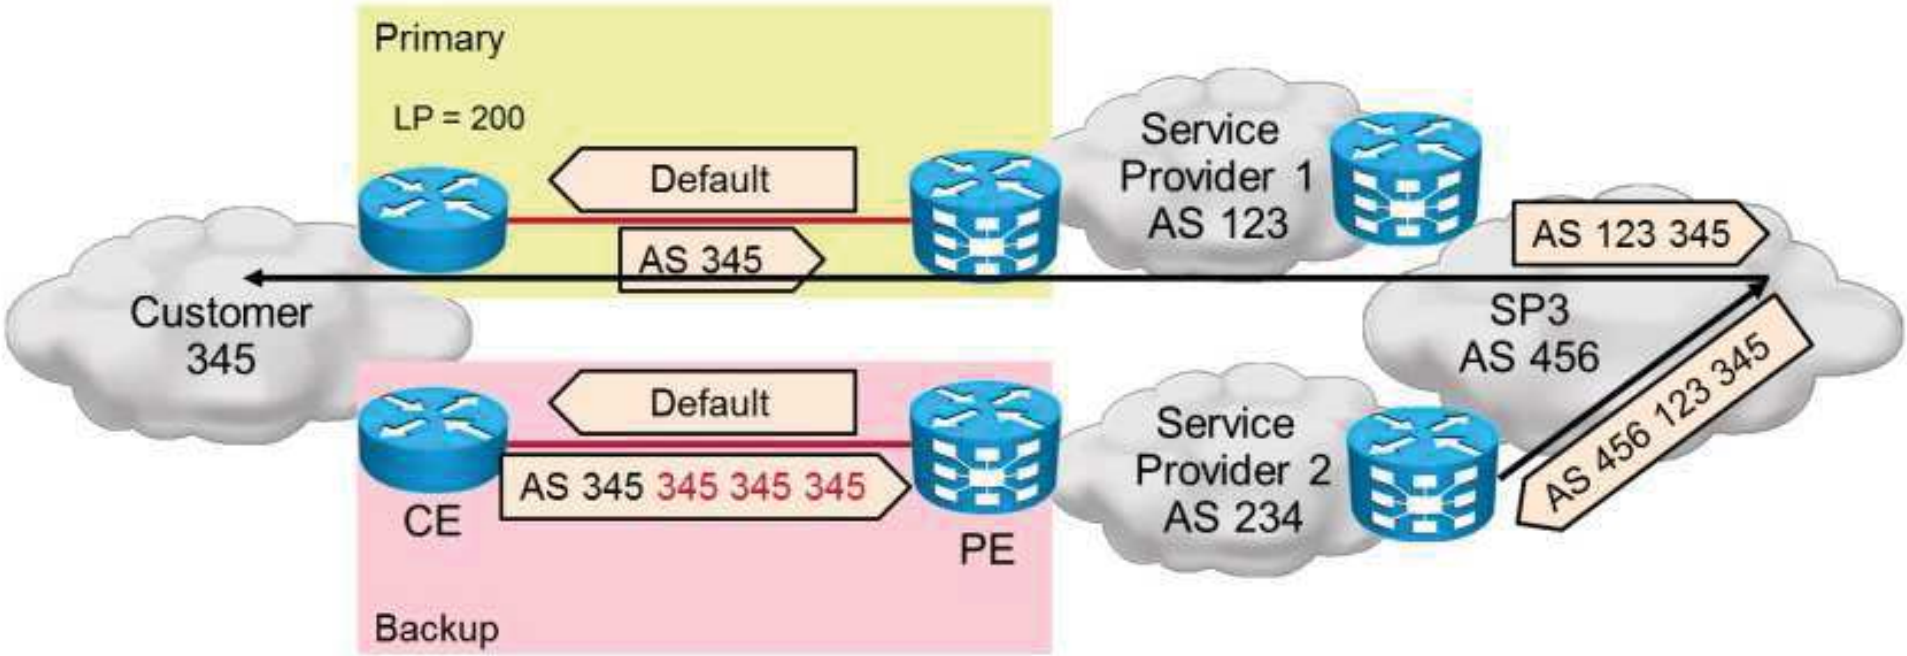

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-Implemented BGP Routing Policies in a Primary and Backup Scenario (Cont.)

- Outbound path:
  - A higher local preference for the default route comes from the primary service provider.
- Return path:
  - AS-path prepending of the customer route is sent to the backup service provider.
- Only the default route is required from both service providers.

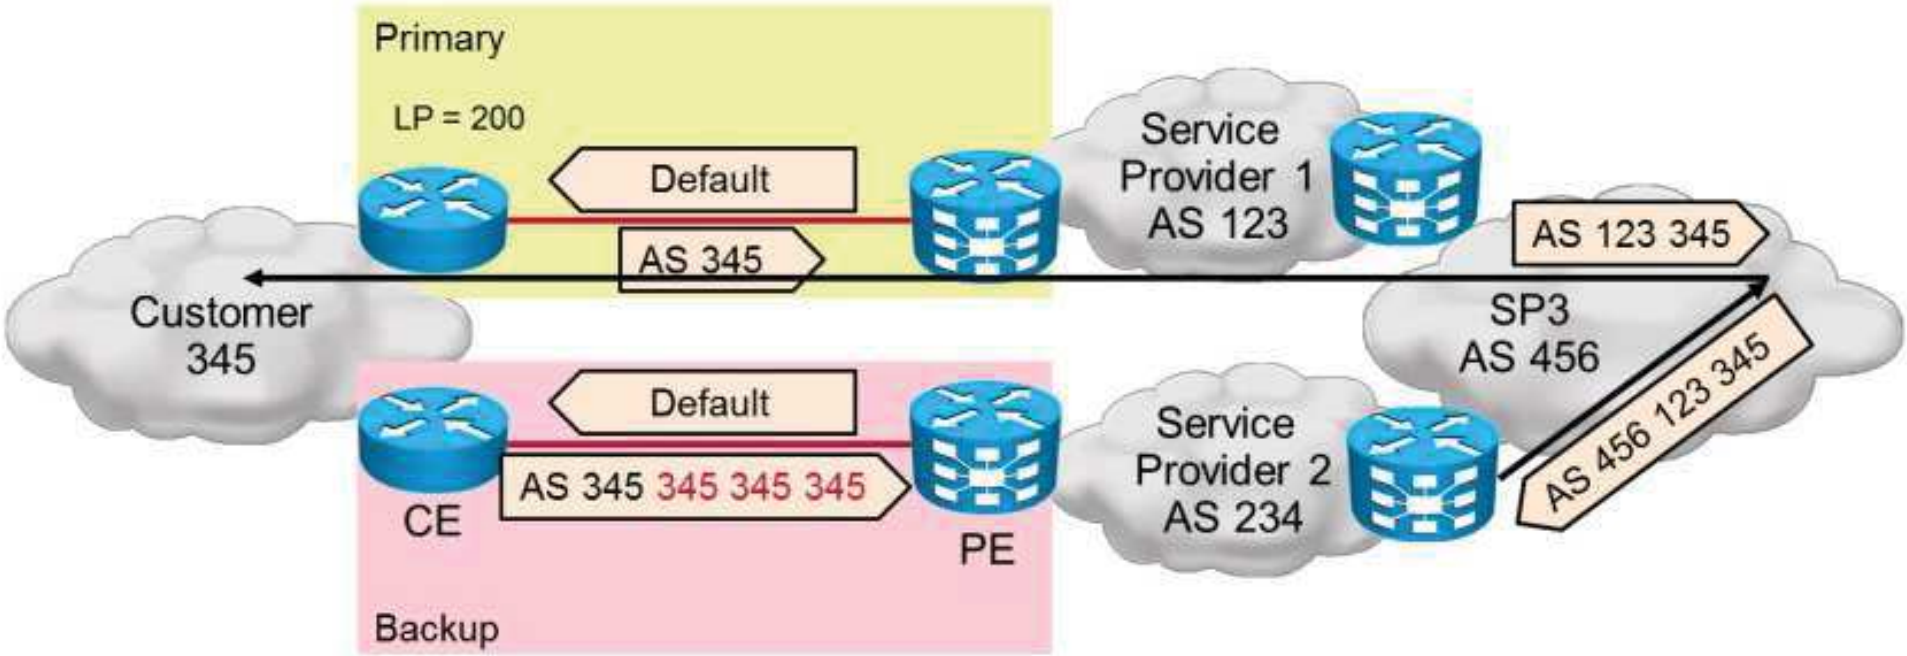

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-Implemented BGP Routing Policies in a Load-Balancing Scenario

Outbound path:

- A higher local preference for half of the routes comes from one service provider.
- A higher local preference for the other half of the routes comes from the other service provider.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-Implemented BGP Routing Policies in a Load-Balancing Scenario (Cont.)

Return path:

- The customer announces portions of the customer address space to each service provider.
- A full routing table is required from both service providers.

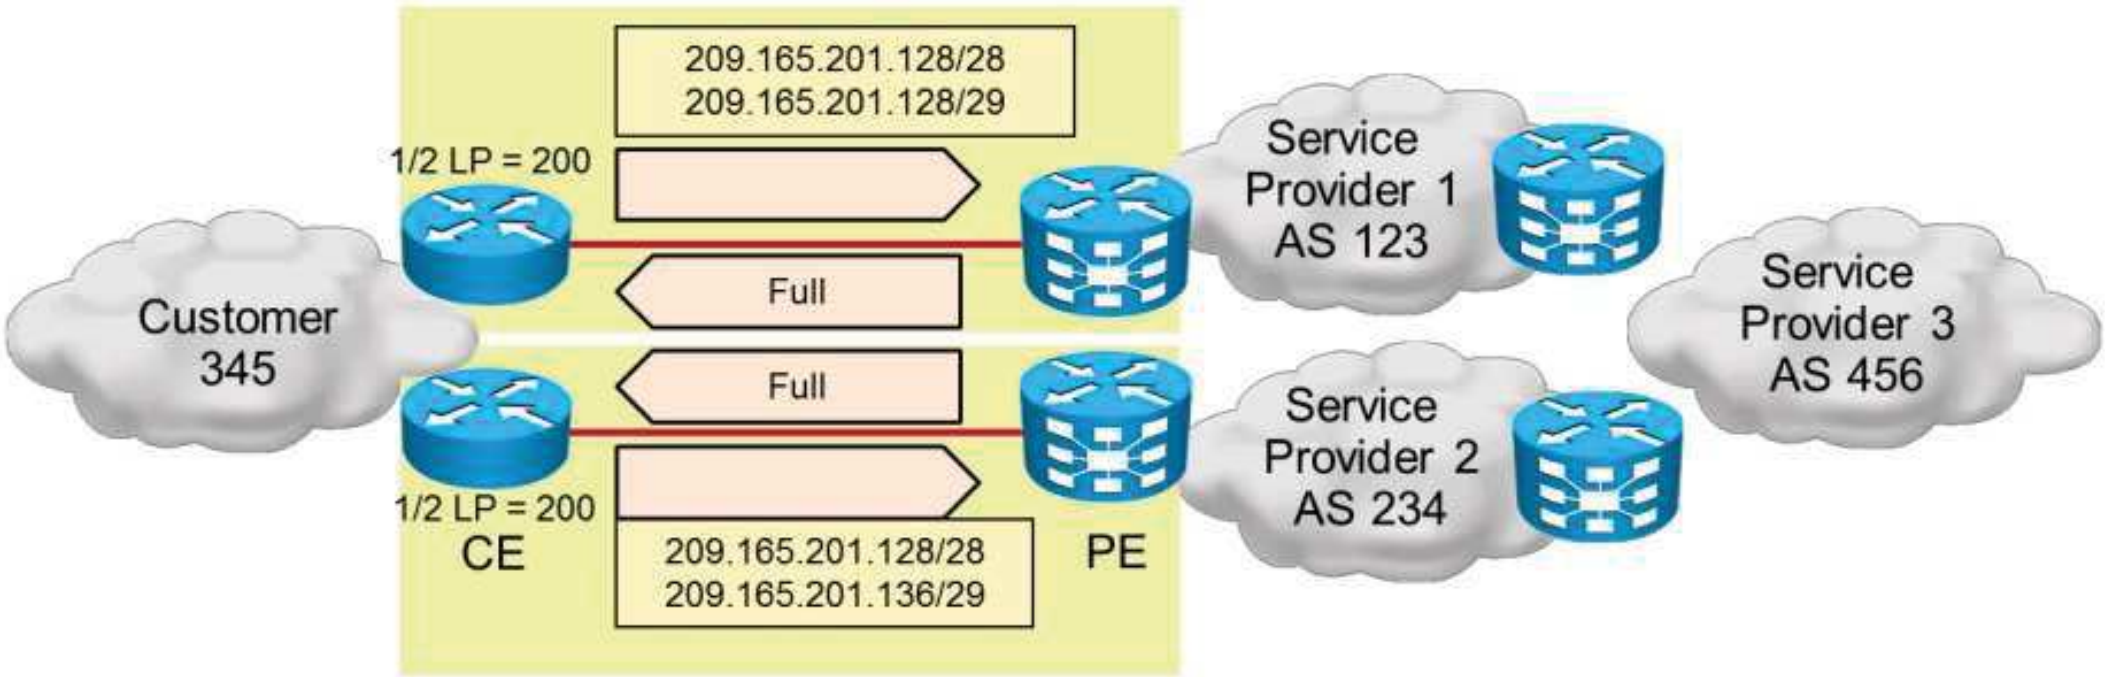

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-Implemented BGP Routing Policies in a Load-Balancing Scenario (Cont.)

Configuration of CE:

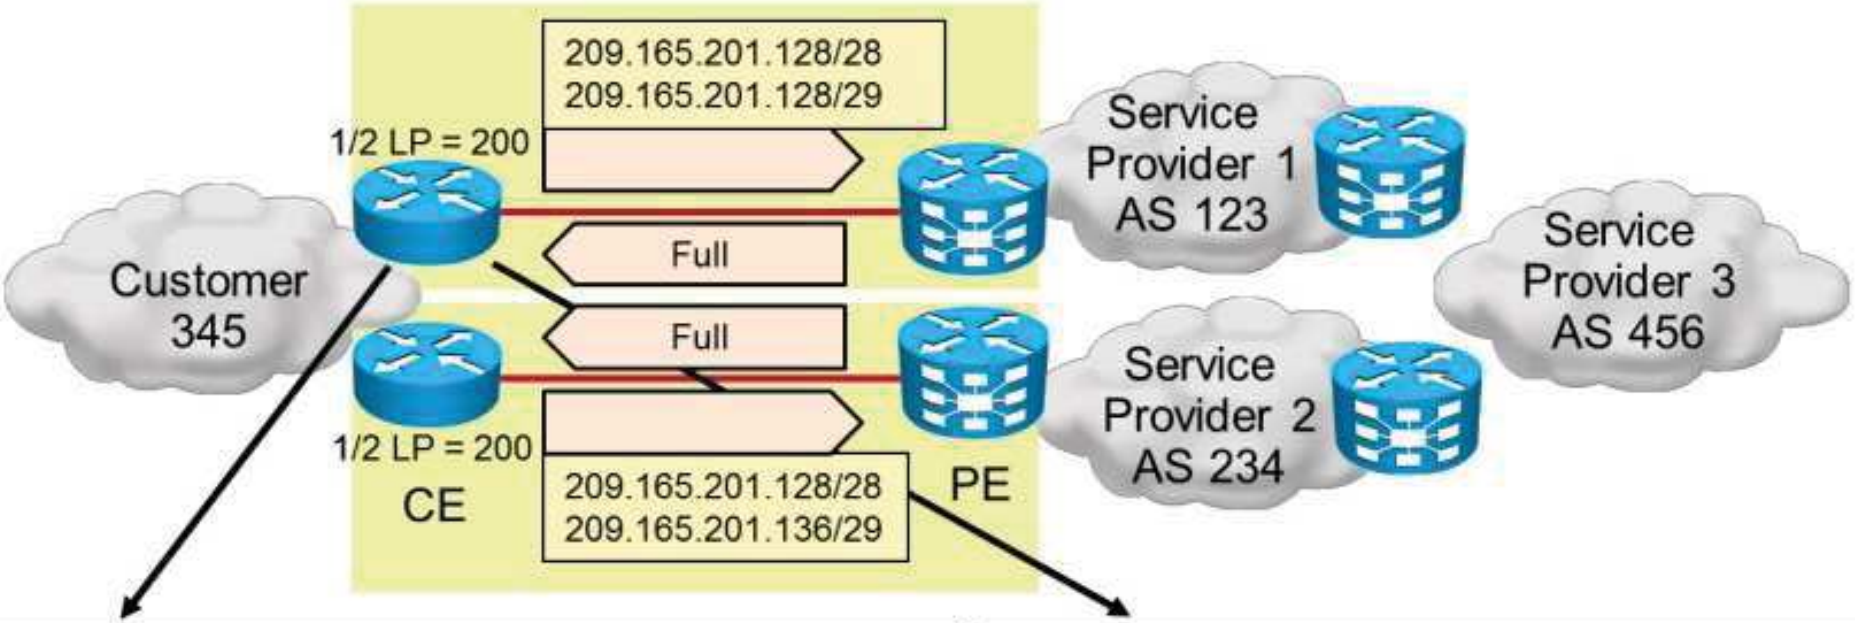


```
ip as-path access-list 1 permit [02468]$
!
route-map FROM_AS123 permit 10
 match as-path 1
 set local-preference 200
!
route-map FROM_AS123 permit 20
!
ip prefix-list TO_AS123 permit 209.165.201.128/28
ip prefix-list TO_AS123 permit 209.165.201.128/29
```

```
router bgp 345
 network 209.165.201.128 mask 255.255.255.240
 network 209.165.201.128 mask 255.255.255.248
!
 neighbor 209.165.201.1 remote-as 123
 neighbor 209.165.201.1 route-map FROM_AS123 in
 neighbor 209.165.201.1 prefix-list TO_AS123 out
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Customer-Implemented BGP Routing Policies in a Load-Balancing Scenario (Cont.)

Configuration of CE (Cont.):

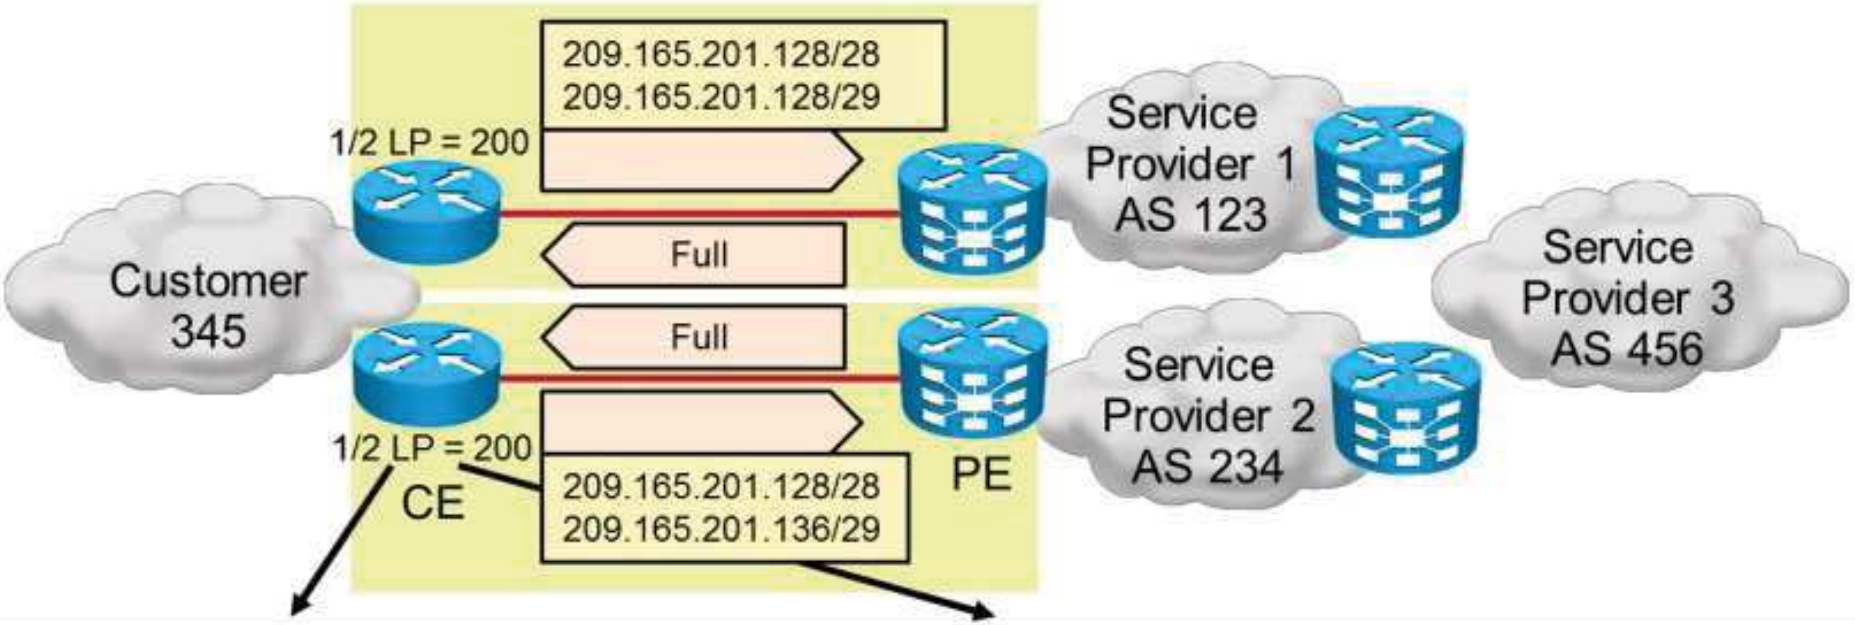


```
ip as-path access-list 1 permit [13579]$
!
route-map FROM_AS234 permit 10
 match as-path 1
 set local-preference 200
!
route-map FROM_AS234 permit 20
!
ip prefix-list TO_AS234 permit 209.165.201.128/28
ip prefix-list TO_AS234 permit 209.165.201.136/29
```

```
router bgp 345
 network 209.165.201.128 mask 255.255.255.240
 network 209.165.201.136 mask 255.255.255.248
!
 neighbor 209.165.201.5 remote-as 234
 neighbor 209.165.201.5 route-map FROM_AS234 in
 neighbor 209.165.201.5 prefix-list TO_AS234 out
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider-Aided Routing Policy: Primary and Backup

Service provider-aided routing policy using signaling:

- The customer signals “backup” service using the BGP community.
- The backup ISP sets a lower local preference on the customer routes that are based on the BGP community.
- The backup ISP does AS-path prepending toward the Internet, which is based on the BGP community.

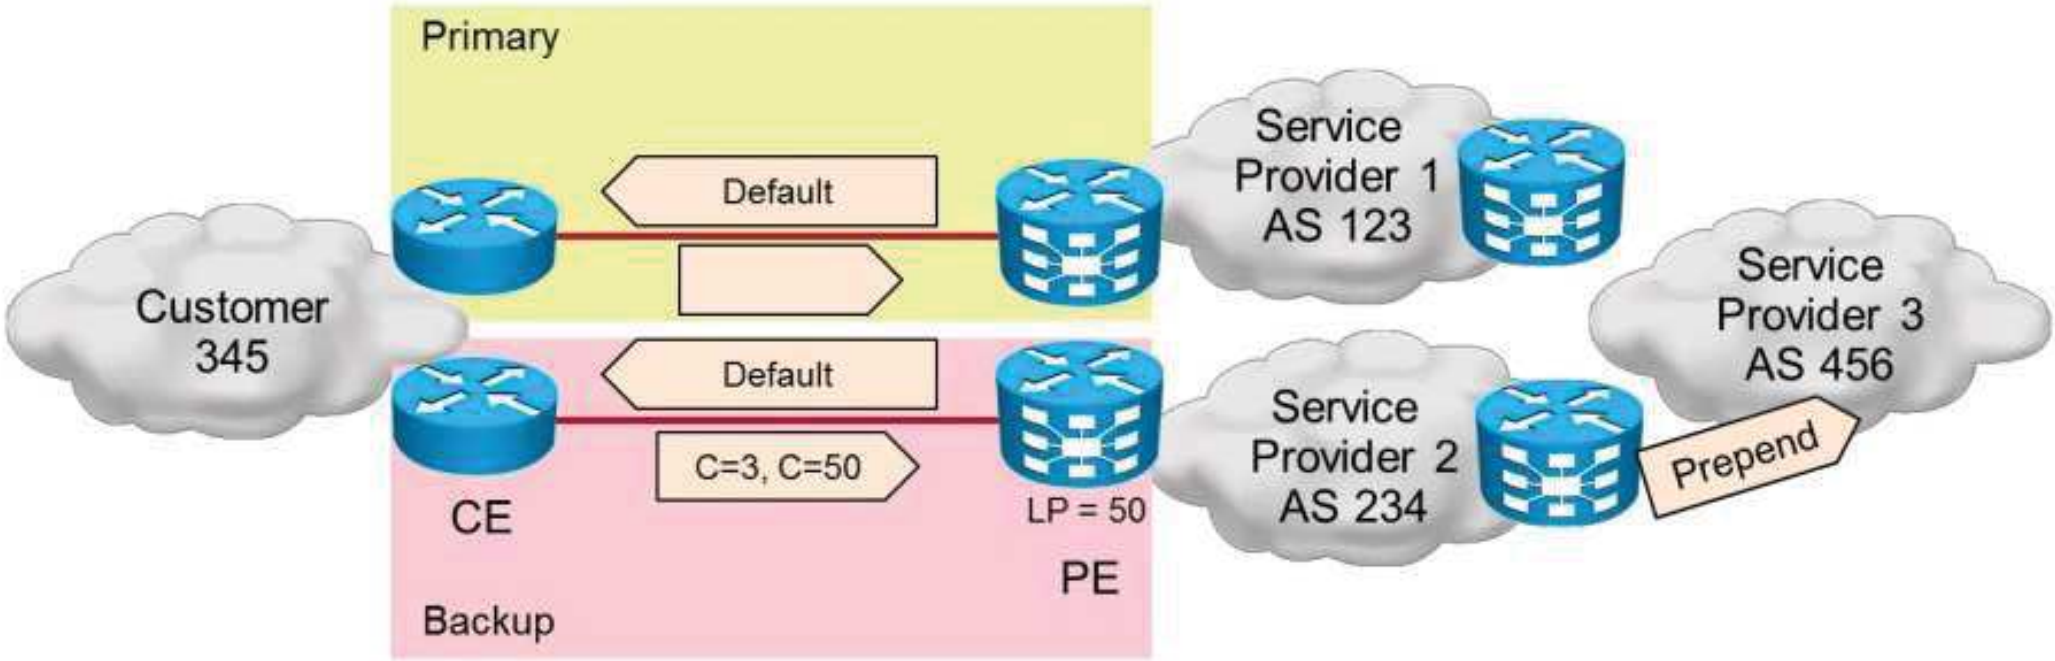

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider-Aided Routing Policy: Primary and Backup (Cont.)

Backup service provider services:

- C = 1: prepend AS number once
- C = 2: prepend AS number twice
- C = 3: prepend AS number three times
- C = 50: local preference equals 50
- C = 200: local preference equals 200

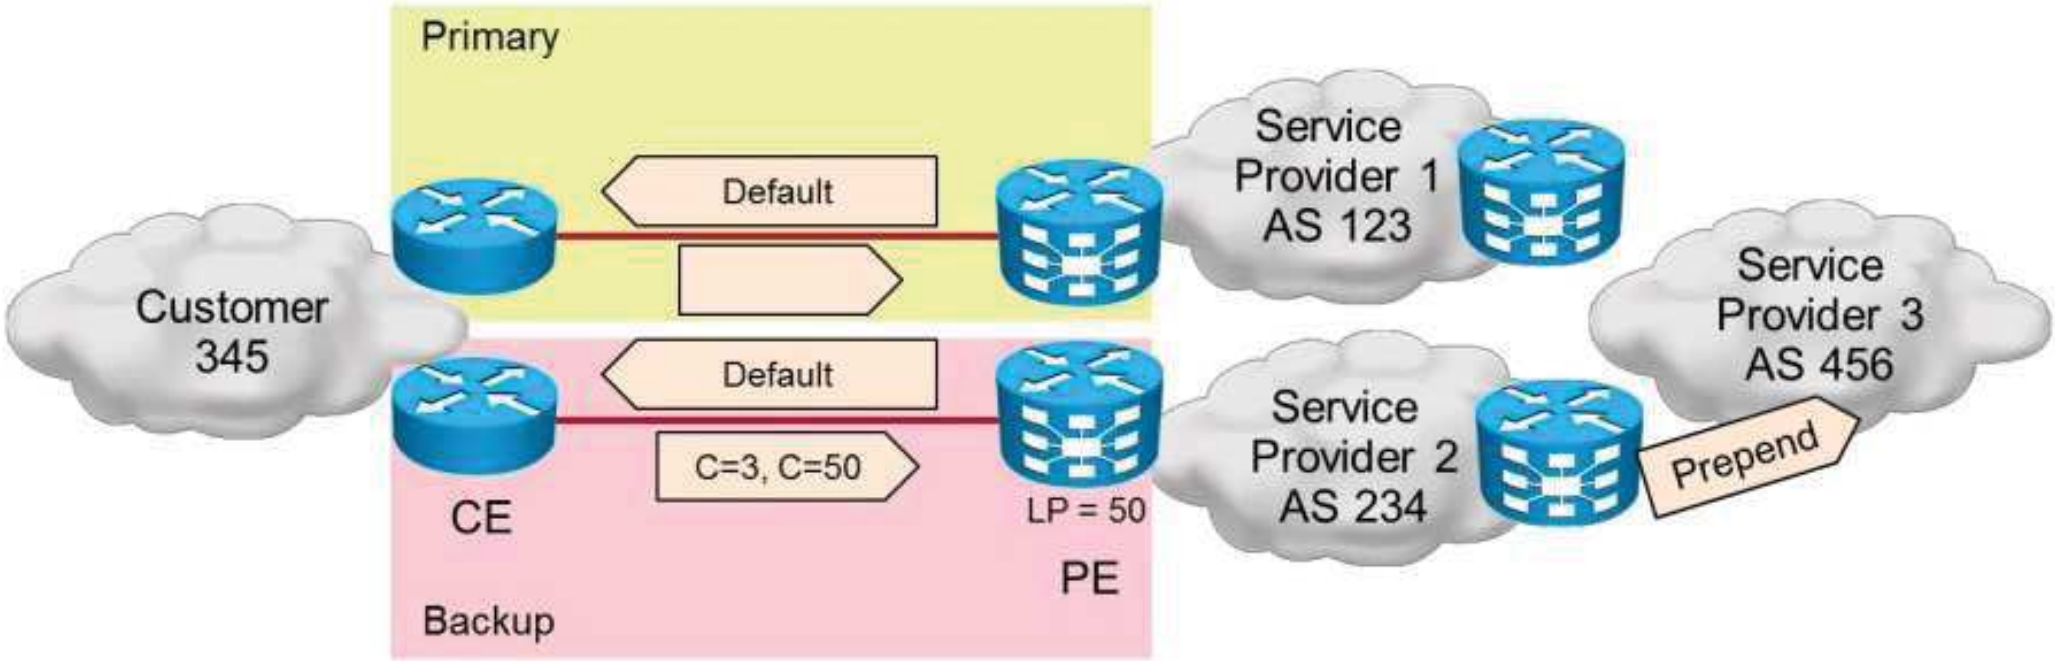

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider-Aided Routing Policy: Primary and Backup (Cont.)

Configuration of PE:

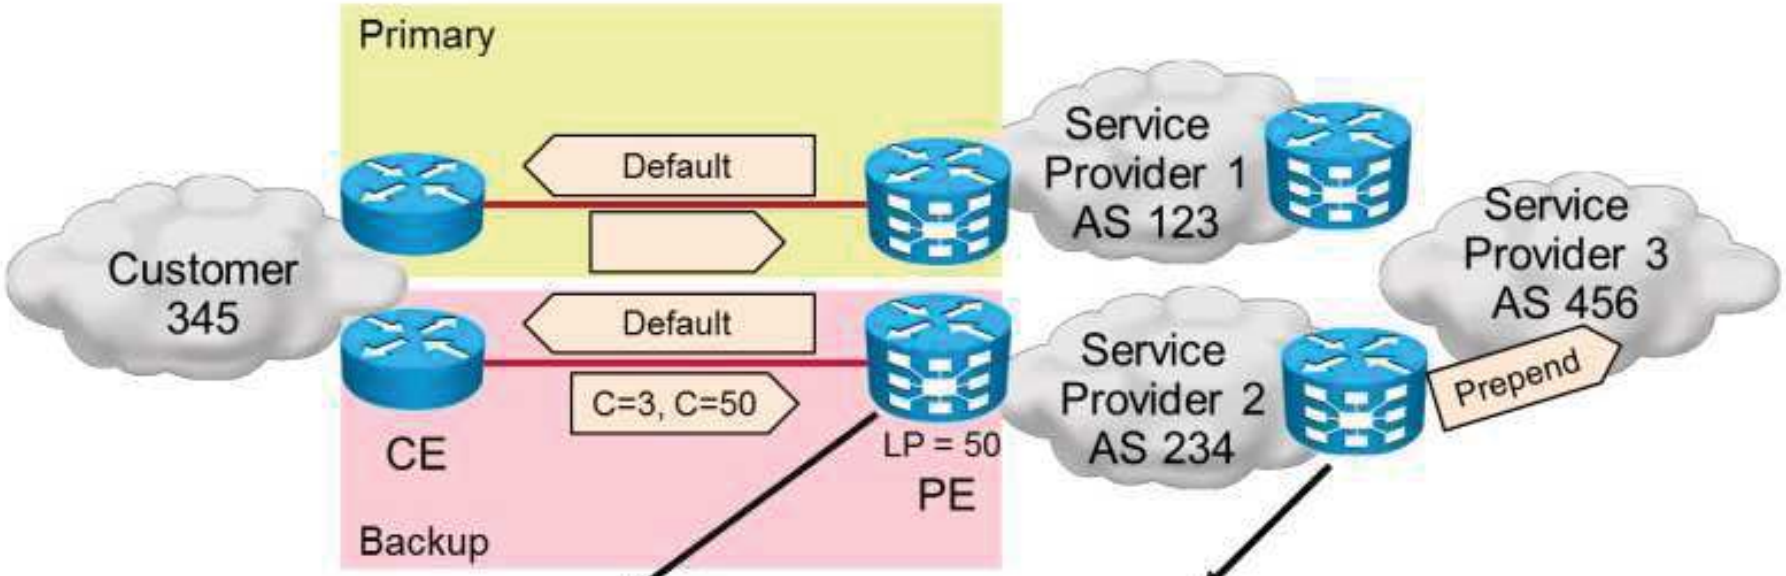


```
route-policy CUST_NET_LP
 if community matches-any (50) then
  set local-preference 50
 endif
end-policy
!
router bgp 234
 neighbor 209.165.201.6
  remote-as 345
  address-family ipv4 unicast
   route-policy CUST_NET_LP in
```

```
route-policy CUST_NET_PREPEND
 if community matches-any (3) then
  prepend as-path 234 3
 endif
end-policy
!
router bgp 234
 neighbor 209.165.201.12
  remote-as 456
  address-family ipv4 unicast
   route-policy CUST_NET_PREPEND out
```

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- Static routing can be used with:
  - Single-attached customers using a single IP address.
  - Single-attached customers using multiple IP addresses.
  - Dual-attached customers using a primary and a backup path.
  - Dual-attached customers using load balancing.
- BGP can be used in a dual-attached scenario for conditional advertising on a CE router.
- BGP is used on PE routers by service providers in dual-attached scenarios.
- BGP must remove private AS numbers in dual-attached scenarios.
- BGP is run on both CE and PE routers in dual-attached scenarios.
- The local AS feature is used to mask the real AS number.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- BGP should be used with:
  - Dual-attached customers using a primary and a backup path.
  - Dual-attached customers using load balancing.
  - Multihomed customers.
  - Multihomed customers using a primary and a backup path.
  - Multihomed customers using load balancing.
- Some service providers allow signalization of BGP services using BGP communities to aid customers using primary and backup paths.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

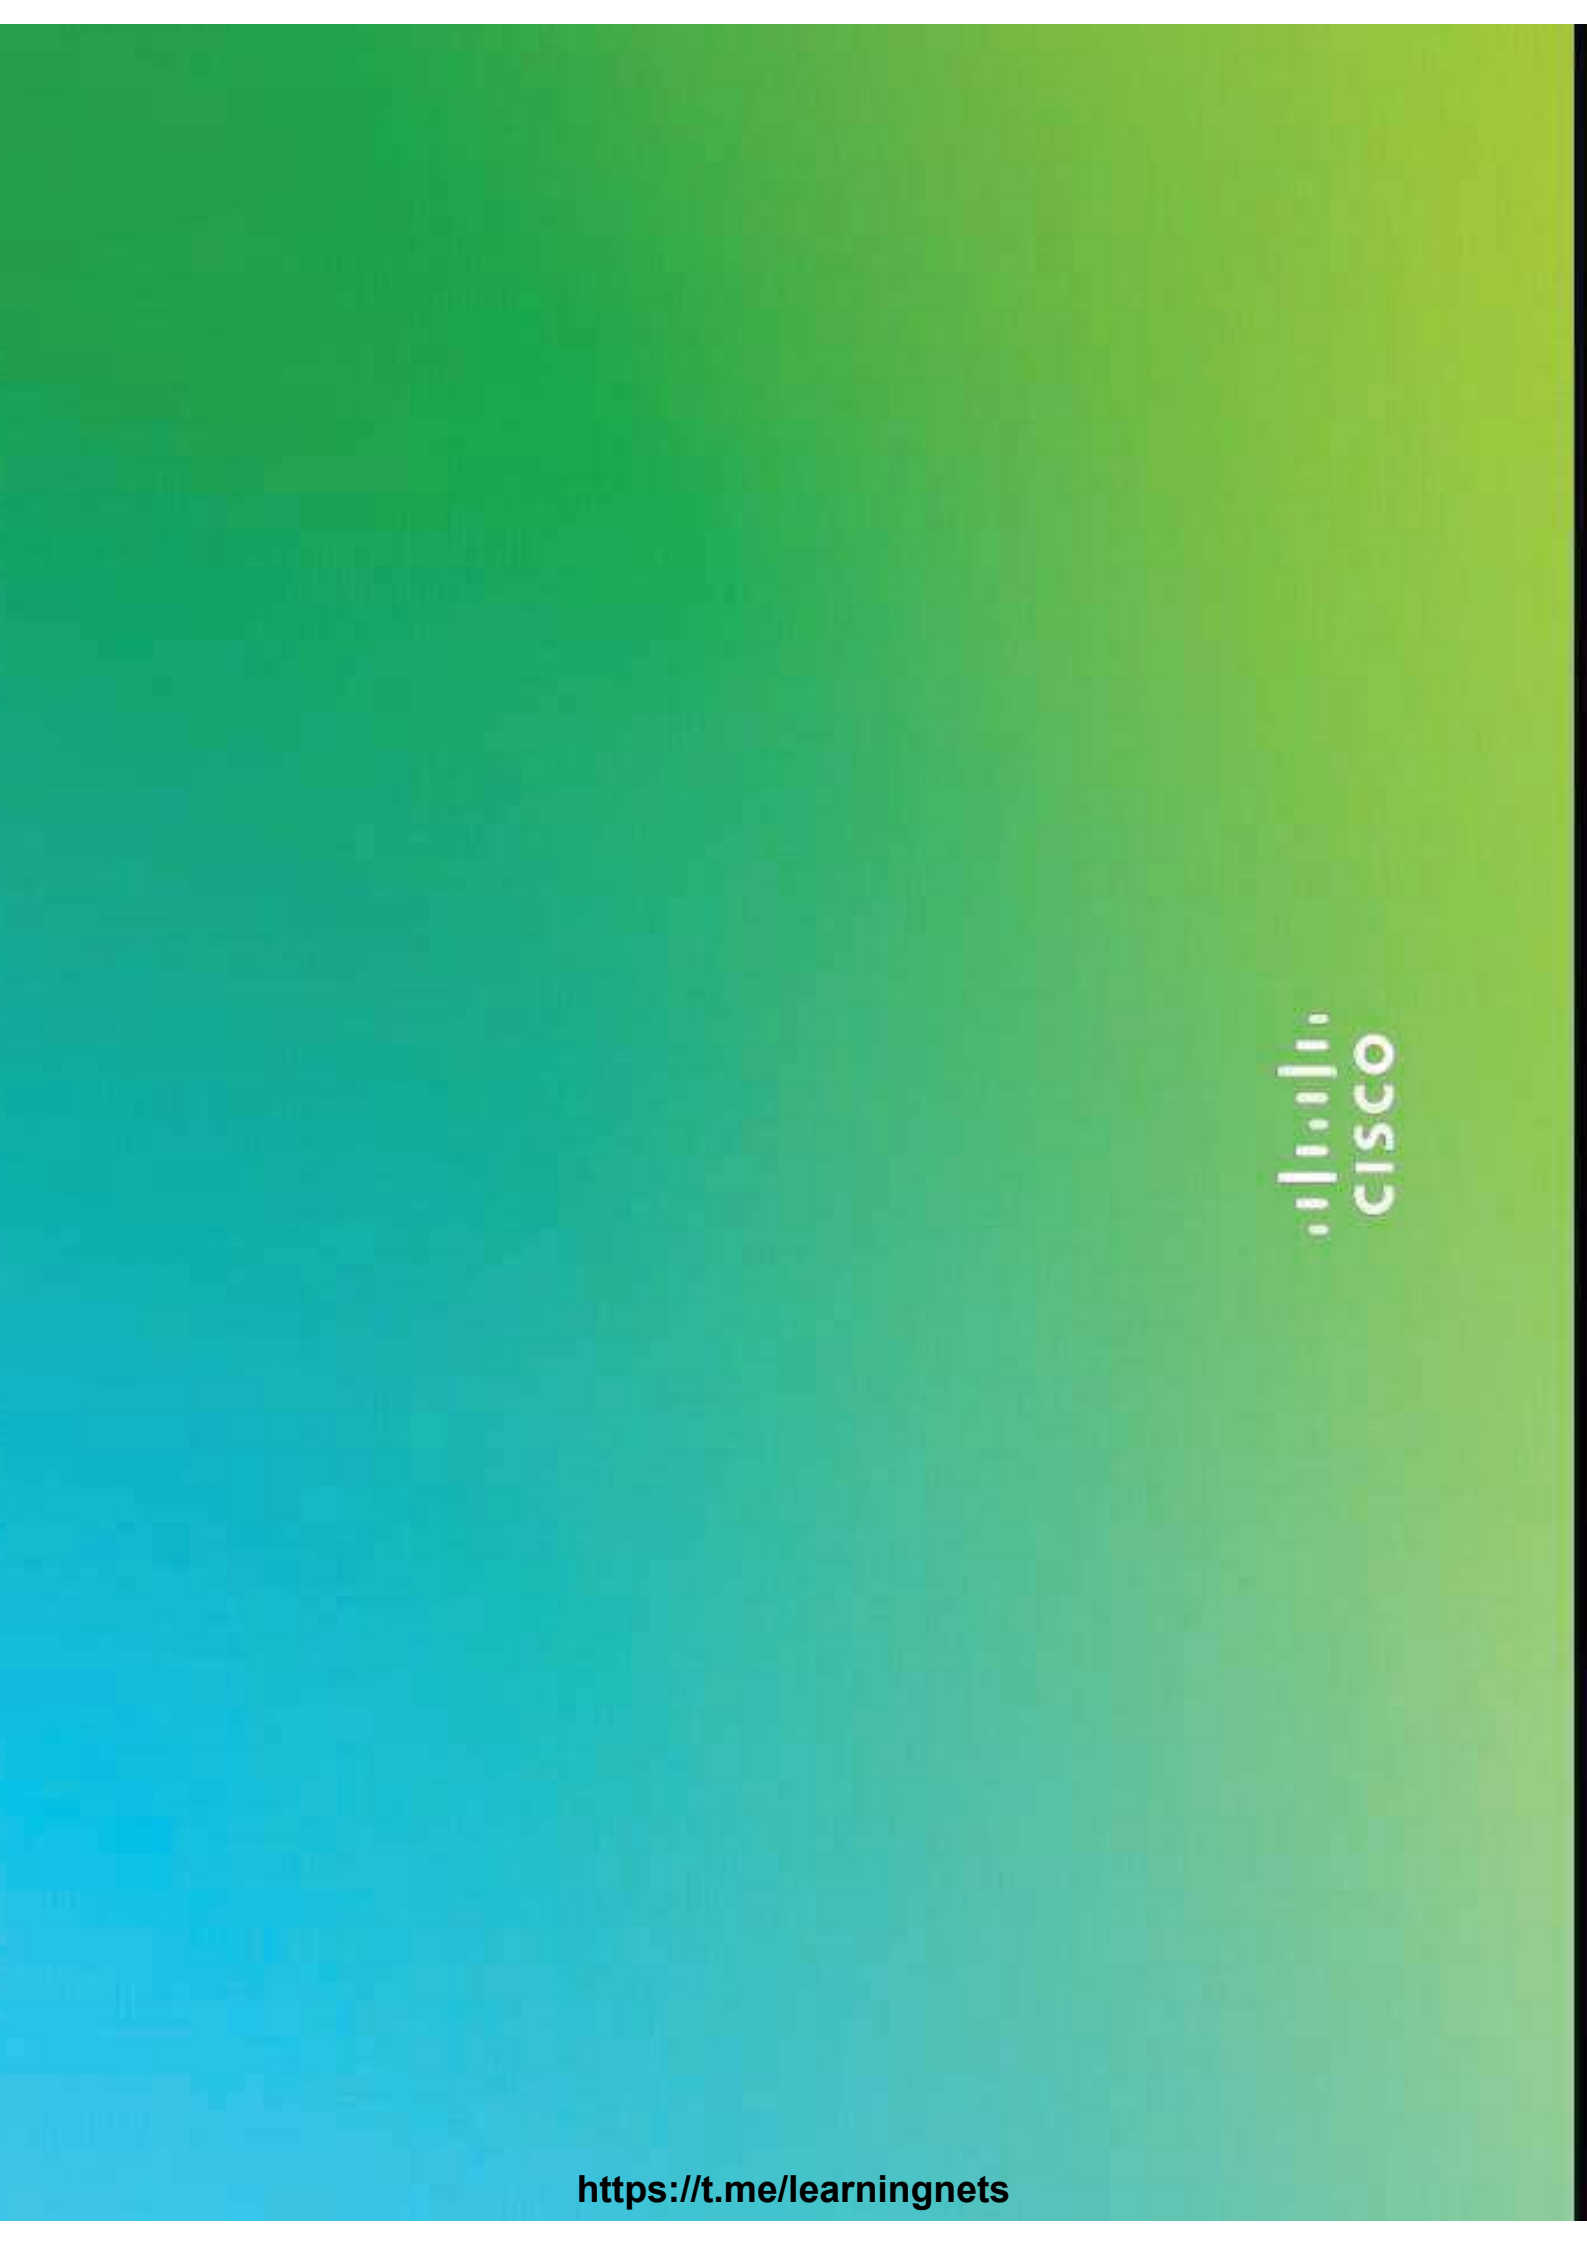

https://t.me/learningnets

Module Summary

- Customer requirements dictate the use of different connectivity types. The routing, IP addressing, and AS numbering used between a customer and service provider depend on the selected connectivity type.
- When using dual links and BGP, BGP attribute manipulation is required to achieve the required traffic distribution pattern.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

https://t.me/learningnets

https://t.me/learningnets

Scaling BGP in Service Provider Networks

Scale Service Provider Network

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco IP NGN Infrastructure Layer

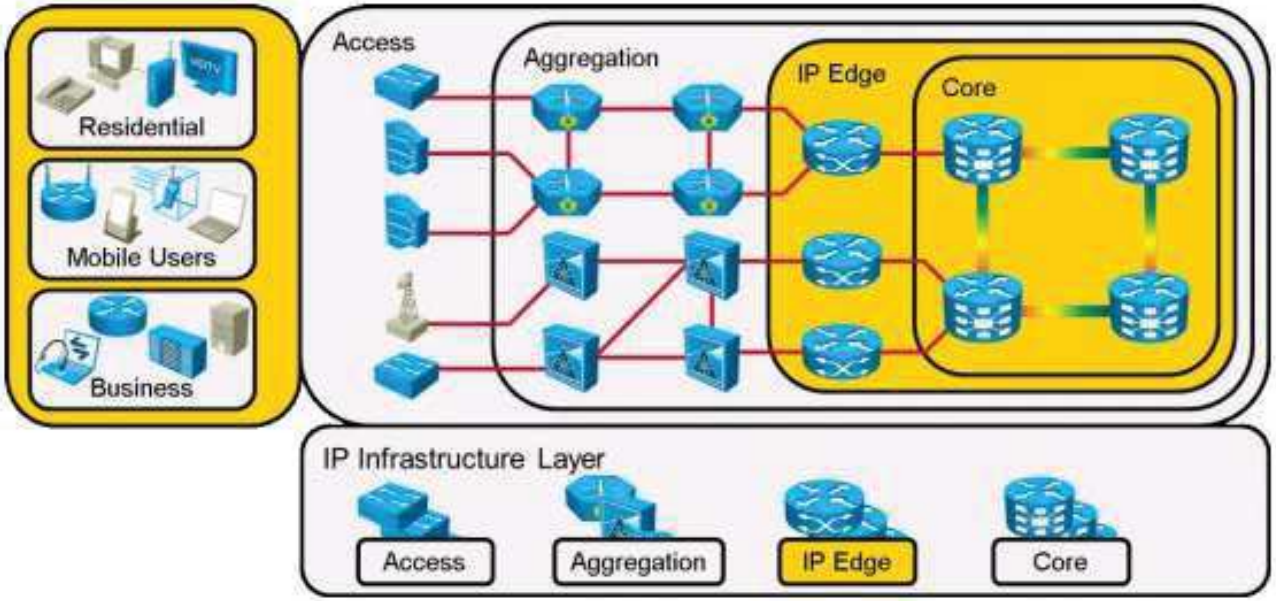

- Route propagation focuses on the IP infrastructure layer of the Cisco IP NGN.
- Route propagation focuses on the core and edge devices of the service provider and on CE devices.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider Network Routing Protocols

Service P-network routing protocols features:

- Runs BGP or static routing with customer.
- Exchanges routes with upstream service providers via BGP.
- Runs full-mesh IBGP between its own BGP routers (unless MPLS, BGP reflectors, or BGP confederations are used).
- Runs one instance of IGP (OSPF or IS-IS).
  - IGP used for internal routes only

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider Network Routing Protocols (Cont.)

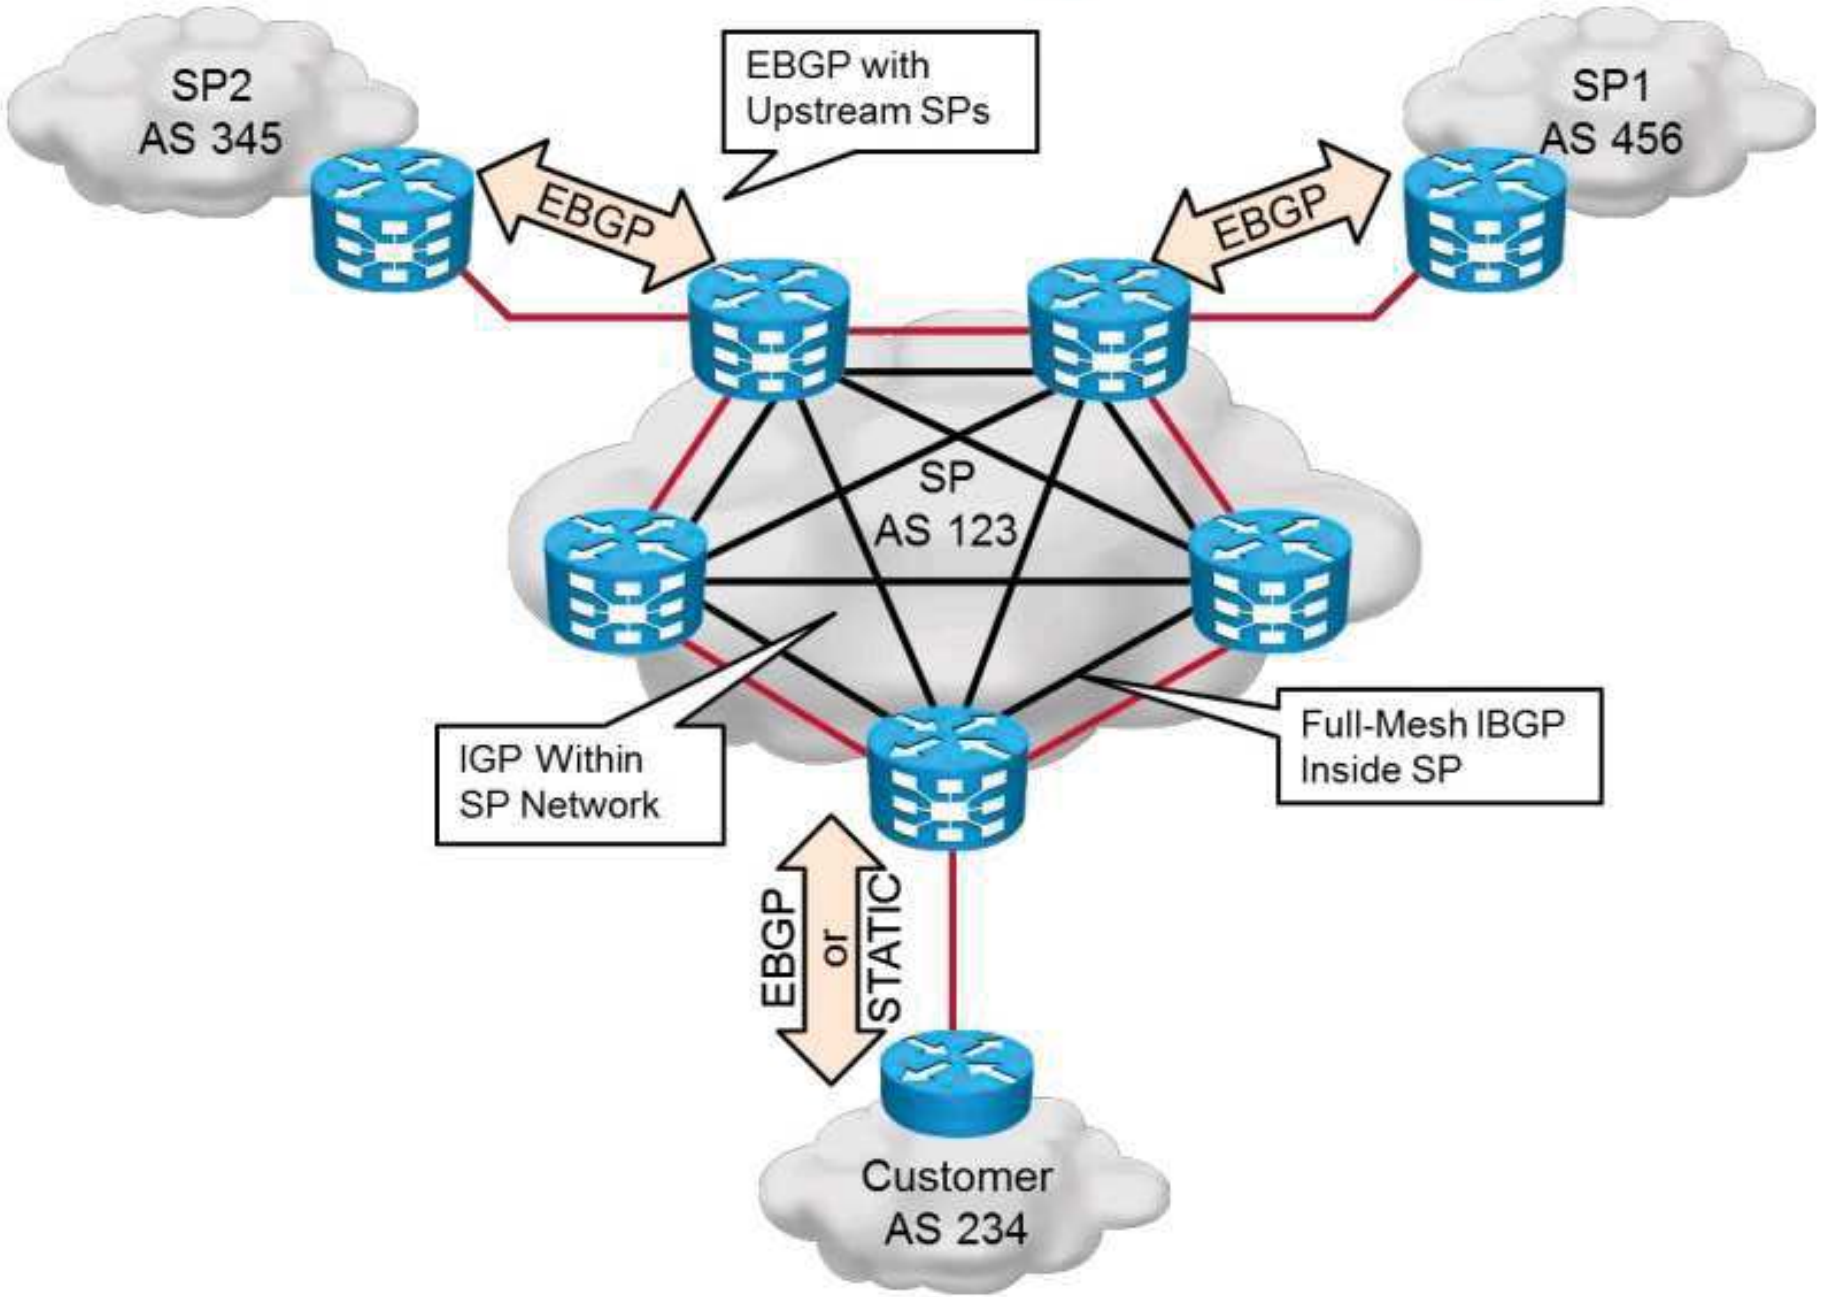

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Service Provider Network Routing Protocols (Cont.)

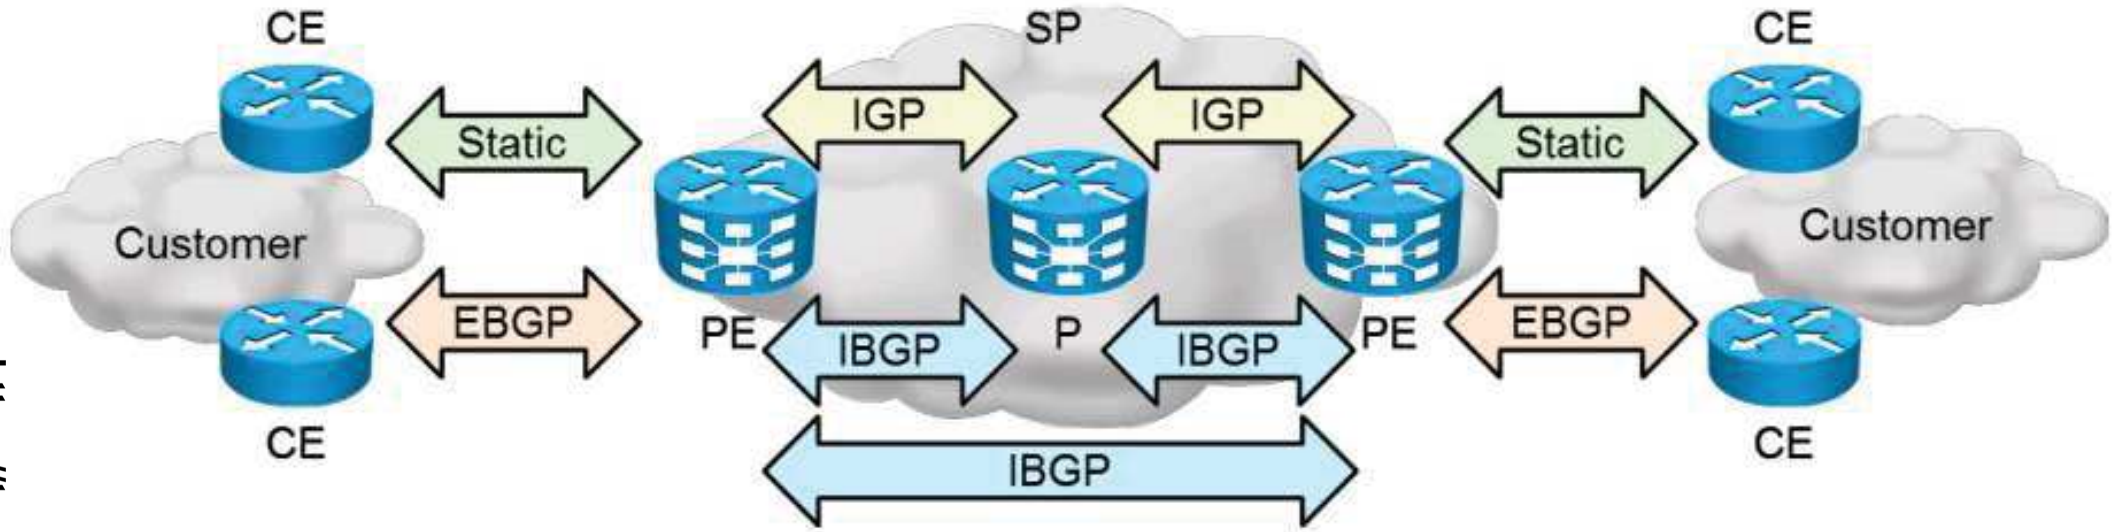

- PE routers use EBGP or static routing with CE routers.
- PE and P routers use full-mesh IBGP routing.
- The provider core IGP is a single instance of IS-IS or OSPF and is used only within the service provider core network.
- Optimal routing between PEs is desired.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Route Propagation in Service Provider Networks

- BGP route propagation:
  - BGP carries other service provider routes.
  - BGP carries customer routes.
- IGP route propagation:
  - IGP is responsible only for the resolution of BGP next hop and internal routes.
- Do not redistribute BGP into IGP:
  - IGP performance and convergence time suffer if a large number of routes are carried.
  - No IGP is capable of carrying full Internet routes.
  - A full Internet IPv4 routing table has exceeded 300,000 routes.
  - A full Internet IPv6 routing table has exceeded 30,000 routes.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Route Information Exchange Between Service Providers

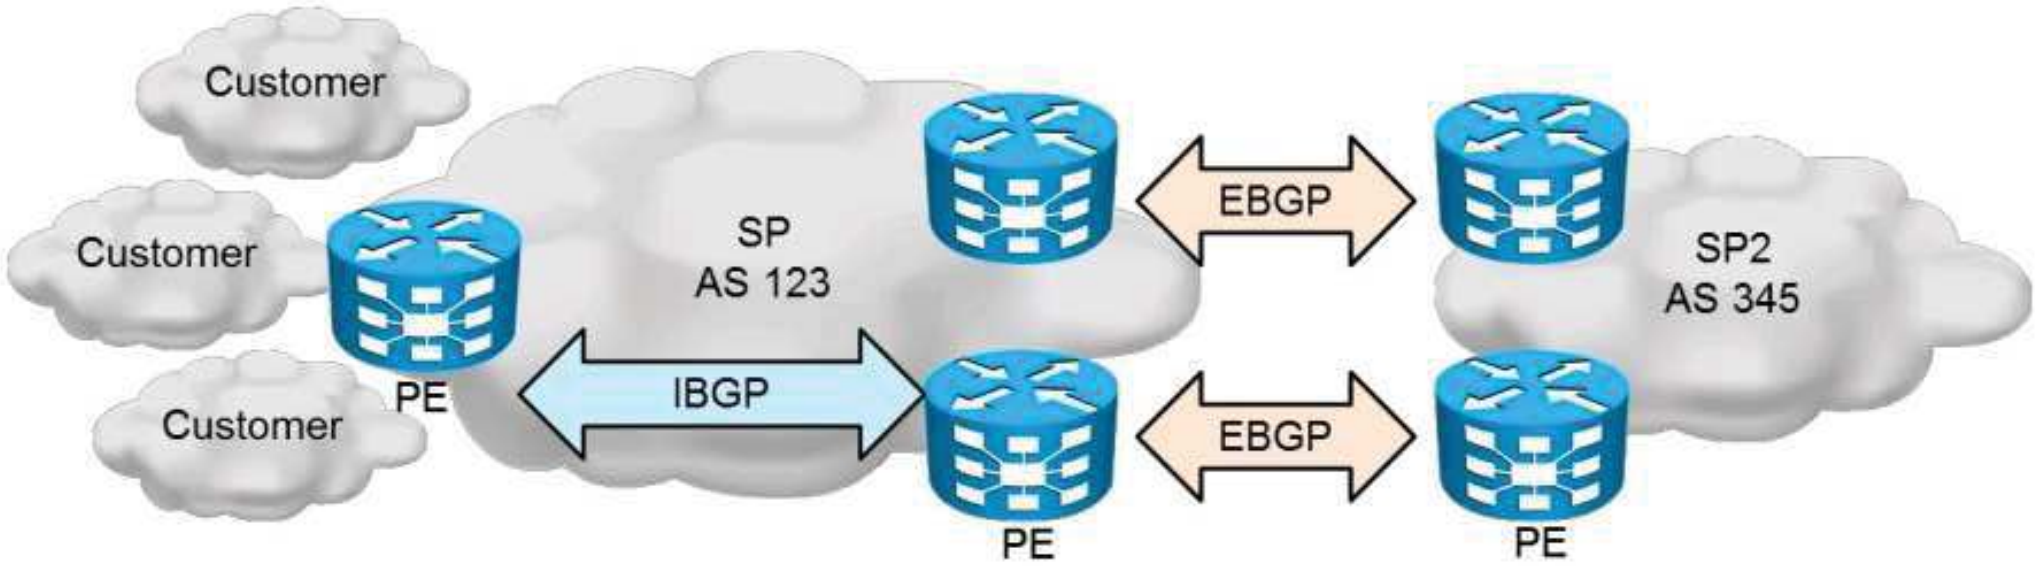

BGP is used to exchange routing information with upstream service providers:

- Service provider sends summary of service provider-owned address space to upstream service provider.
- Service provider sends prefixes owned by customers using independent address space.
- Upstream service provider sends full Internet routing table to the service provider.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Route Information Exchange with Customers

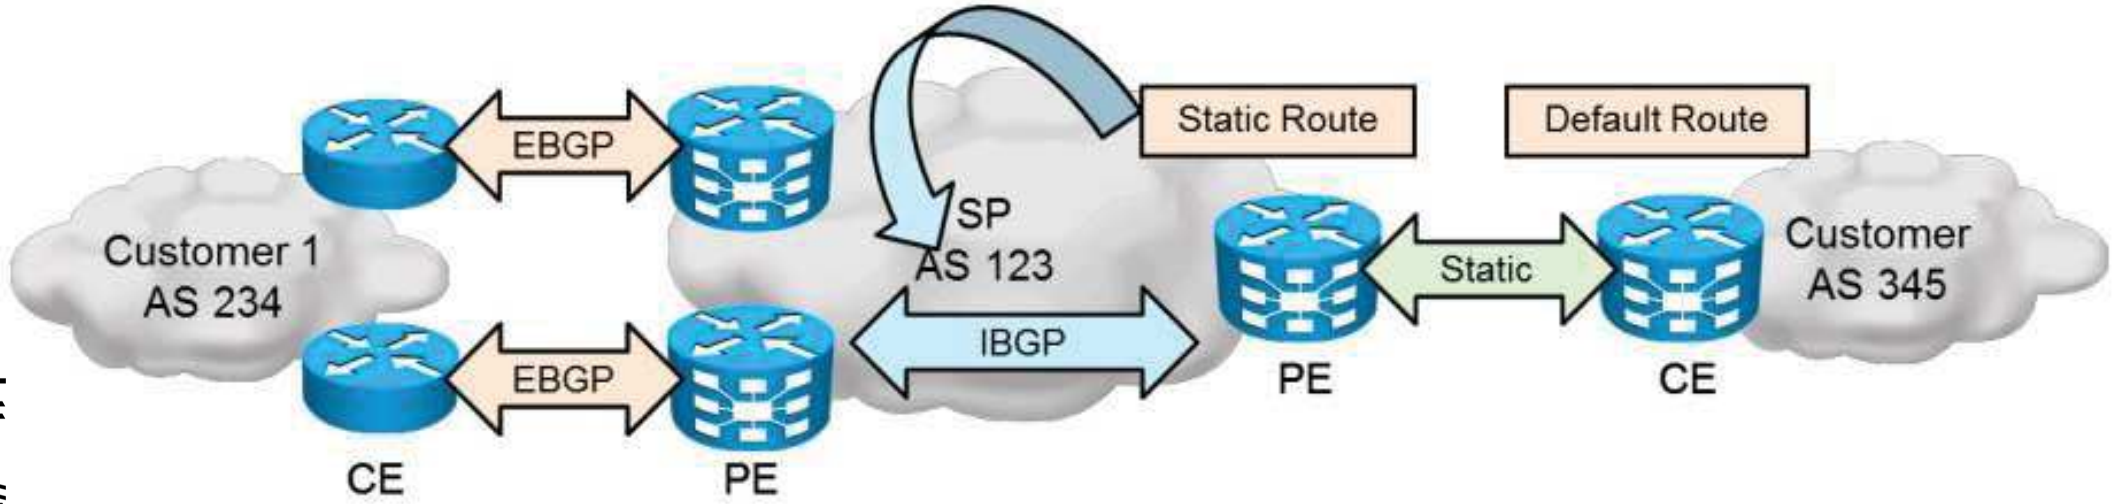

- BGP with customer:
  - Customer advertises its address space.
  - Service provider advertises default route, service provider-owned routes and default route, or full Internet routing table.
- Static routing with customer:
  - Customer uses default route.
  - Service provider uses static route on the PE router for customer address space. Static route is redistributed into BGP on the PE router.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Next-Hop Resolution with IGP

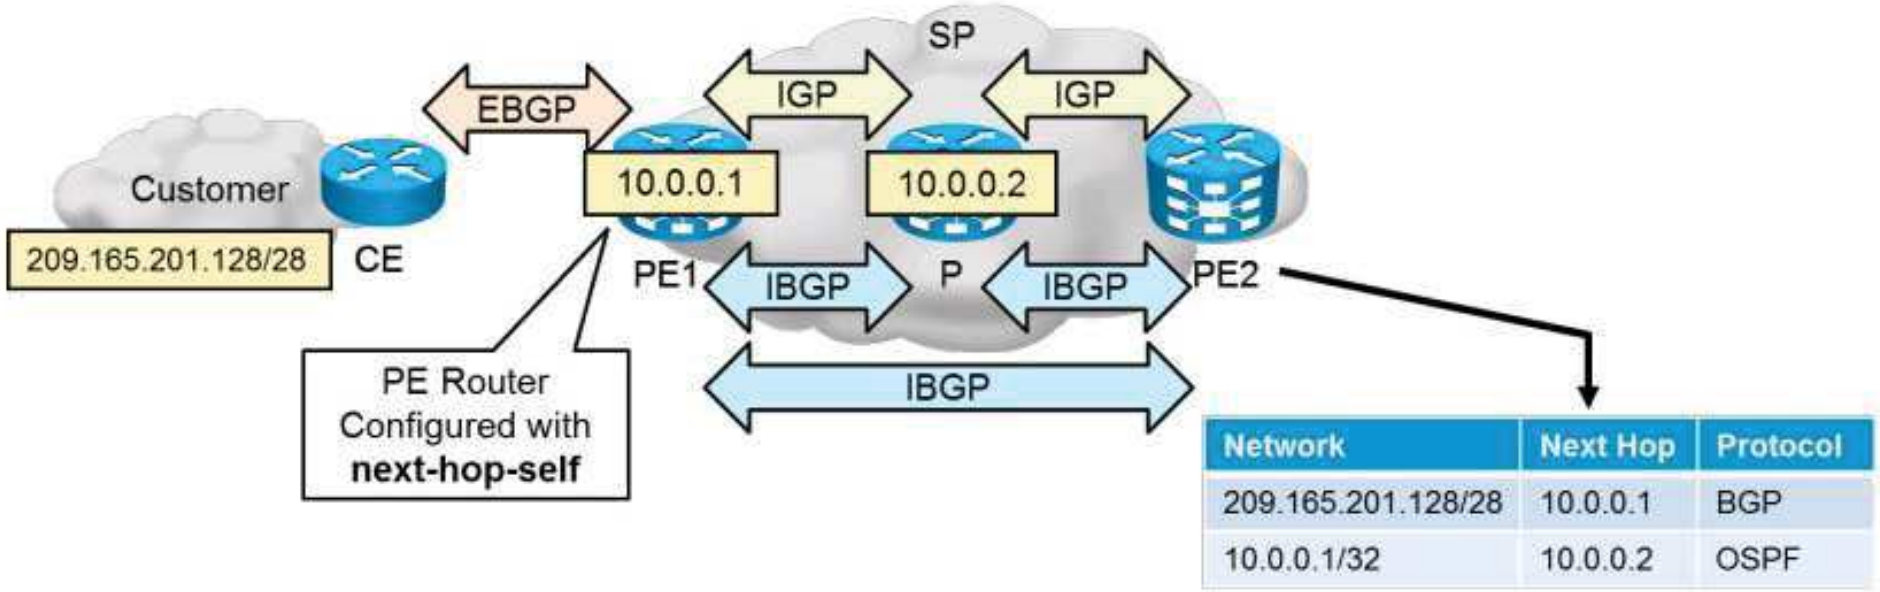

Network	Next Hop	Protocol
209.165.201.128/28	10.0.0.1	BGP
10.0.0.1/32	10.0.0.2	OSPF

- **next-hop-self** on the PE routers removes the need to include access links in IGP, and thus prevents route flapping if access link flaps.
- The service provider core IGP should carry information only about core links and loopback addresses.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Scaling BGP Routing

- BGP policy scaling:
  - The AS routing policy should be uniform and easy to maintain.
  - This goal is achieved by reusing the same configuration in all EBGP routers.
- IBGP scaling:
  - Full-mesh IBGP is not needed since there are other technologies and features available.
- Updates and table size scaling:
  - Route summarization is the key to scalability.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Scaling Addressing in Service Provider Core Network

This list shows how to scale addressing in Service Provider Core Network:

- IPv4:
  - Private or public IP addresses can be used.
  - Private addresses on core links and loopbacks display private IP addresses in a traceroute when run from customers.
    - MPLS with TTL propagation disabled solves the traceroute issue.
  - Private addresses on loopbacks and core links call for careful external routing to prevent advertisement of private addresses to customers or upstream service providers.
  - Otherwise, use public addresses in service provider core networks.
- IPv6:
  - On the core links, only link-local IPv6 addresses can be used.
  - On the loopback interfaces, public IPv6 address should be used.
    - There are no traceroute issues, because transit IPv6-enabled router will always respond from loopback interface.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Policy Accounting

This list shows BGP Policy Accounting characteristics:

- Measures and classifies IP traffic that is sent to, or received from, different peers.
- Accounts for traffic according to the route that it traverses.
- Routes are classified and traffic is measured based on BGP communities, AS number, or AS path.
  - Based on the classification policy, BGP policy accounting assigns each prefix a traffic index (bucket).
- BGP policy accounting can be applied in ingress or egress direction on an interface, where the traffic source IP address, the destination IP address, or both are BGP prefixes.
- Used for:
  - Billing for the traffic routing from customers.
  - Examining and improving design of BGP peering and BGP routing policies.
- Supported for IPv4 only.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Policy Accounting Configuration

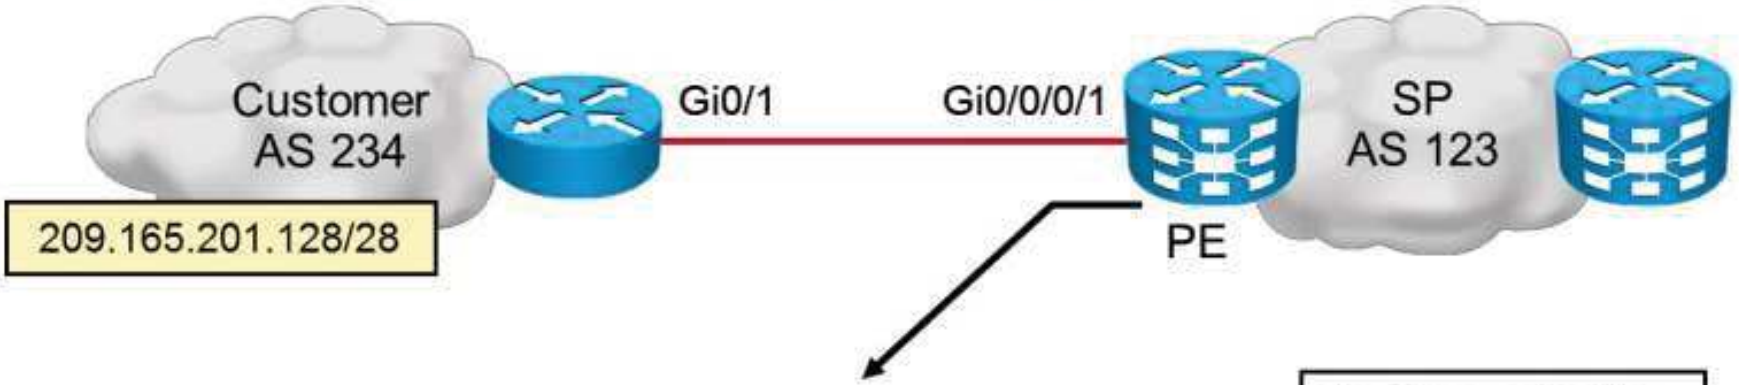


```
route-policy BGP_ACCOUNTING
 if as-path originates-from '234' then
   set traffic-index 11
 endif
end-policy
!
router bgp 123
 address-family ipv4 unicast
  table-policy BGP_ACCOUNTING
!
interface GigabitEthernet0/0/0/1
 ipv4 bgp policy accounting input source-accounting
```

Assigns a traffic index to routes.

Classifies BGP prefixes entered in the routing table.

Enables BGP policy accounting on an interface in ingress direction based on source IP addresses.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Policy Accounting Verification

```
RP/0/RSP0/CPU0:PE# show cef 209.165.201.128/28 detail
Mon Sep 19 13:32:22.655 UTC
209.165.201.128/28, version 4, internal 0x4000001 (ptr 0xad958768) [1], 0x0
(0x0), 0x0 (0x0)
 Updated Sep 19 13:30:11.717
 Prefix Len 24, traffic index 11, precedence routine (0)
```

- Displays assigned traffic index for a prefix.

```
RP/0/RSP0/CPU0:PE# show cef interface GigabitEthernet0/0/0/1 bgp-policy-
statistics
GigabitEthernet0/0/0/1 is UP
Input BGP policy accounting on src IP address enabled
  buckets           packets       bytes
  11                17406     2088447
```

- Displays per-interface traffic statistics.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- Route propagation focuses on the IP infrastructure layer of the Cisco IP NGN.
- Service providers most commonly use integrated IS-IS and OSPF as interior gateway protocols and BGP as the exterior gateway protocol.
- BGP is used to carry customer routes while IGPs are used to carry service provider internal prefix reachability information.
- BGP allows ISP clients to acquire information about all or some networks reachable through the ISP.
- Static routing or BGP can be used by the ISP to direct traffic going to the customers to the correct links.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Next-hop-self can be used to avoid redistributing transit segments into IGP on iBGP neighbors.
- When BGP networks grow, various actions must be taken to make them scalable, for iBGP scalability use route reflectors or confederations.
- When IP networks grow, several aspects of addressing need to be considered to reduce sizes of routing tables and to avoid consuming too many addresses.
- BGP accounting feature can be used when an overview of BGP's use of resources or detailed statistical analysis are required.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

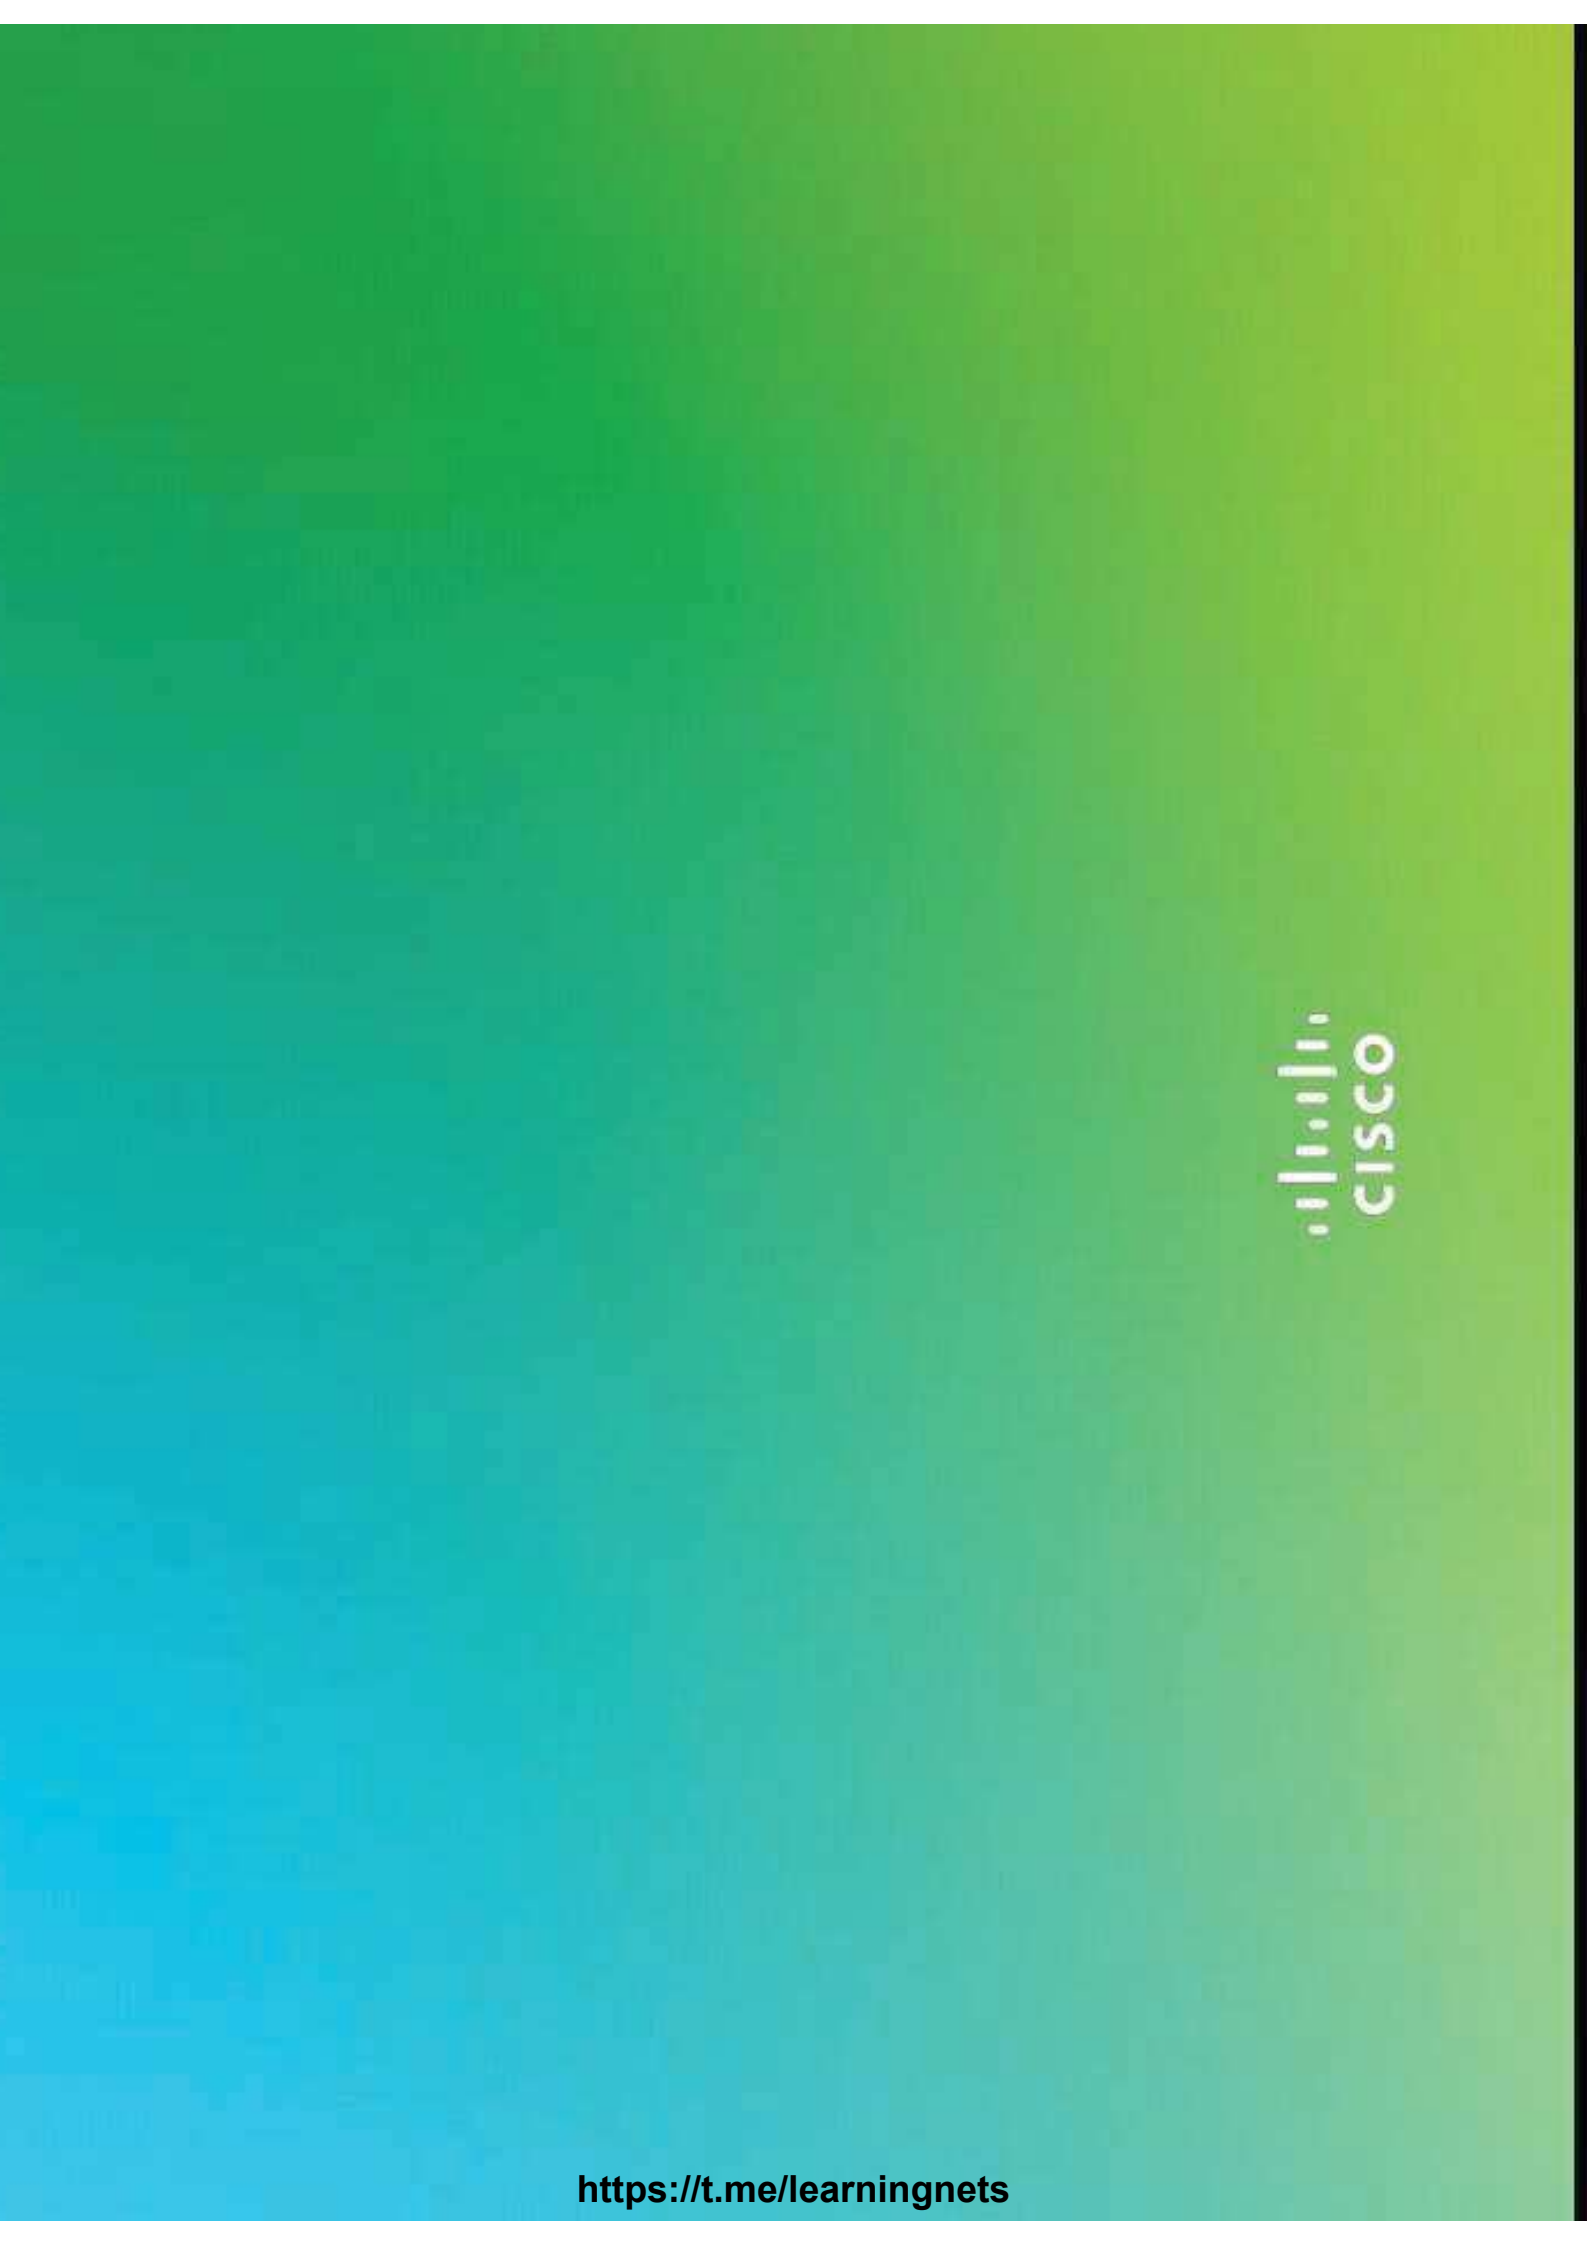

https://t.me/learningnets

https://t.me/learningnets

Introducing BGP Route Reflectors and Confederations

Scale Service Provider Network

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Reflectors and BGP Confederations

IBGP Scalability Issues

- IBGP requires a full mesh between all BGP-speaking routers:
  - Large number of TCP sessions.
  - Unnecessary duplicate routing traffic.
  - Configuration overhead.
- Solutions:
  - Route reflectors modify IBGP split-horizon rules.
  - BGP confederations modify IBGP AS path processing.

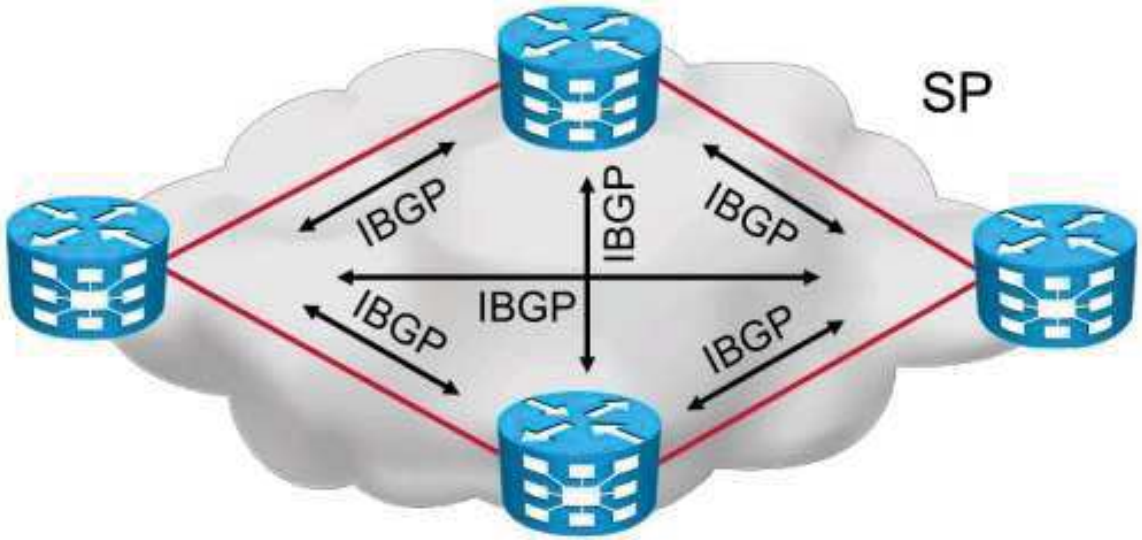

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Split-Horizon Rule

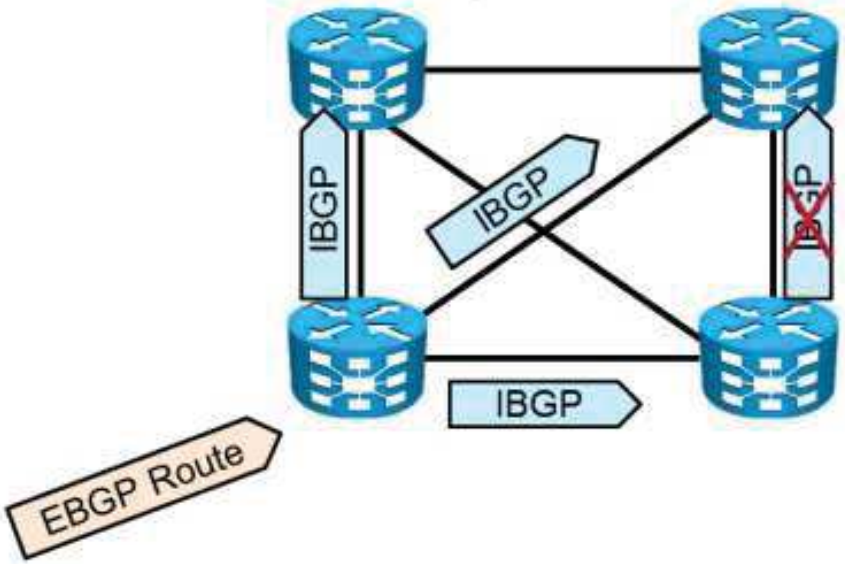

- Classic IBGP:
  - IBGP routes are not propagated to other IBGP peers.
- Full mesh of IBGP peers is therefore required.

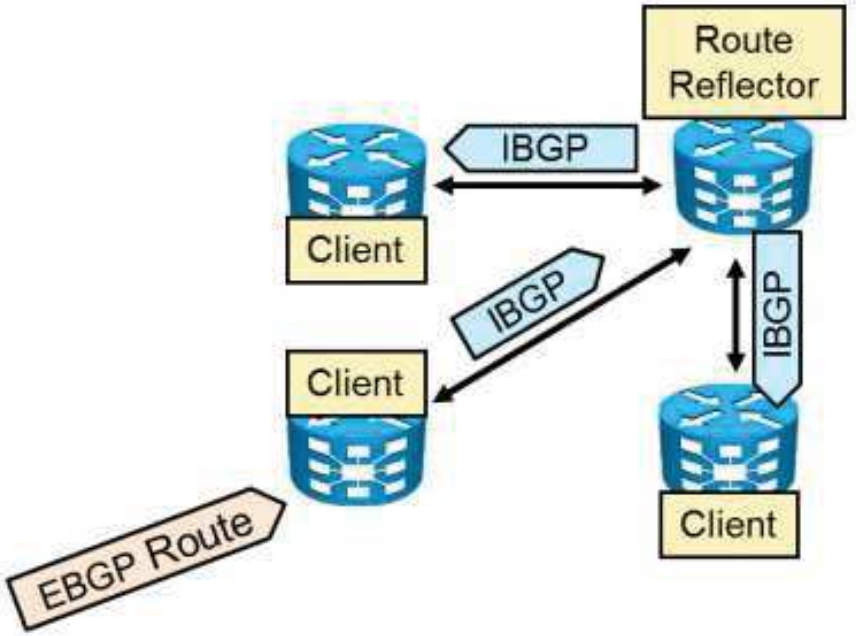

- Route reflector can propagate IBGP routes to other IBGP peers.
- Full mesh of IBGP peers is not required.
- Route reflector-based network includes route reflectors and clients.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Split-Horizon Rule (Cont.)

Route Reflector Split-Horizon Rule

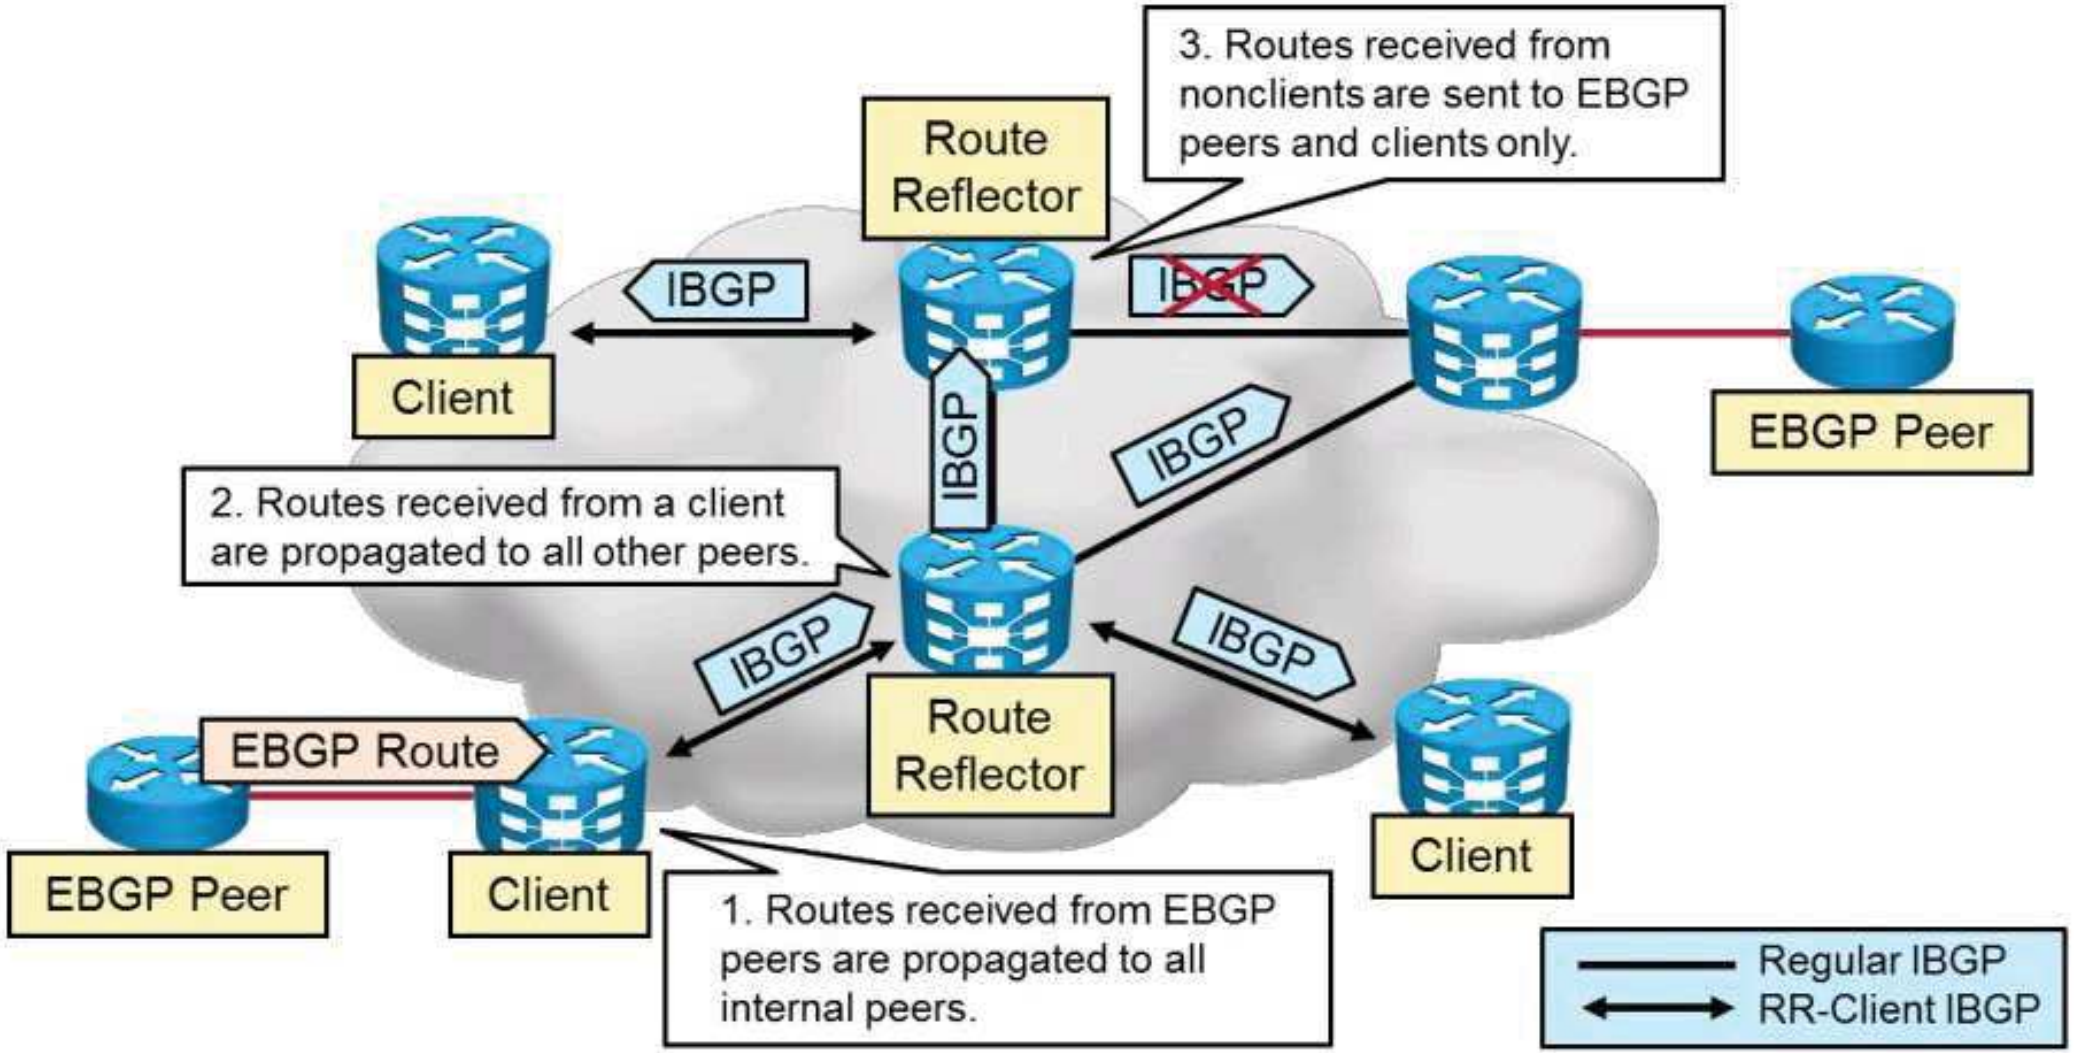

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Steps of Route Propagation in a Route Reflector-Enabled Network (Cont.)

Route Reflector Split-Horizon Rule

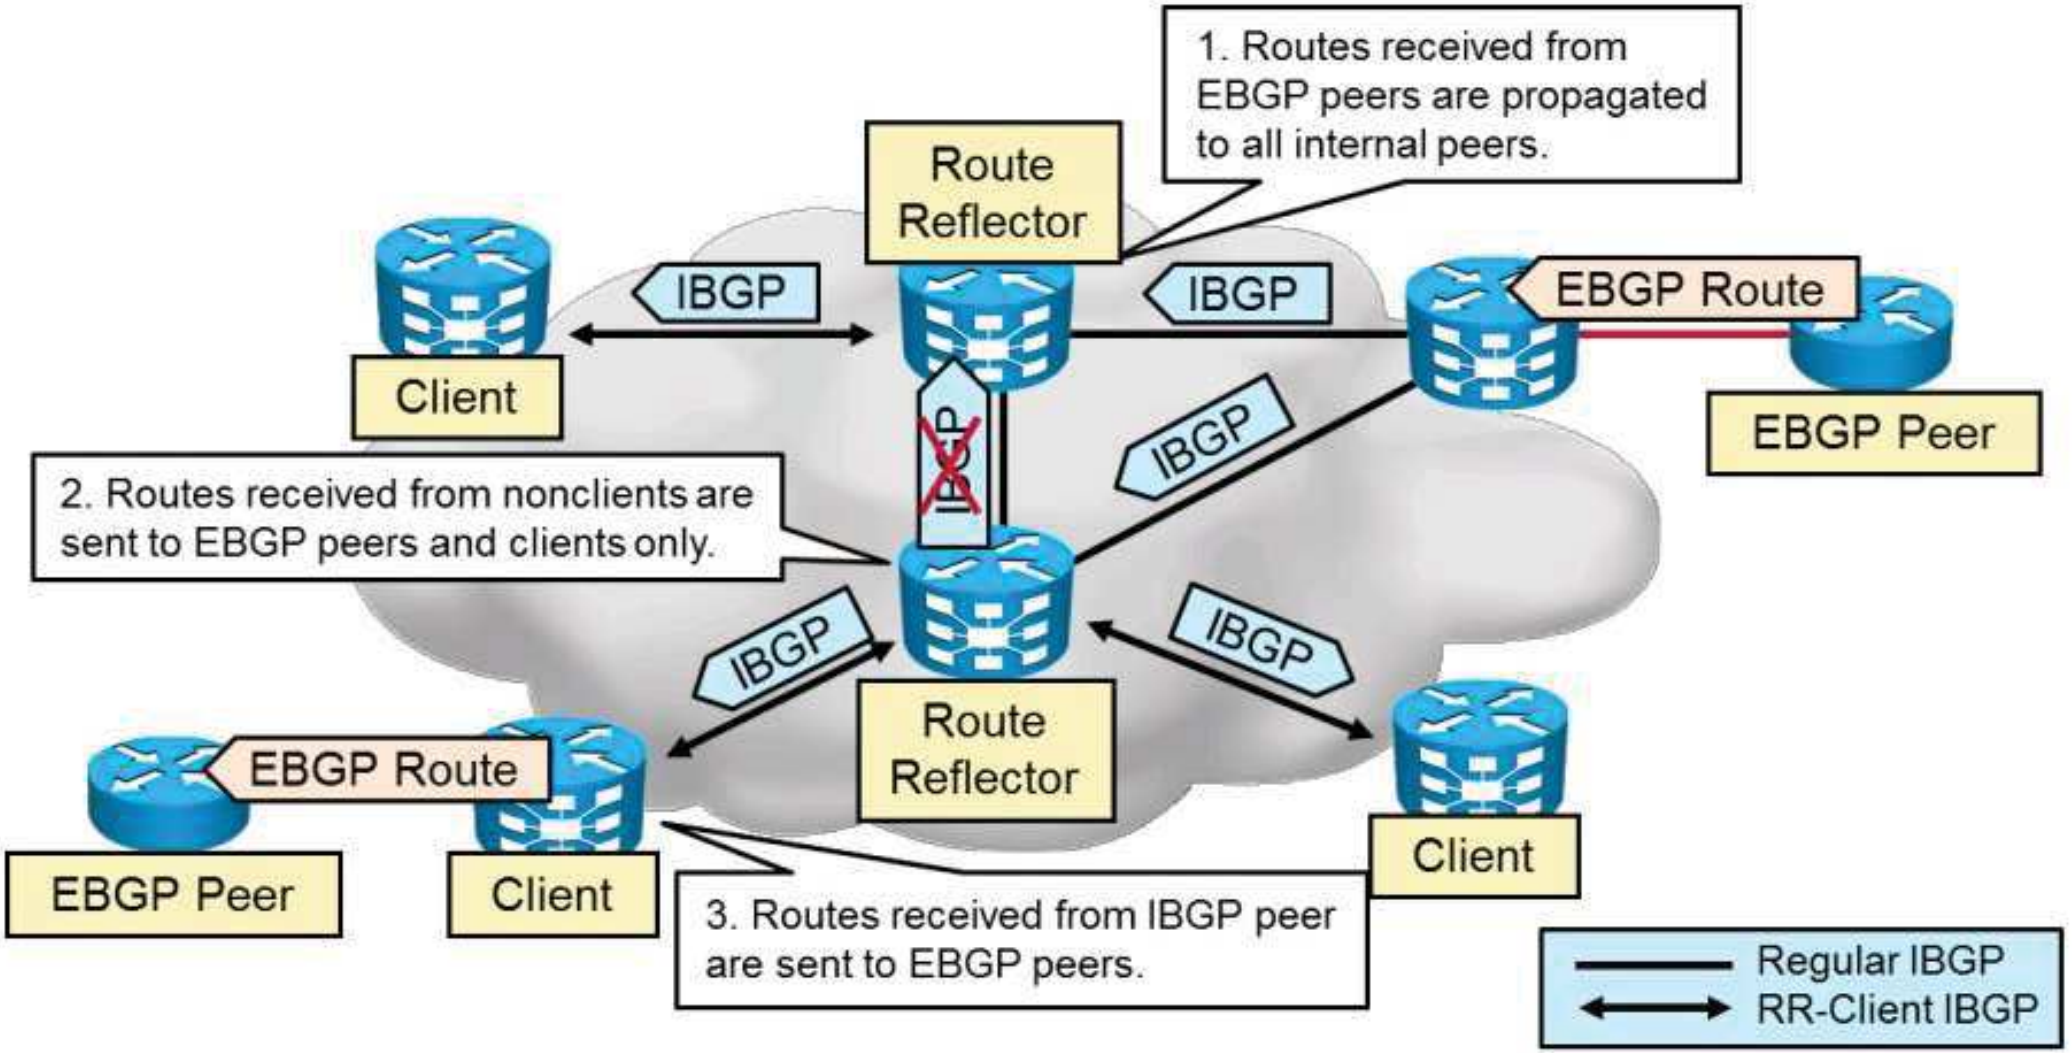

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Redundant Route Reflectors

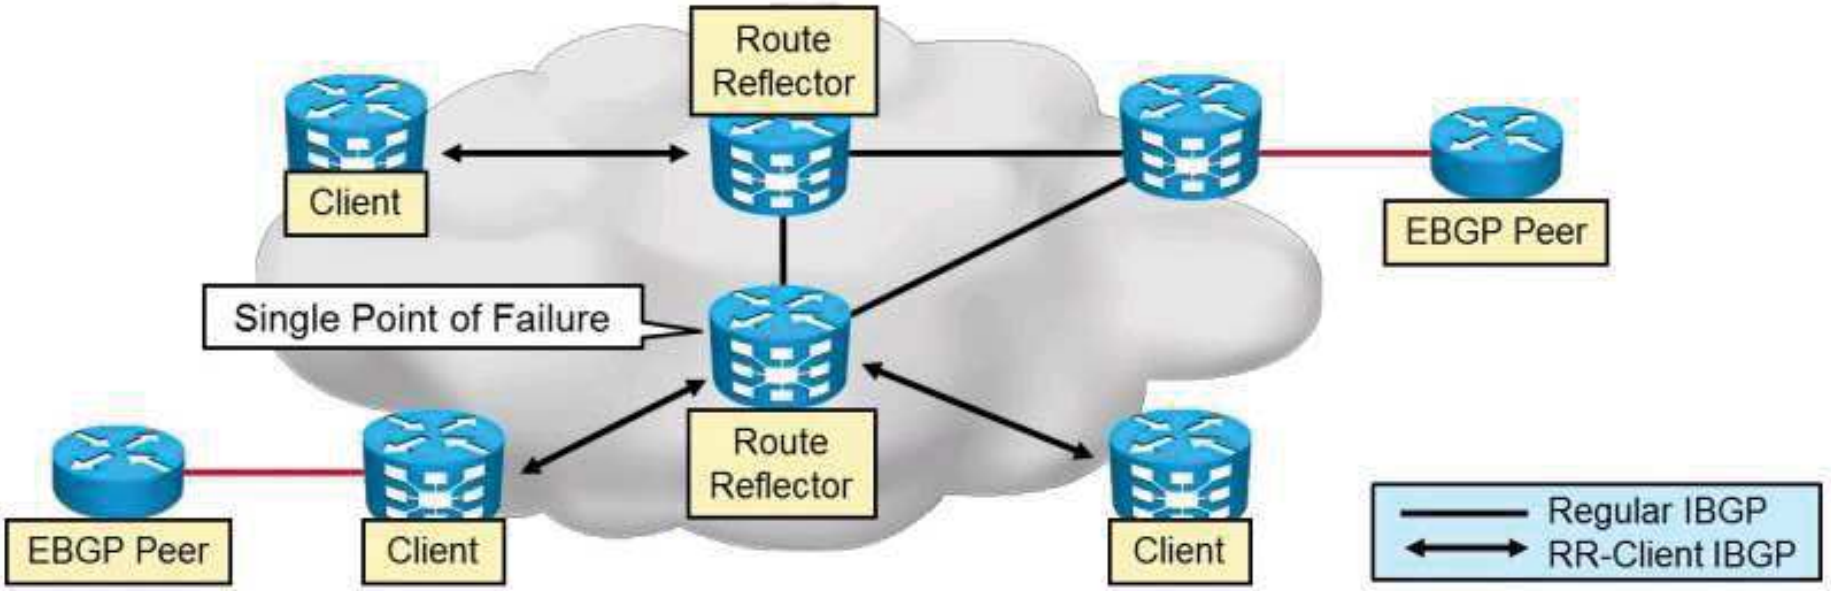

- Clients that have an IBGP session with only one route reflector will not be able to send any BGP updates if the route reflector fails.
- Clients should establish an IBGP session with at least two route reflectors using different physical connections.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Redundant Route Reflectors (Cont.)

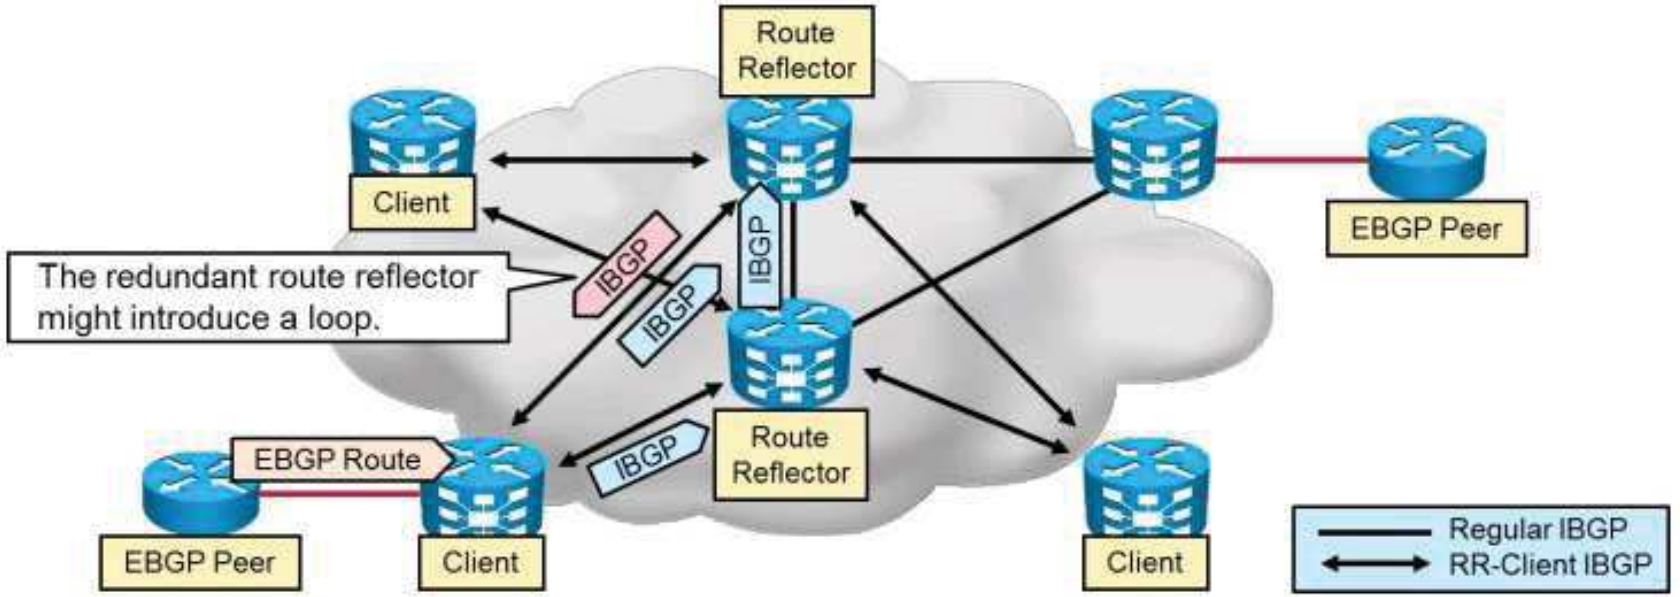

- Redundant reflectors solve the high-availability requirement.
- The concept of clusters is introduced to prevent IBGP routing loops between route reflectors.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Route Reflector Clusters

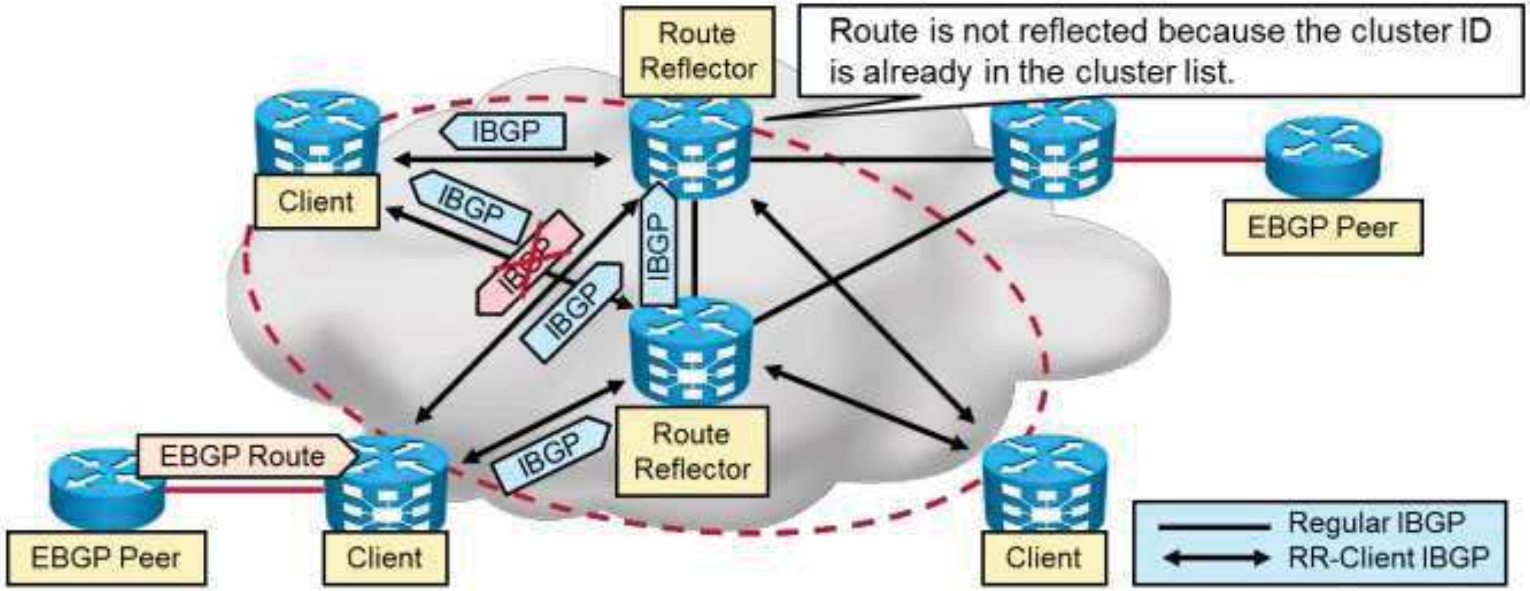

- A group of redundant route reflectors and their clients form a cluster.
- Each cluster must have a unique cluster ID.
- Each time a route is reflected, the cluster ID is added to the cluster-list BGP attribute.
- The route that already contains the local cluster ID in the cluster list is not reflected.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Additional Loop-Prevention Mechanism

This list represents additional loop-prevention mechanism:

- Every time a route is reflected, the router ID of the originating IBGP router is stored in the originator-ID BGP attribute.
- A router receiving an IBGP route with the originator ID set to its own router ID ignores that route.
- The BGP path selection procedure is modified to take into account both the cluster list and the originator ID.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Network Design with BGP Route Reflectors

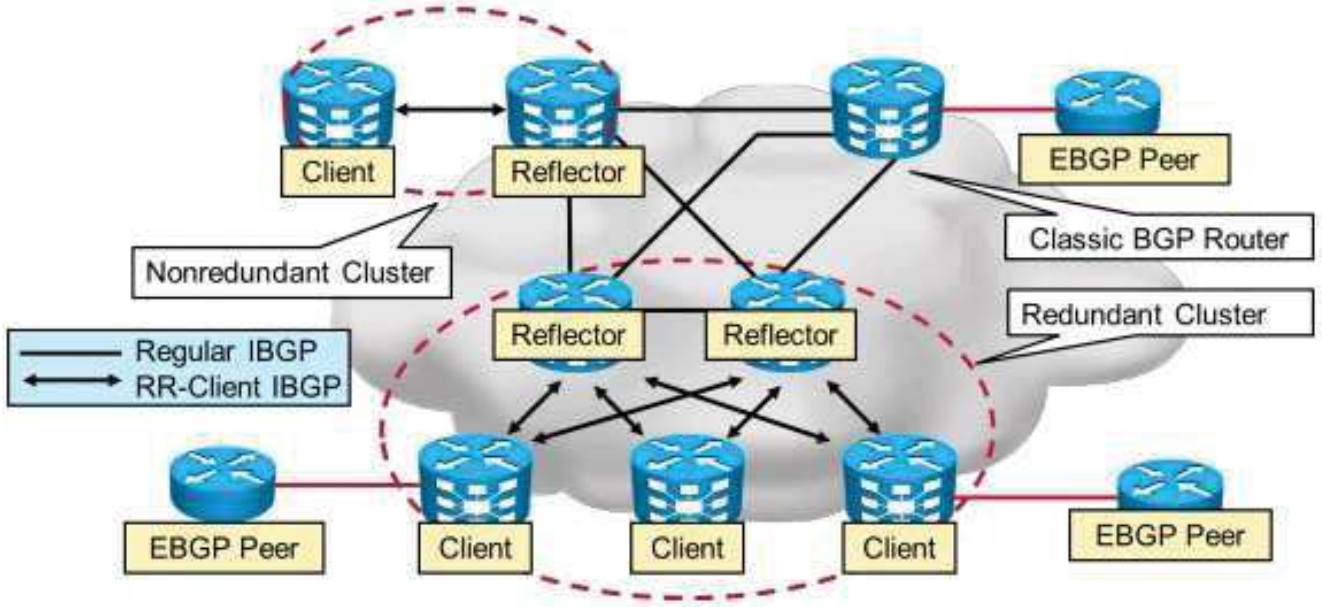

These are route reflector characteristics:

- Route reflector rules divide a transit AS into smaller areas (called clusters).
- Each cluster contains route reflectors and route reflector clients.
- Routers that do not support route reflector functionality act as a one-router cluster or as a route reflector client. These routers have to be fully meshed with route reflectors from all clusters.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Network Design with BGP Route Reflectors (Cont.)

Potential problems that can occur when you deviate from the route reflector network design rules.

Issue:

- Clients do not have sessions with all reflectors in a cluster.
- Clients have sessions with reflectors in several clusters.
- Clients have IBGP sessions with other clients.

Result:

- Clients will not receive all IBGP routes.
- Clients will receive duplicate copies of the same route.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Hierarchical Route Reflectors

These are hierarchical route reflectors characteristics:

- In very large networks, a single layer of route reflectors might not be enough.
- A hierarchy of route reflectors can be established.
  - A route reflector can be a client of another route reflector.
  - The hierarchy can be as deep as needed.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Hierarchical Route Reflectors (Cont.)

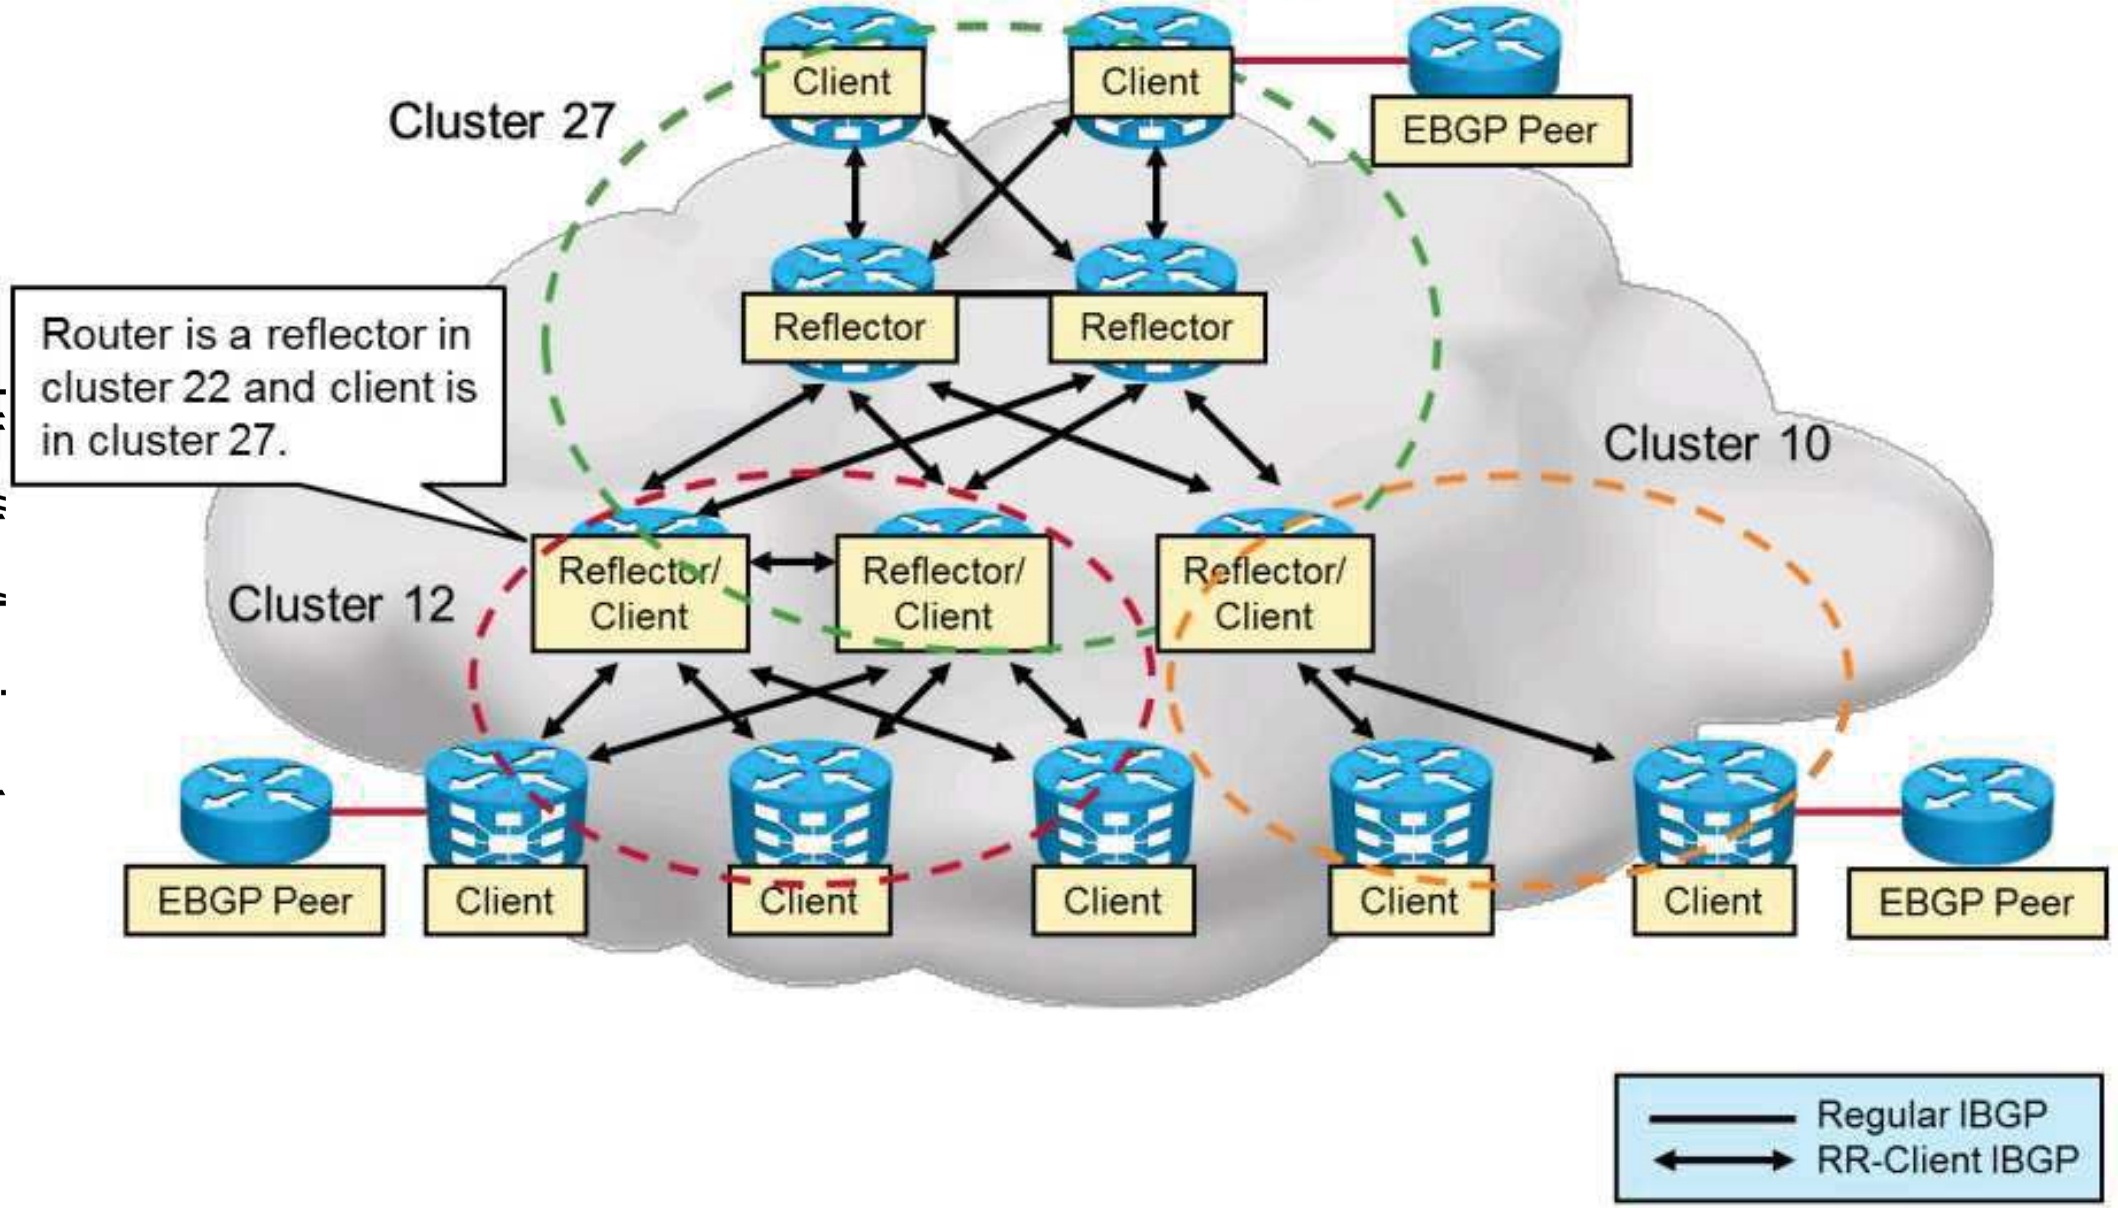

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Implementing BGP Route Reflectors

Route reflector backbone migration steps:

- Divide the AS into areas (clusters):
  - Assign a cluster ID to each area.
- On route reflectors, retain only IBGP sessions with clients in their cluster and with other route reflectors:
  - Configure the cluster ID on every route reflector.
  - Configure clients on every route reflector.
- On route reflector clients, retain only IBGP sessions with route reflectors in their cluster.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Reflectors Configuration

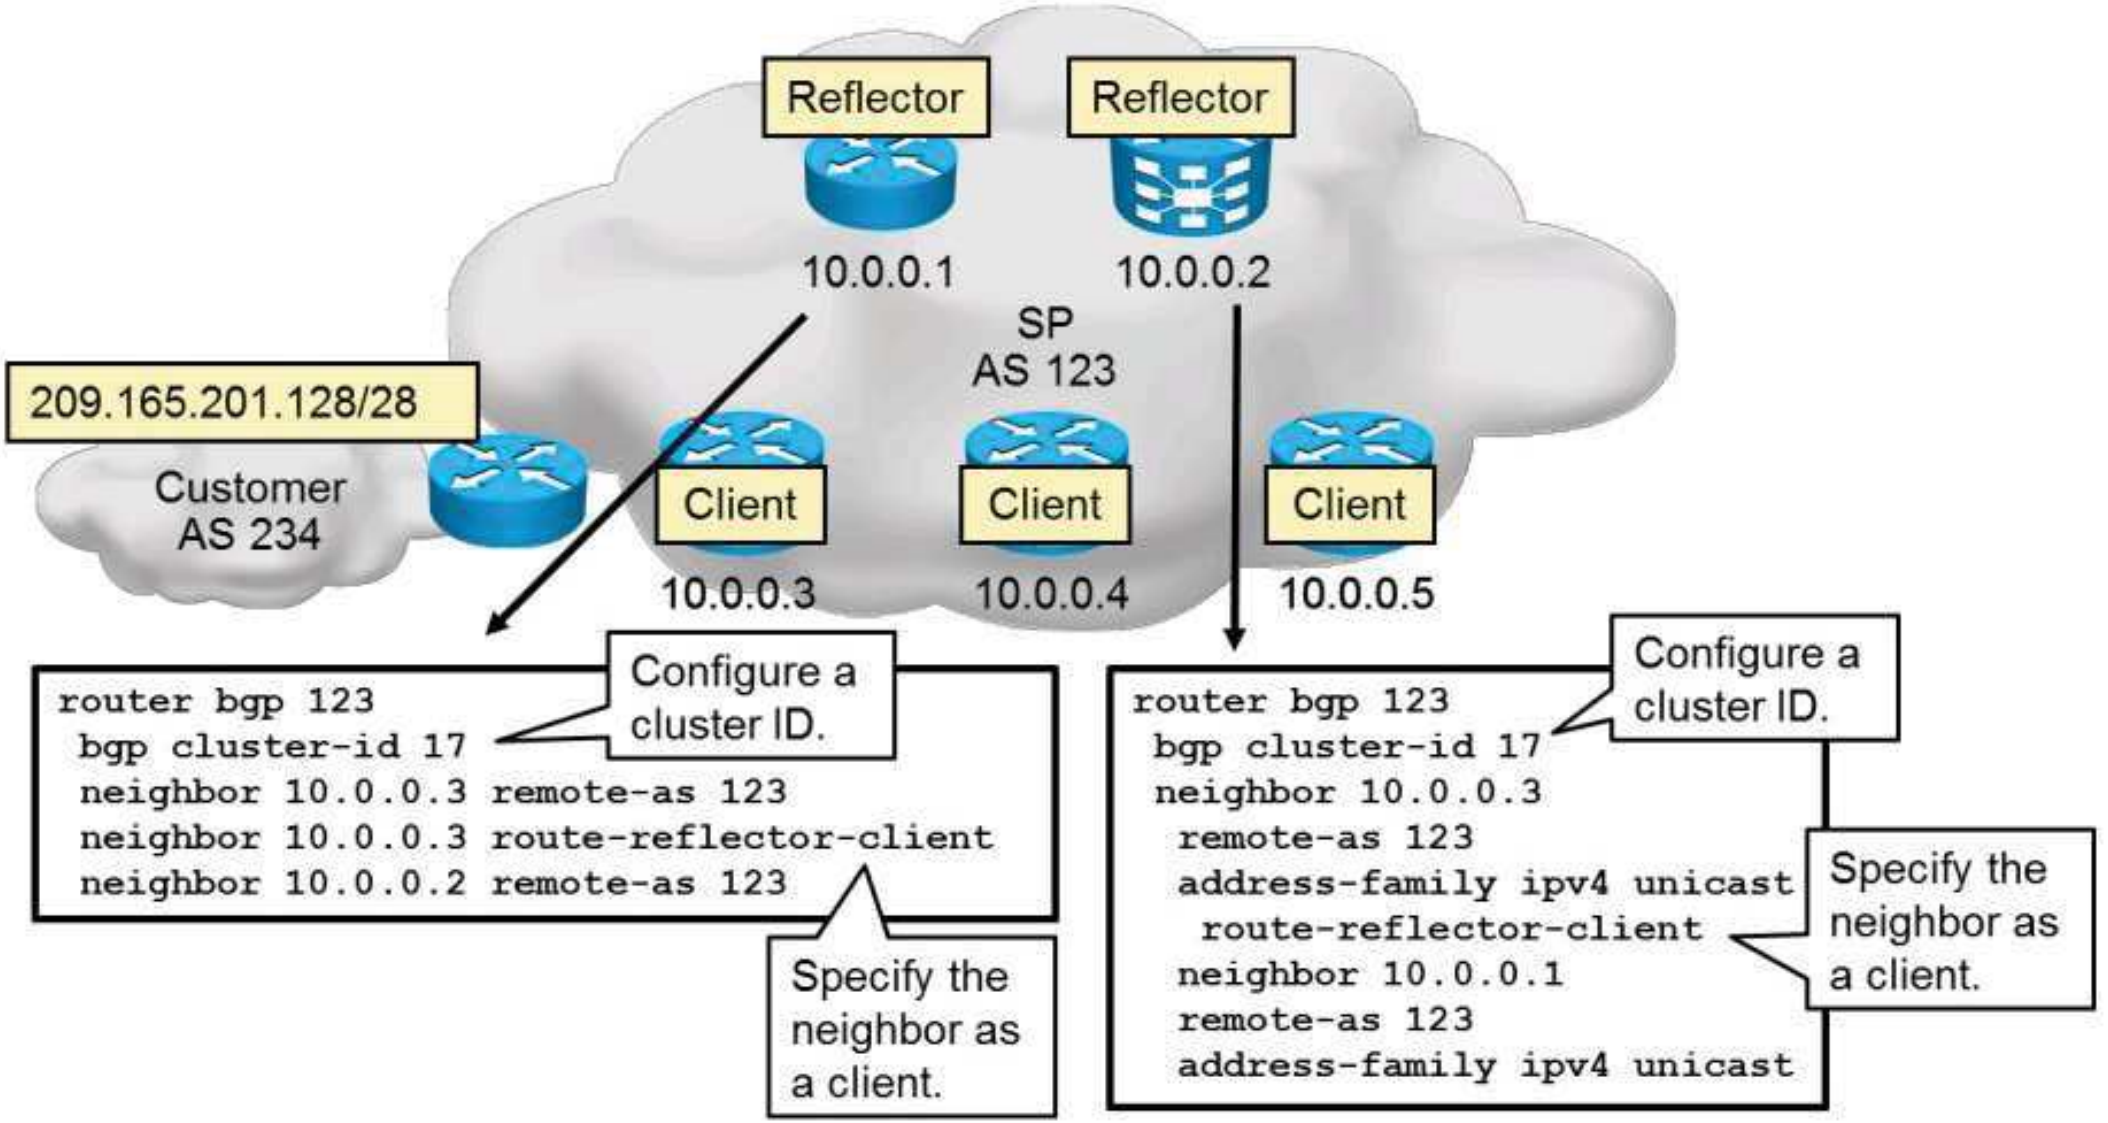

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Reflectors Verification

```
RP/0/RSP0/CPU0:P# show bgp neighbors 10.0.0.3
BGP neighbor is 10.0.0.3
 Remote AS 64500, local AS 64500, internal link
 Remote router ID 10.0.0.3
 Cluster ID 17
  BGP state = Established, up for 00:00:07
<… output omitted …>
For Address Family: IPv4 Unicast
  BGP neighbor version 17
  Update group: 0.1 Filter-group: 0.1 No Refresh request being processed
  Route-Reflector Client
  NEXT_HOP is always this router
<… output omitted …>
```

- Displays information about the BGP session with the neighbor.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Reflectors Verification (Cont.)

```
RP/0/RSP0/CPU0:P# show bgp 209.165.201.128
BGP routing table entry for 209.165.201.128/28
<… output omitted …>
234, (Received from a RR-client)
   10.0.0.3 (metric 2) from 10.0.0.3 (10.0.0.3)
        Origin IGP, metric 0, localpref 100, valid, internal, best, group-best
        Received Path ID 0, Local Path ID 1, version 3
```

- Displays routes received from the client as seen on the reflector

```
RP/0/RSP0/CPU0:PE# show bgp 209.165.201.128
BGP routing table entry for 209.165.201.128/28
<… output omitted …>
234
   10.0.0.1 (metric 2) from 10.0.0.1(10.0.0.1)
         Origin IGP, metric 0, localpref 100, valid, internal
         Received Path ID 0, Local Path ID 0, version 0
         Originator: 10.0.0.1, Cluster list: 0.0.0.17
```

- Displays reflected routes as seen on the client

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Confederations Overview

BGP confederations characteristics:

- Splitting the AS into smaller autonomous systems would reduce the number of BGP sessions, but extra AS numbers are not available.
- Confederations enable internal AS numbers to be hidden and announce only one (external) AS number to EBGP neighbors.
- Inside a confederation, full-mesh IBGP is required.

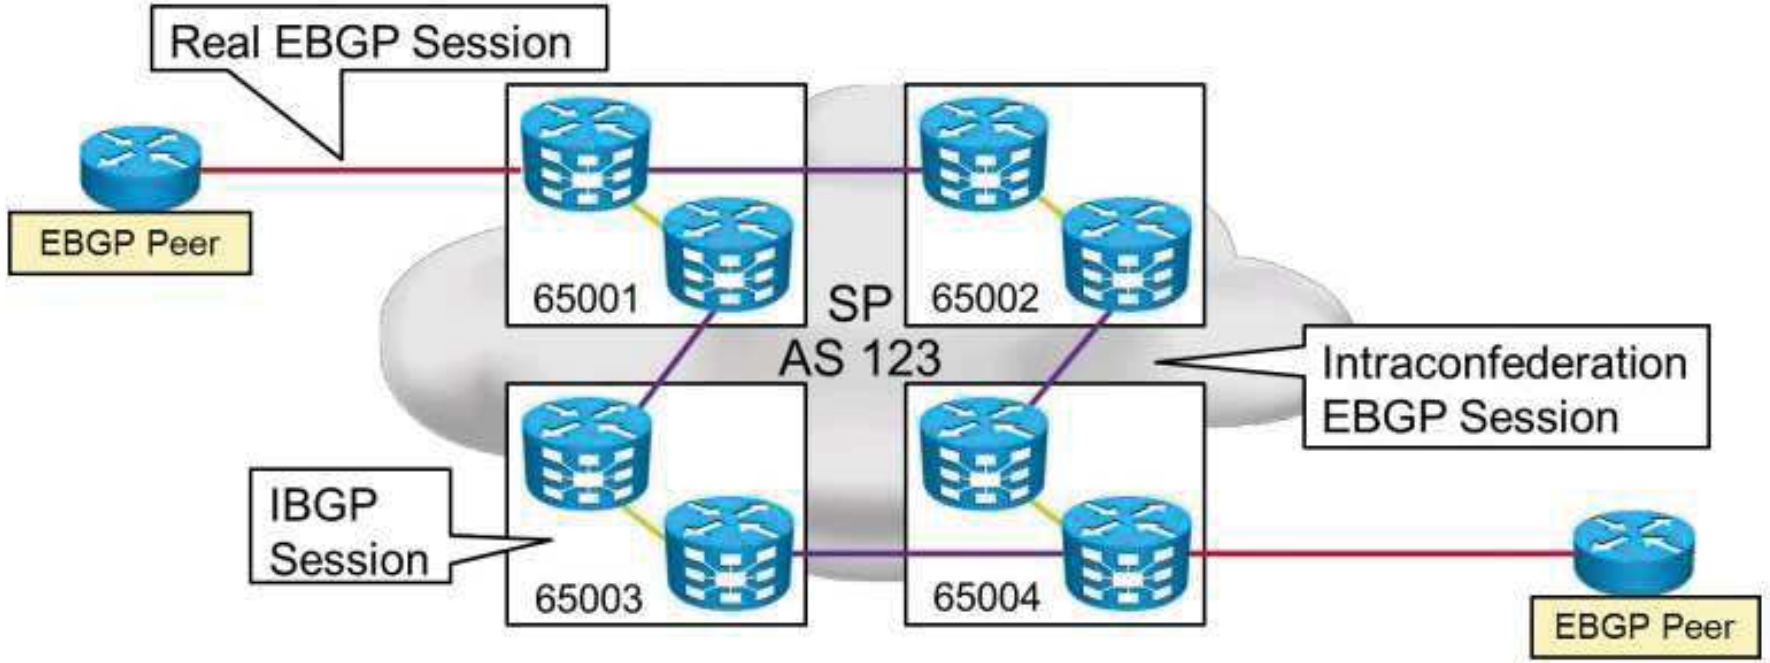

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

AS Path Propagation Within BGP Confederation

AS Path propagation within BGP confederation characteristics:

- IBGP session:
  - The AS path is not changed.
- Intraconfederation EBGP session:
  - The intraconfederation AS number is prepended to the AS path.
  - The intraconfederation AS path is encoded as a separate segment of the AS path.
  - The intraconfederation AS path is displayed in parentheses when you are using **show** commands.
  - A router that does not support BGP confederations will reject an AS path with unknown segment type.
- EBGP session with external peer:
  - Intraconfederation AS numbers are removed from the AS path.
  - The external AS number is prepended to the AS path.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

AS Path Propagation Example

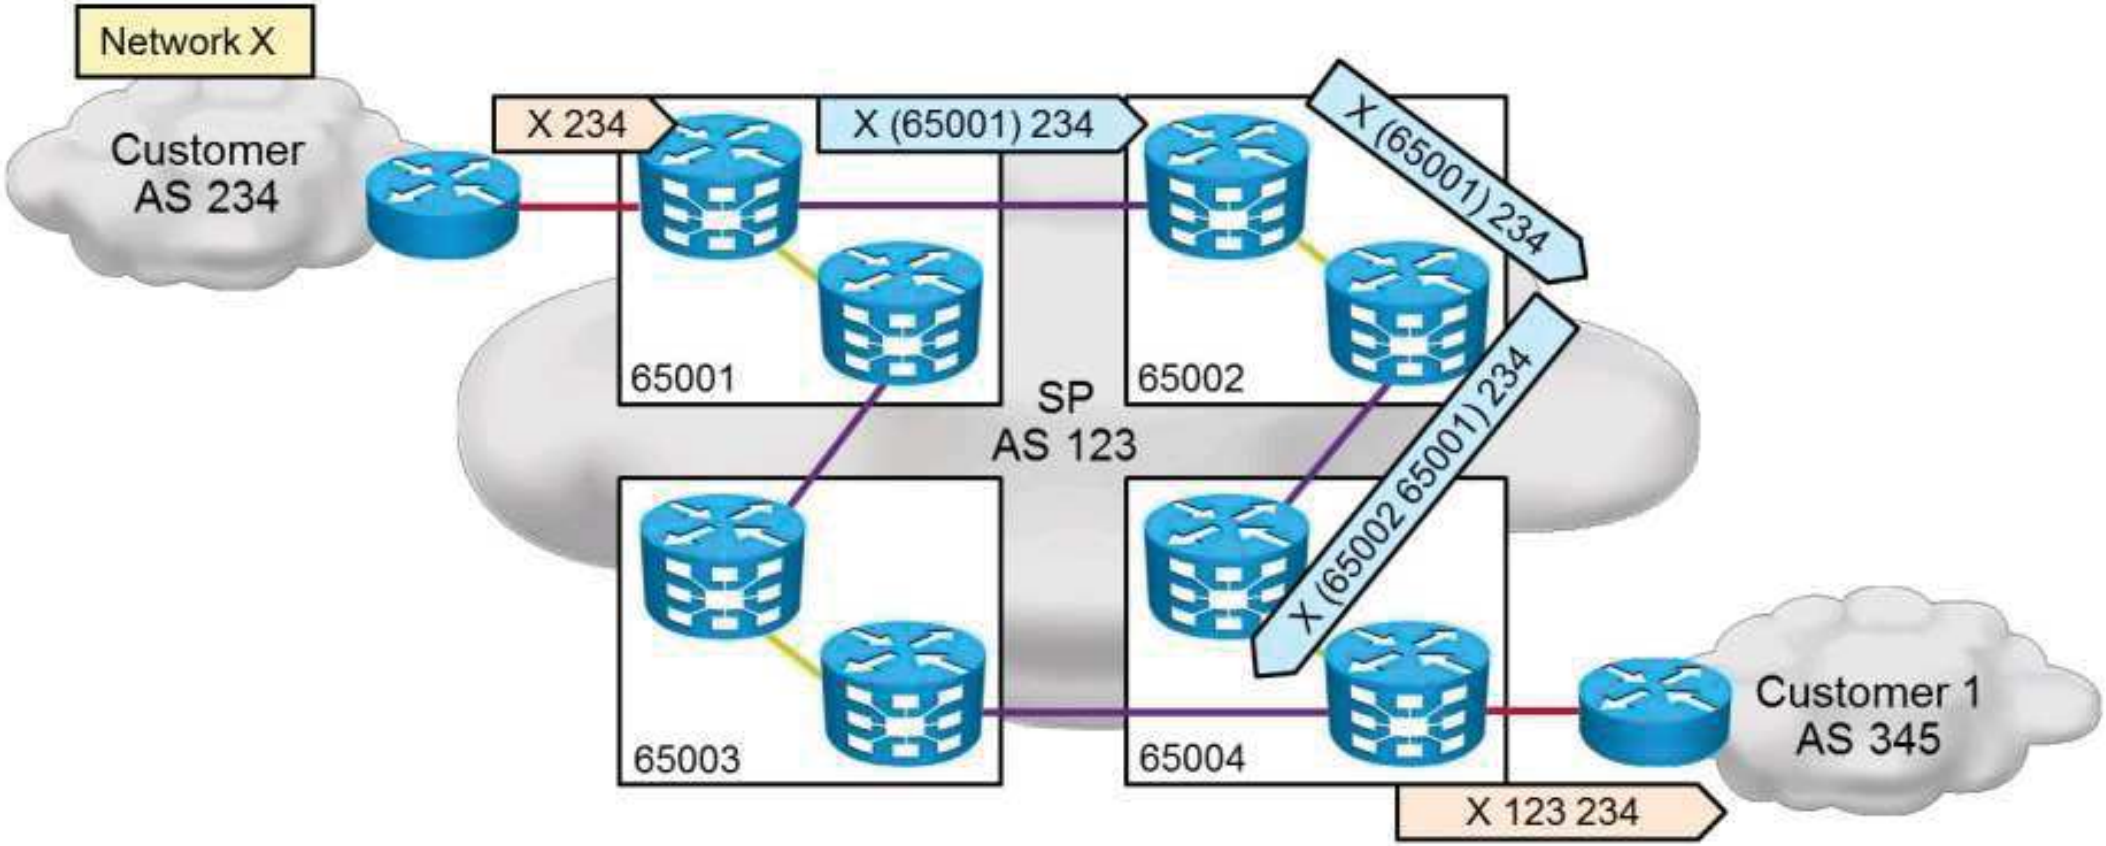

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Intraconfederation EBGP Session Properties

Intraconfederation EBGP session characteristics:

- Behaves like EBGP session during session establishment:
  - The EBGP neighbor has to be directly connected, or you have to configure EBGP multihop on the neighbor.
- Behaves like IBGP session when propagating routing updates:
  - The local preference, MED, and next-hop attributes are retained.
  - The whole confederation can run one IGP, providing optimal routing based on the next-hop attribute in the BGP routing table.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- BGP reflectors are one of two ways to increase IBGP network scalability.
- The IBGP split-horizon rule prevents loops in a fully meshed IBGP cloud.
- A single route reflect is a single point of failure, therefore redundant configuration is essential.
- Route reflectors can be combined into clusters for redundancy.
- The originator-ID BGP attribute is used to prevent routes from being reflected.
- BGP route reflectors allow constructions of flexible network designs.
- BGP route reflector clusters can be nested to many levels.
- BGP route reflectors are configured on reflectors only; clients do not know they are clients.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- IGP confederations allow the creation of AS domains within AS domains.
- The intraconfederation AS number is prepended to the AS path and is displayed in parentheses.
- BGP intraconfederation interneighbor peering sessions combine the properties of both eBGP and IBGP peering sessions at the same time.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

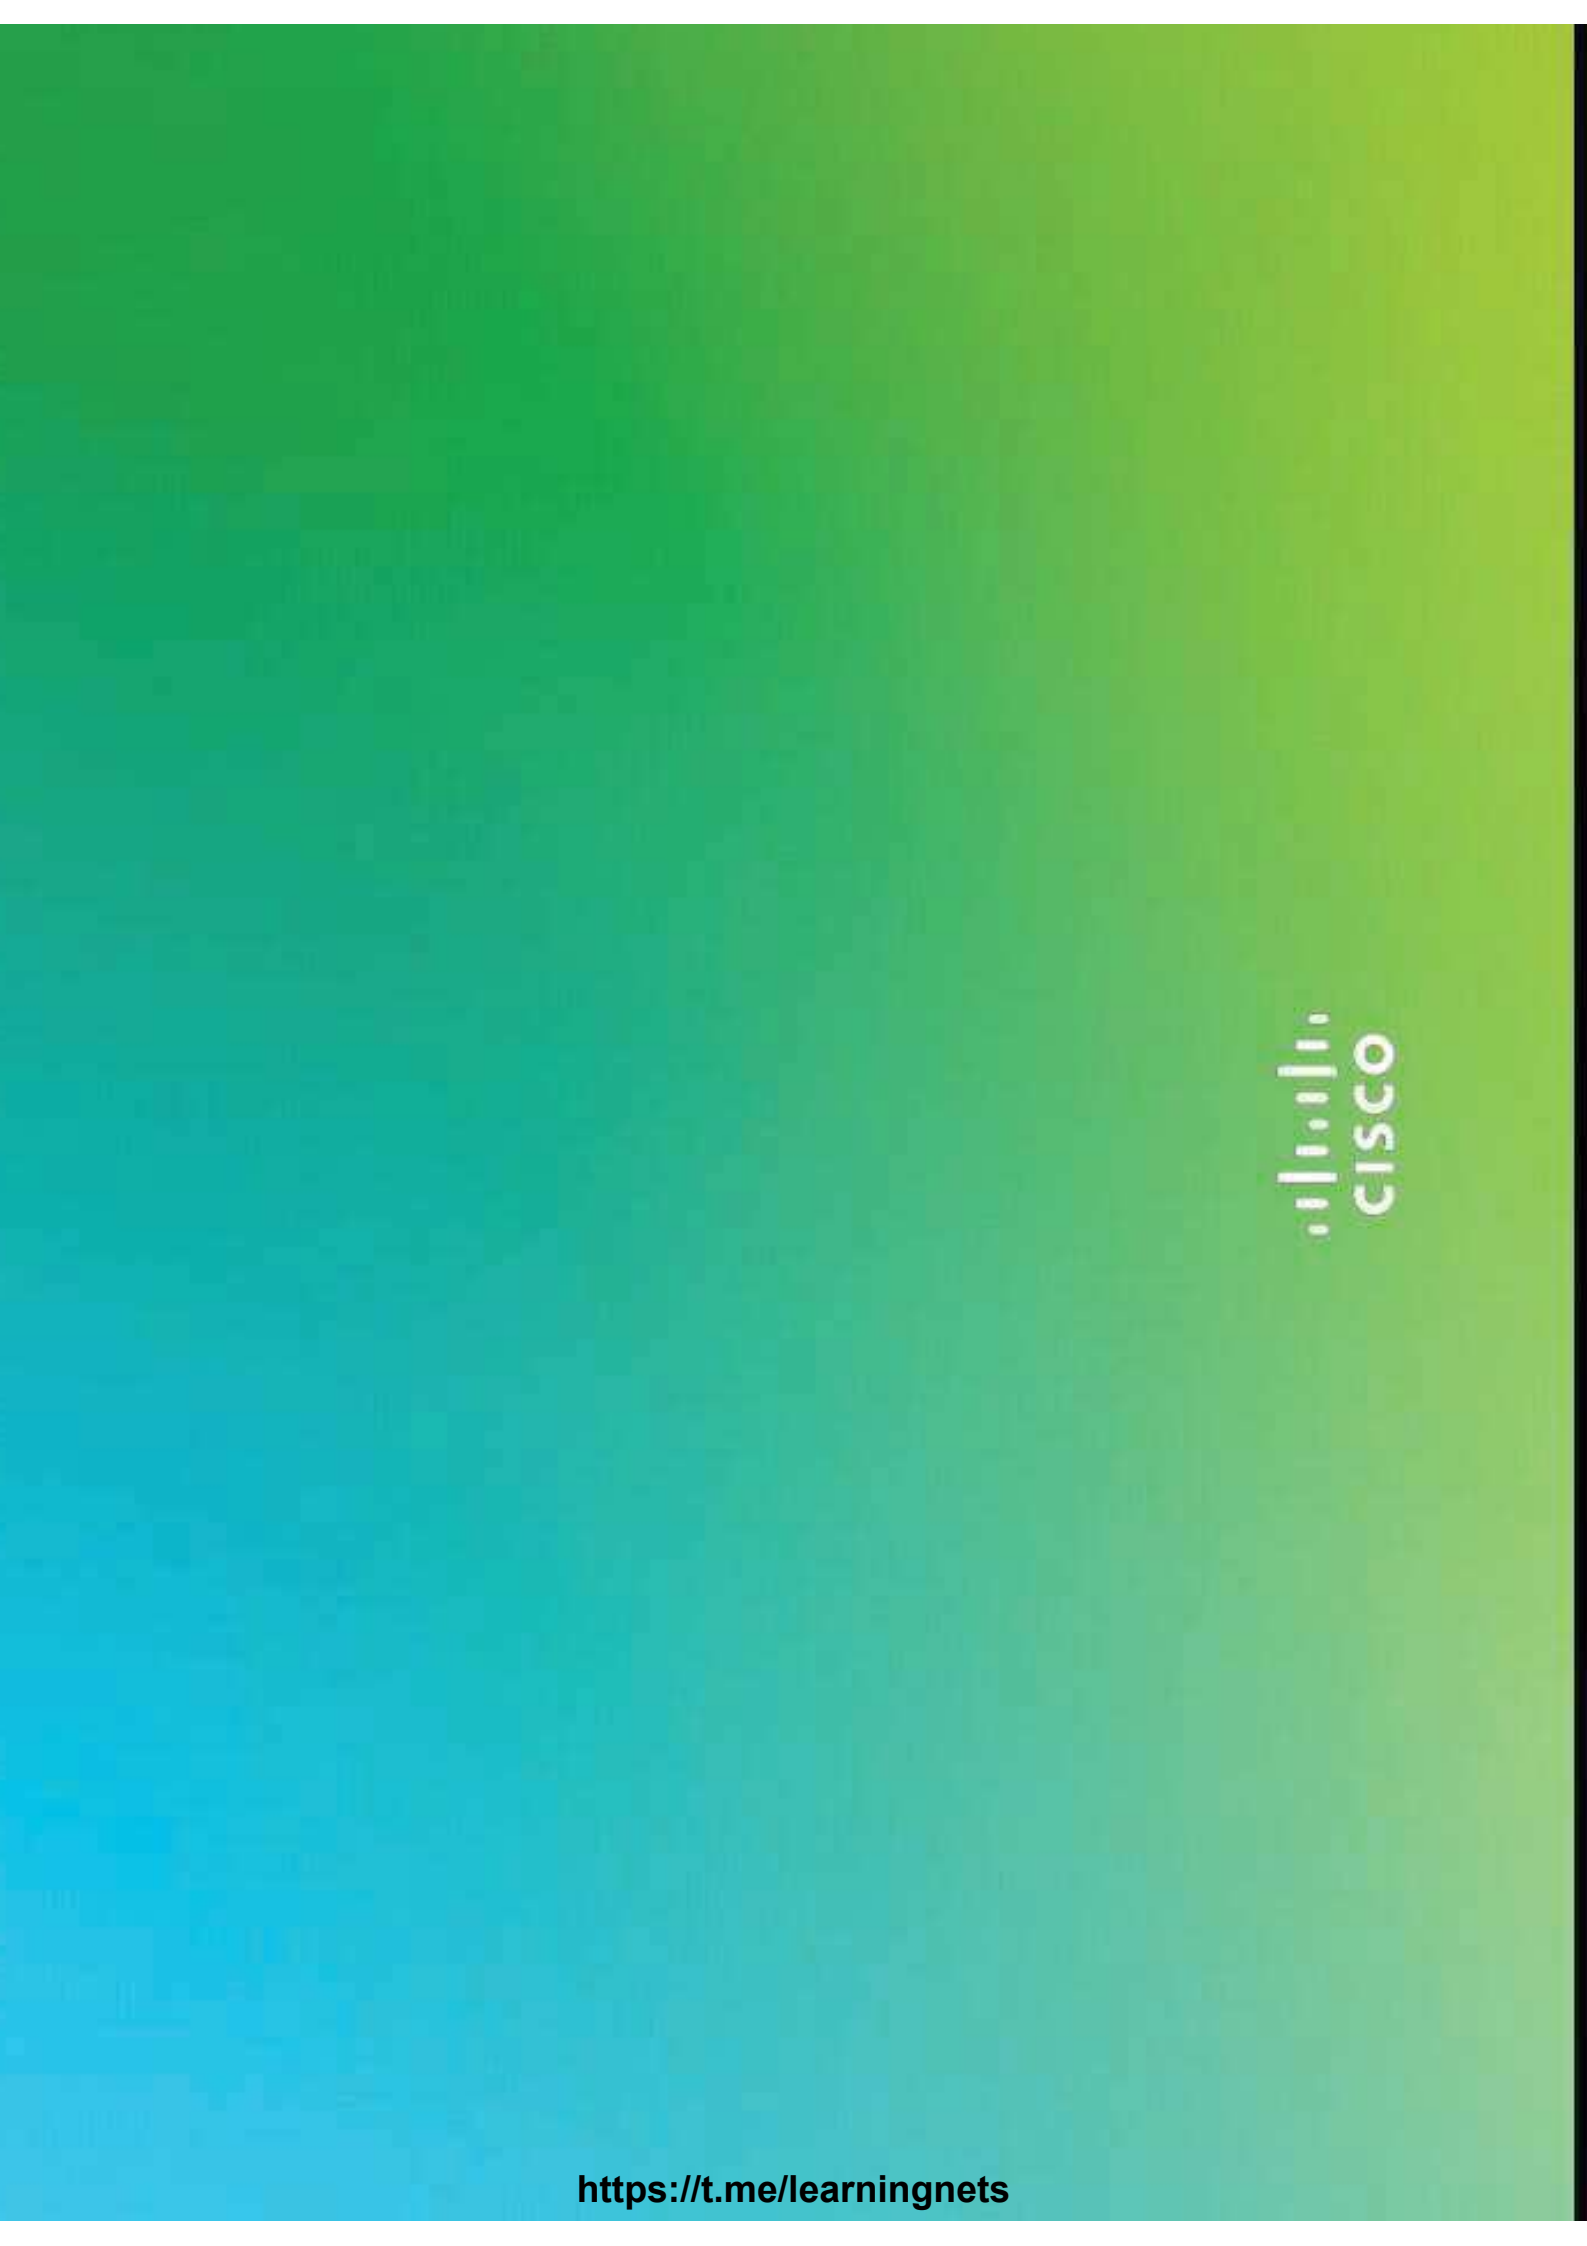

https://t.me/learningnets

Module Summary

- Several routing protocols are used in service provider networks. BGP scalability and IP addressing inside the service provider core network are important factors in service provider scalability.
- BGP route reflectors and BGP confederations do not require full-mesh IBGP, and this significantly improves BGP scalability.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

https://t.me/learningnets

https://t.me/learningnets

Implementing Advanced BGP Operations

Secure and Optimize BGP

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco IP NGN Infrastructure Layer

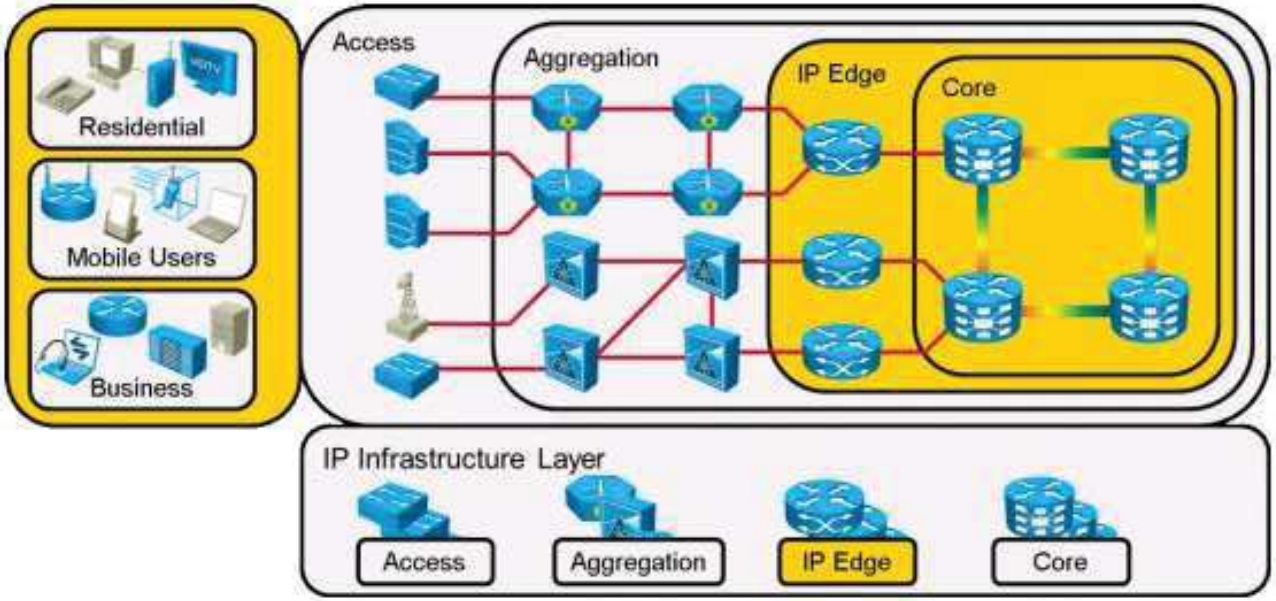

BGP security and optimization options are used in the following Cisco IP NGN infrastructure layers:

- Core and IP edge portions of the service provider network
- Customer edge devices

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Threats in Service Provider Environments

These are threats to the BGP of the service provider:

- BGP relies on TCP as its transport protocol.
- BGP is susceptible to the same attacks that apply to any TCP-based protocol (DoS attacks).
- BGP is the most frequently targeted routing protocol because it is used across the Internet.
- Service providers should take extreme caution to mitigate risks of exploiting BGP routing protocol.
- Inadvertent mistakes during BGP configuration can be serious and can affect networks worldwide.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Threats in Service Provider Environments (Cont.)

Attacks can be performed against a customer, usually using a form of DDoS:

- Such attacks can cause collateral damage to the infrastructure of the service provider, due to large amounts of traffic.
- Countermeasures should be taken to prevent damage to the service provider.

These are the BGP threats:

- BGP routing table manipulation
- BGP route spoofing
- BGP DoS

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Countermeasures Overview

Countermeasure	BGP Table Manipulation	BGP Route Spoofing*	BGP DoS
BGP Neighbor Authentication	Yes	No	No
BGP TTL Security Check	Yes	No	Yes
BGP Maximum Prefix	No	No	Yes
CoPP	Yes	No	Yes

*BGP route spoofing can be prevented using filtering based on prefixes and the AS path.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Limiting

BGP route limiting characteristics:

- All filtering mechanisms specify only what you are willing to accept, but not how much.
- A misconfigured BGP neighbor can send a large number of prefixes, which can exhaust the memory of a router or overload the CPU (several Internet-wide incidents have already occurred).
- BGP maximum prefixes limiting is used to establish a hard limit on the number of prefixes received from a neighbor.
- It is enabled by default on Cisco IOS XR Software to impose limitations for different address families.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Limiting (Cont.)

BGP router terminates peering when a number of maximum prefixes is exceeded.

A router can be configured to:

- Generate a logging message when a specified percentage of the maximum prefixes is reached
- Re-establish BGP peering after a specified time (from 1 to 65535 minutes)
- Generate a logging message when the maximum prefix limit is exceeded, instead of terminating BGP peering

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Limiting Configuration

Enable BGP route limiting for IPv4.

```
router bgp 123
 neighbor 209.165.201.130 remote-as 234
 neighbor 209.165.201.130 maximum-prefix 10
70 restart 5
 !
 neighbor 2001:db8:128::130 remote-as 234
address family ipv6 unicast
 neighbor 2001:db8:128::130 activate
 neighbor 2001:db8:128::130 maximum-prefix
10 70 restart 5
```

Enable BGP route limiting for IPv6.

Enable BGP route limiting for IPv4.

```
router bgp 123
 neighbor 209.165.201.134
  remote-as 345
  address-family ipv4 unicast
   maximum-prefix 100000 80
restart 5
  !
 neighbor 2001:db8:132::134
  remote-as 345
  address-family ipv6 unicast
   maximum-prefix 10000 80
restart 5
```

Enable BGP route limiting for IPv6.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route Limiting Verification

```
RP/0/RSP0/CPU0:PE2# show bgp neighbors 209.165.201.134
BGP neighbor is 209.165.201.134
 Remote AS 345, local AS 123, external link
<… output omitted …>
   Maximum prefixes allowed 100000
   Threshold for warning message 80%, restart interval 5 min
```

https://t.me/learningnets

- Displays BGP neighbor information.

```
RP/0/RSP0/CPU0:PE5# RP/0/RSP0/CPU0:Oct 11 13:31:23.697 : bgp[1048]: %ROUTING-
BGP-4-MAXPFXEXCEED : No. of IPv4 Unicast prefixes received from 209.165.201.134
: 100001 exceed limit 100000
RP/0/RSP0/CPU0:Oct 11 13:31:23.697 : bgp[1048]: %ROUTING-BGP-5-ADJCHANGE :
neighbor 209.165.201.134 Down - Peer exceeding maximum prefix limit (CEASE
notification sent - maximum number of prefixes reached) (VRF: default)
```

- Logging messages displayed when limit is exceeded.

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Neighbor Authentication

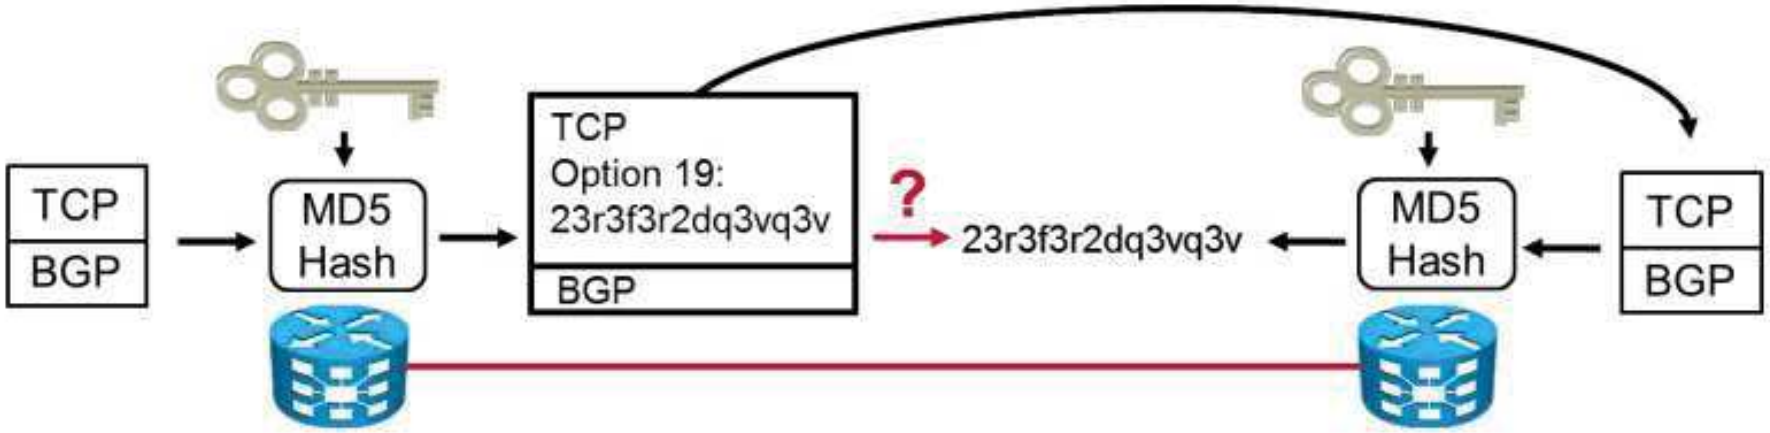

BGP neighbor authentication characteristics:

- BGP neighbors can be authenticated before establishing a TCP session:
  - HMAC-MD5 is used.
  - Cisco IOS XR supports HMAC-SHA1 with key chains.
- To calculate a hash, part of an IP and an entire TCP header with data are used together with a preshared key.
- Every TCP segment is authenticated and the hash is prepended as TCP option 19.
- The hash is calculated on the receiving BGP router and compared with the received hash.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP TTL Security Check

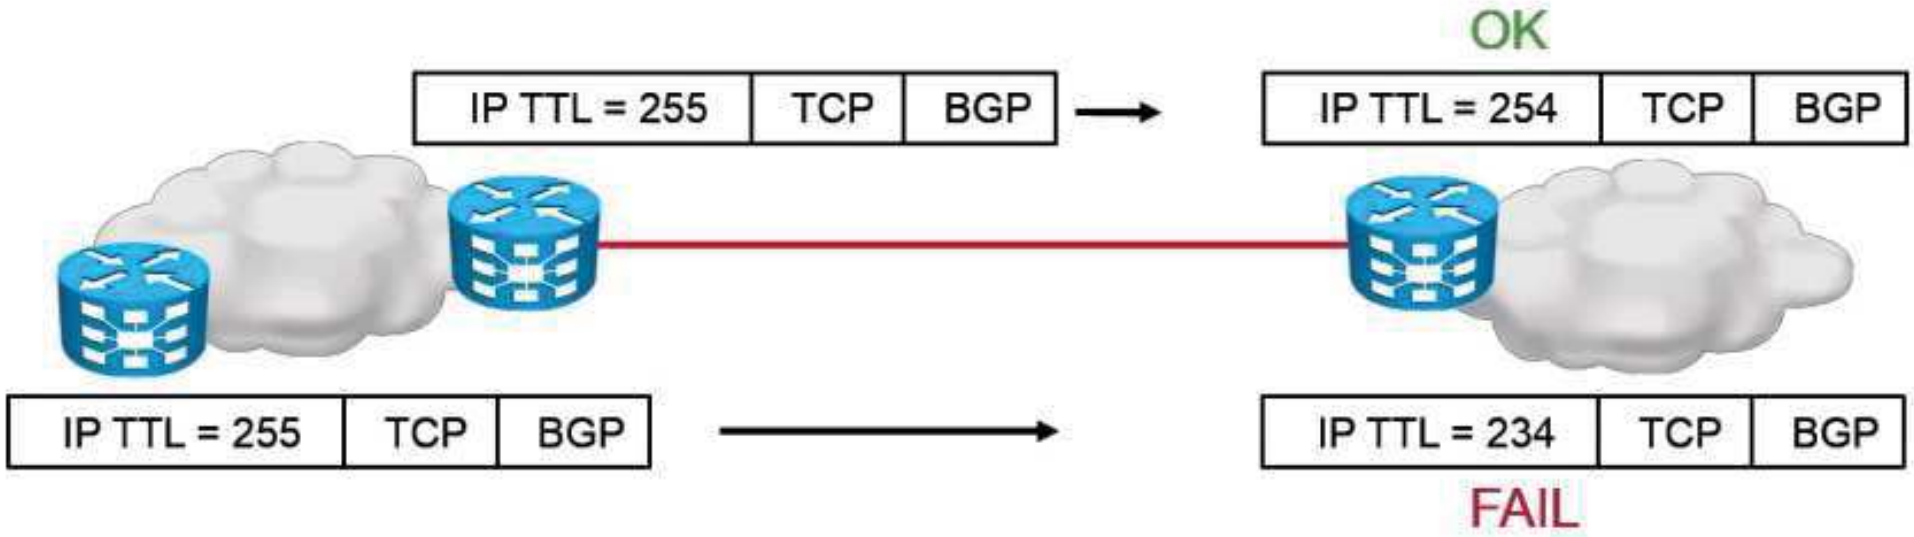

BGP TTL security check characteristics:

- Assumes that valid EBGP sessions are established between connected interfaces or loopback interfaces.
- Enforces that received TTL should match a specified value.
- Can prevent a DoS attack from nondirectly connected neighbors by setting the received TTL to 254 or 253.
- When enabled, sets the TTL of outgoing BGP packets to 255.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Control Plane Policing

CoPP (Cisco IOS and IOS XE Software only) can be used to control traffic to the control plane of a router:

- Permits, denies, and rate limits access to the control plane
- Configured as service policy, applied to virtual control plane interface
- Can be used to filter and rate-limit BGP traffic to the router

LPTS (Cisco IOS XR Software):

- LPTS policers are responsible for policing traffic to the RPs on the incoming line cards.
- Policer values can be changed for each line card separately.
- They can be used to rate-limit BGP traffic to the router.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Neighbor Authentication, TTL Security, and CoPP Configuration

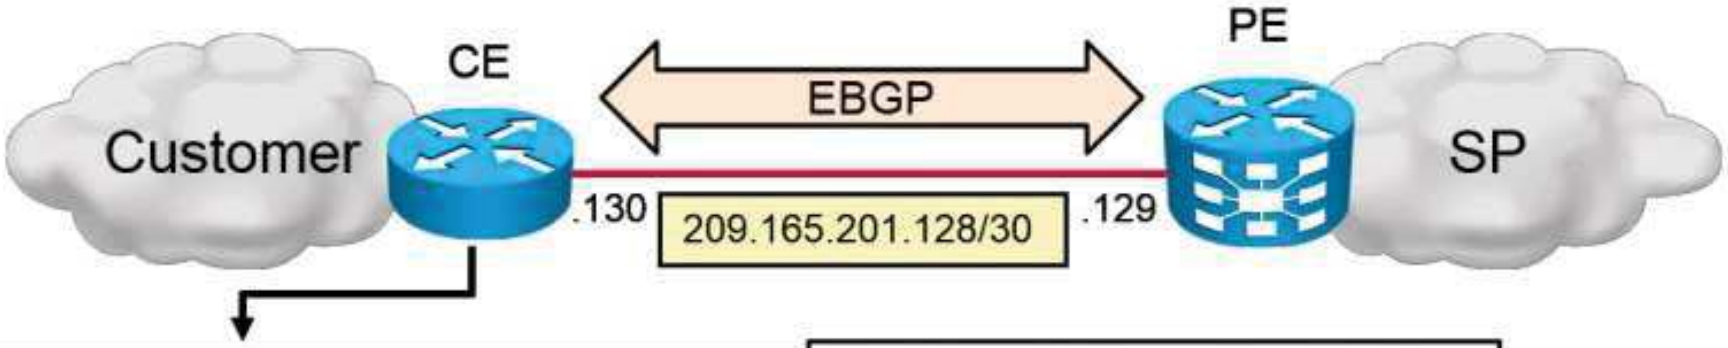


```
router bgp 123
 neighbor 209.165.201.129 password C!sc()
 neighbor 209.165.201.129 ttl-security hops 1
!
ip access-list extended BGP
 permit tcp host 209.165.201.129 host 209.165.201.130 eq bgp
 permit tcp host 209.165.201.129 eq bgp host 209.165.201.130
 deny ip any any
!
class-map BGP_CLASS
  match access-group name BGP
!
policy-map COPP_POLICY
  class BGP_CLASS
    police rate 200 pps conform-action transmit exceed-action drop
!
control-plane
  service-policy input COPP_POLICY
```

Enable neighbor authentication

Enable TTL security

Configure CoPP

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Neighbor Authentication, TTL Security, and CoPP Configuration (Cont.)

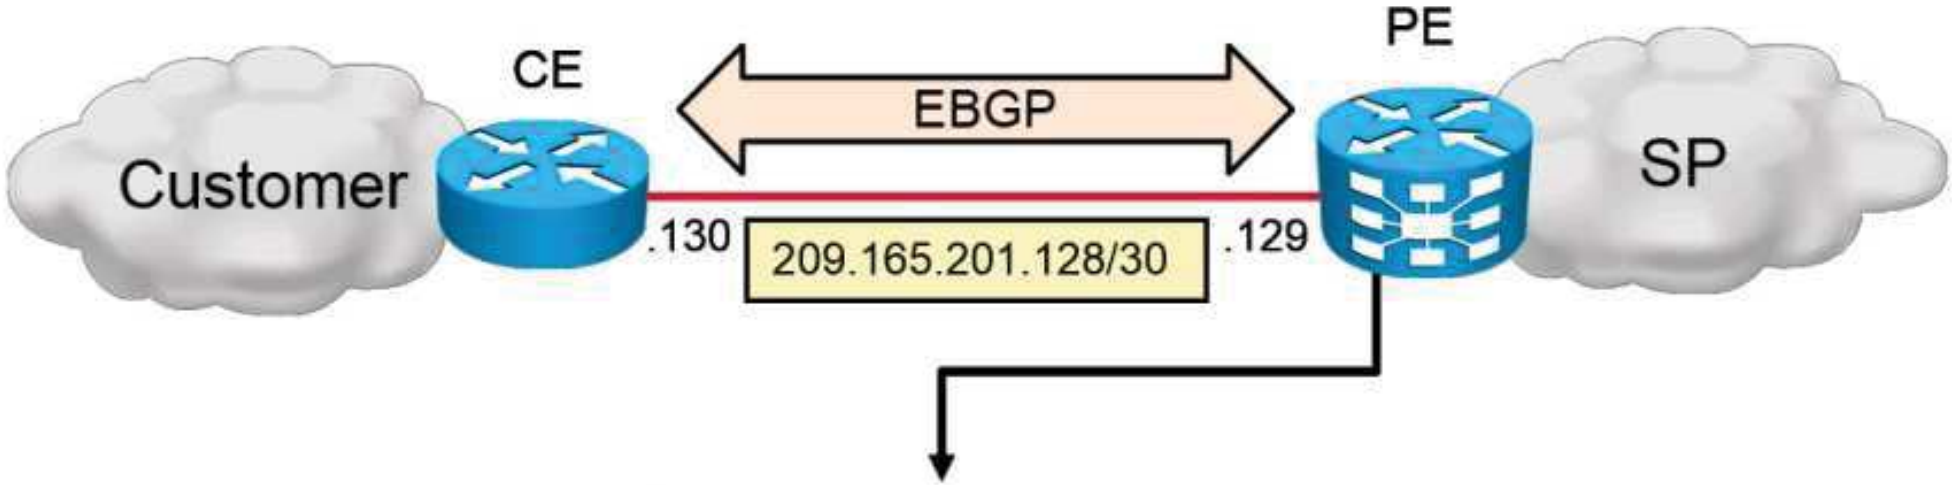

Enable neighbor authentication

Enable TTL security

```
router bgp 123
 neighbor 209.165.201.130
  password C!sc()
  ttl-security
!
lpts pifib hardware police
 flow bgp configured rate 200
 flow bgp default rate 200
 flow bgp known rate 200
```

Configure LPTS

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Remote-Triggered Black-Hole Filtering

Remote-triggered black-hole filtering characteristics:

- When a customer is under DDoS attack, the vast amount of traffic can also cause collateral damage to the infrastructure of the service provider.
- Once the attack has been detected, traffic related to the DDoS should be discarded on the edge of the service provider network.
- One BGP router should be designated as the signaling router:
  - The router signals over BGP to the edge routers that traffic causing DoS should be discarded.
- Destination-based RTBH:
  - Traffic going to the IP addresses of the customer is discarded on the edge.
- Source-based RTBH:
  - Traffic coming from the IP addresses of the attacker is discarded on the edge.
  - Uses strict uRPF with BGP signaling.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Destination-Based RTBH

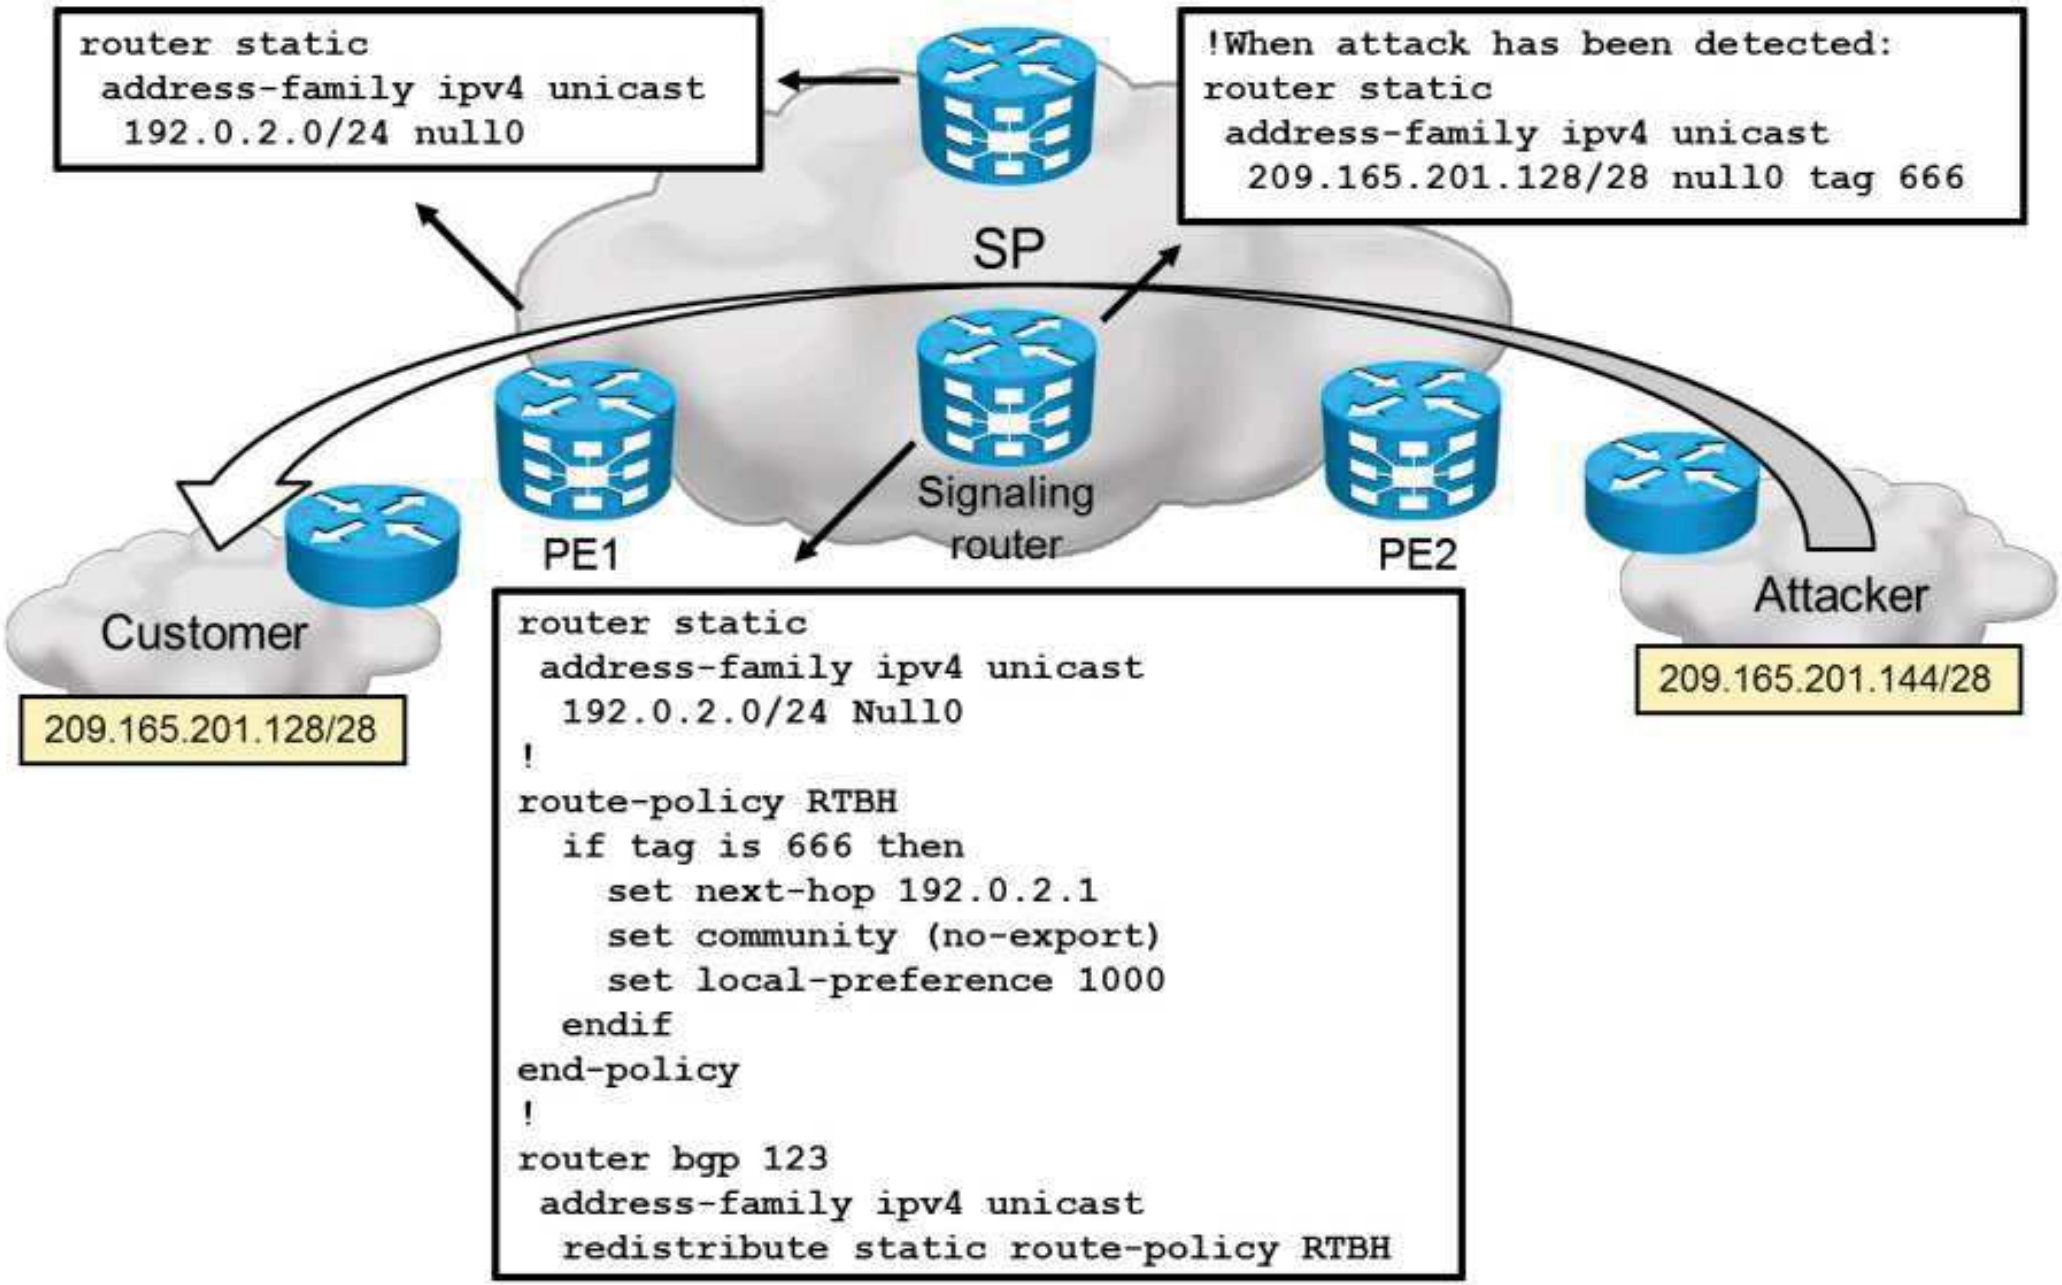

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

RTBH Verification

```
RP/0/RSP0/CPU0:PE1# show bgp
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
*>i209.165.201.128/28 192.0.2.1                0   1000      0 I
<… output omitted …>
```

- Displays BGP table.

```
RP/0/RSP0/CPU0:PE1# ping 209.165.201.129
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.201.129, timeout is 2 seconds:
UUUUU
Success rate is 0 percent (0/5)
```

- Pings the target.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Source-Based RTBH

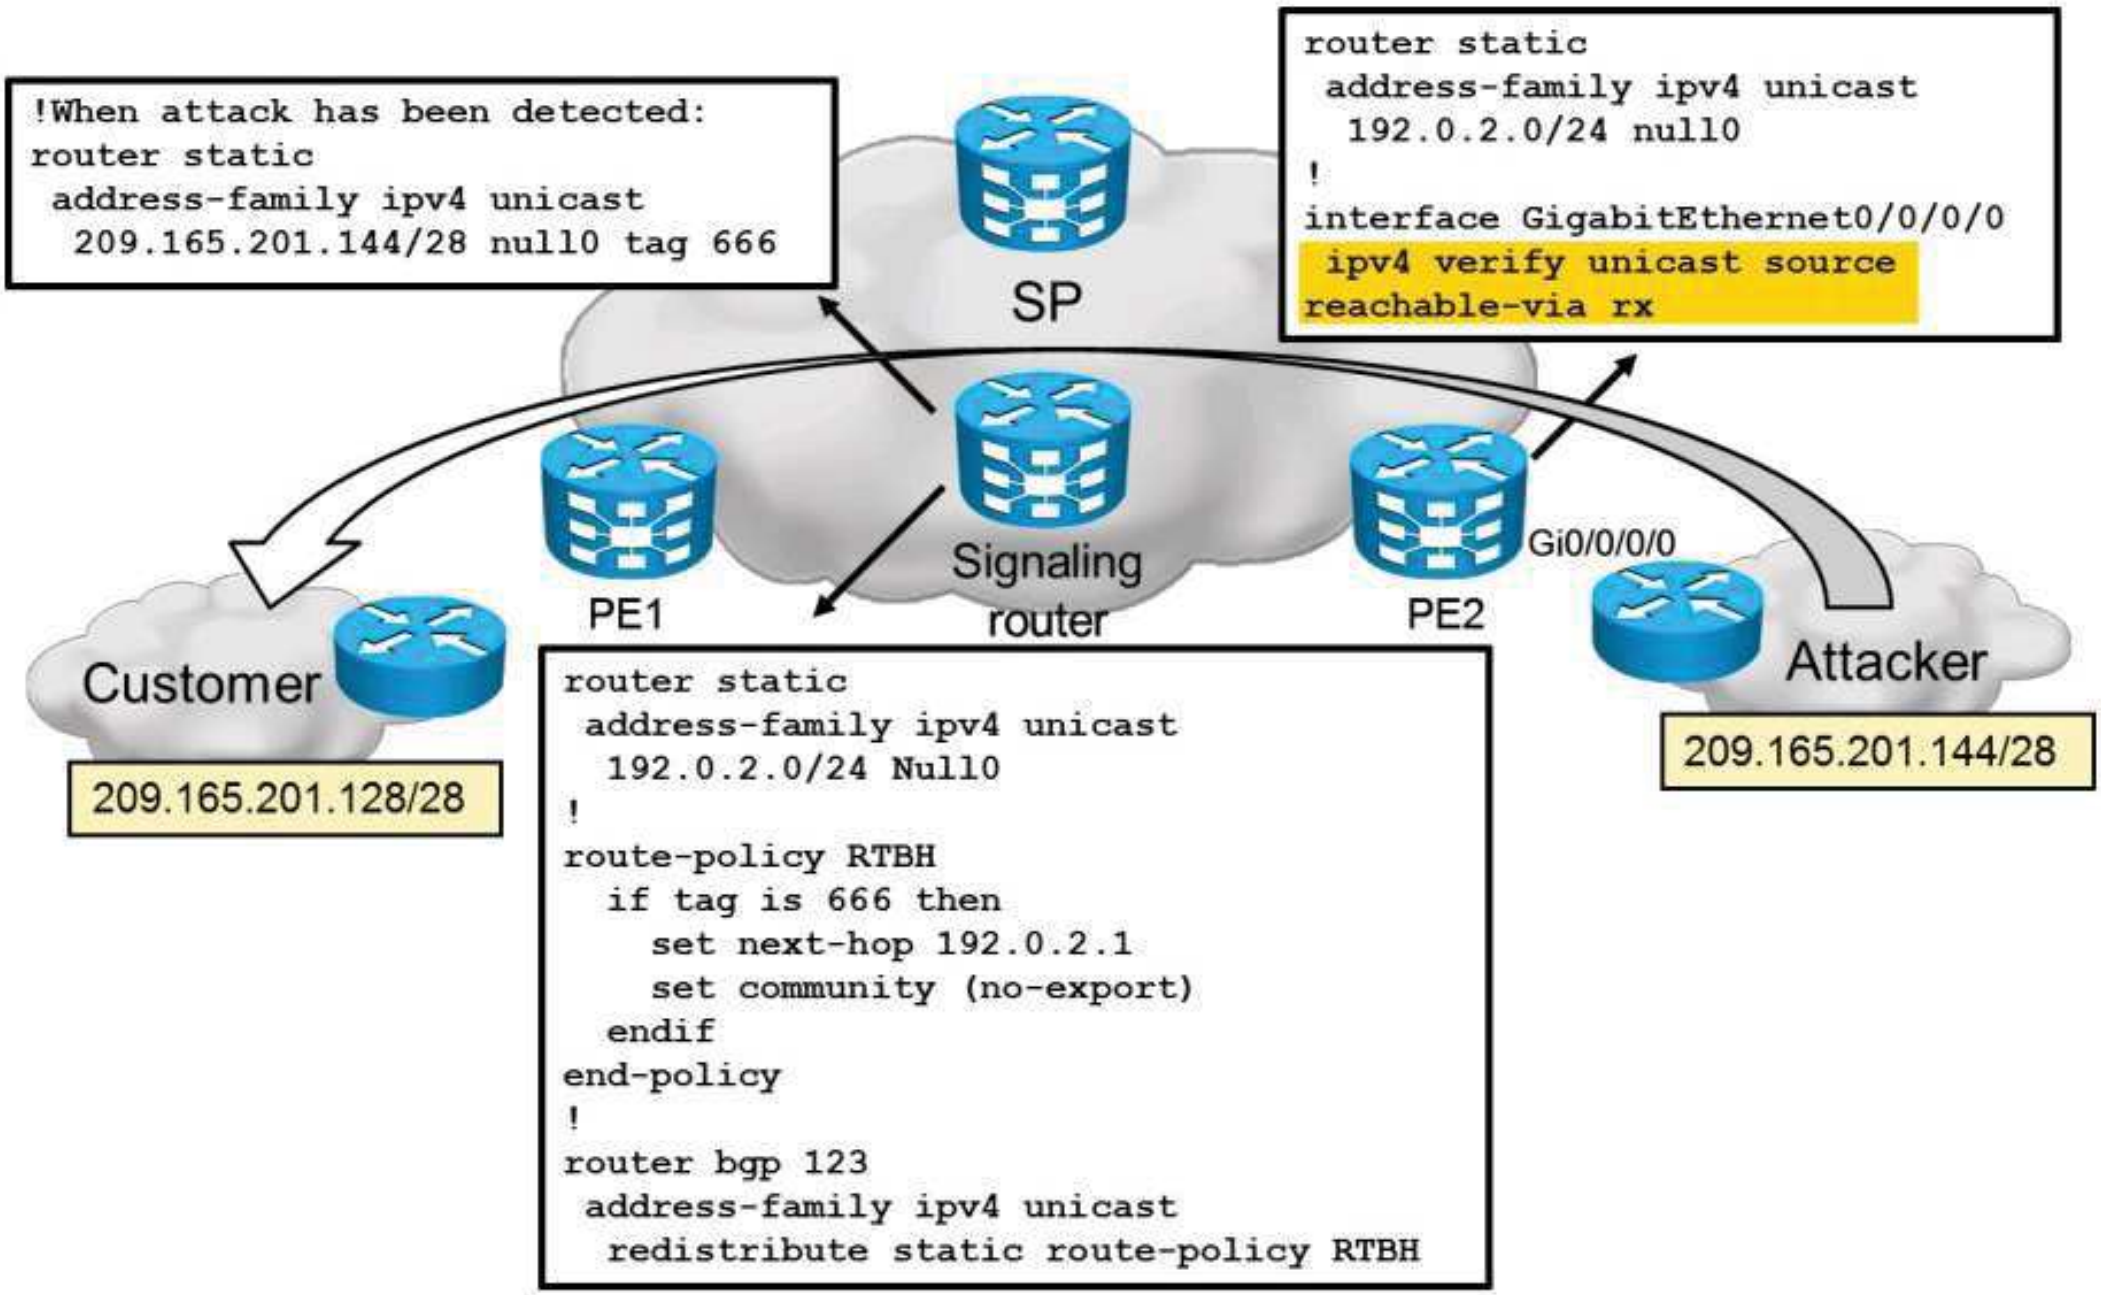

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Nonstop Forwarding

Cisco NSF characteristics:

- Cisco NSF is applicable in platforms with dual RPs and works together with SSO.
- Cisco NSF allows:
  - Routing neighbor relationships remain established during SSO.
  - Routes on neighboring routers remain valid.
  - Forwarding of data packets continues while the routing process on the new RP converges.
- Cisco NSF is supported by:
  - Routing protocols (OSPF, IS-IS, EIGRP, BGP)
  - Forwarding operation (Cisco Express Forwarding)
- The device must be Cisco NSF-capable.
- The neighboring device must be Cisco NSF-aware.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Nonstop Forwarding (Cont.)

Cisco NSF overview:

- One RP is active, one is standby.
- Cisco Express Forwarding on the active RP synchronizes the FIB and adjacency table to the standby RP.
- Upon switchover, the new active RP uses the old FIB and adjacency table to forward packets while the routing protocol reconverges.
- BGP must:
  - Establish neighbor relationship without causing a reset of neighbor relationship
  - Learn routing information
- As the routing protocol starts to repopulate the RIB, it updates Cisco Express Forwarding.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Nonstop Forwarding (Cont.)

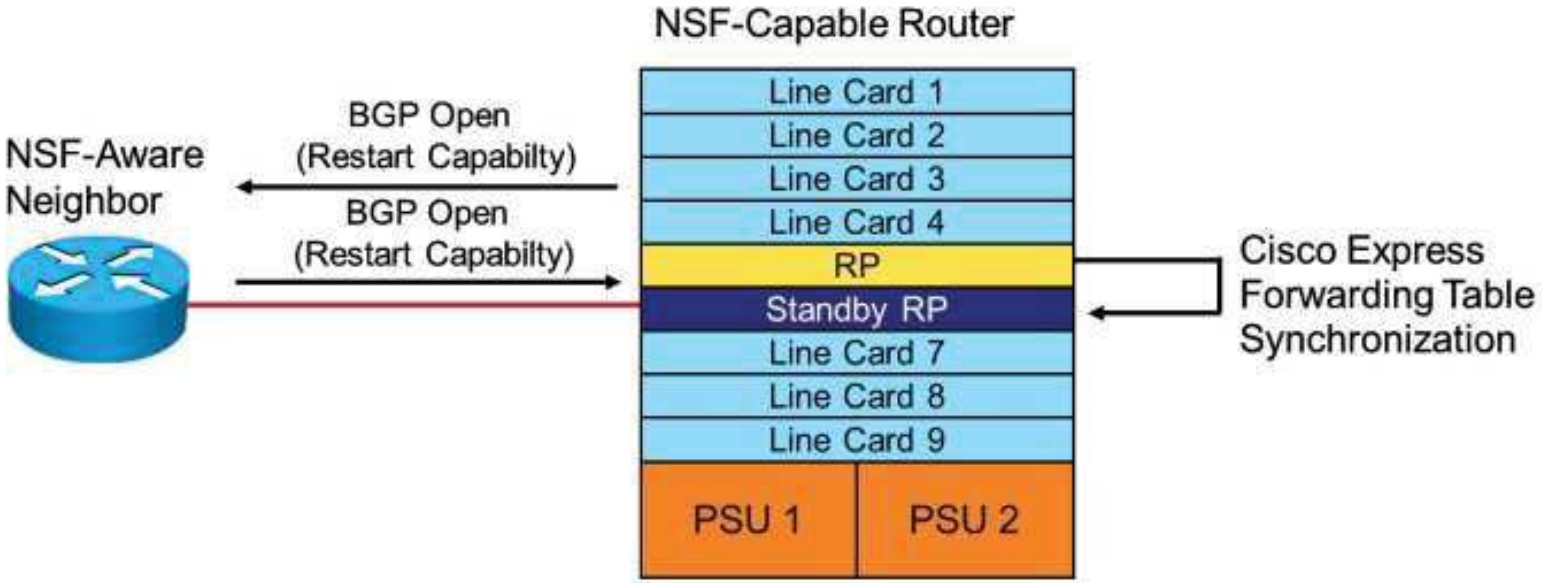

Cisco NSF-aware neighbors:

- Do not reconverge
- Help the NSF-capable router restart
- Continue forwarding traffic to the restarting router

Cisco NSF-capable router:

- Data is forwarded in the hardware base.
- Rebuilds the Layer 3 routing protocol database from the neighbor.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Cisco Nonstop Routing

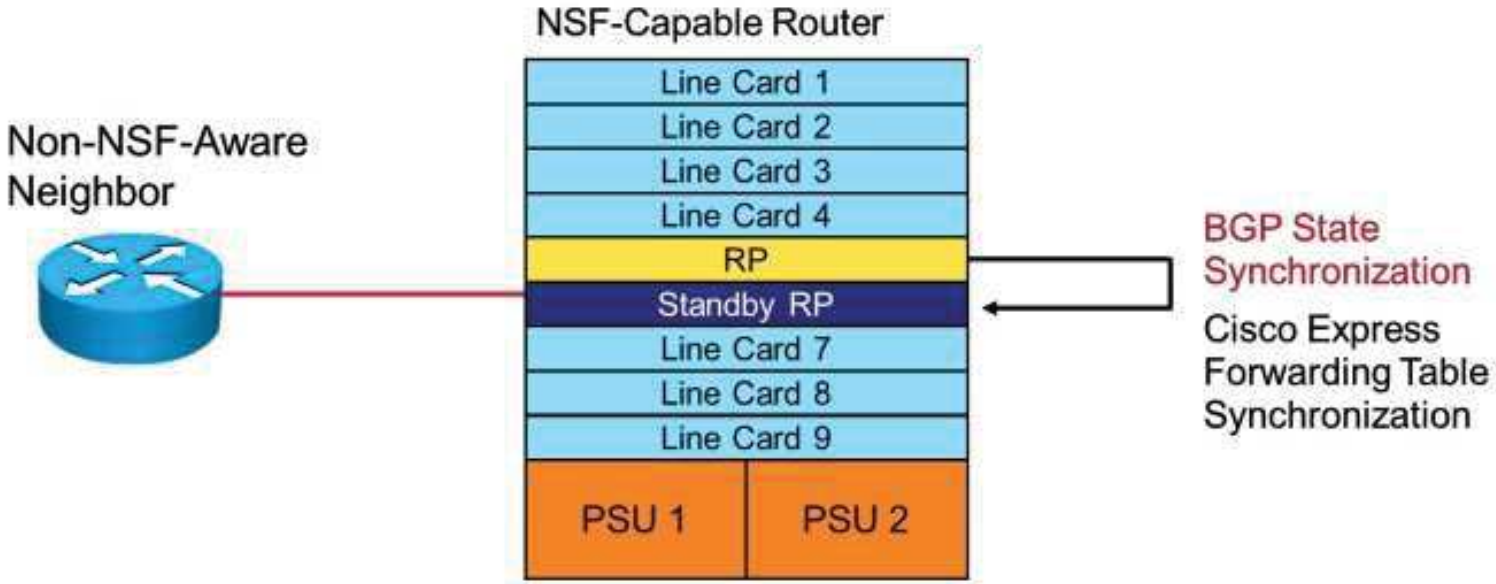

Cisco NSR characteristics:

- Used between CE and PE routers running BGP.
- Cisco NSR brings NSF operations to CE routers that are not NSF-aware:
  - BGP NSR uses SSO to maintain BGP state for EBGP connections between RPs.
  - NSR detects NSF-aware neighbors and runs NSF with them to save resources.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP NSF and NSR Configuration

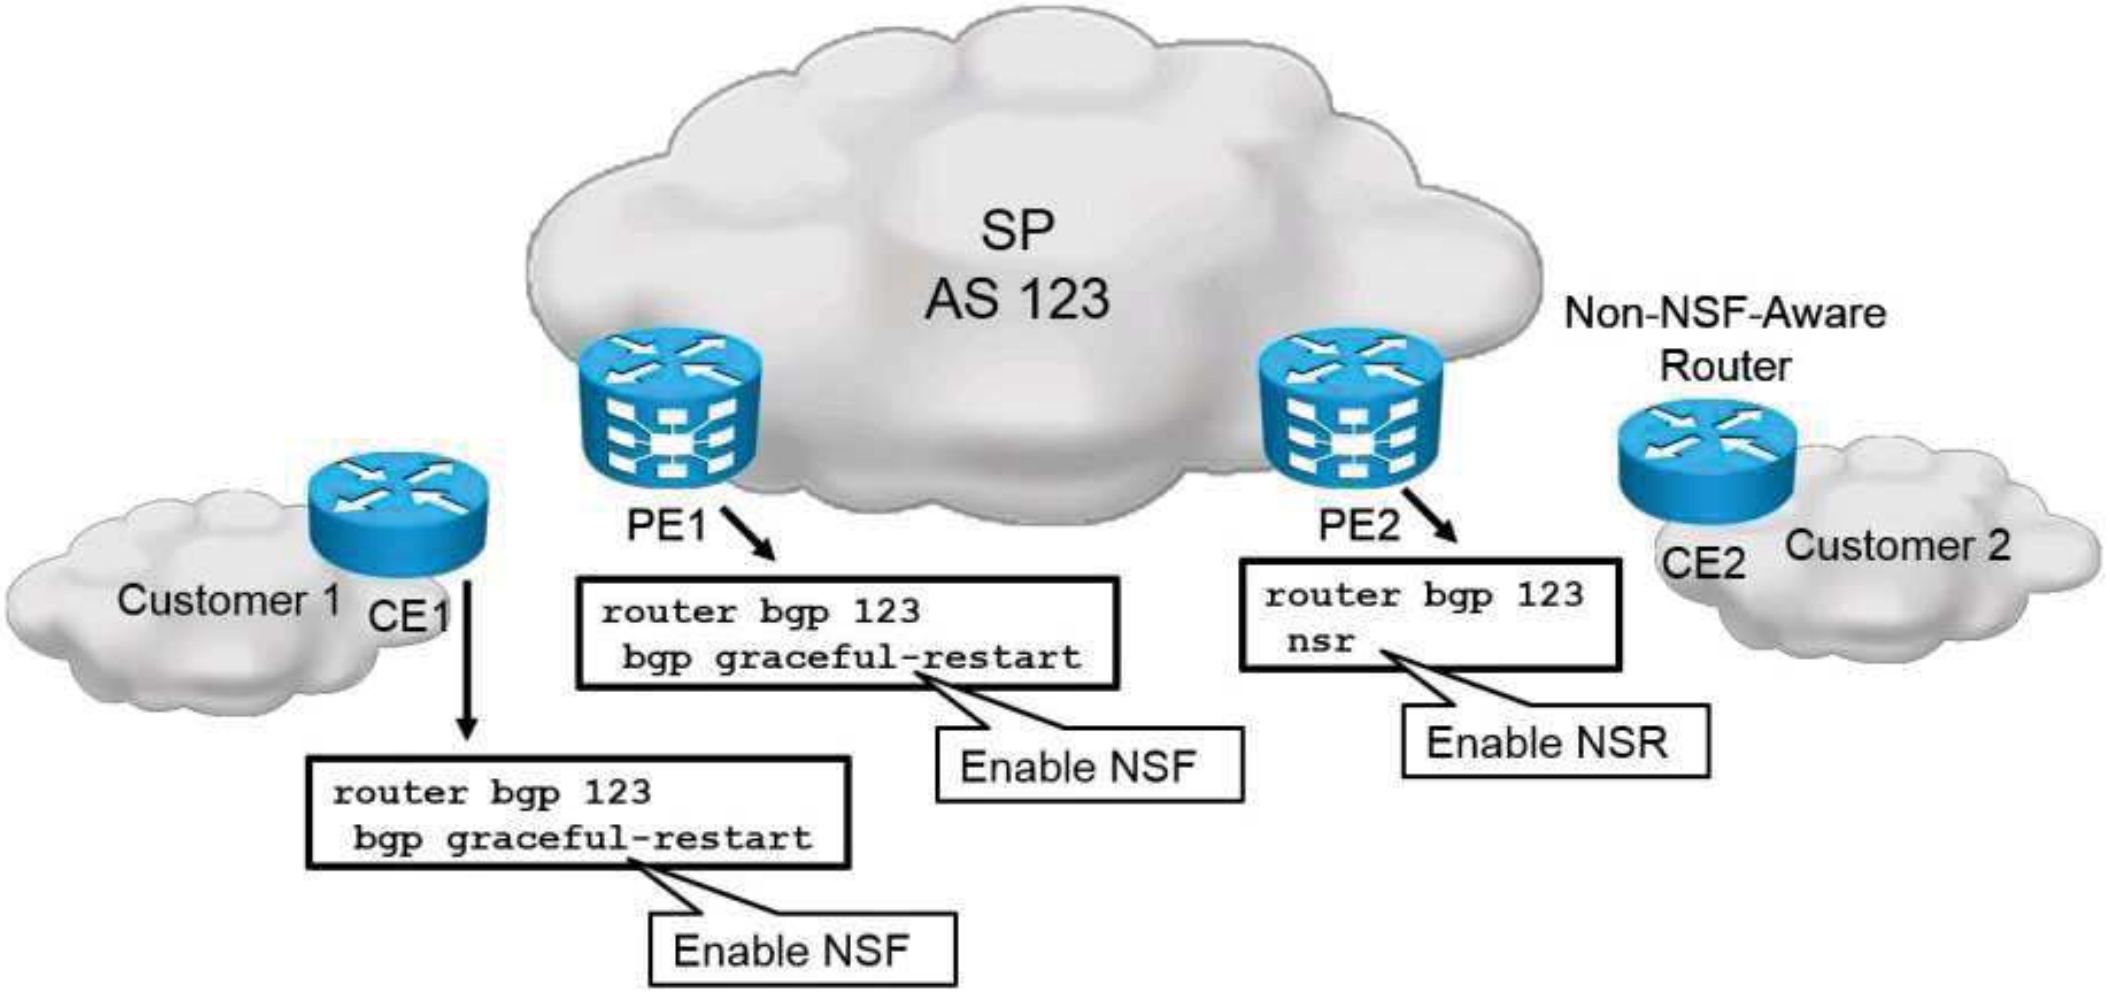

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Process Restart

Cisco IOS XR microkernel architecture enables restart of most processes.

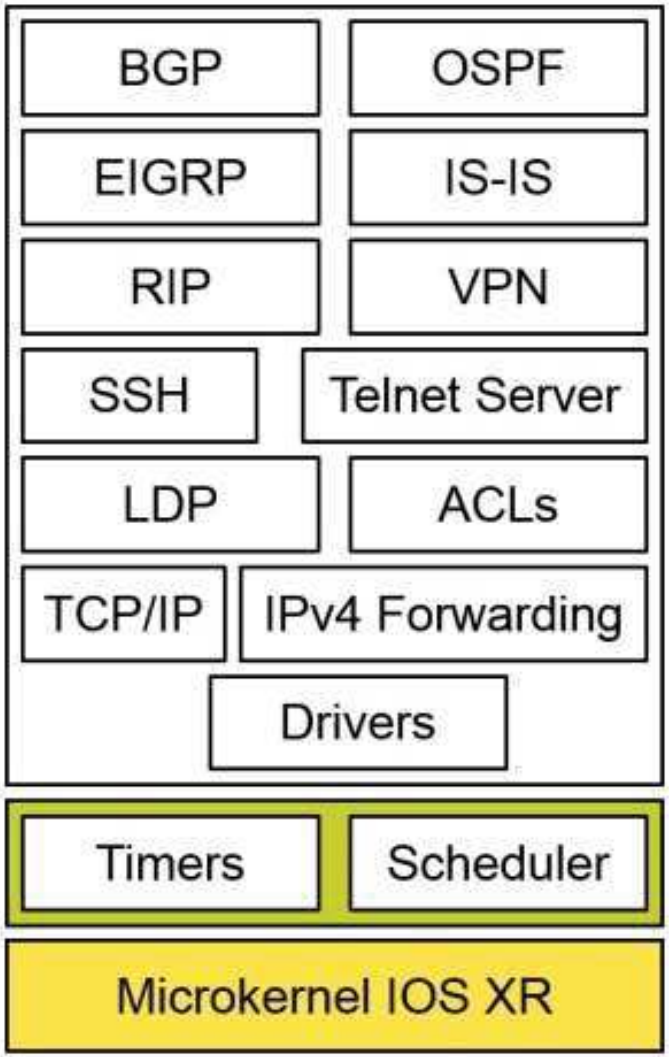

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- BGP security and optimization options are used in the infrastructure layer of the Cisco IP NGN.
- BGP as any service can be a target of malicious attacks.
- Protection mechanisms can reduce the risk of an attack if implemented.
- Maximum prefix limitation will shut down peering session if too many prefixes are received.
- BGP passwords add security mechanism that prevents arbitrary connections even if other parameters match.
- TTL security check can be used to verify hop distance of the peer.
- Control plane policing can limit the amount of traffic sent to the control plane of the router and thus reduce the possibility of an attacker overwhelming the CPU.
- All these features are configured on per-neighbor basis.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- RTBH makes it possible for the network to signal border routers which traffic to discard to reduce the effects of DDoS attacks based on destination or source addresses.
- With the destination-based RTBH implementation, traffic going to the IP address of the victim is discarded on the edge of the service provider.
- With the source-based RTBH implementation, traffic coming from the IP addresses of the attacker is discarded on the service provider edge.
- NSF allows forwarding to continue in the event of supervisor failover.
- NSR allows NSF capabilities with routers that are not NSF-aware.
- BGP supports both NSF and NSR.
- In case of a software failure IOS XR supports automatic recovery by restarting BGP processes.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

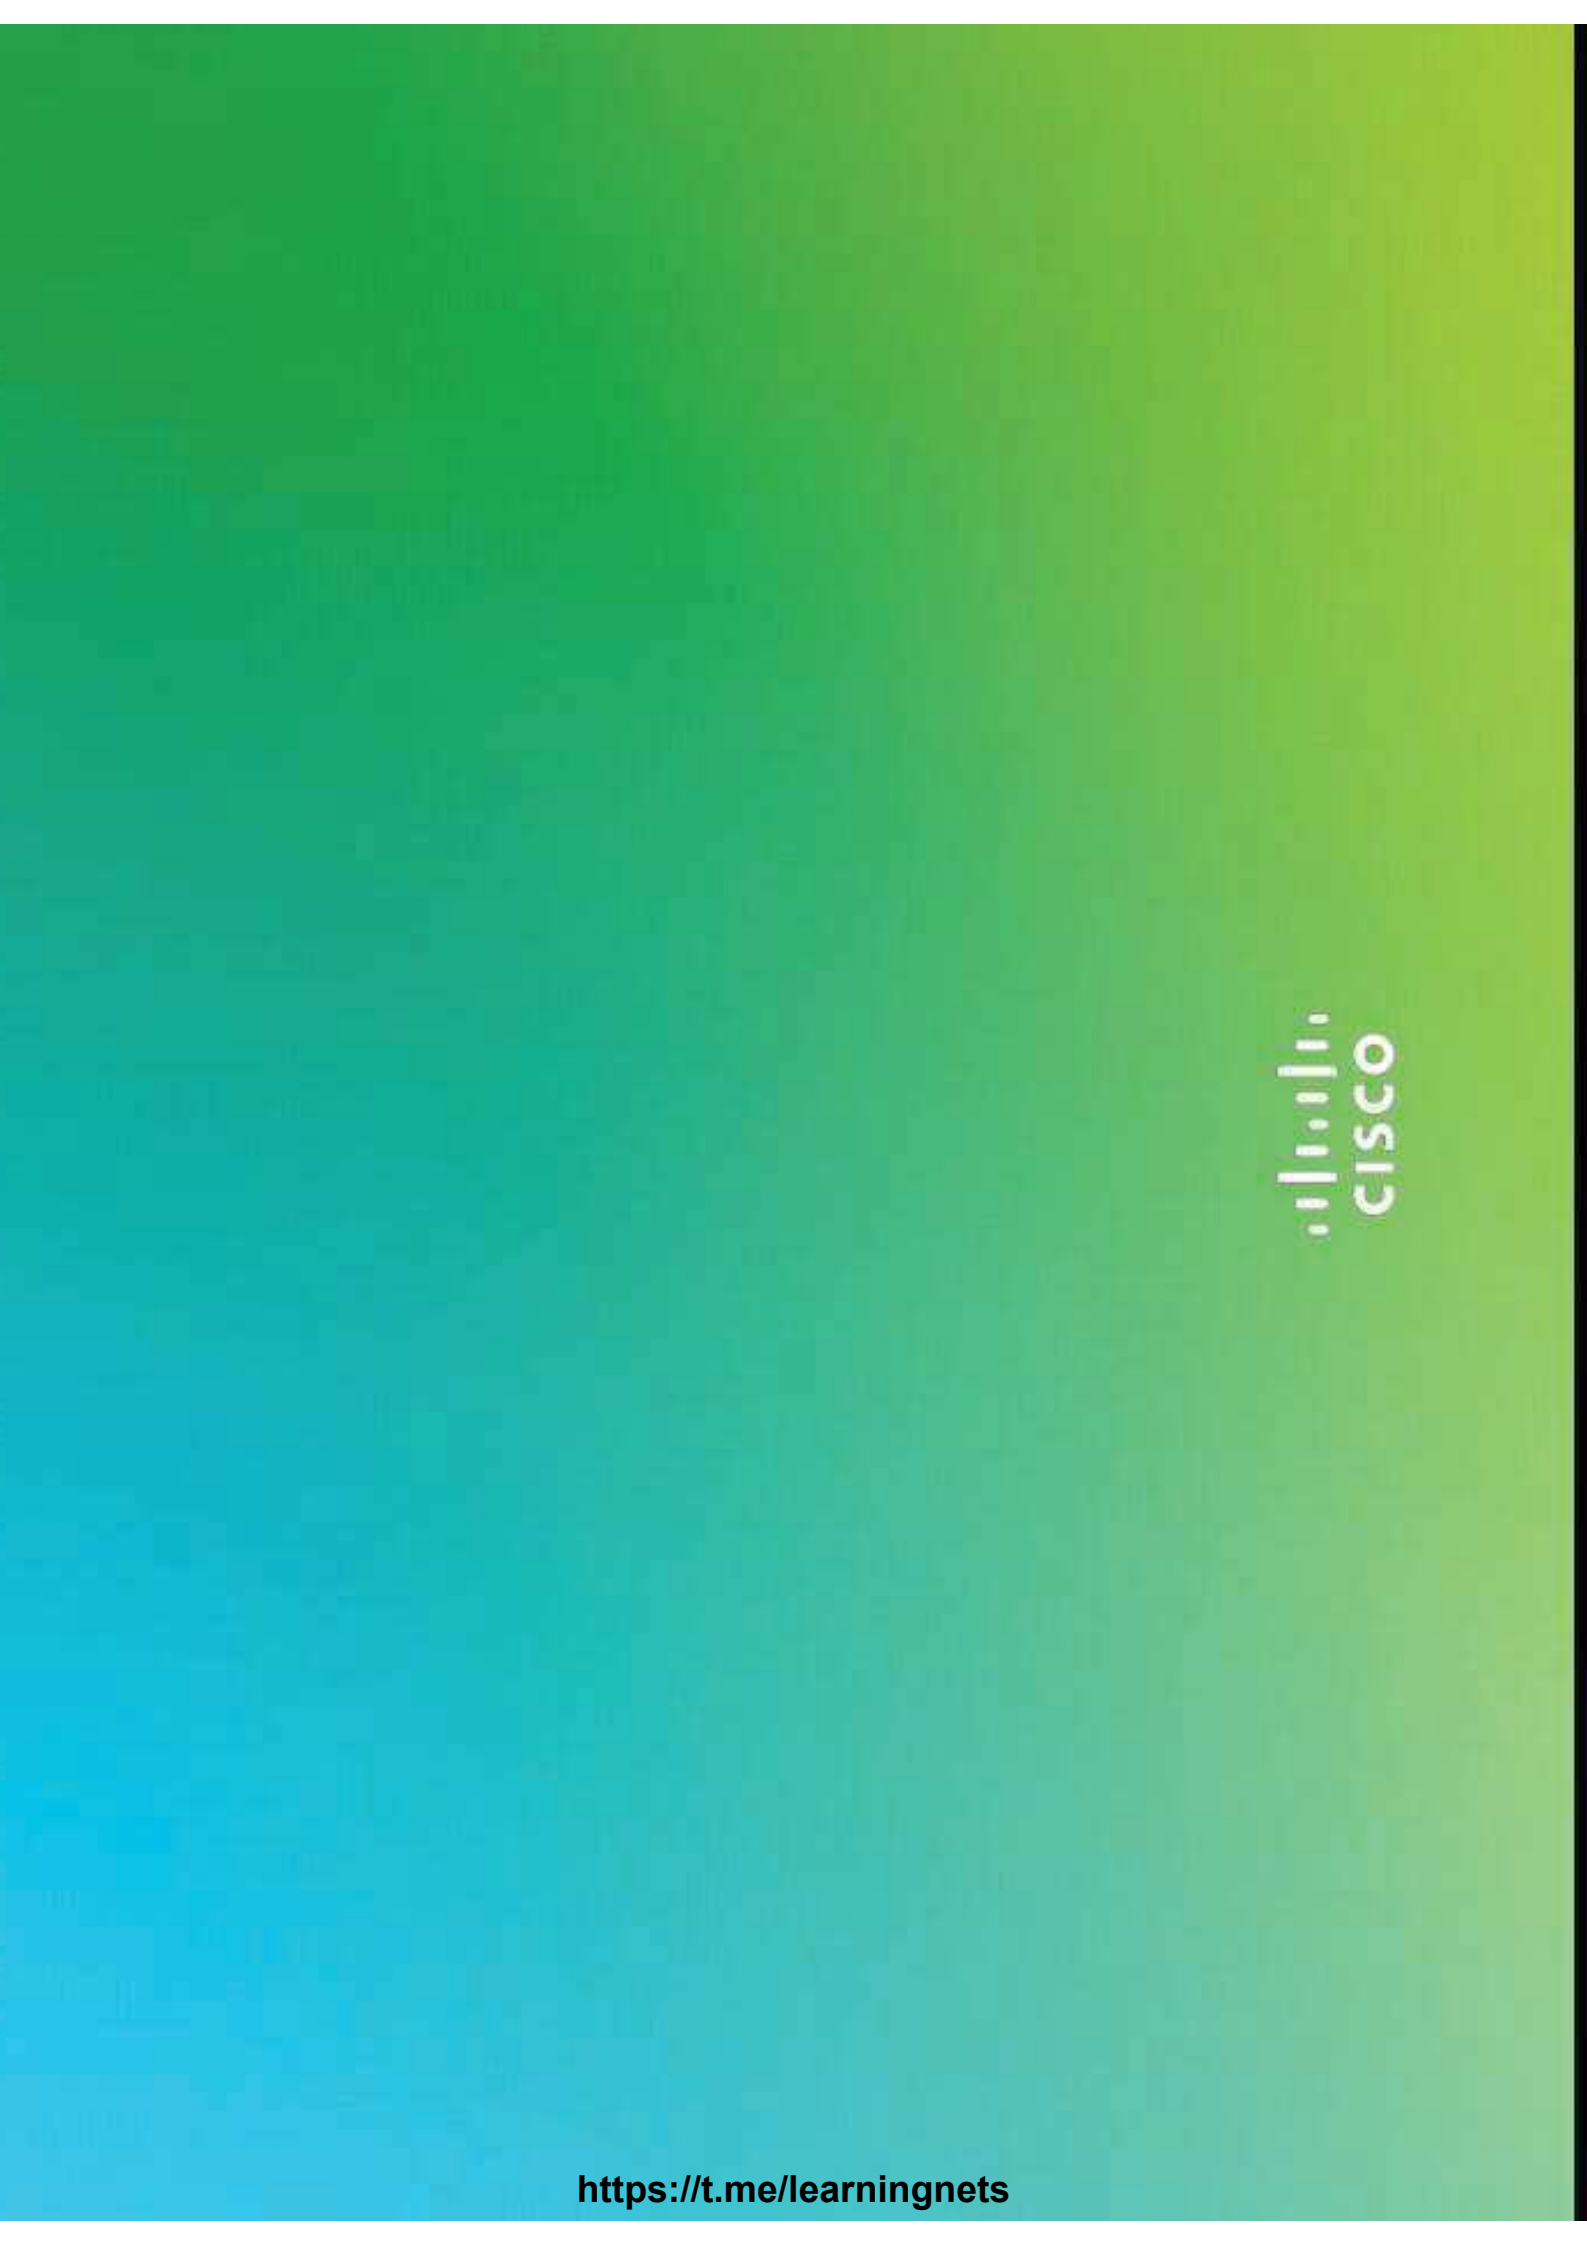

https://t.me/learningnets

https://t.me/learningnets

Improving BGP Convergence

Secure and Optimize BGP

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route-Dampening Overview

BGP route-dampening characteristics:

- Designed to reduce router processing load caused by unstable routes
- Minimizes the amount of BGP update processing in the Internet by suppressing unstable (flapping) routes
- Does not suppress routes that occasionally flap
- Suppresses routes that are likely to flap in the future, based on the history of their behavior

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route-Dampening Overview (Cont.)

Route-dampening operation:

- Each time an EBGP route flaps, it gets 1000 penalty points.
- The penalty placed on a route decays according to the exponential decay algorithm.
- When the penalty exceeds the suppress limit, the route is dampened.
- A penalty is applied to the individual path in the BGP table, not to the IP prefix.

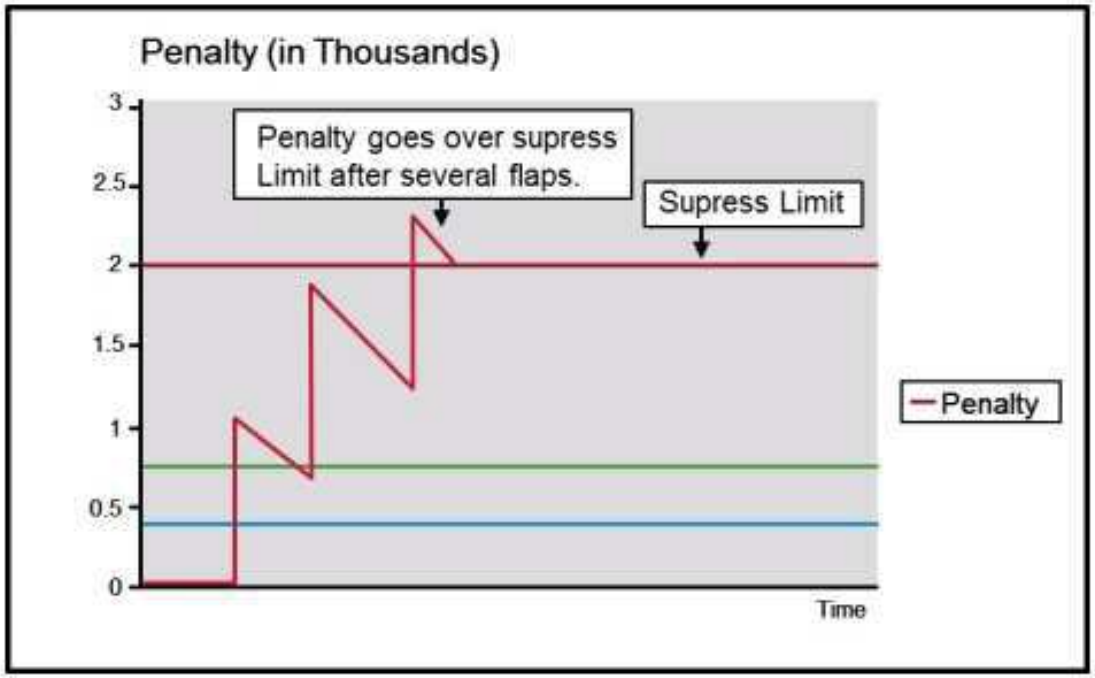

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Route-Dampening Overview (Cont.)

Route-dampening operation continuation:

- The route is never dampened for longer than the maximum suppression time limit.
- An unreachable route with a flap history is put in the history state.
- A dampened route is propagated when the penalty drops below the reuse limit.
- The flap history is forgotten when the penalty drops below half of the reuse limit.

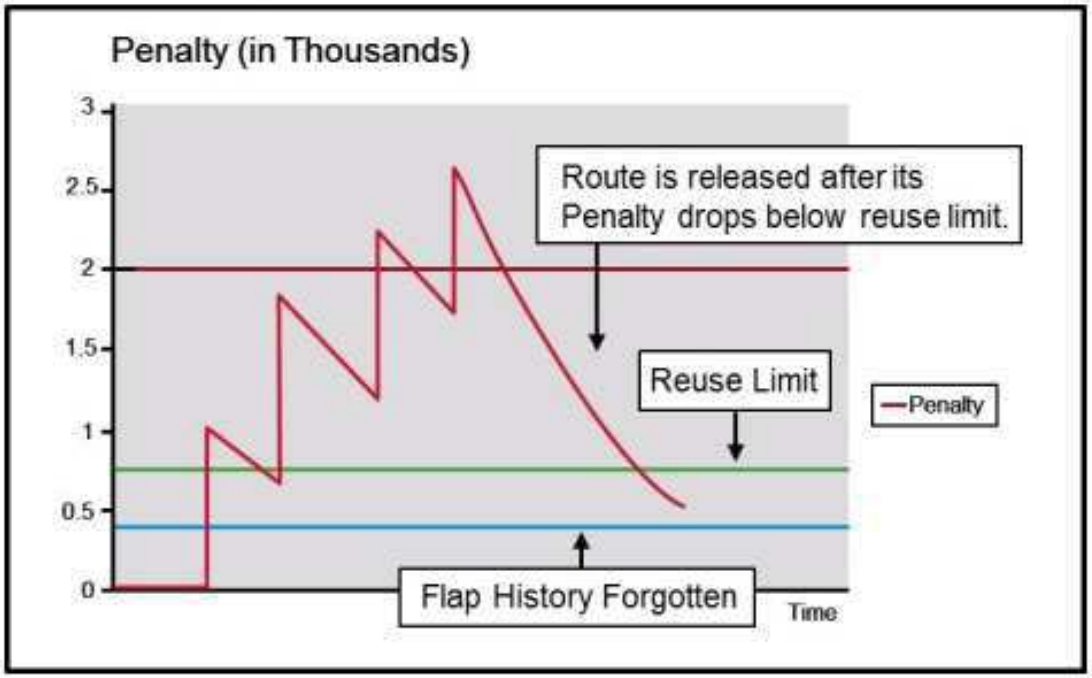

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Configuring BGP Route Dampening

BGP route-dampening configuration characteristics:

- Configured globally
- BGP dampening parameters:
  - **half-life**: Decay time in which the penalty is halved
  - **suppress**: Value when the route starts dampening
  - **reuse**: Value when the dampened route is reused
  - **max-suppress-time**: Maximum time to suppress the route
- Default values that are used by most service providers:
  - **half-life**: 15 minutes
  - **suppress**: 2000
  - **reuse**: 750
  - **max-suppress-time**: 60 minutes (4x half-life)
- Can also be enabled selectively using a route policy or a route map

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Configuring BGP Route Dampening (Cont.)


```
router bgp 123
 bgp dampening 10 1000 3000 40
```

Enable BGP dampening for all IPv4 routes.

```
route-policy BGP_DAMP
 if destination in (209.165.201.144/28) then
  set dampening halflife 10 suppress 3000
reuse 1000 max-suppress 40
  endif
end-policy
!
router bgp 123
 address-family ipv4 unicast
  bgp dampening route-policy BGP_DAMP
```

Create a route policy to match only specific routes and set dampening parameters.

Enable selective BGP dampening

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Verifying BGP Route Dampening

```
RP/0/RSP0/CPU0:PE1# show bgp 209.165.201.144/28
<… output omitted …>
234, (suppressed due to dampening)
      209.165.201.144 from 192.168.105.51 (10.5.100.1)
         Origin IGP, metric 0, localpref 100, valid, external
         Received Path ID 0, Local Path ID 0, version 0
         Dampinfo: penalty 3620, flapped 4 times in 00:04:14, reuse in 00:27:00
                   half life 00:10:00, suppress value 3000, reuse value 1000
                   Maximum suppress time 00:40:00
```

- Displays a BGP route.

```
RP/0/RSP0/CPU0:PE1# show bgp dampened-paths
<… output omitted …>
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            From             Reuse    Path
*d 209.165.201.144/28 192.168.105.51   00:28:10 64505 i
```

- Displays dampened BGP routes.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Verifying BGP Route Dampening (Cont.)

```
RP/0/RSP0/CPU0:PE1#
debug bgp [address-family] dampening
```

- Displays the BGP dampening events.

```
RP/0/RSP0/CPU0:PE1#
show [address-family] bgp flap-statistics
```

- Displays flap statistics for all routes with dampening history.

```
RP/0/RSP0/CPU0:PE1#
clear bgp [address-family] dampening [ip-address/prefix]
```

- Releases all the dampened routes or just the specified network.

```
RP/0/RSP0/CPU0:PE1#
clear bgp [address-family] flap-statistics [ip-address/prefix]
```

- Clears BGP flap statistics for all routes or just the specified network.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Convergence

BGP convergence characteristics:

- As the number of routes in the Internet routing table grows, the time required for BGP to converge increases.
- The Internet currently contains more than 300,000 prefixes.
- Network convergence times can range from 10 minutes to more than one hour.
- BGP is considered converged when:
  - All routes have been accepted.
  - All routes have been installed in the routing table.
  - The input queue and output queue for all peers is 0.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Processes

Process	Description	Interval
BGP open	Performs BGP peer establishment.	At initialization, when establishing a TCP connection with a BGP peer
BGP I/O	Handles queuing and processing of BGP packets (updates and keepalives).	As BGP control packets are received
BGP scanner	Walks the BGP table and confirms reachability of the next hops. BGP scanner also checks conditional advertisement to determine whether BGP should advertise condition prefixes. Performs route dampening.	Every 60 seconds
BGP router	Calculates the best BGP path and processes any route changes. Also sends and receives routes, establishes peers, and interacts with the routing information base.	Once per second and when adding, removing, or soft-reconfiguring a BGP peer

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Processes (Cont.)

CPU effects of BGP processes:

- BGP scanner process:
  - High CPU utilization stemming from the BGP scanner process can be expected for short durations on a router carrying a large Internet routing table.
  - While the BGP scanner runs, low-priority processes need to wait a longer time to access the CPU.
- BGP router process:
  - The BGP router process runs about once per second to check for work.
  - The BGP router consumes all free CPU cycles.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Improving BGP Convergence

You can reduce BGP convergence time and high CPU utilization caused by BGP processes in the following ways:

- Implement distributed BGP:
  - Reduces CPU utilization and increases stability
- Implement BFD:
  - Reduces BGP convergence by fast detection of neighbor failure
- Implement BGP PIC:
  - Reduces convergence by storing BGP backup/alternate path in RIB and FIB
- Enable the path MTU feature:
  - Improves efficiency by dynamically determining the largest MTU that you can use without creating packets that need to be fragmented
- Increase interface input queues:
  - Improves convergence by reducing dropped TCP ACKs

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Distributed BGP

Distributed BGP characteristics:

- Supported on Cisco IOS XR Software only
- Splits BGP functionality into three process types with several instances:
  - BGP process manager (one instance)
  - BGP RIB process (one instance per address family)
  - BGP speaker process (up to 15 instances)
- Used to reduce the impact that a fault in one address family has on another address family
- Must be enabled

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

PMTU Discovery

PMTU discovery characteristics:

- Used to automatically determine TCP MSS used for TCP connections from a router
- Default TCP MSS value for BGP is 536 bytes
- Small TCP MSS affects BGP convergence:
  - Higher TCP MSS can improve BGP convergence.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

PMTU Increasing Input Queue Depth

PMTU increasing input queue depth characteristics:

- Available on Cisco IOS and IOS XE Software only.
- Input queue on an interface specifies how many packets can be queued before dropping the packets.
- BGP routers with several peers might experience packet drops on an interface due to a large number of TCP ACK segments.
- The default input hold queue is platform-dependent.
- A length of 1000 will normally resolve problems caused by input queue drops of TCP ACKs.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

PMTU Discovery, Hold Queue, and Distributed BGP Configurations

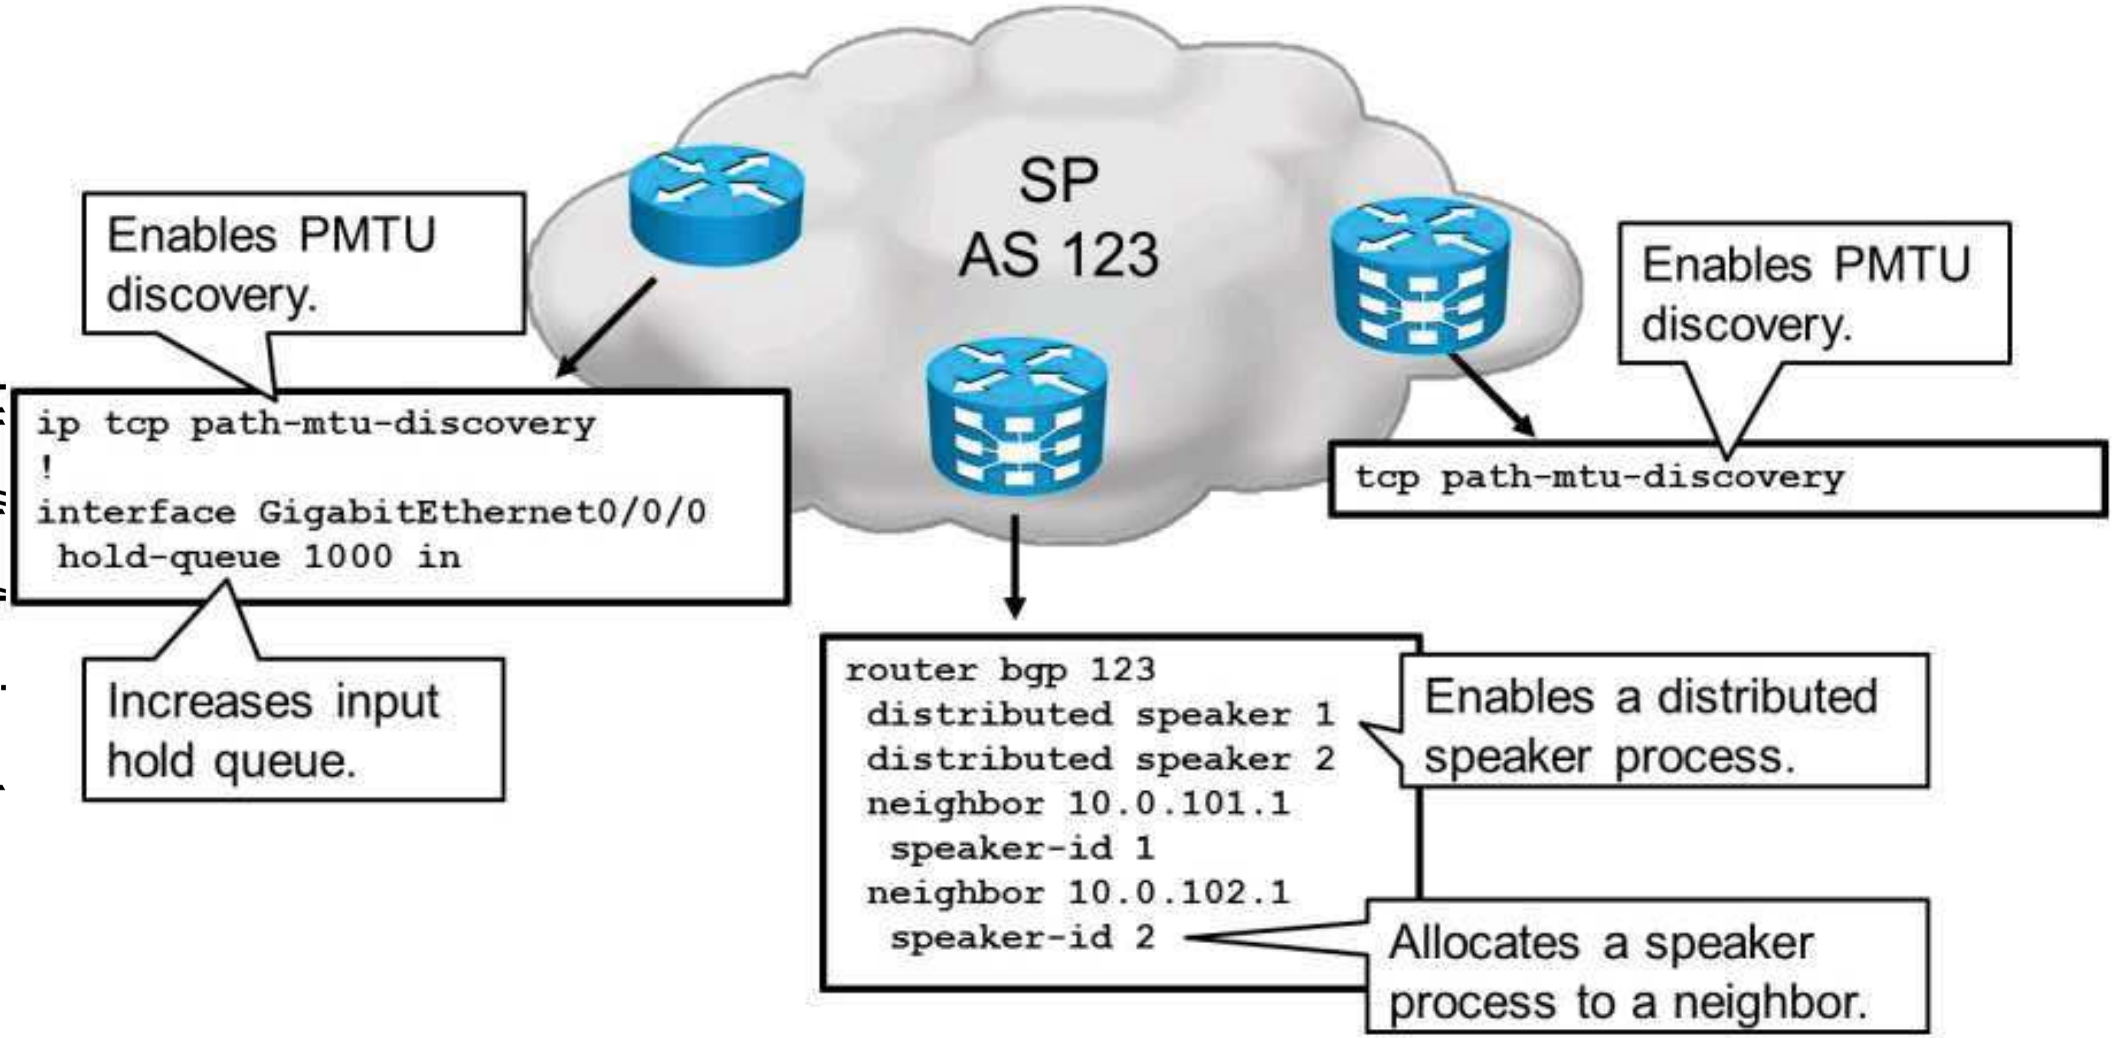

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

PMTU Discovery, Hold Queue, and Distributed BGP Verifications

```
RP/0/RSP0/CPU0:P# show bgp process
BGP Process Information:
BGP is operating in DISTRIBUTED mode
Autonomous System number format: ASPLAIN
Autonomous System: 64500
Router ID: 10.5.1.1
Default Cluster ID: 10.5.1.1
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60
<… output omitted …>
```

- Displays BGP process information.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

PMTU Discovery, Hold Queue, and Distributed BGP Verifications (Cont.)

```
PE1# show ip bgp neighbors | include Datagrams
Datagrams (max data segment is 1460 bytes):
<… output omitted …>
```

- Displays BGP neighbor information, including TCP MSS

```
PE1# show interface GigabitEthernet0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Hardware is ASR1001, address is e8b7.48fb.7100 (bia e8b7.48fb.7100)
  Internet address is 192.168.106.60/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 100Mbps, link type is auto, media type is T
  output flow-control is on, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:13, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/1000/0/0 (size/max/drops/flushes); Total output drops: 0
<… output omitted …>
```

- Displays interface information, including input queue depth.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Prefix Independent Convergence

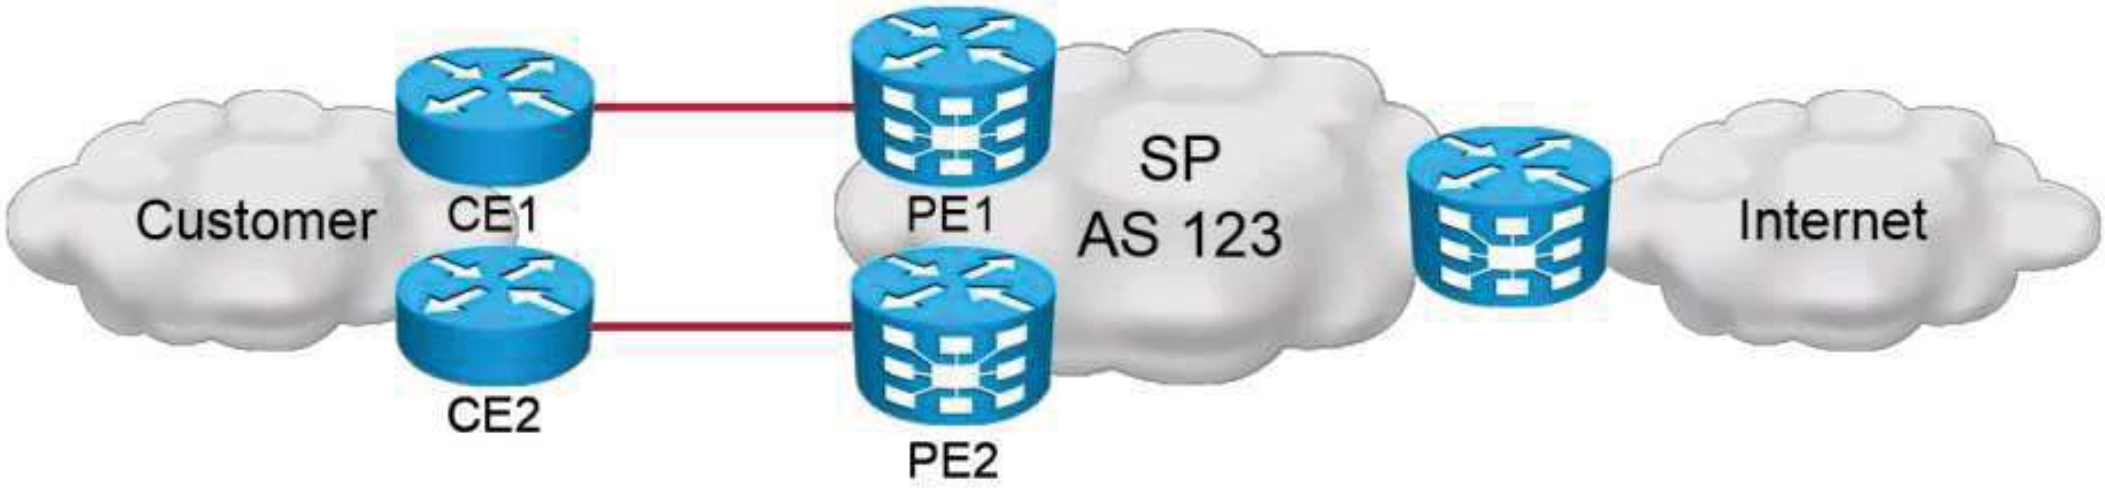

https://t.me/learningnets

BGP prefix independent convergence characteristics:

- PIC enhances BGP convergence, regardless of the number of BGP prefixes.
- PIC stores the BGP backup/alternate path for each prefix in BGP, RIB, and FIB tables.
- When the primary goes down, Cisco Express Forwarding quickly selects a different egress port for the affected destination.

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP PIC Configuration

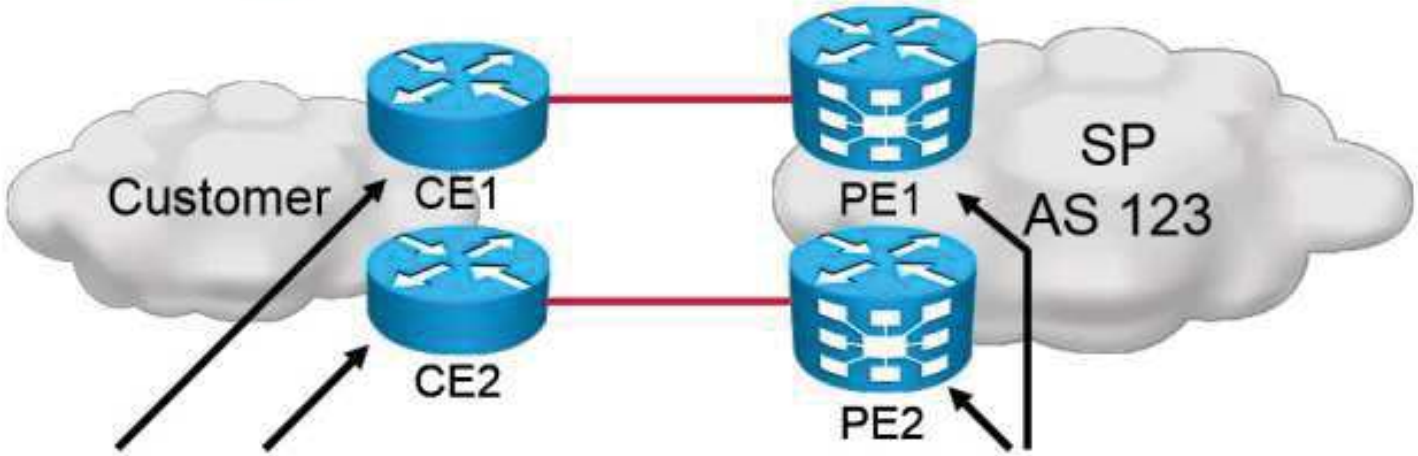


```
router bgp 234
 address-family ipv4 unicast
  bgp additional-paths install
 address-family ipv6 unicast
  bgp additional-paths install
```

Enable BGP PIC

```
route-policy ALL
  pass
end-policy
!
router bgp 234
 address-family ipv4 unicast
  additional-paths selection route-policy ALL
 address-family ipv6 unicast
  additional-paths selection route-policy ALL
```

Enable BGP PIC

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Bidirectional Forwarding Detection for BGP

BFD for BGP characteristics:

- Extremely lightweight hello protocol that uses UDP to test bidirectional communication
- Used to detect failures in the forwarding path between two adjacent routers
- Millisecond resolution of forwarding plane failure
- Relies on routing protocols to detect neighbors

https://t.me/learningnets

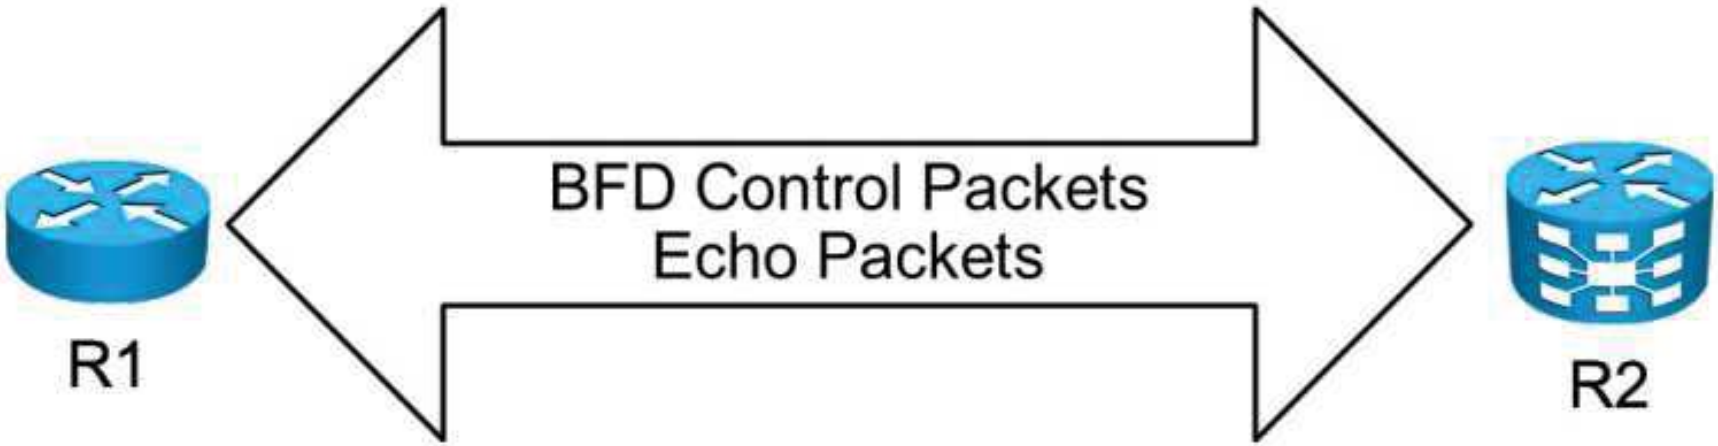

© 2014 Cisco and/or its affiliates. All rights reserved.

Bidirectional Forwarding Detection for BGP (Cont.)

BFD operation:

- Routing protocol (BFD client) bootstraps BFD to create a BFD session to a neighbor:
  - BFD client receives link status change notification.
  - Receive and transmit intervals are negotiated and configurable.
- The two systems agree on a method to detect failure.
- In case of failure, BFD notifies the BFD client:
  - The BFD client independently decides on the action.

https://t.me/learningnets

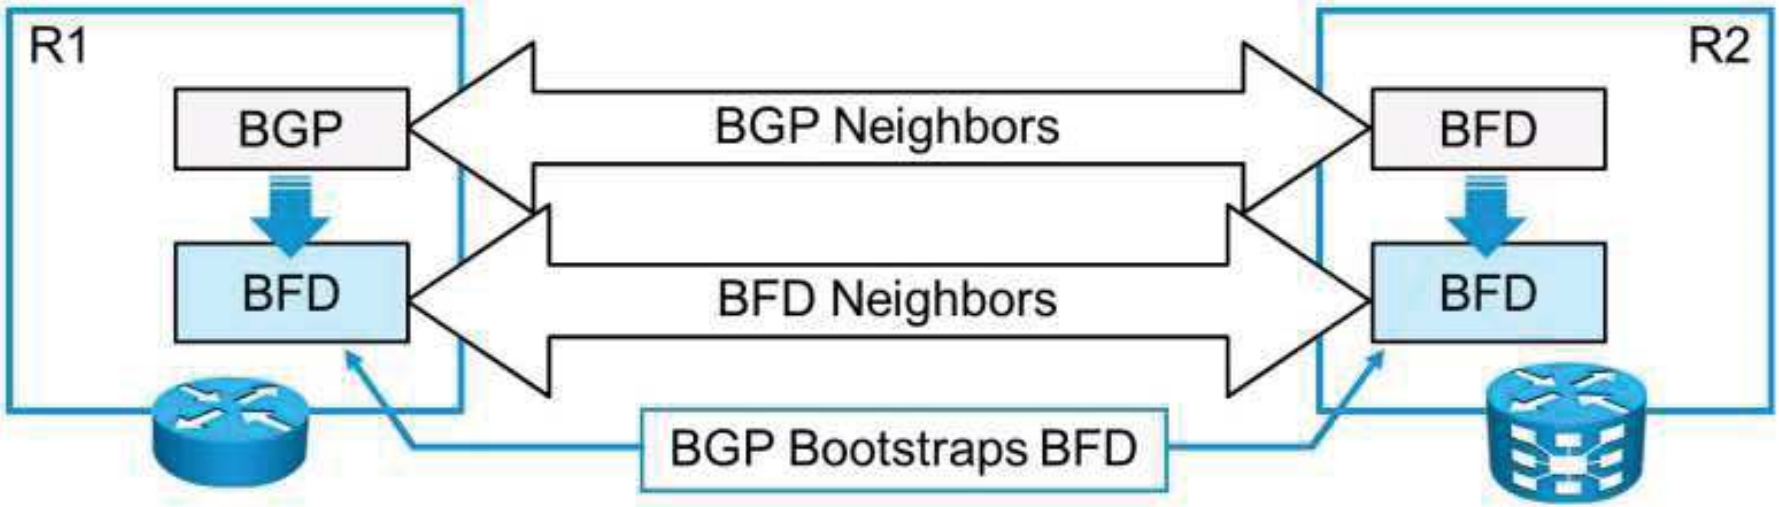

© 2014 Cisco and/or its affiliates. All rights reserved.

BFD Configuration

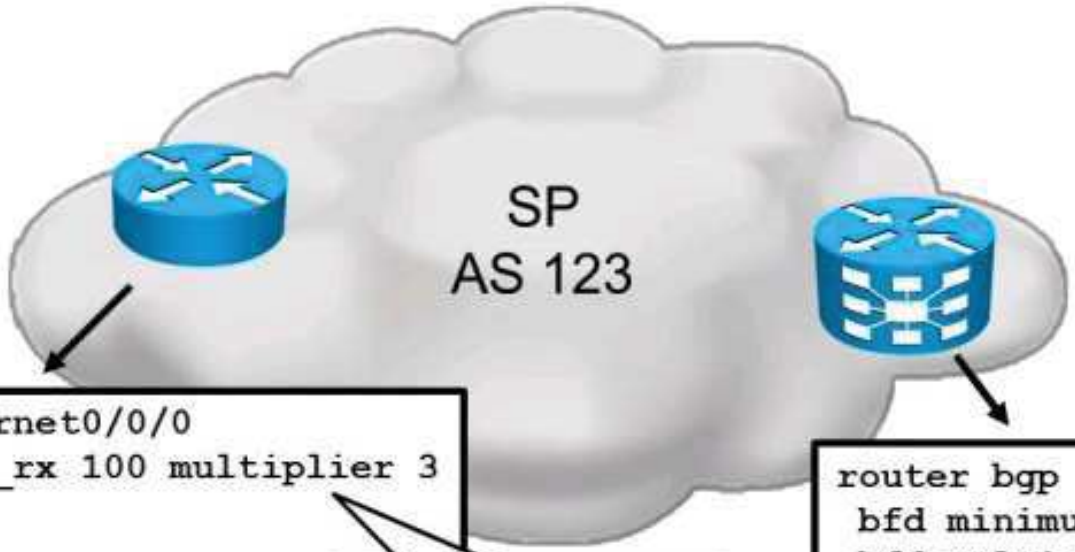


```
interface GigabitEthernet0/0/0
 bfd interval 100 min_rx 100 multiplier 3
!
router bgp 123
 neighbor 10.0.0.6 fall-over bfd
```

Enable BFD on an interface.

Enable BFD support for BGP.

```
router bgp 123
 bfd minimum-interval 100
 bfd multiplier 3
 neighbor 10.0.101.1
  bfd fast-detect
```

Enable BFD support for BGP.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Timers and Intervals

BGP scan time characteristics:

- Defines how often the BGP scanner process scans the BGP table.
- Needed to confirm that next hops are still available.
- The BGP scanner process is also responsible for advanced features such as conditional advertisement check and performing route dampening.
- Set to 60 seconds by default.

BGP advertisement interval characteristics:

- Defines a time that must elapse between two successive updates about the same destination that are sent to a neighbor.
- Default values are different for IBGP and EBGP neighbors.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Timers and Intervals (Cont.)

BGP keepalive timer characteristics:

- Defines a time between successive keepalive messages.
- Set to 60 seconds by default.

BGP hold-down timer characteristics:

- Defines how long a router will wait from the last received keepalive or update message before declaring the session dead.
- Set to 3x keepalive = 180 seconds by default.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Timers and Intervals (Cont.)

Improving BGP Convergence

- BGP convergence can also be improved to some extent by:
  - Lowering the scan time interval for the BGP scanner process.
  - Lowering the advertisement interval between BGP neighbors.
  - Lowering the keepalive and hold-down timers.
- Limitations:
  - Not recommended in routers dealing with large BGP tables.
  - Could lead to CPU or memory exhaustion.
  - Lower hold-down timers could lead to undesired session terminations.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Timers and Intervals Configuration

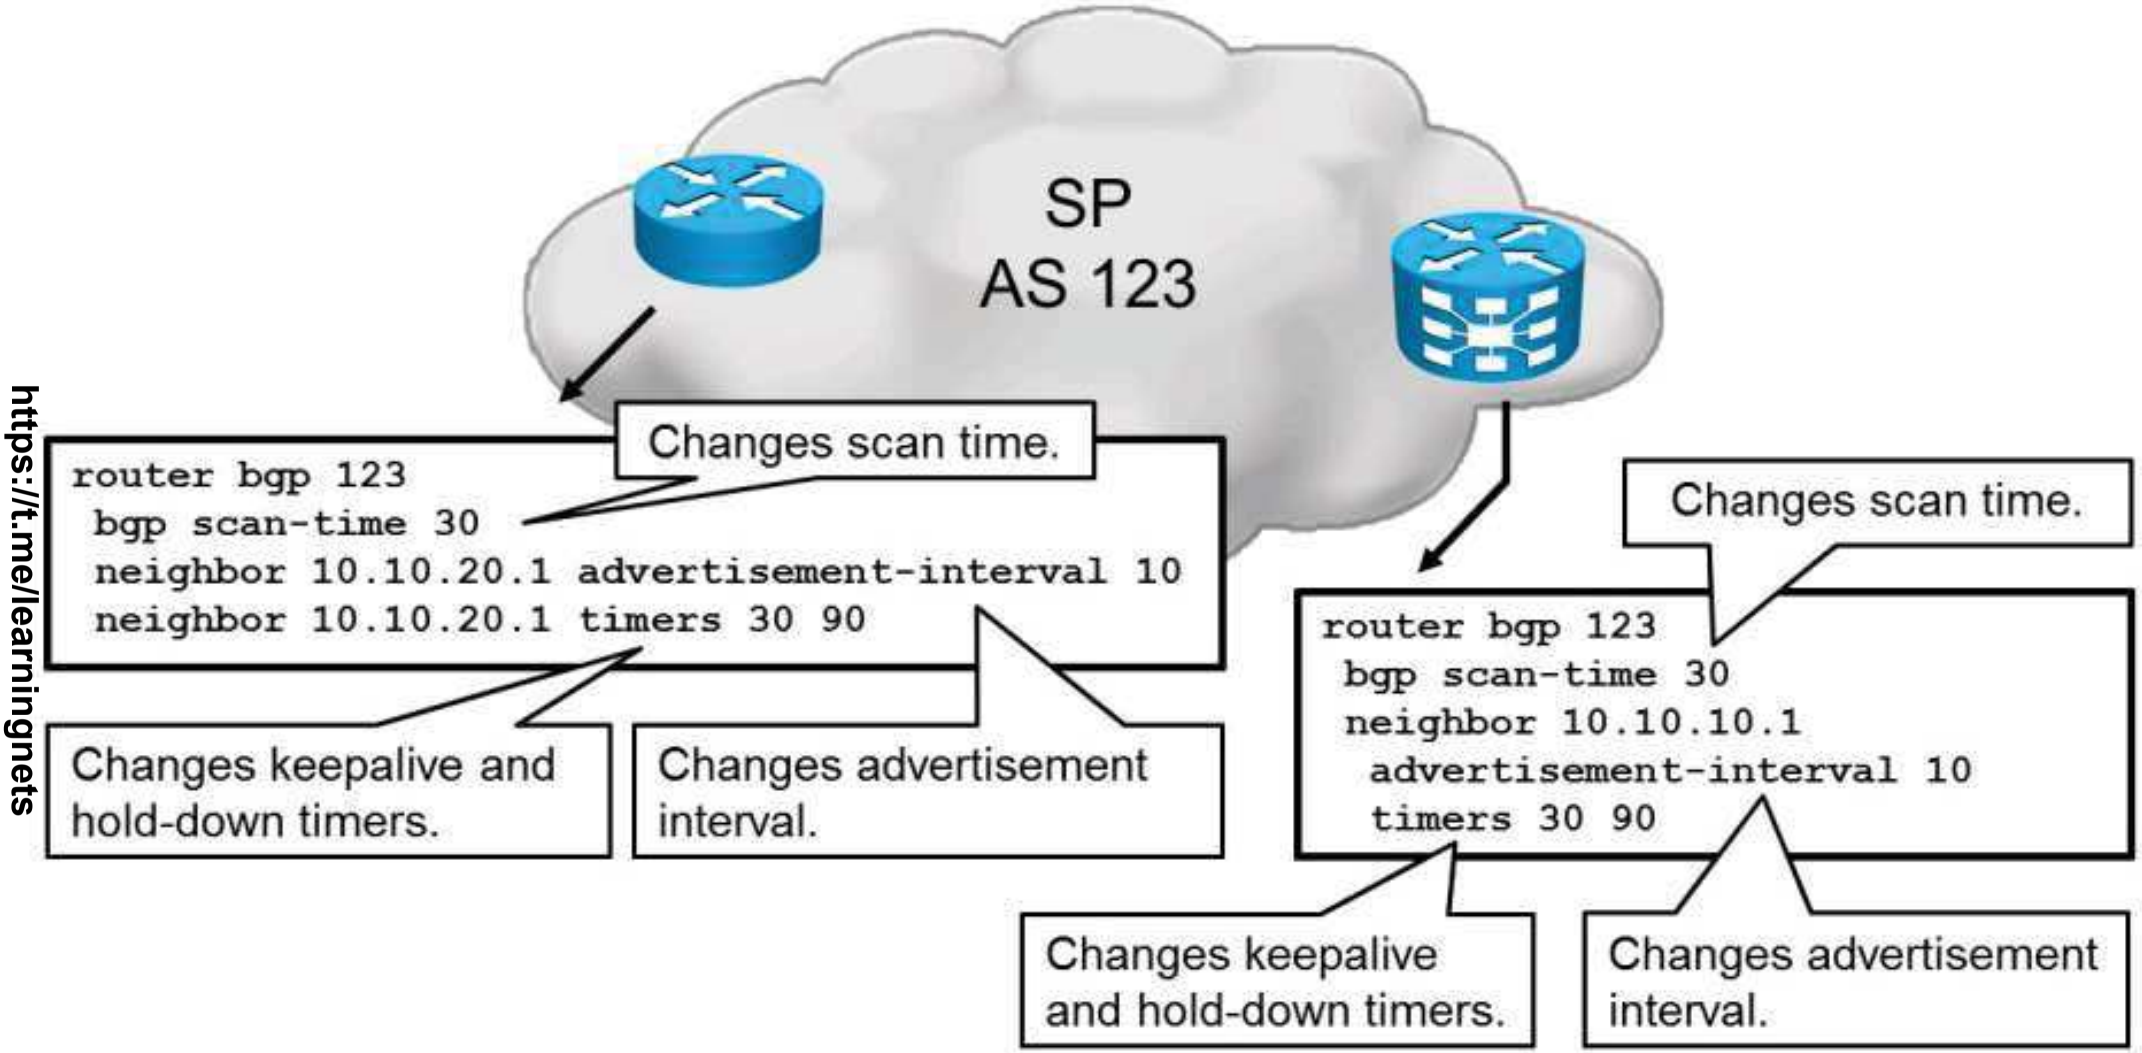

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Timers and Intervals Verification

```
RP/0/RSP0/CPU0:P# show bgp process
BGP Process Information:
BGP is operating in DISTRIBUTED mode
Autonomous System number format: ASPLAIN
Autonomous System: 64500
Router ID: 10.5.1.1
Default Cluster ID: 10.5.1.1
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 30
<… output omitted …>
```

- Displays BGP process information.

```
RP/0/RSP0/CPU0:PE5# show bgp neighbor 10.0.1.1 | include advertisement
    Minimum time between advertisement runs is 10 secs
```

- Displays BGP neighbor information.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- BGP route dampening reduces the effects of route flaps by removing offending prefixes from the routing table for a period of time.
- To enable BGP route dampening, use the **bgp dampening** command under BGP router configuration mode.
- To verify BGP route dampening, use the **show bgp** and **show bgp dampened-paths** commands.
- BGP is considered converged when no outstanding updates are to be sent to neighbors.
- BGP comprises several processes of equal importance for the functioning of BGP as a complete protocol.
- BGP provides several features to improve convergence times.
- Distributed BGP splits processing of different address families to reduce impact in case of a failure.
- PMTU discovery benefits BGP because it benefits any transmission over IP.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Input queue depth can be increased to improve performance.
- PMTUD, hold queue, and distributed BGP need to be configured for the BGP process.
- PIC enhances BGP convergence regardless of the number of BGP prefixes.
- BFD reduces conversion time by providing rapid detection of failed peers.
- BFD needs to be configured for directly connected eBGP peers or else its capabilities are available through the use of an IGP in conjunction with BFD.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

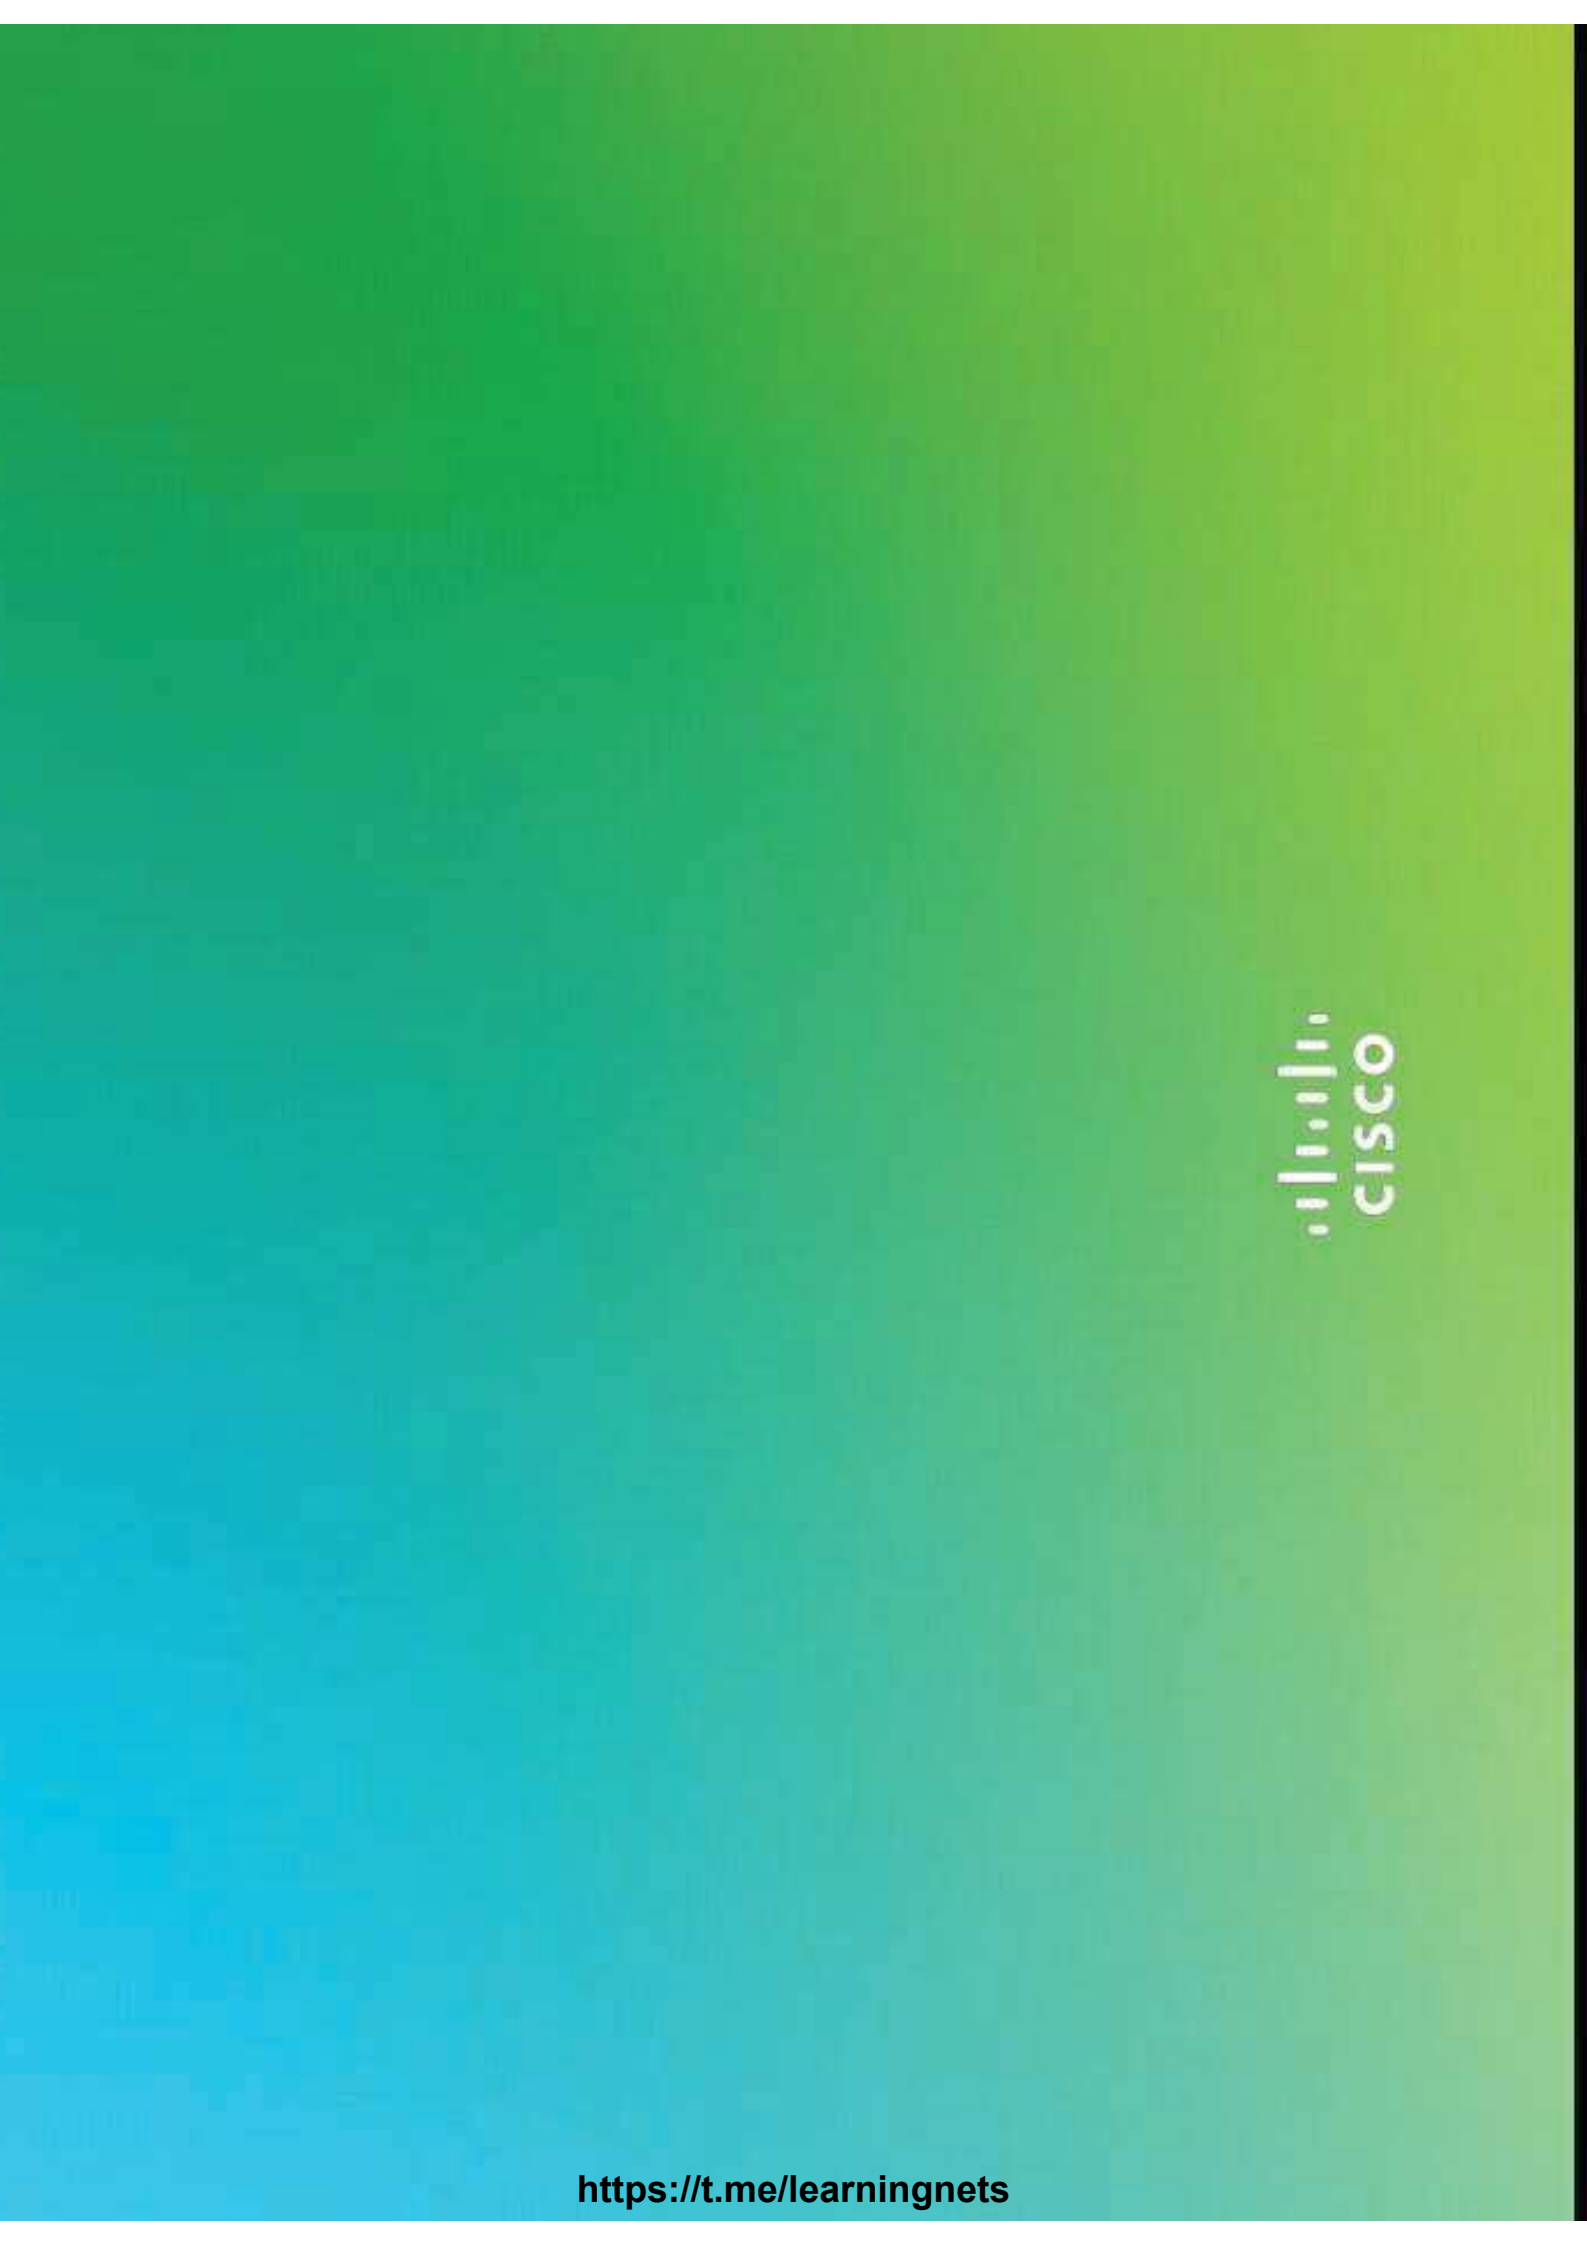

https://t.me/learningnets

https://t.me/learningnets

Improving BGP Configuration Scalability

Secure and Optimize BGP

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Groups Overview

These are the BGP peer groups characteristics:

- BGP routers could have a large number of neighbors with similar requirements:
  - Provider edge router with many customer connections.
  - BGP route reflector with many IBGP peers.
  - Provider edge router at an exchange point.
  - Most of the parameters specified for the BGP neighbors are identical, with a few exceptions.
- The solution is to group common parameters in a BGP peer group.
- Available only on Cisco IOS and IOS XE Software.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Example: CE Connections

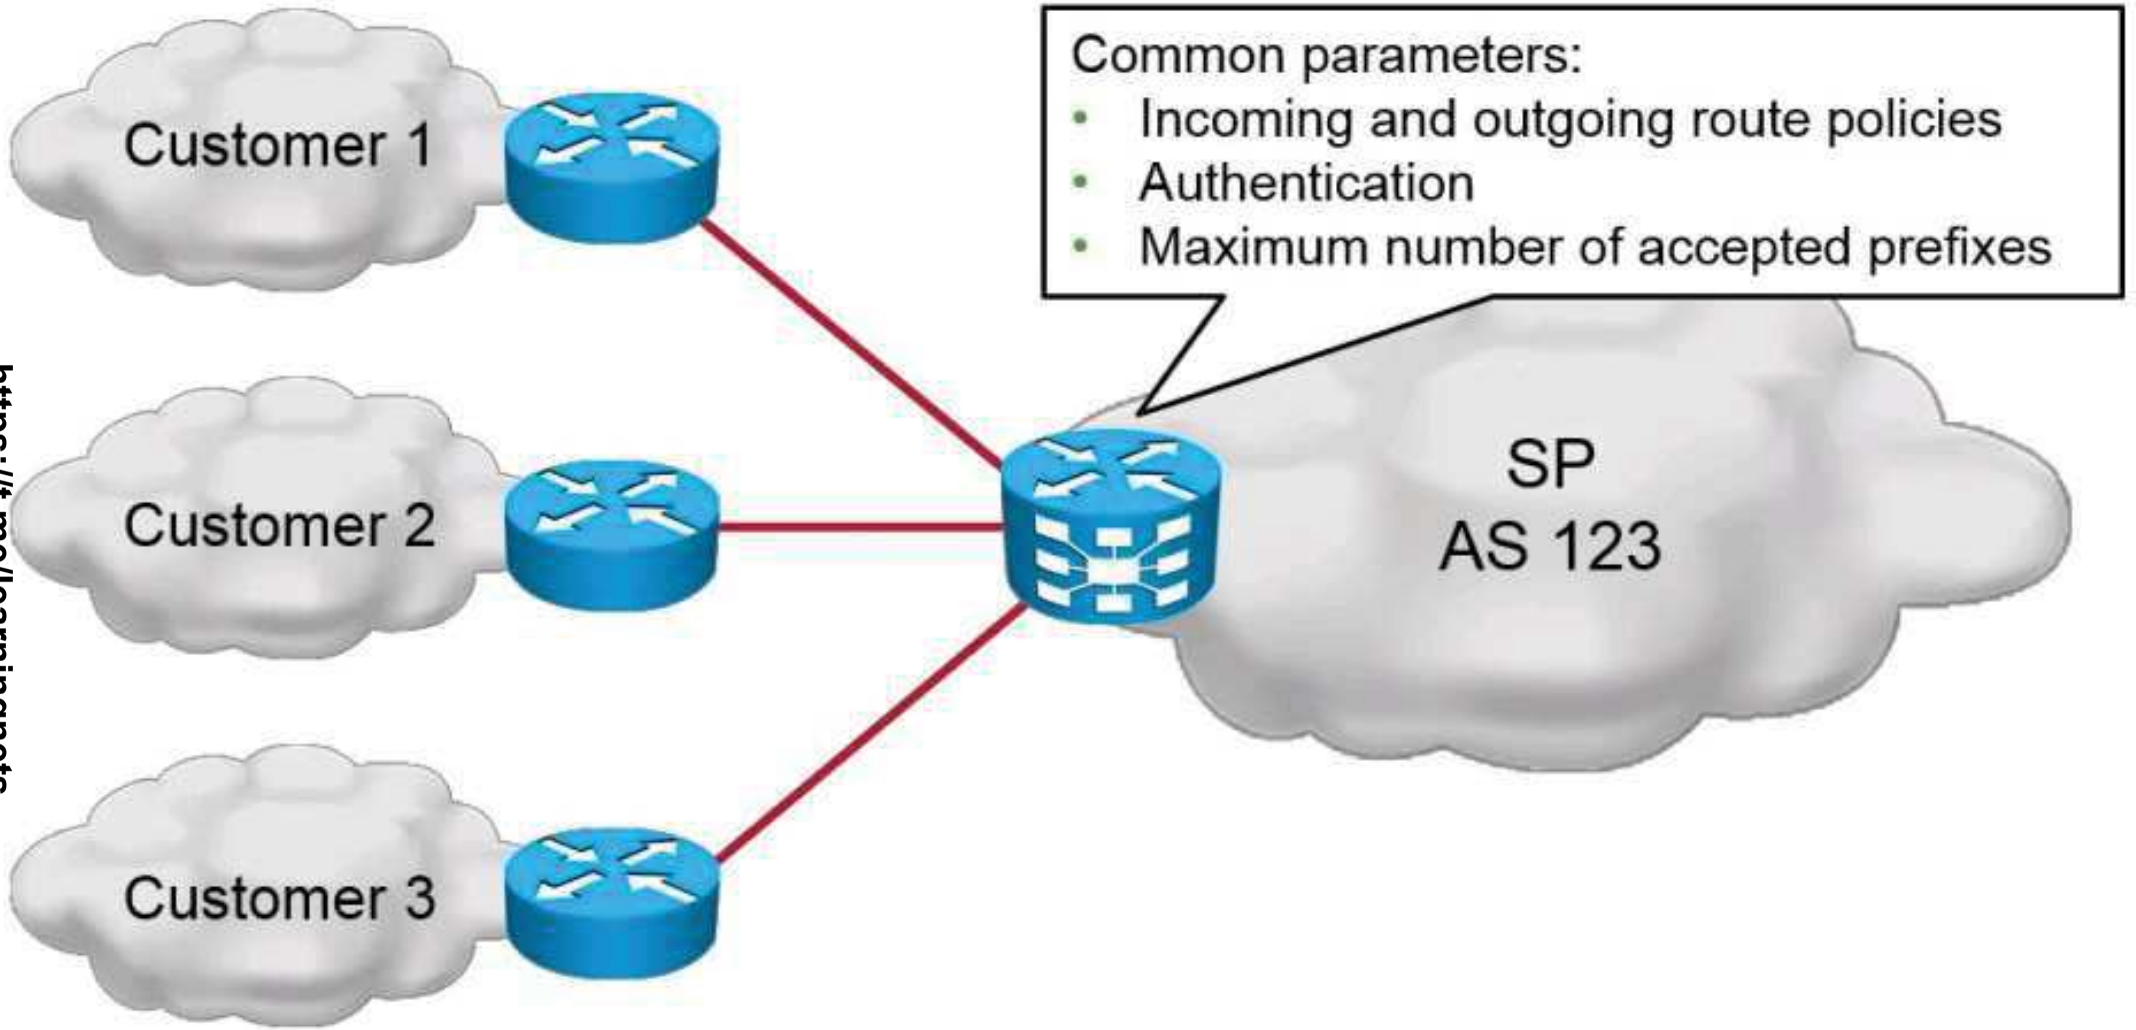

CE = Customer edge

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Example: BGP Route Reflector

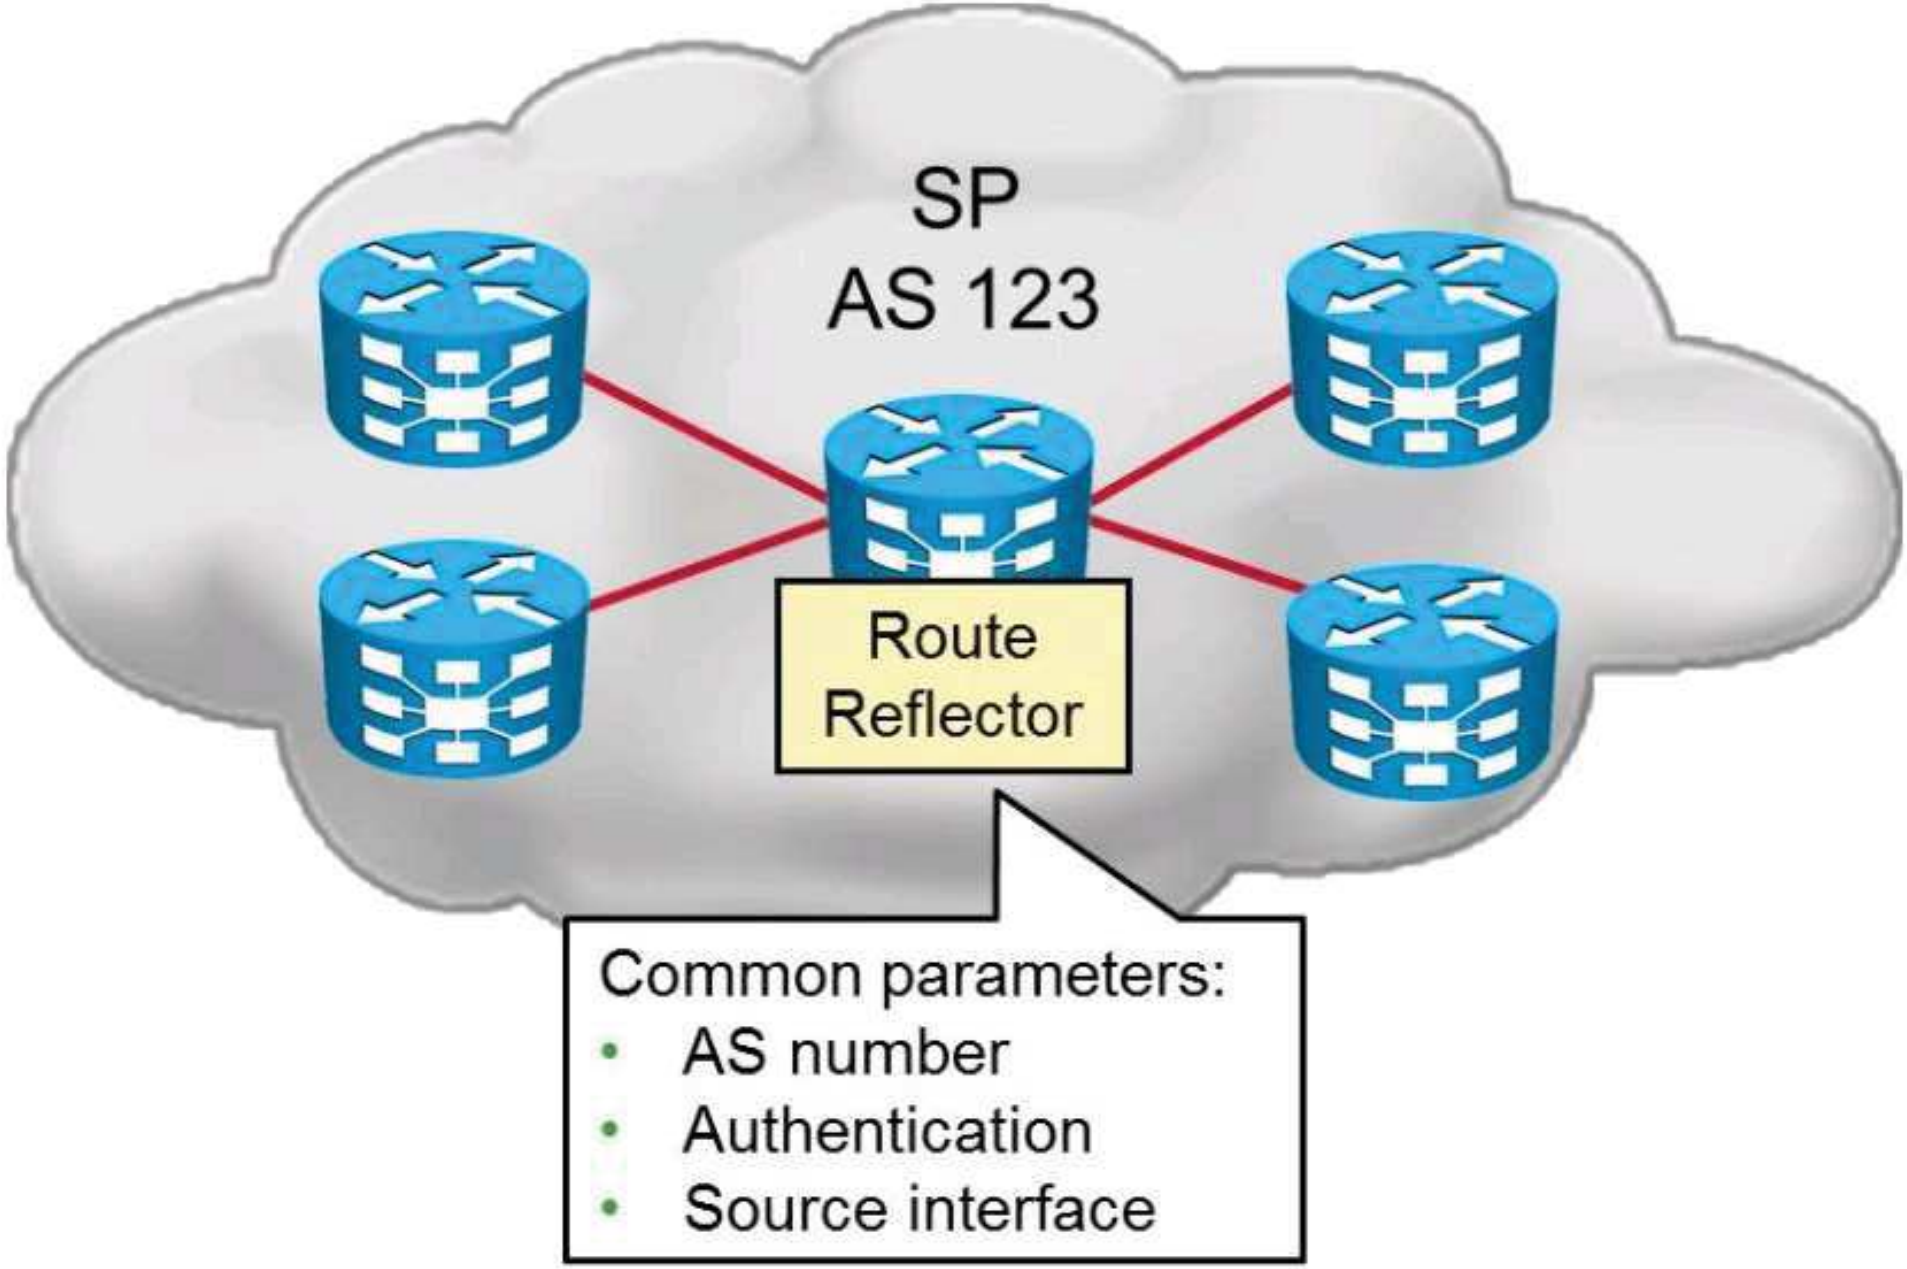

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Groups as Performance Tool

These are the BGP peer groups characteristics:

- Combine common configuration into a peer group.
- Neighbors are configured by assigning them to the peer group.
- A single BGP update is built for all members of a BGP peer group:
  - The CPU load does not increase linearly with the increased number of neighbors.
  - Use peer groups wherever possible to reduce the CPU load of the BGP process.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Groups Configuration

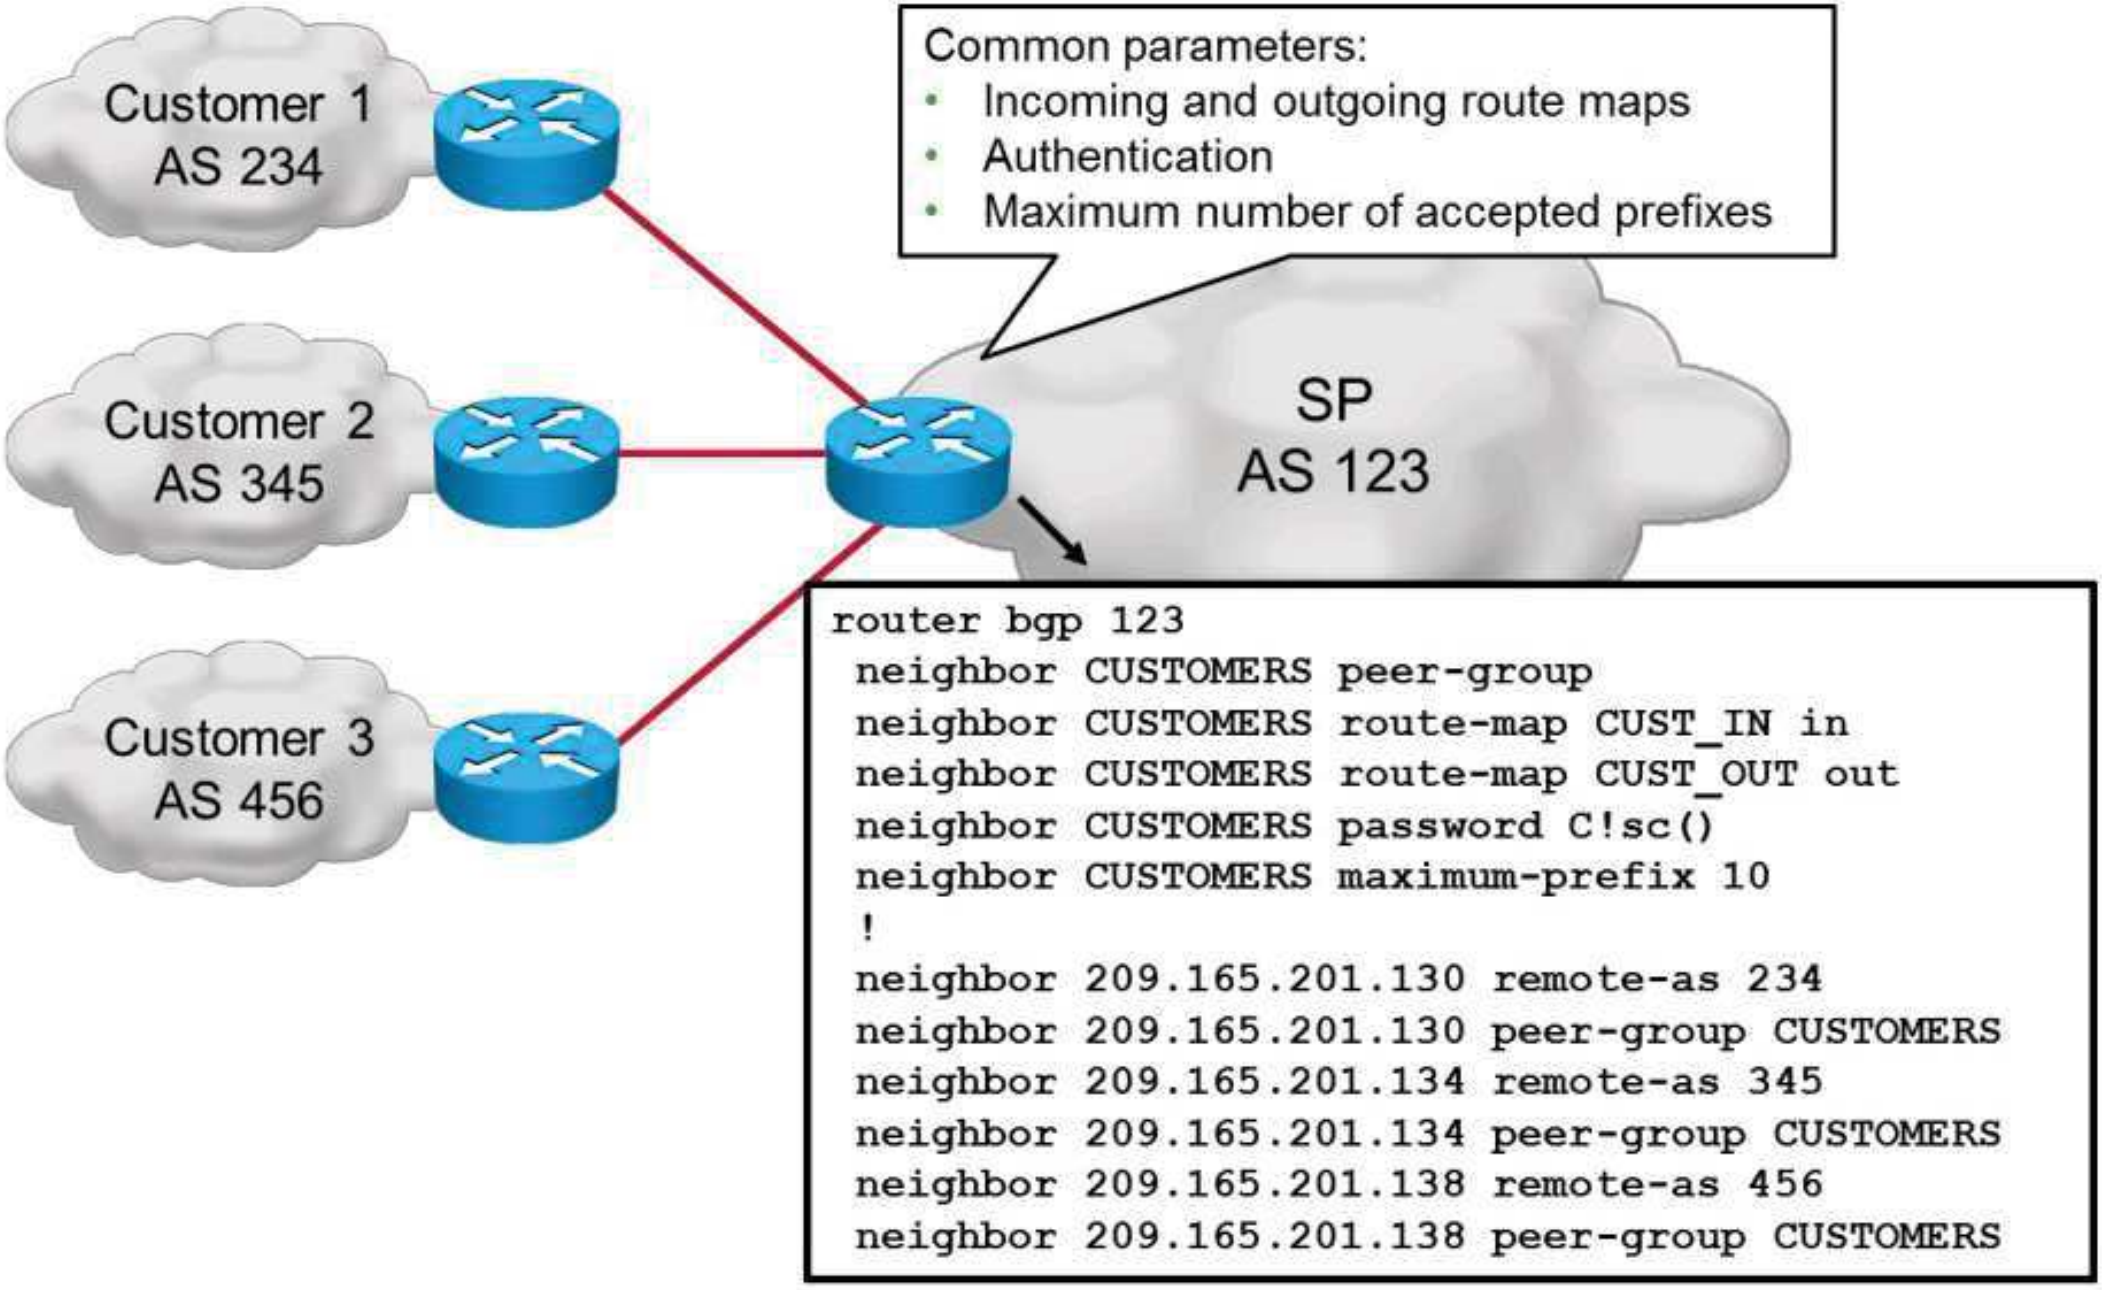

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Groups Limitation

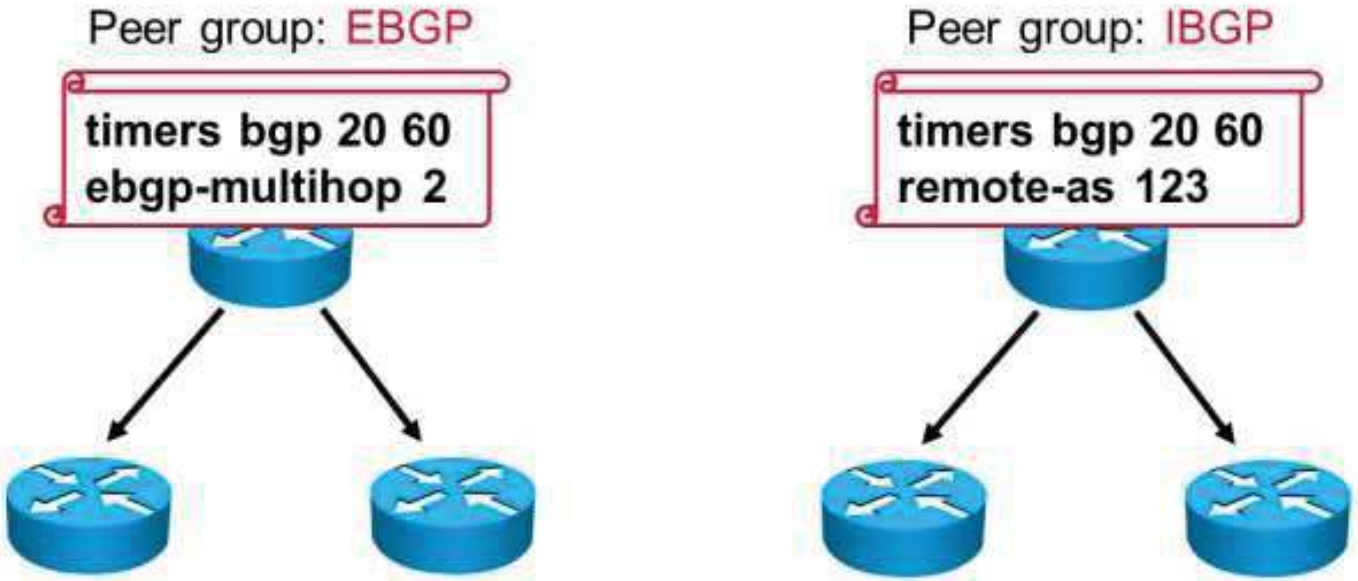

These are the BGP peer groups limitation characteristics:

- Peer groups were intended to be used only for CPU optimization.
- Awkward with similar but not identical configuration policies.
- IBGP and EBGP neighbors cannot be mixed in a peer group.
- Per-neighbor BGP parameters that affect outbound updates cannot be changed for peer group members.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Dynamic Update Groups

These are the BGP dynamic update groups characteristics:

- Separates BGP update generation from neighbor configuration.
- Introduces a new algorithm that dynamically calculates and optimizes update groups of neighbors that share the same outbound policies.
- Requires no configuration.
- Optimal BGP update message generation occurs automatically and independently.
- Available on Cisco IOS, IOS XR, and IOS XE Software platforms.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Templates

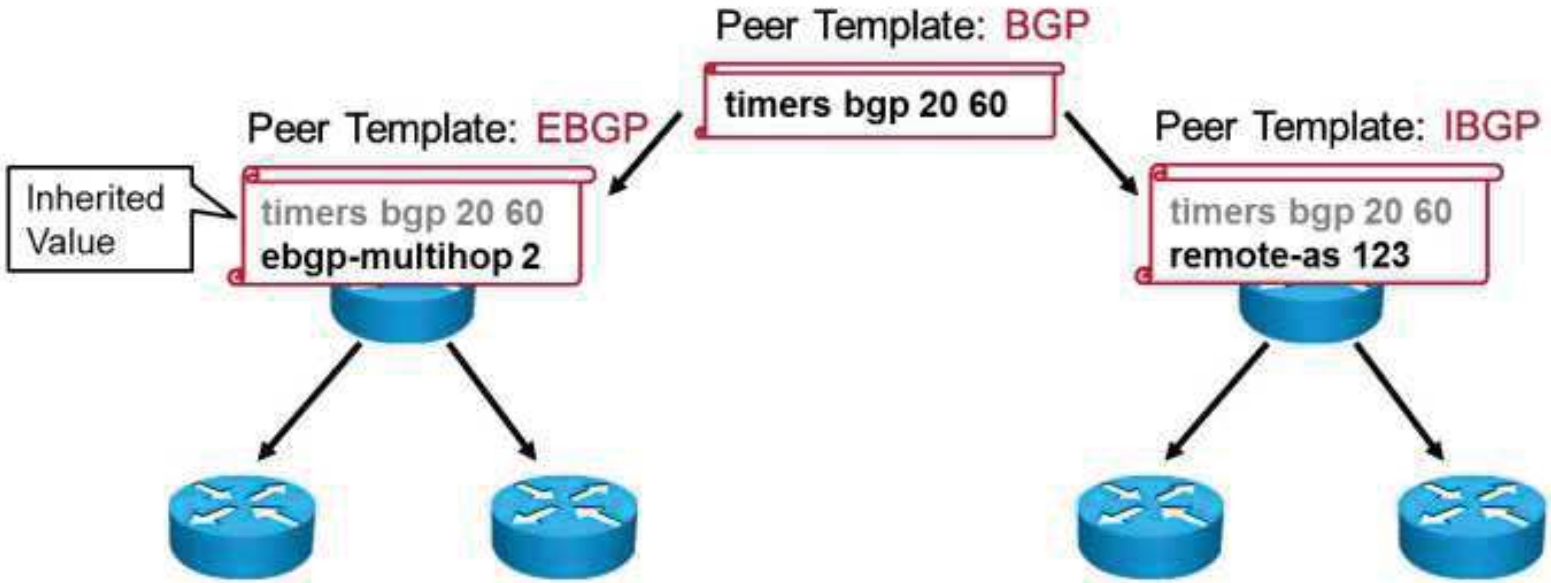

These are the BGP peer templates characteristics:

- Peer templates contain configuration patterns that can be applied to neighbors that share common policies.
- They are reusable and support inheritance, which allows you to group and apply distinct neighbor configurations for BGP neighbors.
- You can define complex configuration patterns through the use of inheritance.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configuration Templates

These are the BGP configuration templates characteristics:

- Available on Cisco IOS XR Software.
- Three types of configuration templates available:
  - Address family group: used to group address family-specific commands:
    - The same commands as in the neighbor address family configuration submode.
  - Session group: used to group address family-independent commands:
    - The same commands as in the neighbor configuration submode.
  - Neighbor group: used to group all commands:
    - All commands under neighbor and neighbor address family configuration submodes.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configuration Templates Inheritance

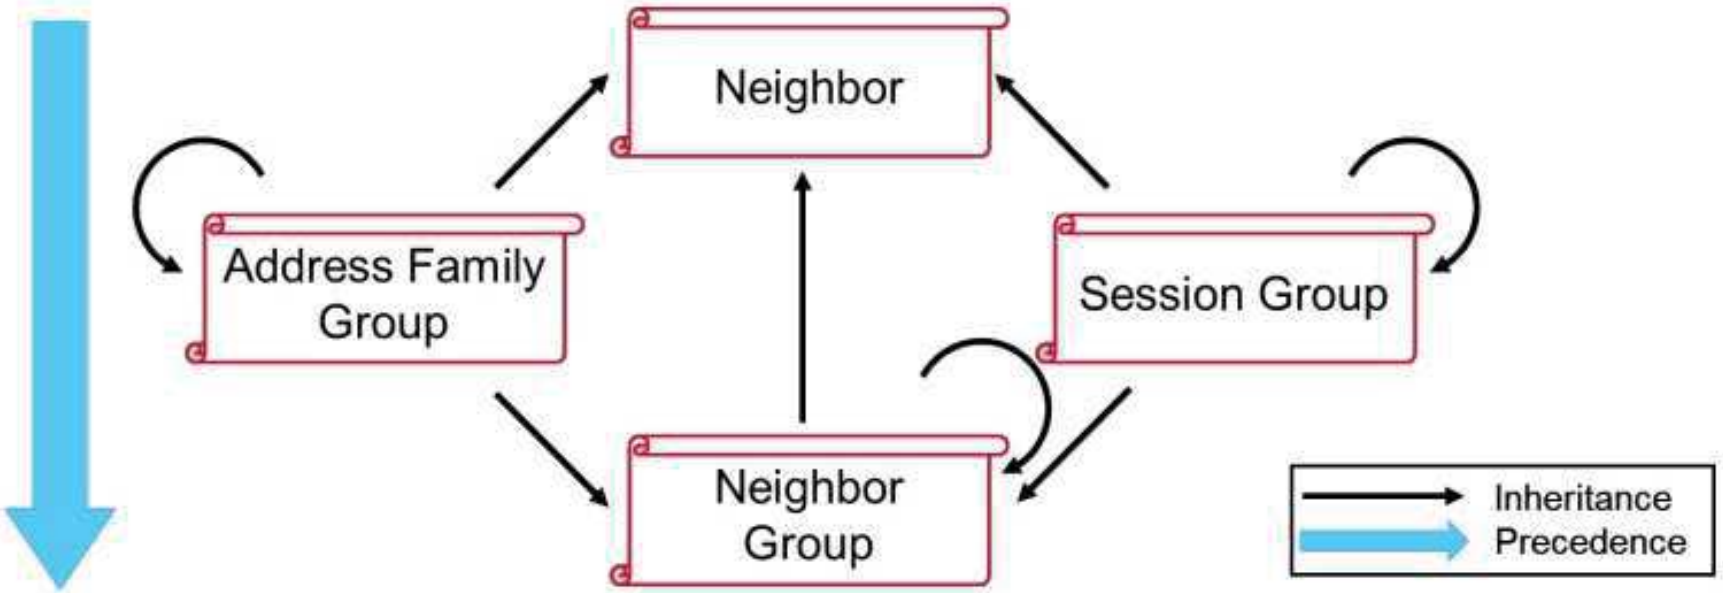

These are the BGP configuration templates inheritance characteristics:

- Address family groups can inherit from other address family groups.
- Session groups can inherit from other session groups.
- Neighbor groups can inherit from address family groups, session groups, and other neighbor groups.
- Neighbors can inherit from address family groups, session groups, and neighbor groups.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Configuring BGP Configuration Templates

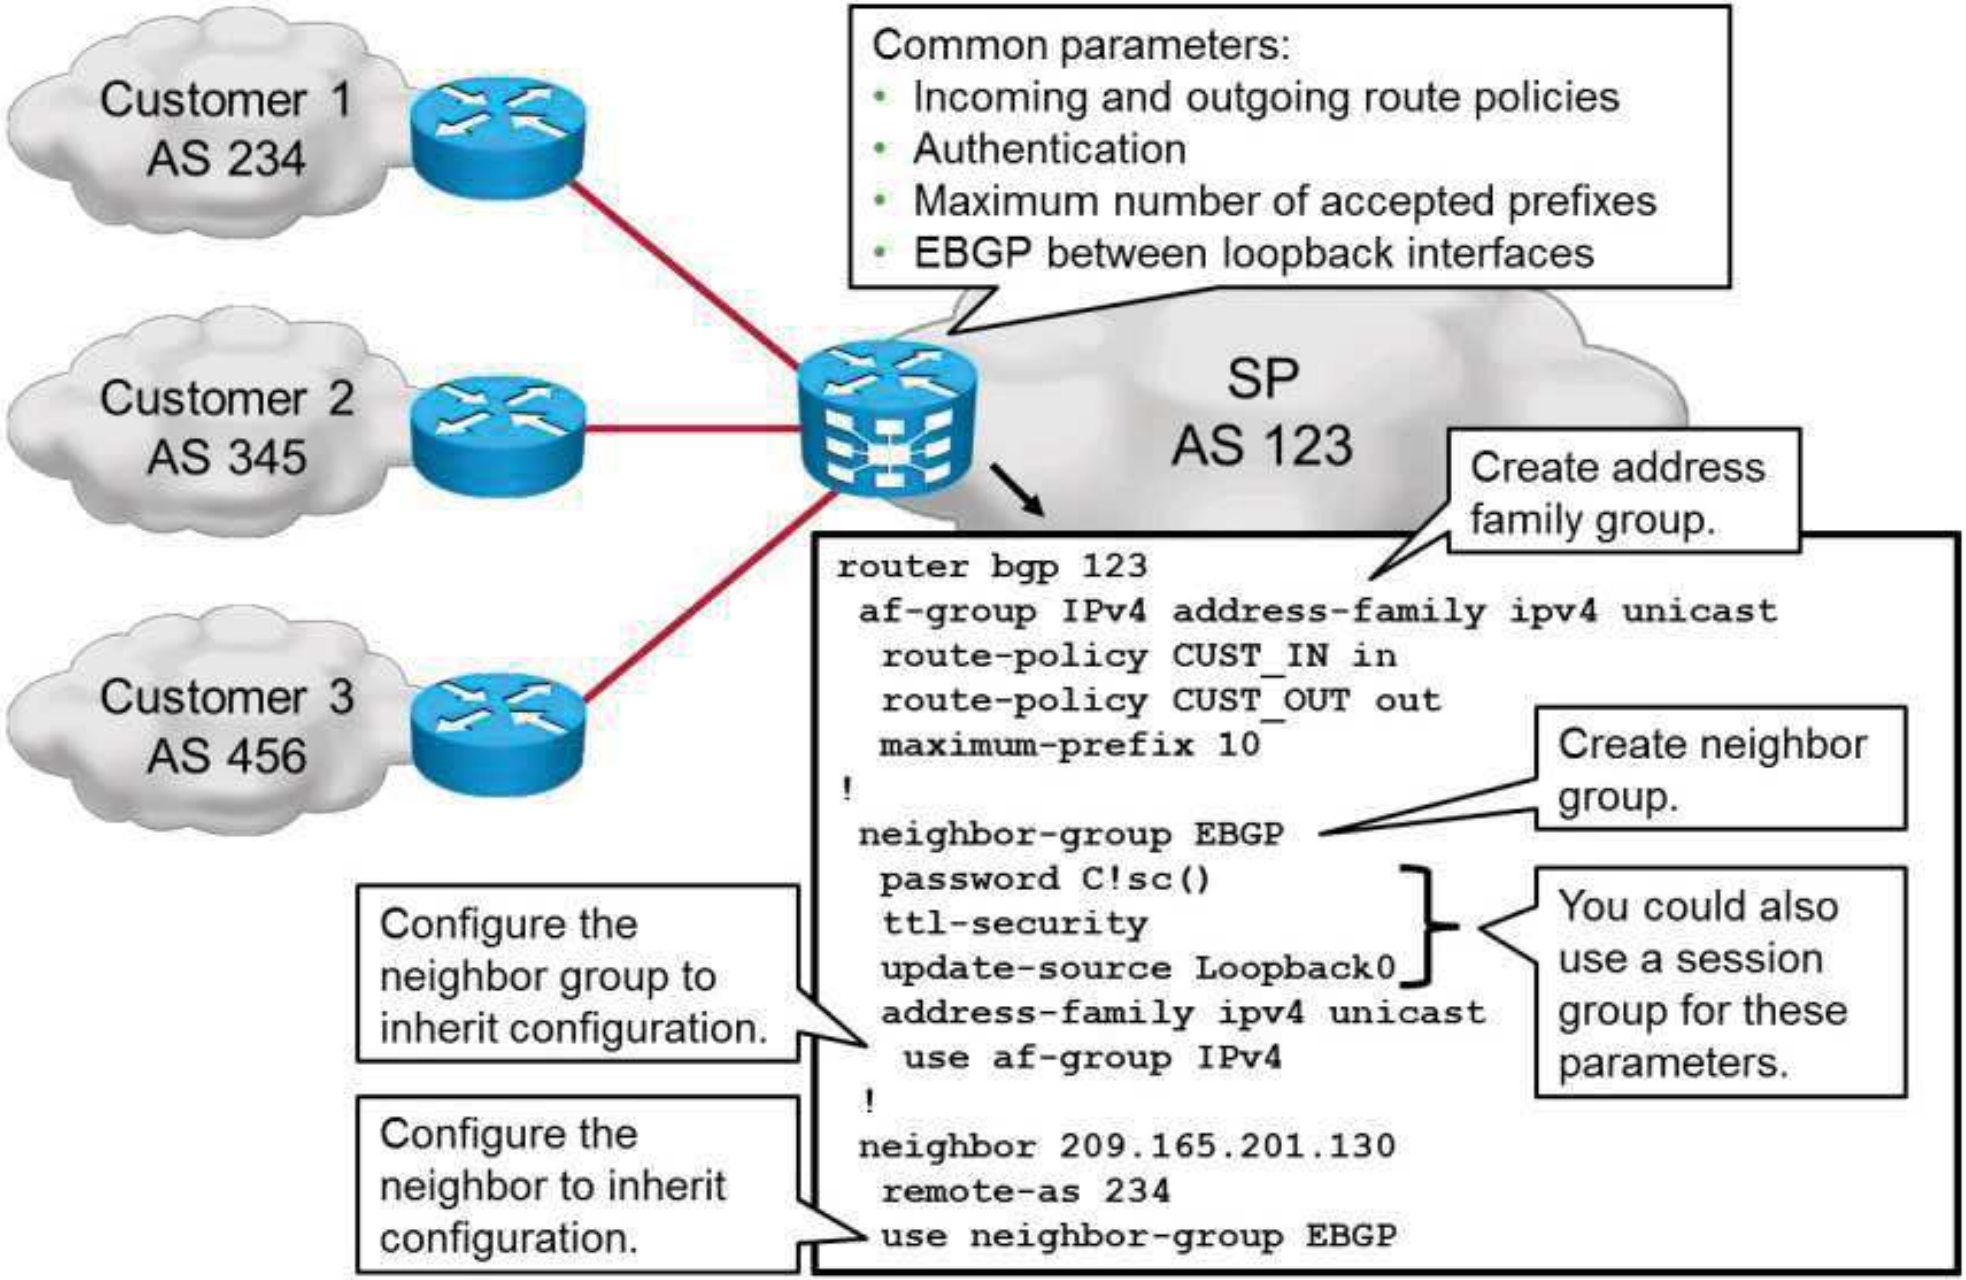

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configuration Templates Verification

```
RP/0/RSP0/CPU0:PE1# show bgp af-group IPv4 configuration
af-group IPv4 address-family IPv4 Unicast
  maximum-prefix 10 75                          []
  policy CUST_IN in                             []
  policy CUST_OUT out                           []
```

- Displays BGP address family group configuration.

```
RP/0/RSP0/CPU0:PE5# show bgp af-group IPv4 users
IPv4 Unicast: 209.165.201.130 n:EBGP
```

- Displays the neighbors, neighbor groups, and address family groups that inherit configuration from this address family group.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configuration Templates Verification (Cont.)

```
RP/0/RSP0/CPU0:PE1# show bgp neighbor-group EBGP configuration
neighbor-group EBGP
 password encrypted 143453180F4C63 []
 update-source Loopback0           []
 ttl-security                      []
 address-family IPv4 Unicast       []
  maximum-prefix 10 75             [a:IPv4]
  policy CUST_IN in                [a:IPv4]
  policy CUST_OUT out              [a:IPv4]
```

- Displays BGP neighbor group configuration.

```
RP/0/RSP0/CPU0:PE1# show bgp neighbor-group EBGP users
Session:       209.165.201.130
IPv4 Unicast: 209.165.201.130
```

- Displays the neighbors and neighbor groups that inherit configuration from this neighbor group.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configuration Templates Verification (Cont.)

```
RP/0/RSP0/CPU0:PE1# show bgp neighbors 209.165.201.130 configuration
neighbor 209.165.201.130
 remote-as 234                          []
 password encrypted 143453180F4C63      [n:EBGP]
 update-source Loopback0                [n:EBGP]
 ttl-security                           [n:EBGP]
 address-family IPv4 Unicast            [n:EBGP]
  maximum-prefix 10 75                  [n:EBGP a:IPv4]
  policy CUST_IN in                     [n:EBGP a:IPv4]
  policy CUST_OUT out                   [n:EBGP a:IPv4]
```

- Displays the effective configuration for the neighbor, including any settings that have been inherited from groups.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Templates

Two types of peer templates are available on Cisco IOS and IOS XE Software:

- **Peer session template:** Used to group and apply the configuration of general session commands that are common to all address families.
- **Peer policy template:** Used to group and apply the configuration of commands that are applied within specific address-family configuration modes.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Templates Inheritance

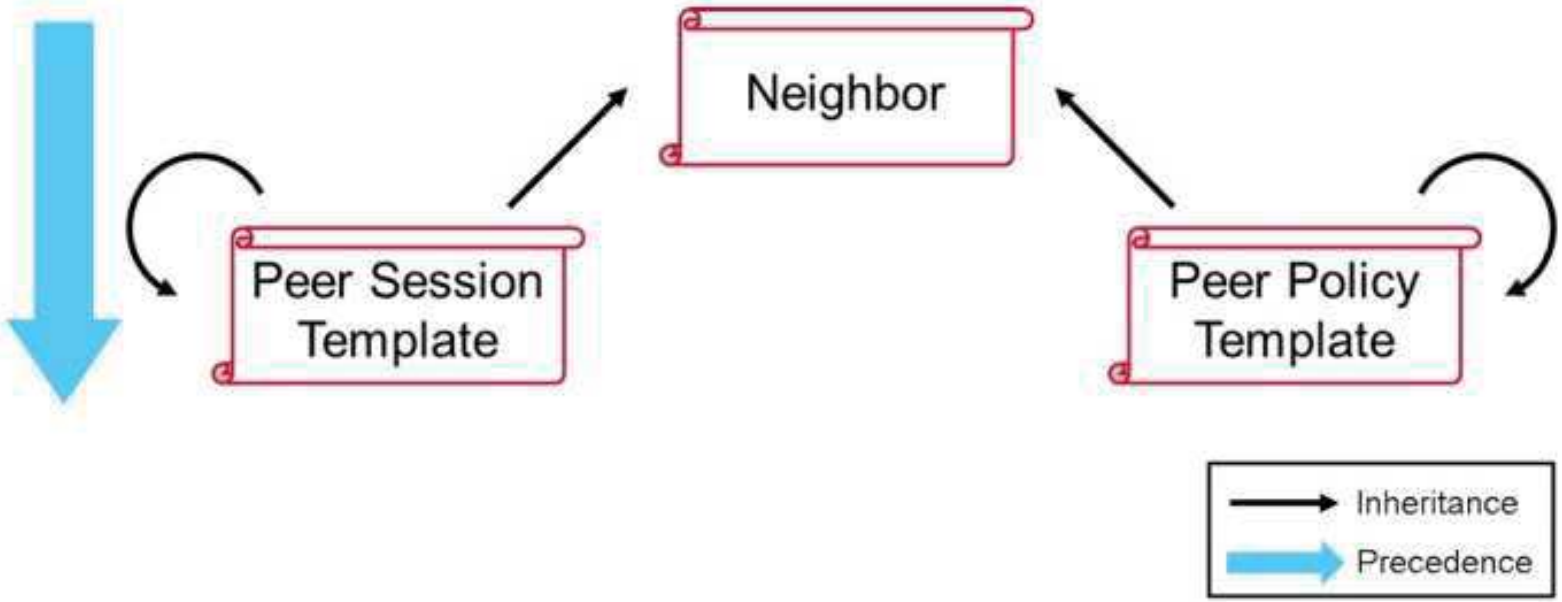

These are the BGP peer templates inheritance characteristics:

- A session template can inherit configuration from another session template.
- A policy template can inherit configuration from another policy template.
- Neighbors can inherit from a session and a policy template.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Templates Configuration

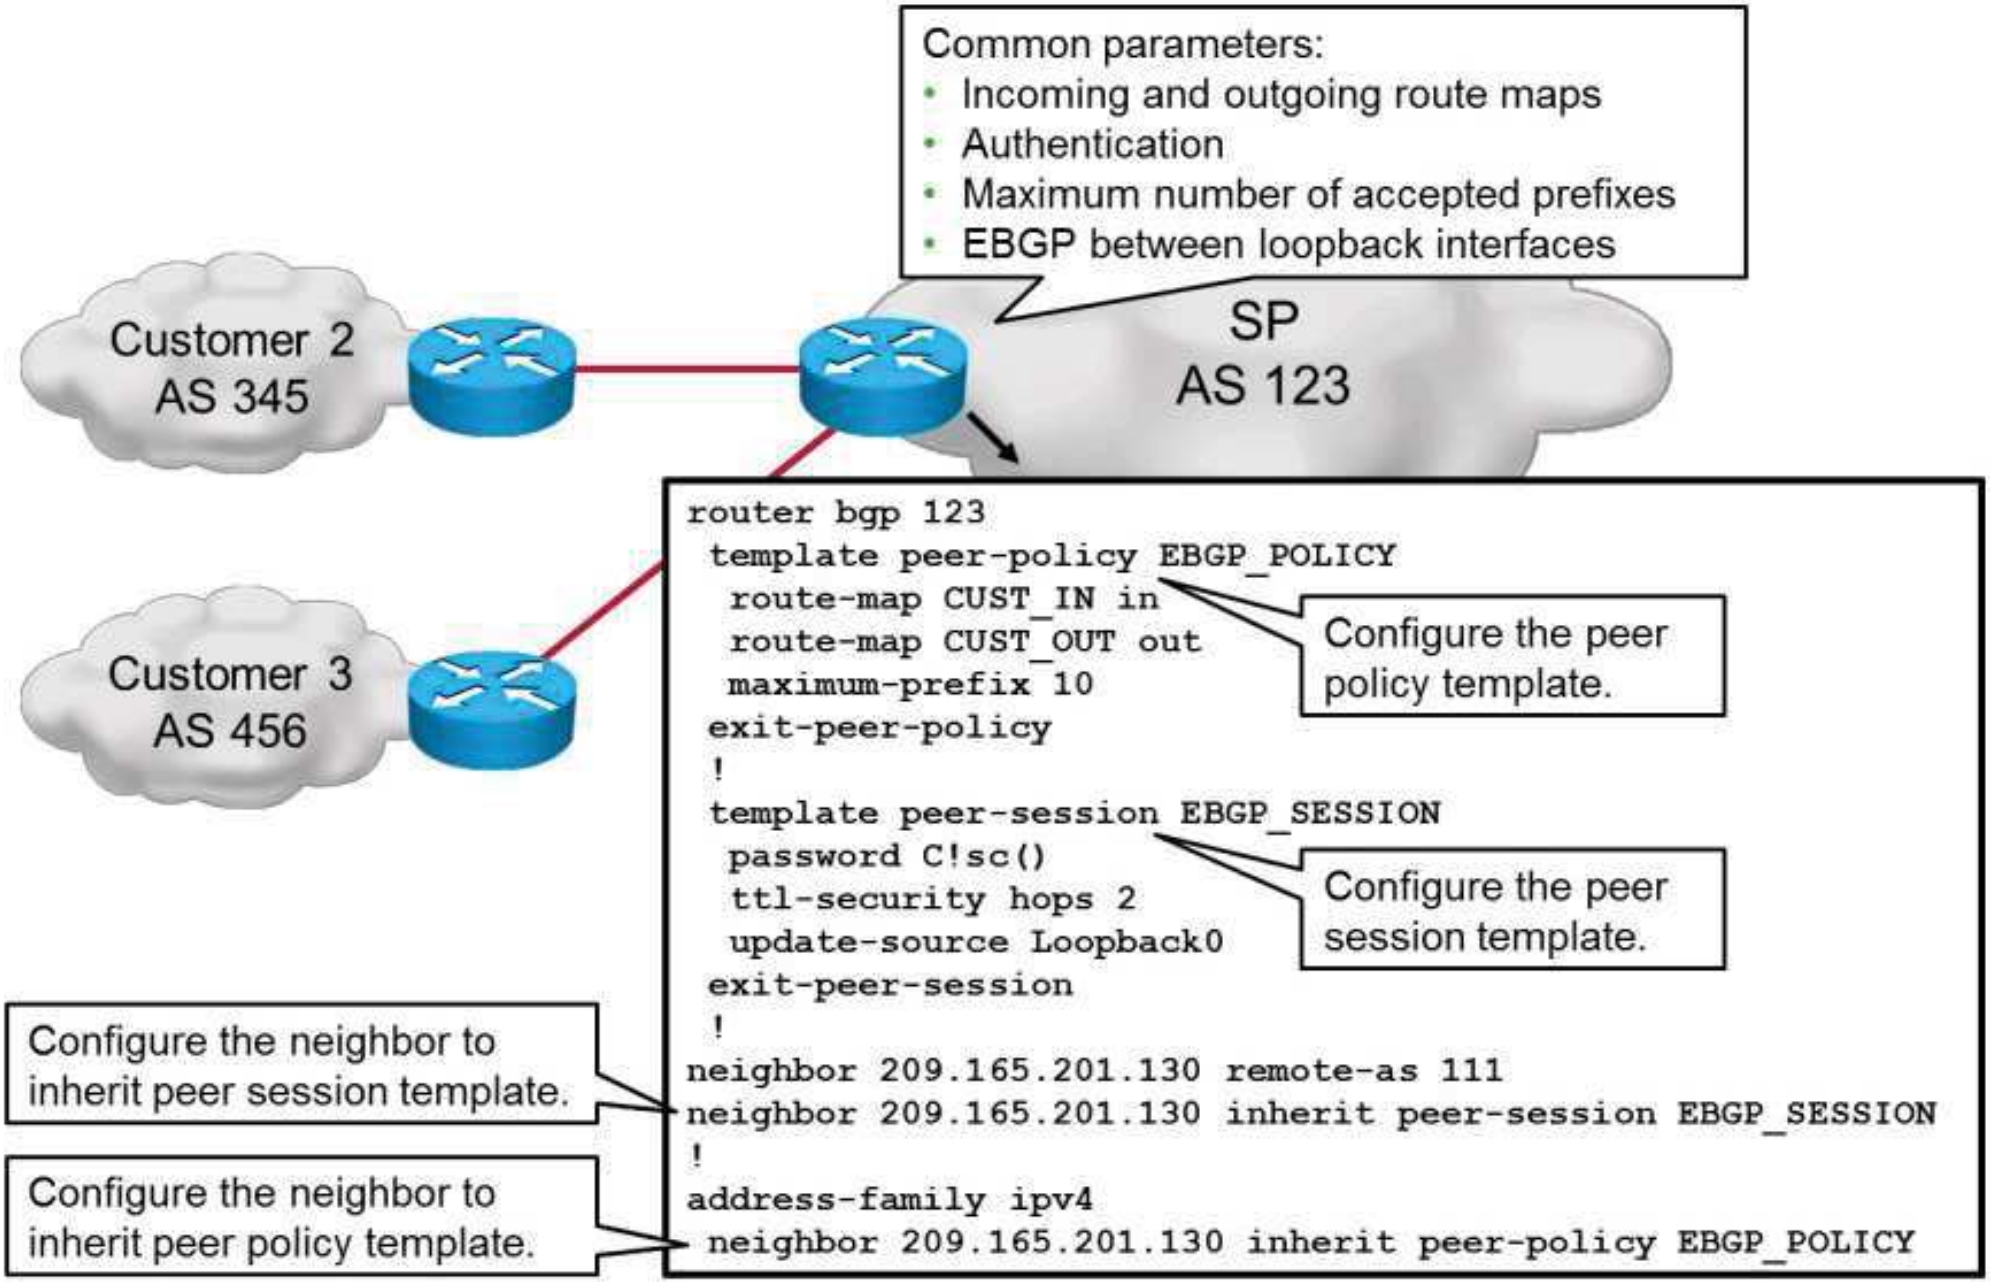

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Templates Verification

```
PE1# show ip bgp template peer-session
Template:EBGP_SESSION, index:1
Local policies:0x890, Inherited polices:0x0
Locally configured session commands:
 password is configured
 ttl-security hops 2
 update-source Loopback0
Inherited session commands:
```

- Displays locally configured peer session templates.

```
PE1# show ip bgp template peer-policy
Template:EBGP_POLICY, index:1.
Local policies:0x80003, Inherited polices:0x0
Local disable policies:0x0, Inherited disable policies:0x0
Locally configured policies:
    route-map CUST_IN in
    route-map CUST_OUT out
    maximum-prefix 10
Inherited policies:
```

- Displays locally configured peer policy templates.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Peer Templates Verification (Cont.)

```
PE1# show ip bgp neighbors 209.165.201.130 policy
  Neighbor: 209.165.201.130, Address-Family: IPv4 Unicast
  Inherited polices:
    route-map CUST_IN in
    route-map CUST_OUT out
    maximum-prefix 10
```

- Displays the policies applied to a neighbor per address family.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary

- BGP peer groups (Cisco IOS and IOS XE Software) were designed primarily for CPU optimization.
- BGP peer groups (Cisco IOS and IOS XE Software) can also be used for configuration optimization.
- To use BGP peer groups, use the **neighbor** *name* **peer-group** command to create a peer group. Then, assign specific BGP parameters to the peer group.
- BGP peer groups are used on a router with several neighbors that have similar but not completely identical policies.
- BGP dynamic update groups separate BGP update generation from neighbor configuration.
- BGP configuration (Cisco IOS XR Software) and peer templates (Cisco IOS and IOS XE Software) are reusable configuration patterns that support inheritance.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- IOS XR supports three types of BGP templates:
  - AF groups: contain address family-dependent information
  - session groups: contain address family-independent information
  - neighbor groups: universal
- BGP configuration template inheritance occurs between different template groups.
- To configure the neighbor to inherit configuration from the neighbor group, enter neighbor configuration mode, and use the **use neighbor-group** command, followed by the neighbor group name.
- To display the neighbors and neighbor groups that inherit configuration from a neighbor group, use the **show bgp neighbor-group users** command.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Cisco IOS and IOS XE Software support two types of BGP templates:
  - Peer session templates: contain configuration common to all address families
  - Peer policy templates: contain configuration applied within a specific address family
- BGP peer session templates support direct and indirect inheritance.
- To create a session BGP peer session template, use the **template peer-session** command, followed by the session template name.
- To display BGP peer policy template configurations, use the **show ip bgp template peer-session** command.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

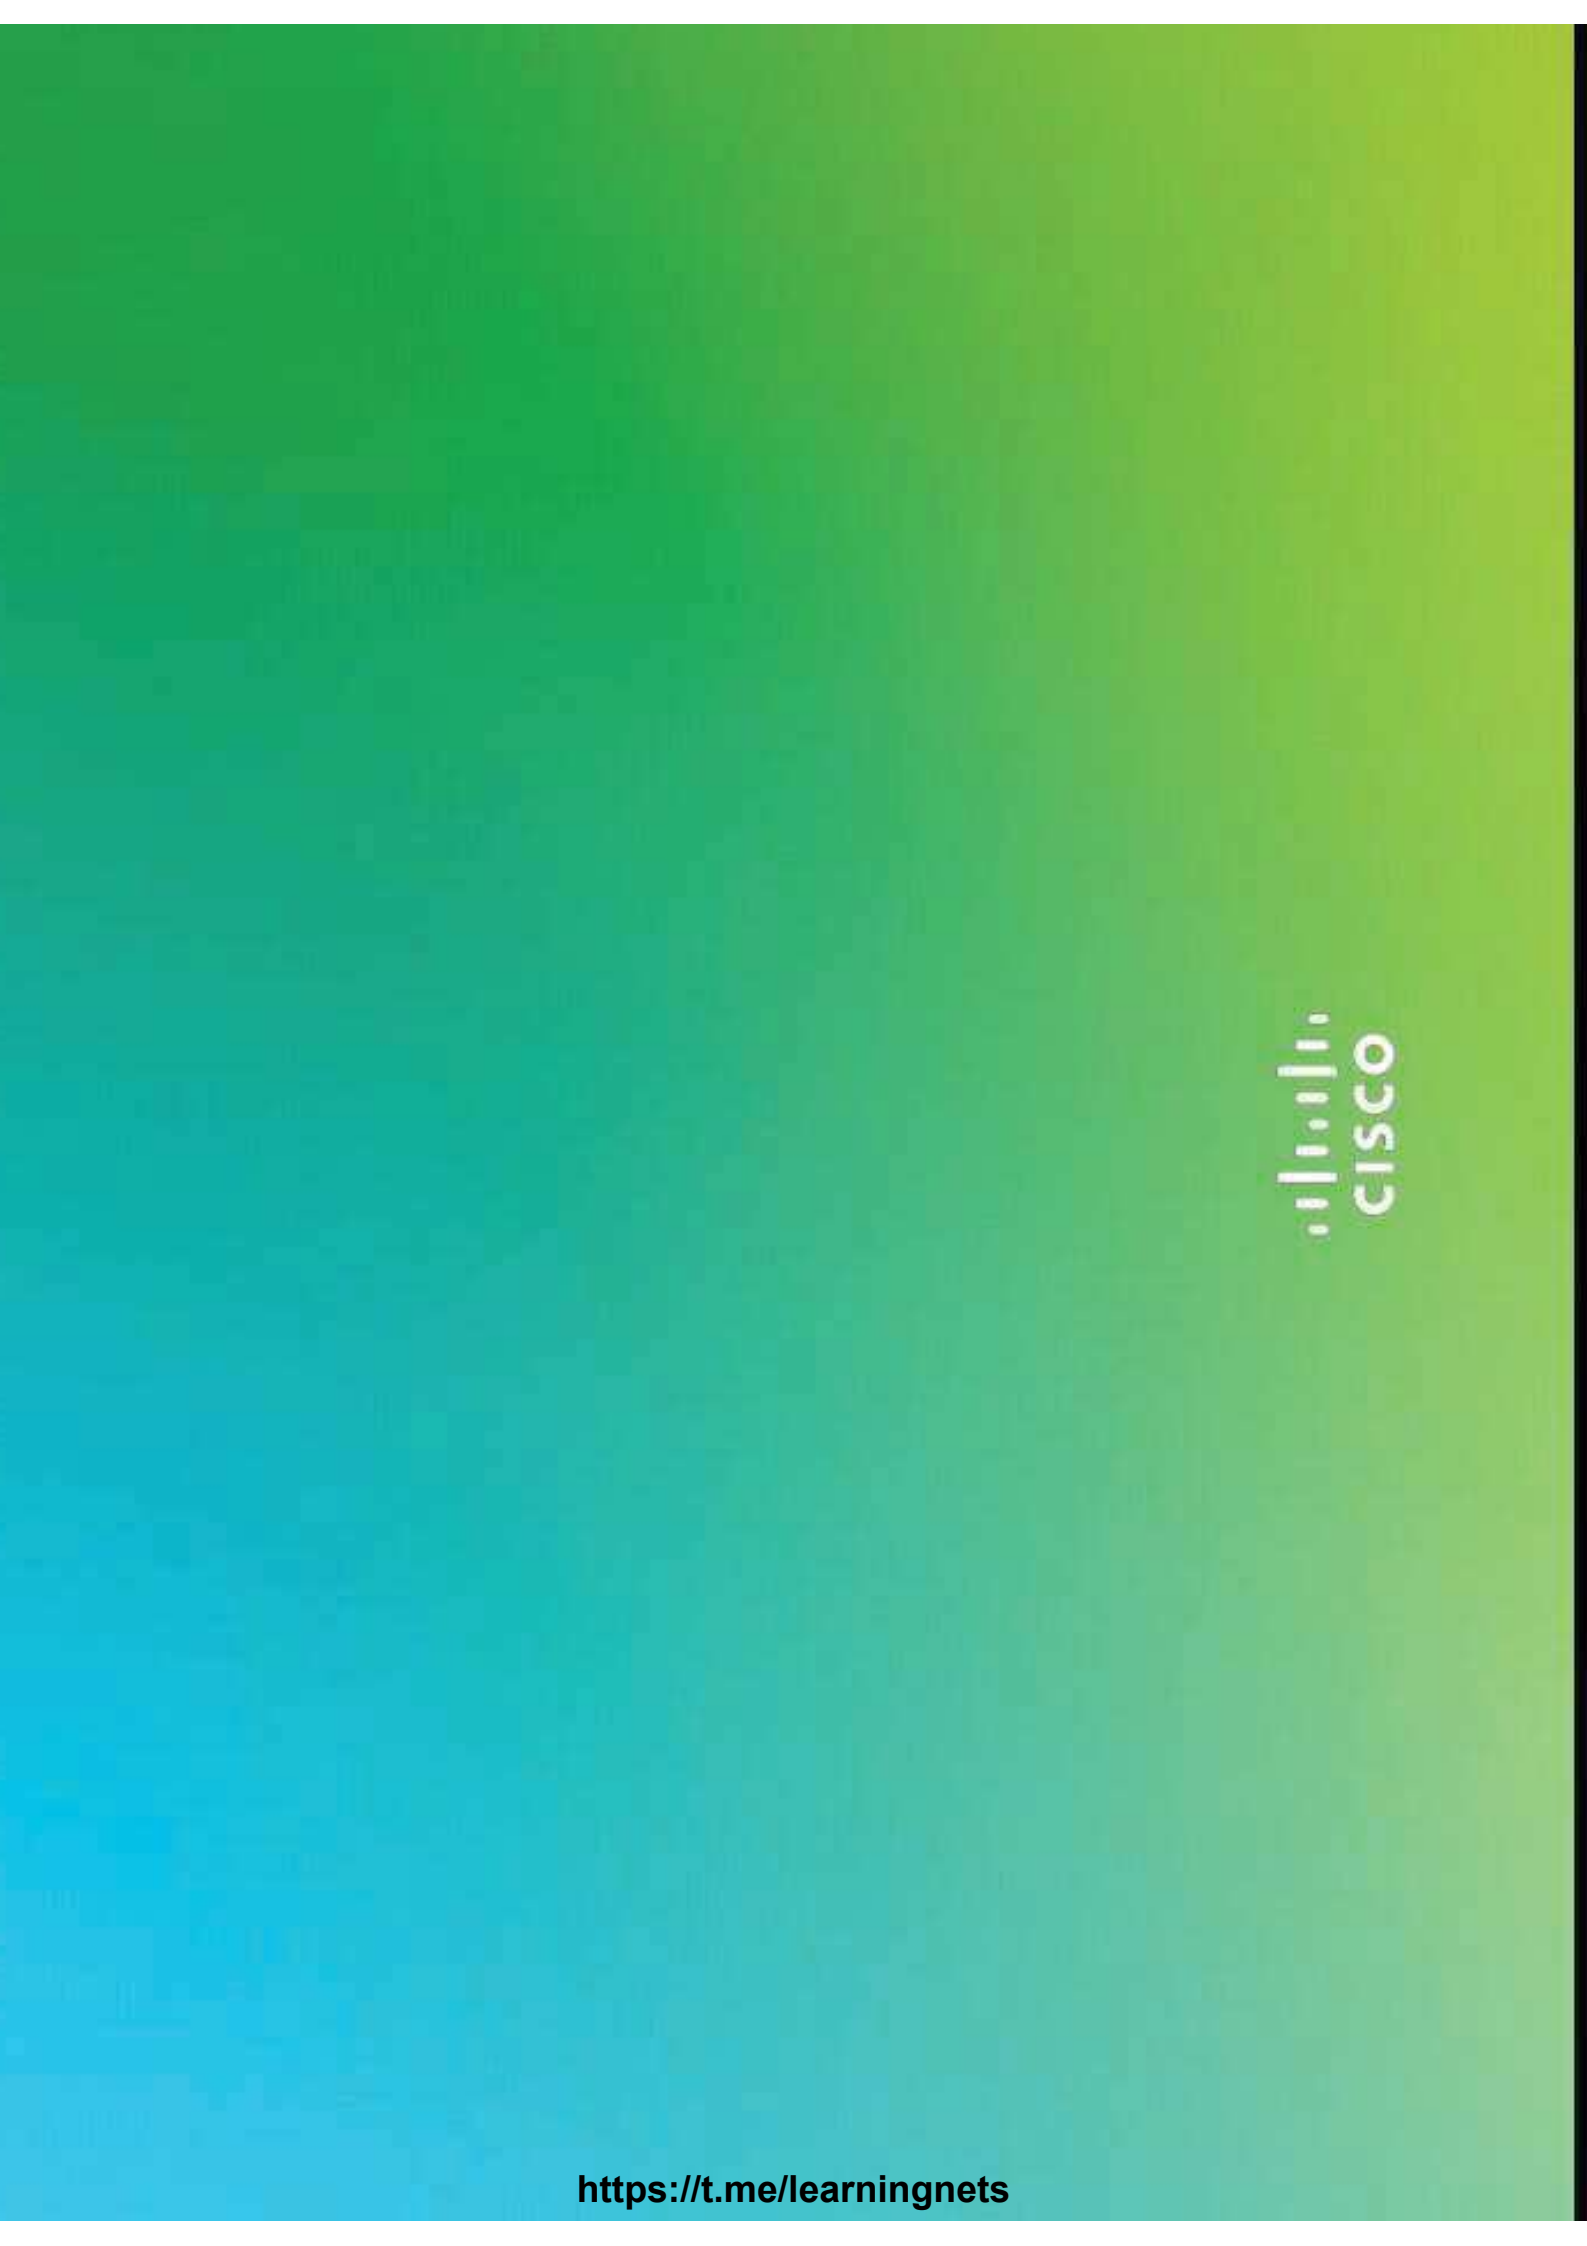

https://t.me/learningnets

Module Summary

- Service providers can implement various BGP security options to prevent attacks to service providers and customer networks.
- Several features are available to improve BGP convergence and reduce CPU utilization: distributed BGP, BGP peer groups, the PMTU feature, and interface input queues.
- The features that are available to improve BGP scalability are the maximum prefix feature, dynamic update groups, configuration and peer templates, and BGP route dampening.

https://t.me/learningnets

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

https://t.me/learningnets