Introducing BGP Routing

Implement BGP in the Service Provider Network

© 2012 Cisco and/or its affiliates. All rights reserved.

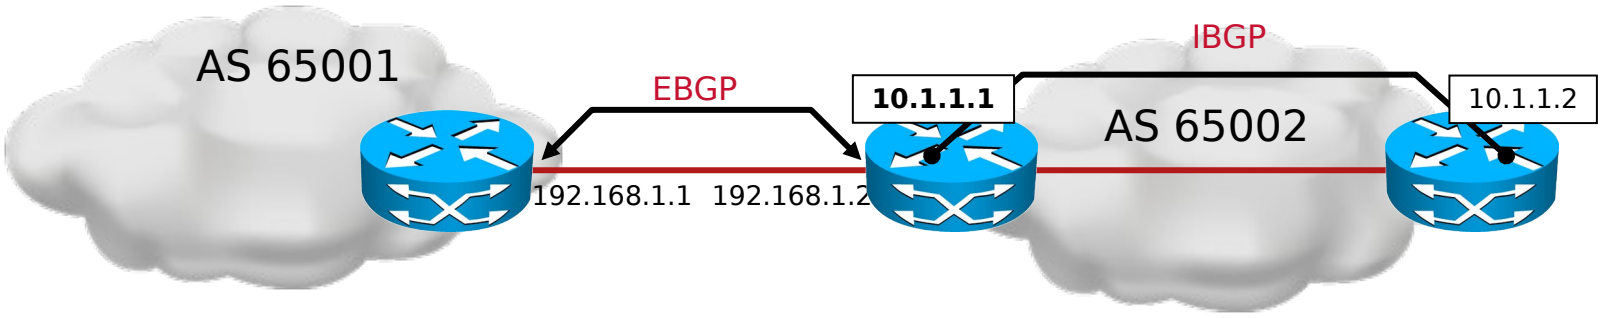
Objectives

- Describe BGP terminology
- Describe autonomous systems in BGP networks
- Describe BGP routing between autonomous systems
- Describe BGP path vectors
- Describe BGP routing policies
- Describe features of BGP
- Describe the tables used by a router to store BGP information
- Describe the four BGP message types
- Describe the multiprotocol extensions add to BGP to support IPv6
- Describe MPLS VPNs

© 2012 Cisco and/or its affiliates. All rights reserved.

BGP Terminology

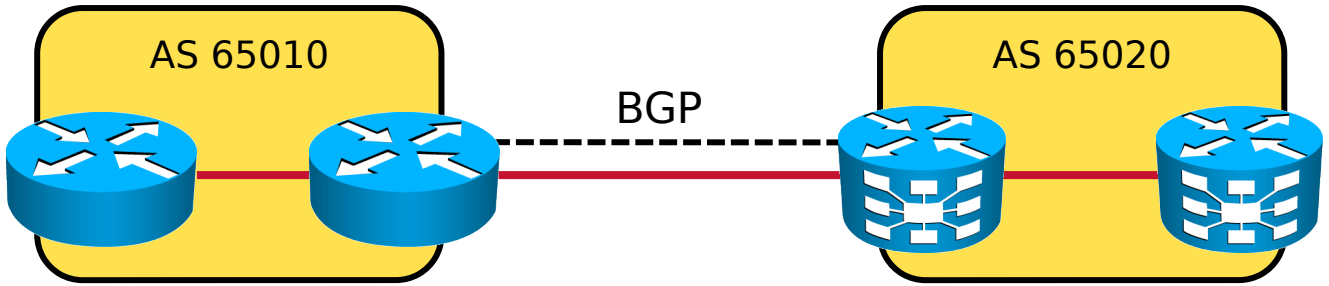
- **Autonomous system**: A collection of networks under a single administrative domain
- **Interdomain routing:** Routing between the customer and the service provider
- **Internal routing:** Uses IGP protocol (RIP, OSPF, IS-IS, and so on) to exchange routing information inside the autonomous system
- **External routing:** Uses EGP protocol (BGP) to exchange routes between autonomous systems
- Two BGP implementations:
  - **Internal BGP (IBGP):** When BGP is used inside an AS
  - **External BGP (EBGP):** When BGP is used between autonomous systems



© 2012 Cisco and/or its affiliates. All rights reserved.

Autonomous System

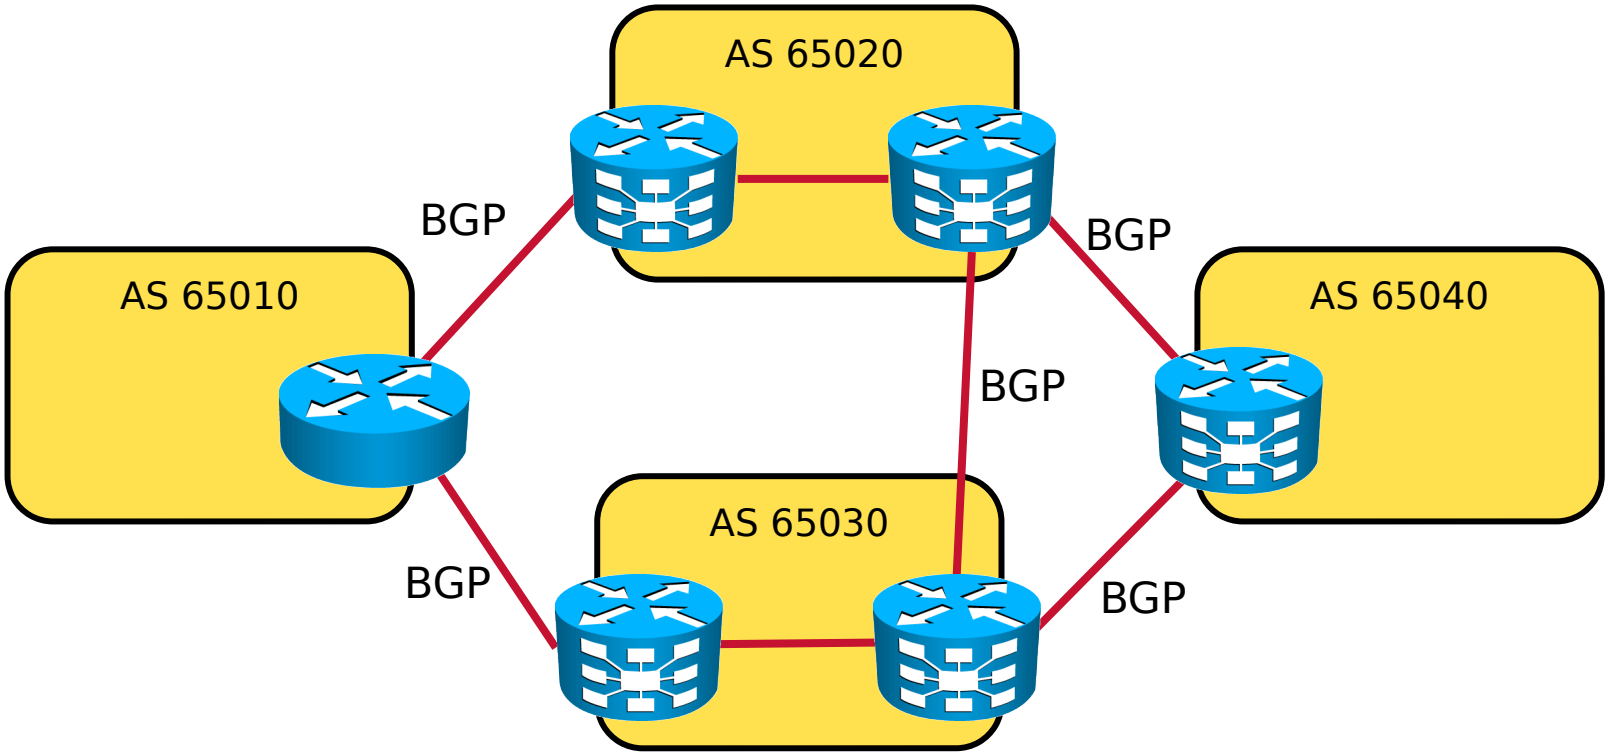
- An autonomous system (AS) is a collection of networks under a single technical administration.
  - 16-bit numbers (as of January 2009, 32-bit numbers are available)
  - Ranging from 1 to 65535
  - Private AS: 64512–65535
- The IANA allocates AS numbers.
- IGPs operate within an AS.
- BGP is used between autonomous systems.



© 2012 Cisco and/or its affiliates. All rights reserved.

BGP Routing Between Autonomous Systems

- BGP is used to provide an interdomain routing system.
- BGP guarantees the exchange of loop-free routing information.
- BGP works differently than IGPs.
  - BGP is a policy-based routing protocol.
  - BGP controls traffic flow using multiple BGP path attributes.



© 2012 Cisco and/or its affiliates. All rights reserved.

Path Vector Functionality

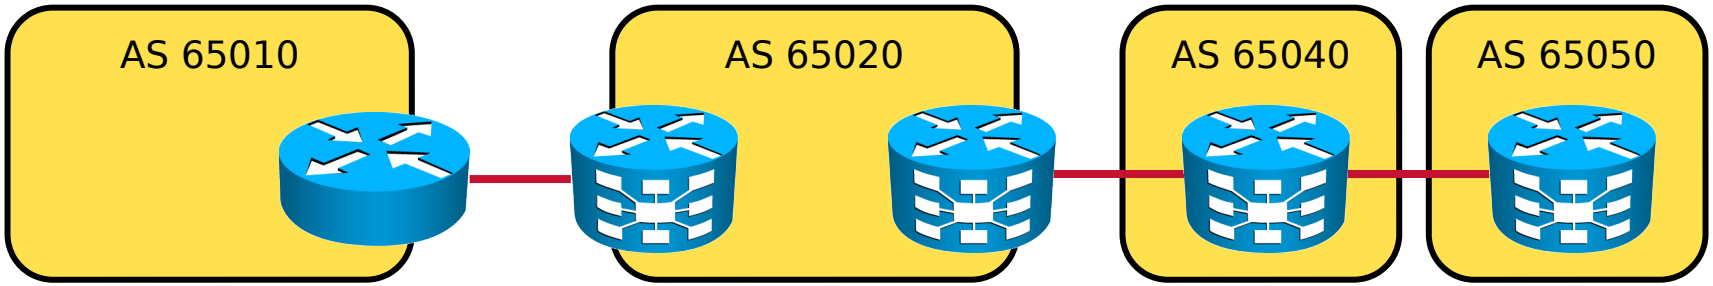
- BGP announces this information:
  - Paths (set of AS numbers)
  - Networks that are reachable at the end of the path
- The path is described by using attributes.
- The administrator can define data flow through autonomous systems



© 2012 Cisco and/or its affiliates. All rights reserved.

BGP Routing Policies

- BGP can support any policy conforming to the hop-by-hop (AS-by-AS) routing paradigm.

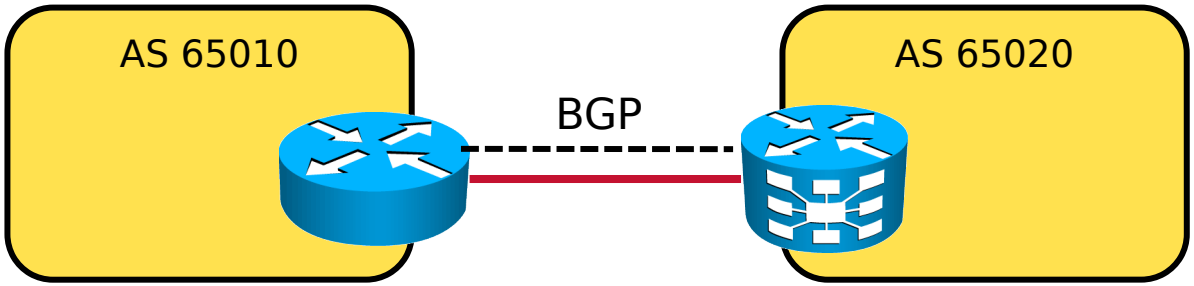


© 2012 Cisco and/or its affiliates. All rights reserved.

Features of BGP

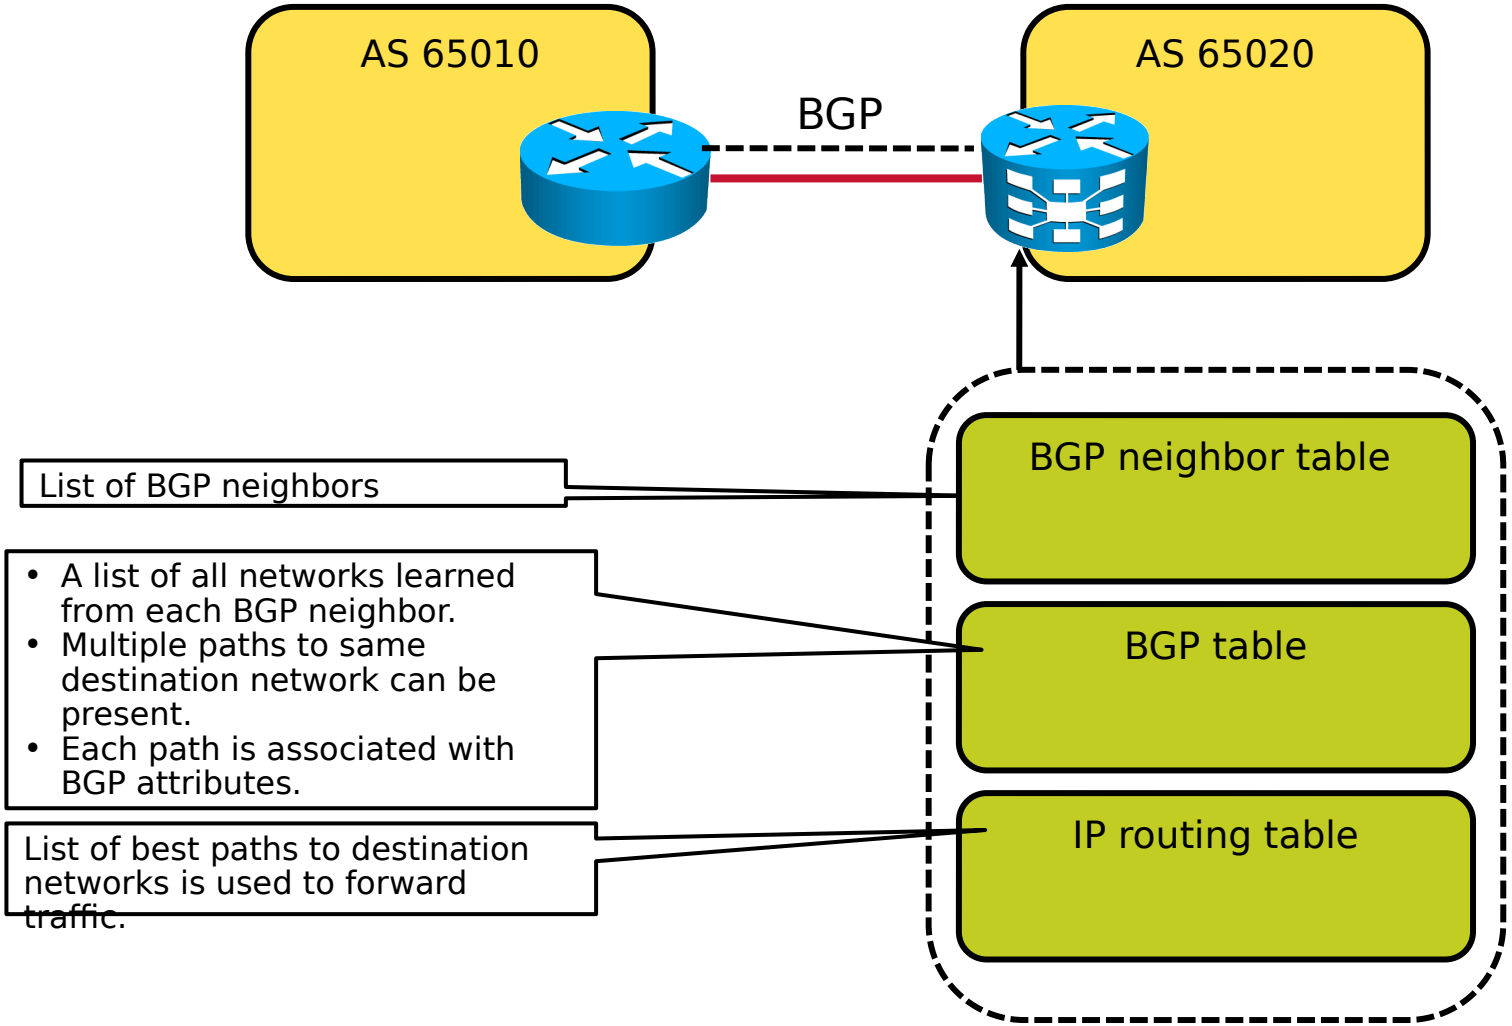
BGP is a path vector protocol with the following properties:

- Reliable updates: BGP runs on top of TCP (port 179)
- Incremental, triggered updates only
- Periodic keepalive messages to verify TCP connectivity
- Rich metrics (called path vectors or attributes)
- Designed to scale to huge internetworks



© 2012 Cisco and/or its affiliates. All rights reserved.

BGP Databases



© 2012 Cisco and/or its affiliates. All rights reserved.

BGP Message Types

BGP defines the following message types:

- Open
  - Includes hold time and BGP router ID
- Keepalive
- Update
  - Information for one path only (could be to multiple networks)
  - Includes path attributes and networks
- Notification
  - When an error is detected
  - BGP connection closed after message is sent

© 2012 Cisco and/or its affiliates. All rights reserved.

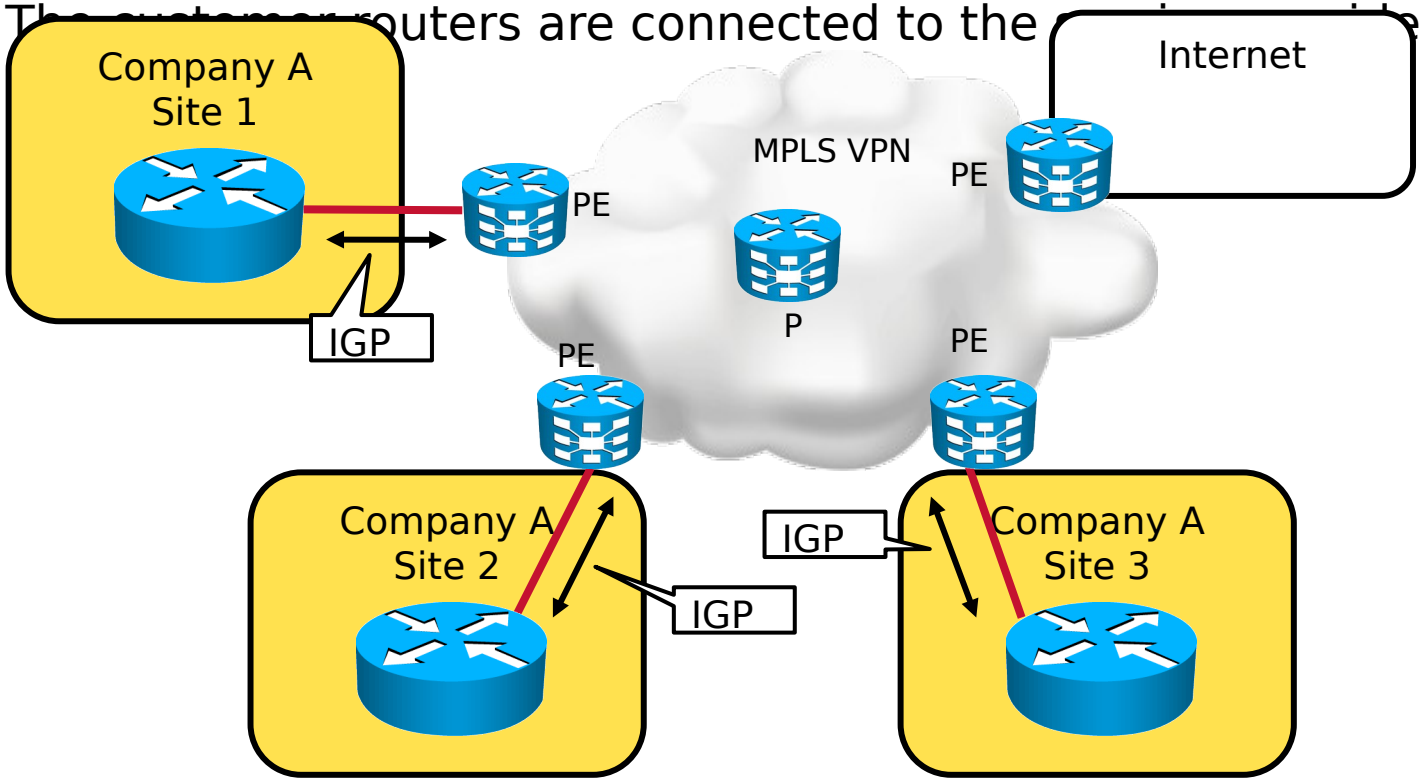
Multiprotocol Extensions for BGP4

- BGP originally designed for IPv4:
  - Carries IPv4 prefix reachability information
  - Uses IPv4 for transport
- Multiprotocol extensions for BGP4:
  - Enables other protocols besides IPv4
  - New identifier for the address family
  - Most often used in MPLS networks for MPLS VPN
- IPv6-specific extensions:
  - Scoped addresses: NEXT_HOP contains a global IPv6 address and potentially a link-local address.
  - NEXT_HOP and NLRI are expressed as IPv6 addresses and prefixes in the multiprotocol attributes.
- Still uses TCP for transport:
  - TCP can run over IPv4, transporting IPv6 information.
  - TCP can run natively over IPv6.

© 2012 Cisco and/or its affiliates. All rights reserved.

MPLS VPN Overview

- It is used to connect multiple customer locations via the common Layer 3 infrastructure of a service provider.
  - A special VPN can be used to provide Internet connectivity.
  - Routing used can be static or dynamic, depending on the service provider.
  - The customer routers are connected to the service provider PE



© 2012 Cisco and/or its affiliates. All rights reserved.

Summary

- BGP is the external routing protocol used between autonomous systems.
- Autonomous system is a collection of networks under a single administration and is represented by 16-bit or 32-bit number.
- Forwarding is based on policies and not on best path. BGP routers exchange network reachability information, called path vectors, which are made up of path attributes.
- BGP announces set of AS numbers and netwroks that are reachable at the end of the path. The path is described by attributes.
- BGP router can advertise to neighboring autonomous systems only those routes that it actually uses.

© 2012 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- BGP routers establish a TCP session and then exchange routing tables. After that, BGP peers send only incremental triggered updates.
- BGP uses three databases: neighbor table, BGP table, and routing table.
- The four BGP message types are open, keepalive, update, and notification.
- BGP4, with multiprotocol extensions, enables the use of many address families. Address families define the type of addressews being carried.
- MPLS VPNs are used by customers with multiple locations that do not want to use expensive Layer 2 technologies.

© 2012 Cisco and/or its affiliates. All rights reserved.