Juniper Networks | Education Services

Advanced Junos Service Provider Routing

STUDENT GUIDE

Revision V19A

Education Services Courseware

Advanced Junos Service Provider Routing

V19A

Student Guide
Volume 1 of 2

Education Services

1133 Innovation Way
Sunnyvale, CA 94089 USA
408-745-2000
www.juniper.net

Course Number: EDU-JUN-AJSPR

This document is produced by Juniper Networks, Inc.

This document or any part thereof may not be reproduced or transmitted in any form under penalty of law, without the prior written permission of Juniper Networks Education Services.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Advanced Junos Service Provider Routing Student Guide, Revision V19A

Copyright © 2020 Juniper Networks, Inc. All rights reserved.

Printed in USA.

Revision History:

Revision 10.a — March 2011

Revision 10.b—September 2011

Revision 11.a—January 2012

Revision 12.a—September 2013

Revision V-17.a—June 2017

Revision V19A—February 2020

The information in this document is current as of the date listed above.

The information in this document has been carefully verified and is believed to be accurate for software Release 19.3. Juniper Networks assumes no responsibilities for any inaccuracies that may appear in this document. In no event will Juniper Networks be liable for direct, indirect, special, exemplary, incidental, or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

YEAR 2000 NOTICE

Juniper Networks hardware and software products do not suffer from Year 2000 problems and hence are Year 2000 compliant. The Junos operating system has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

SOFTWARE LICENSE

The terms and conditions for using Juniper Networks software are described in the software license provided with the software, or to the extent applicable, in an agreement executed between you and Juniper Networks, or Juniper Networks agent. By using Juniper Networks software, you indicate that you understand and agree to be bound by its license terms and conditions. Generally speaking, the software license restricts the manner in which you are permitted to use the Juniper Networks software, may contain prohibitions against certain uses, and may state conditions under which the license is automatically terminated. You should consult the software license for further details.

Contents

Chapter 1: Course Introduction 1-1

Chapter 2: OSPF 2-1
OSPFv2 Review 2-3
Protocol Operations 2-30
OSPF Authentication 2-46
OSPFv3 2-51
Lab: Configuring and Monitoring OSPF 2-60

Chapter 3: Advanced OSPF 3-1
Review of OSPF Areas 3-3
Stub Area Operation 3-9
Stub Area Configuration 3-14
NSSA Operation 3-17
NSSA Configuration 3-23
Route Summarization 3-26
Lab: Configuring and Monitoring OSPF Areas and Route Summarization 3-38

Chapter 4: Advanced OSPF Options 4-1
OSPF Multi-Area Adjacencies 4-3
Virtual Links 4-12
External Reachability 4-22
Lab: Configuring Advanced OSPF Options 4-40

Chapter 5: Troubleshooting OSPF 5-1
Troubleshooting OSPF Adjacency Issues 5-3
Troubleshooting LSDB Consistency Issues 5-16
Troubleshooting OSPF Routing Issues 5-18
Lab: Troubleshooting OSPF 5-35

Chapter 6: IS-IS 6-1
Overview of IS-IS 6-3
IS-IS PDUs 6-12
Neighbors and Adjacencies 6-34
Configuring and Monitoring IS-IS 6-39
Lab: IS-IS Configuration and Monitoring 6-54

Chapter 7: Advanced IS-IS Operations and Configuration Options 7-1
IS-IS Operations 7-3
IS-IS Configuration Options 7-16
IS-IS Routing Policy 7-29
Lab: Advanced IS-IS Configuration Options and Routing Policy 7-41

Acronym List ACR-1

Course Overview

This five-day course is designed to provide students with detailed coverage of OSPF, IS-IS, BGP, and routing policy. Through demonstrations and hands-on labs, students will gain experience in configuring, monitoring, and troubleshooting the Junos operating system and in monitoring device and protocol operations.

This course uses Juniper Networks vMX Series Routers for the hands-on component, but the lab environment does not preclude the course from being applicable to other Juniper hardware platforms running the Junos OS. This course is based on the Junos OS Release 19.3.

Course Level

Advanced Junos Service Provider Routing (AJSPR) is an advanced-level course.

Intended Audience

This course benefits individuals responsible for implementing, monitoring, and troubleshooting Layer 3 components of a service provider's network.

Prerequisites

Students should have intermediate-level networking knowledge and an understanding of the Open Systems Interconnection (OSI) model and the TCP/IP protocol suite. Students should also attend the *Introduction to the Junos Operating System* (IJOS) and *Junos Intermediate Routing* (JIR) courses prior to attending this class.

Objectives

After successfully completing this course, you should be able to:

- Describe the various OSPF link-state advertisement (LSA) types.
- Explain the flooding of LSAs in an OSPF network.
- Describe the shortest-path-first (SPF) algorithm.
- Explain OSPF link metrics.
- Describe the various OSPF authentication methods.
- Explain the key differences between OSPFv2 and OSPFv3.
- Describe OSPF area types and operations.
- Configure various OSPF area types.
- Summarize and restrict routes.
- Configure OSPF multi-area adjacencies
- Configure OSPF virtual links.
- Explain OSPF external reachability.
- List useful commands that are used to troubleshoot and verify OSPF.
- Isolate different OSPF issues.
- Explain the concepts and operation of IS-IS.
- Describe various IS-IS link-state protocol data unit (LSP) types.
- List IS-IS adjacency rules and troubleshoot common adjacency issues.
- Configure and monitor IS-IS.
- Display and interpret the link-state database (LSDB).
- Perform advanced IS-IS configuration options.
- Implement IS-IS routing policy.
- Explain the default operation in multilevel IS-IS.
- Describe address summarization methods used in IS-IS.

- Describe address summarization methods used in IS-IS.
- Configure and monitor a multilevel IS-IS network.
- List useful commands to troubleshoot and verify IS-IS problems.
- Troubleshoot and isolate different IS-IS issues.
- Describe basic BGP operation.
- List common BGP attributes.
- Explain the route selection process for BGP.
- Describe how to alter the route selection process.
- Configure some advanced options for BGP peers.
- Explain how policies function in BGP.
- Describe BGP attributes and explain how these attributes can be used to manipulate traffic.
- Show how Regex can be used in policies to manipulate AS-path.
- Describe the BGP attributes Origin, MED and communities in detail and explain the operation of those attributes.
- Manipulate these BGP attributes using routing policy.
- Describe the operation of BGP route reflection.
- Configure a route reflector.
- Describe the virtual route reflector
- Describe the operation of optimal route reflection
- Describe the operation of a BGP confederation.
- Configure confederations.
- Describe peering relationships in a confederation.
- Describe DDoS attacks.
- Describe DDoS mitigation techniques.
- Describe FlowSpec Operations.
- Configure and Monitor FlowSpec.
- Review common BGP troubleshooting procedures.
- List common BGP troubleshooting commands.
- Identify issues with BGP peering.
- Isolate problems on routing policy structure and configuration.
- Identify common commands for troubleshooting routing policy.
- Explain the causes for route instability.
- Describe the effect of damping on BGP routing.
- Explain the default behavior of damping on links.
- Control damping using routing policy.
- View damped routes using command-line interface (CLI) commands.

Course Agenda

Day 1

Chapter 1: Course Introduction

Chapter 2: OSPF

Lab 1: Configuring OSPF

Chapter 3: OSPF Areas

Lab 2: OSPF Route Summarization

Chapter 4: Advanced OSPF Options

Lab 3: Advanced OSPF

Day 2

Chapter 5: Troubleshooting OSPF

Lab 4: Troubleshooting OSPF

Chapter 6: IS-IS

Lab 5: IS-IS Configuration and Monitoring

Chapter 7: Advanced IS-IS Operations

Lab 6: Advanced IS-IS Configuration Options and Routing Policy

Day 3

Chapter 8: Multilevel IS-IS Networks

Lab 7: Configuring a Multilevel IS-IS Network

Chapter 9: Troubleshooting IS-IS

Lab 8: Troubleshooting IS-IS

Chapter 10: BGP

Lab 9: Configuring BGP

Day 4

Chapter 11: BGP Attributes and Policy—Part 1

Lab 10: BGP Attributes – Part 1

Chapter 12: BGP Attributes and Policy—Part 2

Lab 11: BGP Attributes – Part 2

Chapter 13: Route Reflection and Confederations

Lab 12: Scaling BGP

Day 5

Chapter 14: BGP FlowSpec

Lab 13: BGP FlowSpec

Chapter 15: Troubleshooting BGP

Lab 14: Troubleshooting BGP

Chapter 16: Troubleshooting Policy

Lab 15: Troubleshooting Routing Policy

Appendix A: BGP Route Damping

Lab AA: BGP Route Damping

Document Conventions

CLI and GUI Text

Frequently throughout this course, we refer to text that appears in a command-line interface (CLI) or a graphical user interface (GUI). To make the language of these documents easier to read, we distinguish GUI and CLI text from chapter text according to the following table.

Style	Description	Usage Example
Franklin Gothic	Normal text.	Most of what you read in the Lab Guide and Student Guide.

`Courier New`	Console text:
	• Screen captures
	• Noncommand-related syntax
	GUI text elements:
	• Menu names
	• Text field entry
		
		`commit complete`
		`Exiting configuration mode`
		
		Select `File > Open`, and then click `Configuration.conf` in the `Filename` text box.

Input Text Versus Output Text

You will also frequently see cases where you must enter input text yourself. Often these instances will be shown in the context of where you must enter them. We use bold style to distinguish text that is input versus text that is simply displayed.

Style	Description	Usage Example
`Normal CLI`	No distinguishing variant.	`Physical interface:fxp0, Enabled`
`Normal GUI`		View configuration history by clicking `Configuration > History`.
`CLI Input`	Text that you must enter.	`lab@San_Jose>` **`show route`**
`GUI Input`		Select `File > Save`, and type **`config.ini`** in the `Filename` field.

Undefined Syntax Variables

Finally, this course distinguishes syntax variables, where you must assign the value (undefined variables). Note that these styles can be combined with the input style as well.

Style	Description	Usage Example
`CLI Undefined` *`GUI Undefined`*	Text where the variable's value is the user's discretion or text where the variable's value as shown in the lab guide might differ from the value the user must input according to the lab topology.	Type **`set policy`** ***`policy-name`***. **`ping 10.0.`*****`x.y`*** Select `File > Save`, and type ***`filename`*** in the `Filename` field.

Additional Information

Education Services Offerings

You can obtain information on the latest Education Services offerings, course dates, and class locations from the World Wide Web by pointing your Web browser to: http://www.juniper.net/training/education/.

About This Publication

This course was developed and tested using the software release listed on the copyright page. Previous and later versions of software might behave differently so you should always consult the documentation and release notes for the version of code you are running before reporting errors.

This document is written and maintained by the Juniper Networks Education Services development team. Please send questions and suggestions for improvement to training@juniper.net.

Technical Publications

You can print technical manuals and release notes directly from the Internet in a variety of formats:

- Go to http://www.juniper.net/techpubs/.
- Locate the specific software or hardware release and title you need, and choose the format in which you want to view or print the document.

Documentation sets and CDs are available through your local Juniper Networks sales office or account representative.

Juniper Networks Support

For technical support, contact Juniper Networks at http://www.juniper.net/customers/support/, or at 1-888-314-JTAC (within the United States) or 408-745-2121 (from outside the United States).

Education Services

Advanced Junos Service Provider Routing

Chapter 1: Course Introduction

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - Get to know one another
  - Identify the objectives, prerequisites, facilities, and materials used during this course
  - Identify additional Education Services courses at Juniper Networks
  - Describe the Juniper Networks Certification Program

© 2020 Juniper Networks, Inc. All Rights Reserved.

juniper NETWORKS | 2

We Will Discuss:

- Objectives and course content information;
- Additional Juniper Networks, Inc. courses; and
- The Juniper Networks Certification Program.

Introductions

- Before we get started…
  - What is your name?
  - Where do you work?
  - What is your primary role in your organization?
  - What kind of network experience do you have?
  - Are you certified on Juniper Networks?
  - What is the most important thing for you to learn in this training session?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Introductions

The slide asks several questions for you to answer during class introductions.

Prerequisites

- The prerequisites for this course include:
  - Intermediate-level networking knowledge and familiarity with the Junos OS command-line interface (CLI)
  - Successful completion of the following courses prior to attending this class:
    - *Introduction to the Junos Operating System* (IJOS)
    - *Junos Intermediate Routing* (JIR)
    - *Junos Service Provider Switching* (JSPX)

© 2020 Juniper Networks, Inc. All Rights Reserved.

Prerequisites

The slide lists the prerequisites for this course.

Course Contents (1 of 2)

- Contents:
  - Chapter 1: Course Introduction
  - Chapter 2: OSPF
  - Chapter 3: OSPF Areas
  - Chapter 4: Advanced OSPF Options
  - Chapter 5: Troubleshooting OSPF
  - Chapter 6: IS-IS
  - Chapter 7: Advanced IS-IS Operations and Configuration Options
  - Chapter 8: Multilevel IS-IS Networks
  - Chapter 9: Troubleshooting ISIS

© 2020 Juniper Networks, Inc. All Rights Reserved.

Course Contents

The slide lists the topics we discuss in this course.

Course Contents (2 of 2)

- Contents: (contd.)
  - Chapter 10: BGP
  - Chapter 11: BGP Attributes and Policy—Part 1
  - Chapter 12: BGP Attributes and Policy—Part 2
  - Chapter 13: Route Reflection and Confederations
  - Chapter 14: BGP FlowSpec
  - Chapter 15: Troubleshooting BGP
  - Chapter 16: Troubleshooting Policy

© 2020 Juniper Networks, Inc. All Rights Reserved.

Course Administration

- The basics:
  - Sign-in sheet
  - Schedule
    - Class times
    - Breaks
    - Lunch
  - Break and restroom facilities
  - Fire and safety procedures
  - Communications
    - Telephones and wireless devices
    - Internet access

© 2020 Juniper Networks, Inc. All Rights Reserved.

General Course Administration

The slide documents general aspects of classroom administration.

Education Materials

- Available materials for classroom-based and instructor-led online classes:
  - Lecture material
  - Lab guide
  - Lab equipment
- Self-paced online courses also available
  - http://www.juniper.net/courses

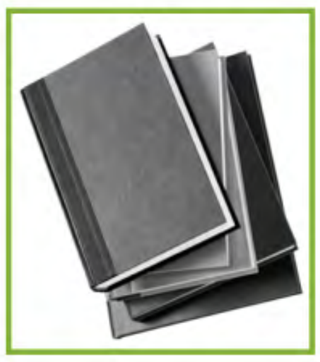

© 2020 Juniper Networks, Inc. All Rights Reserved.

Training and Study Materials

The slide describes Education Services materials that are available for reference both in the classroom and online.

Additional Resources

- For those who want more:
  - Juniper Networks Technical Assistance Center (JTAC)
    - http://www.juniper.net/support/requesting-support.html
  - Juniper Networks books
    - http://www.juniper.net/books
  - Hardware and software technical documentation
    - Online: http://www.juniper.net/techpubs
    - Portable libraries: http://www.juniper.net/techpubs/resources
  - Certification resources
    - http://www.juniper.net/certification

© 2020 Juniper Networks, Inc. All Rights Reserved.

Additional Resources

The slide provides links to additional resources available to assist you in the installation, configuration, and operation of Juniper Networks products.

Satisfaction Feedback

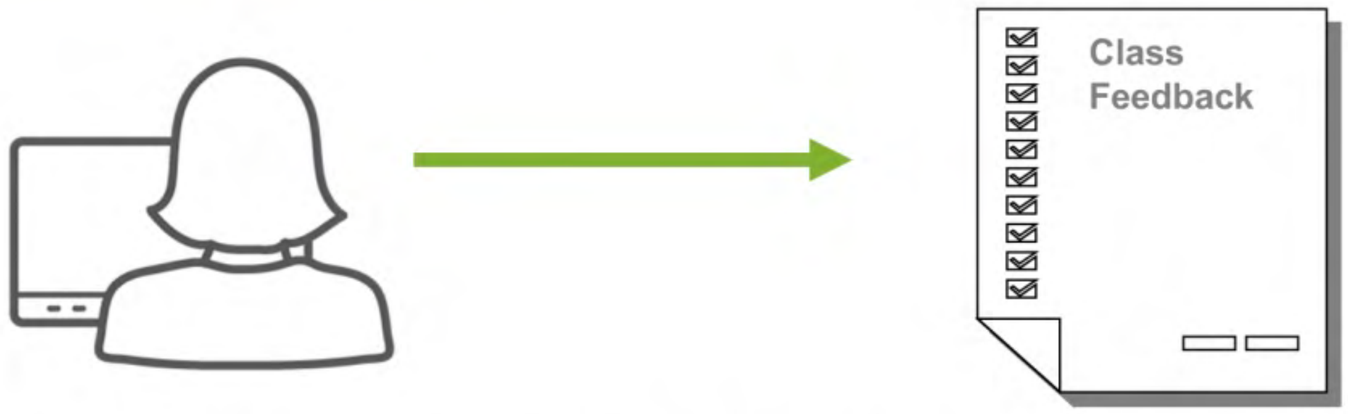

- To receive your certificate, you must complete the survey
  - Either you will receive a survey to complete at the end of class, or we will e-mail it to you within two weeks
  - Completed surveys help us serve you better!

© 2020 Juniper Networks, Inc. All Rights Reserved.

Satisfaction Feedback

Juniper Networks uses an electronic survey system to collect and analyze your comments and feedback. Depending on the class you are taking, please complete the survey at the end of the class, or be sure to look for an e-mail about two weeks from class completion that directs you to complete an online survey form. (Be sure to provide us with your current e-mail address.)

Submitting your feedback entitles you to a certificate of class completion. We thank you in advance for taking the time to help us improve our educational offerings.

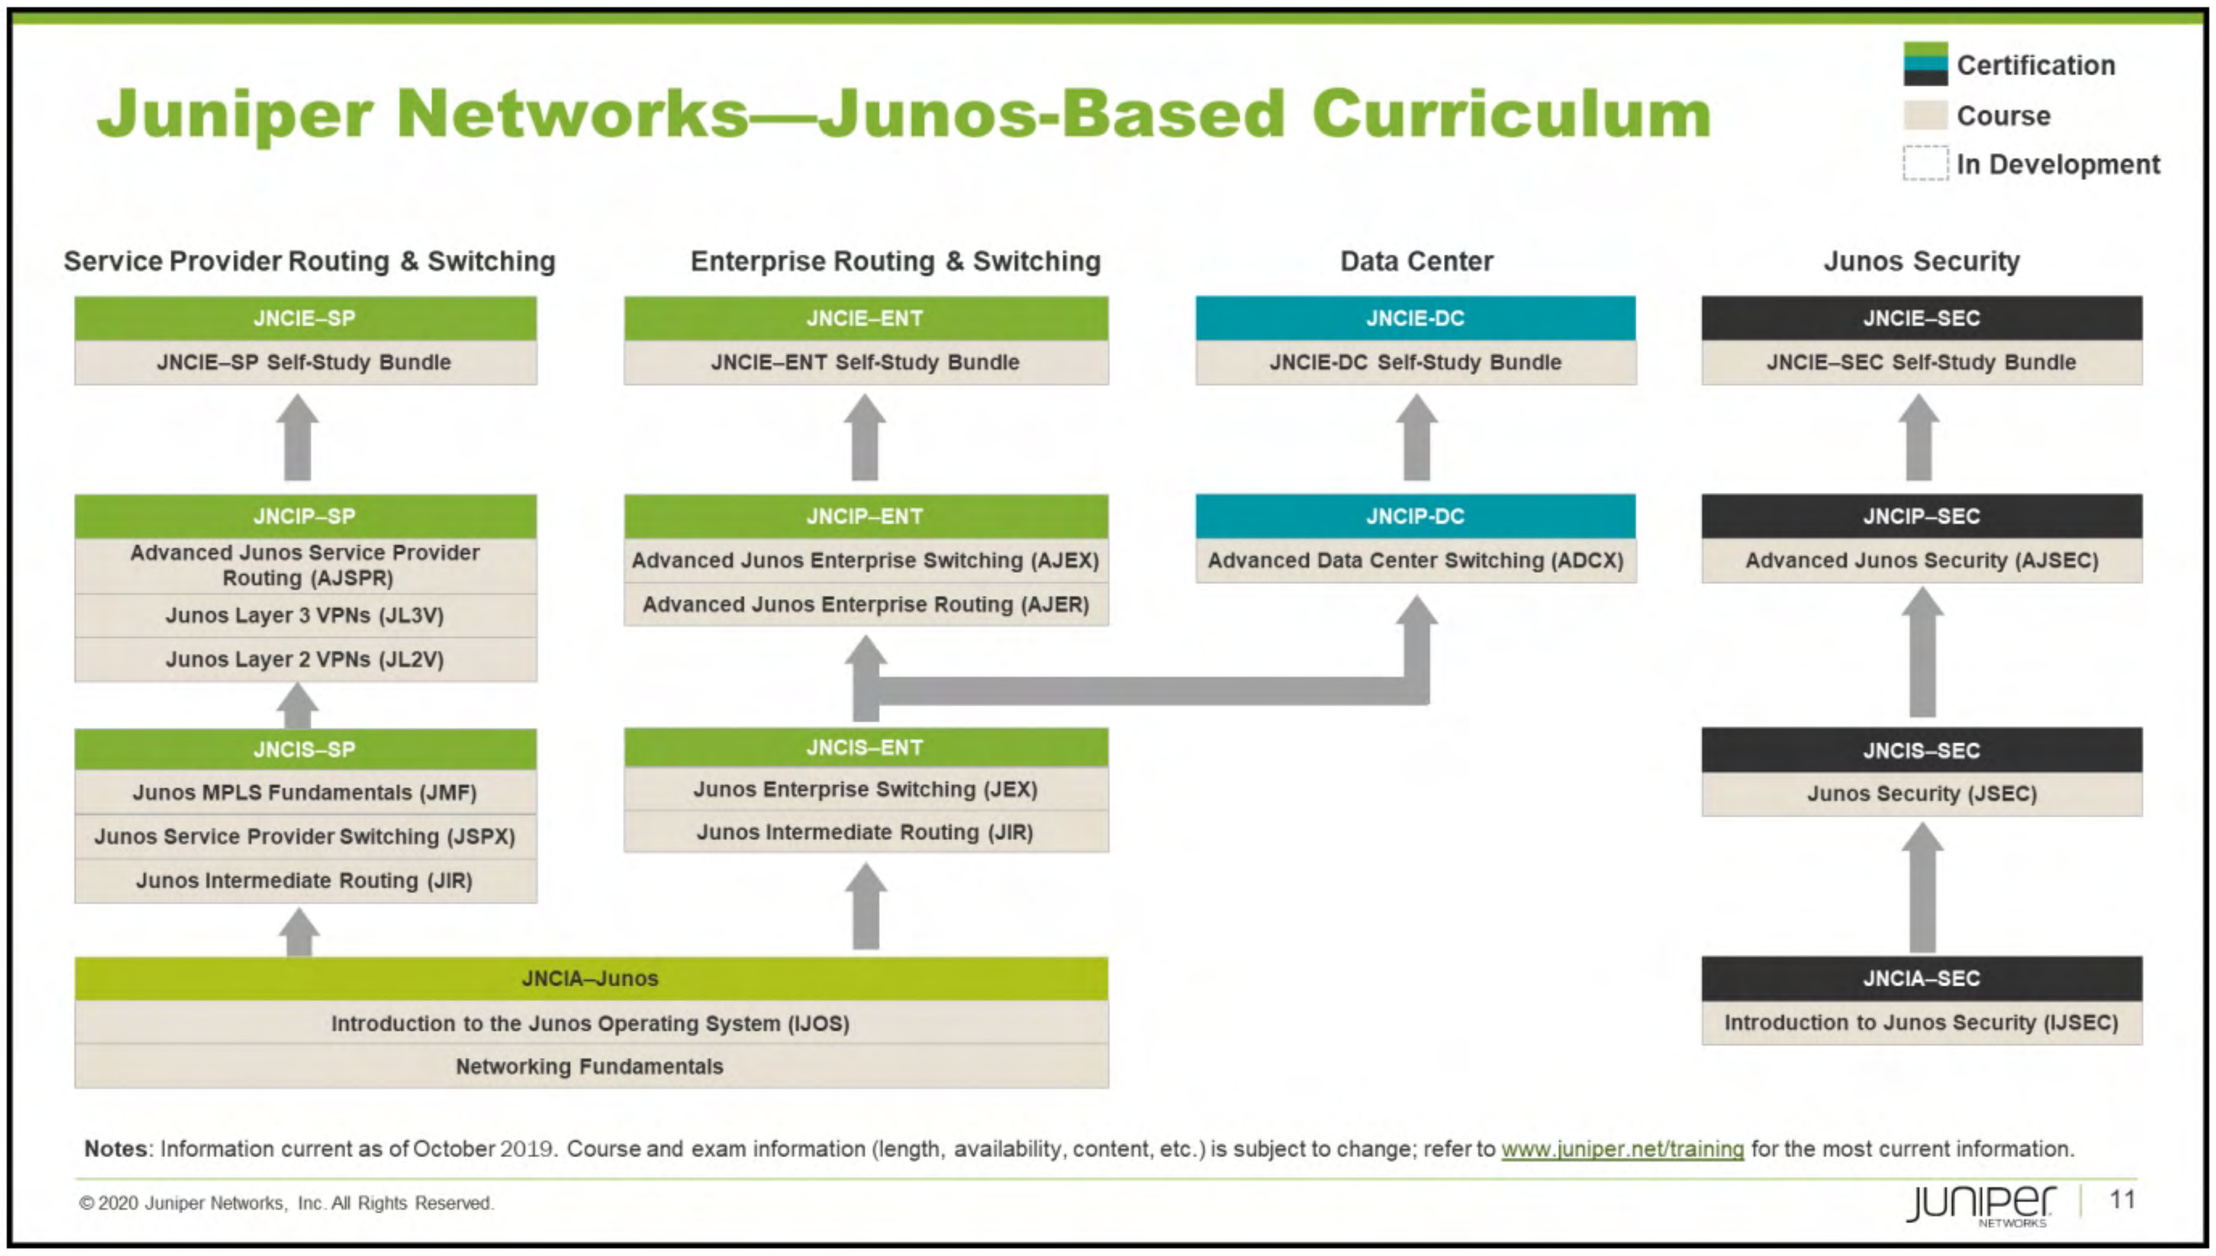

Juniper Networks Education Services Curriculum

Juniper Networks Education Services can help ensure that you have the knowledge and skills to deploy and maintain cost-effective, high-performance networks for both enterprise and service provider environments. We have expert training staff with deep technical and industry knowledge, providing you with instructor-led hands-on courses in the classroom and online, as well as convenient, self-paced eLearning courses. In addition to the courses shown on the slide, Education Services offers training in automation, E-Series, firewall/VPN, IDP, network design, QFabric, support, and wireless LAN.

Courses

Juniper Networks courses are available in the following formats:

- Classroom-based instructor-led technical courses;
- Online instructor-led technical courses;
- Self-paced on-demand training with labs;
- Hardware installation eLearning courses as well as technical eLearning courses;
- Learning bytes: Short, topic-specific, video-based lessons covering Juniper products and technologies.

Find the latest Education Services offerings covering a wide range of platforms at www.juniper.net/training.

Juniper Networks Certification Program

- Why earn a Juniper Networks certification?
  - Juniper Networks certification makes you stand out
    - Demonstrate your solid understanding of networking technologies
    - Distinguish yourself and grow your career
  - Capitalize on the promise of the New Network
    - Develop and deploy the services you need
    - Lead the way and increase your value
  - Unique benefits for certified individuals

Juniper Networks | Education Services CERTIFICATION PROGRAM

© 2020 Juniper Networks, Inc. All Rights Reserved.

Juniper Networks Certification Program

A Juniper Networks certification is the benchmark of skills and competence on Juniper Networks technologies.

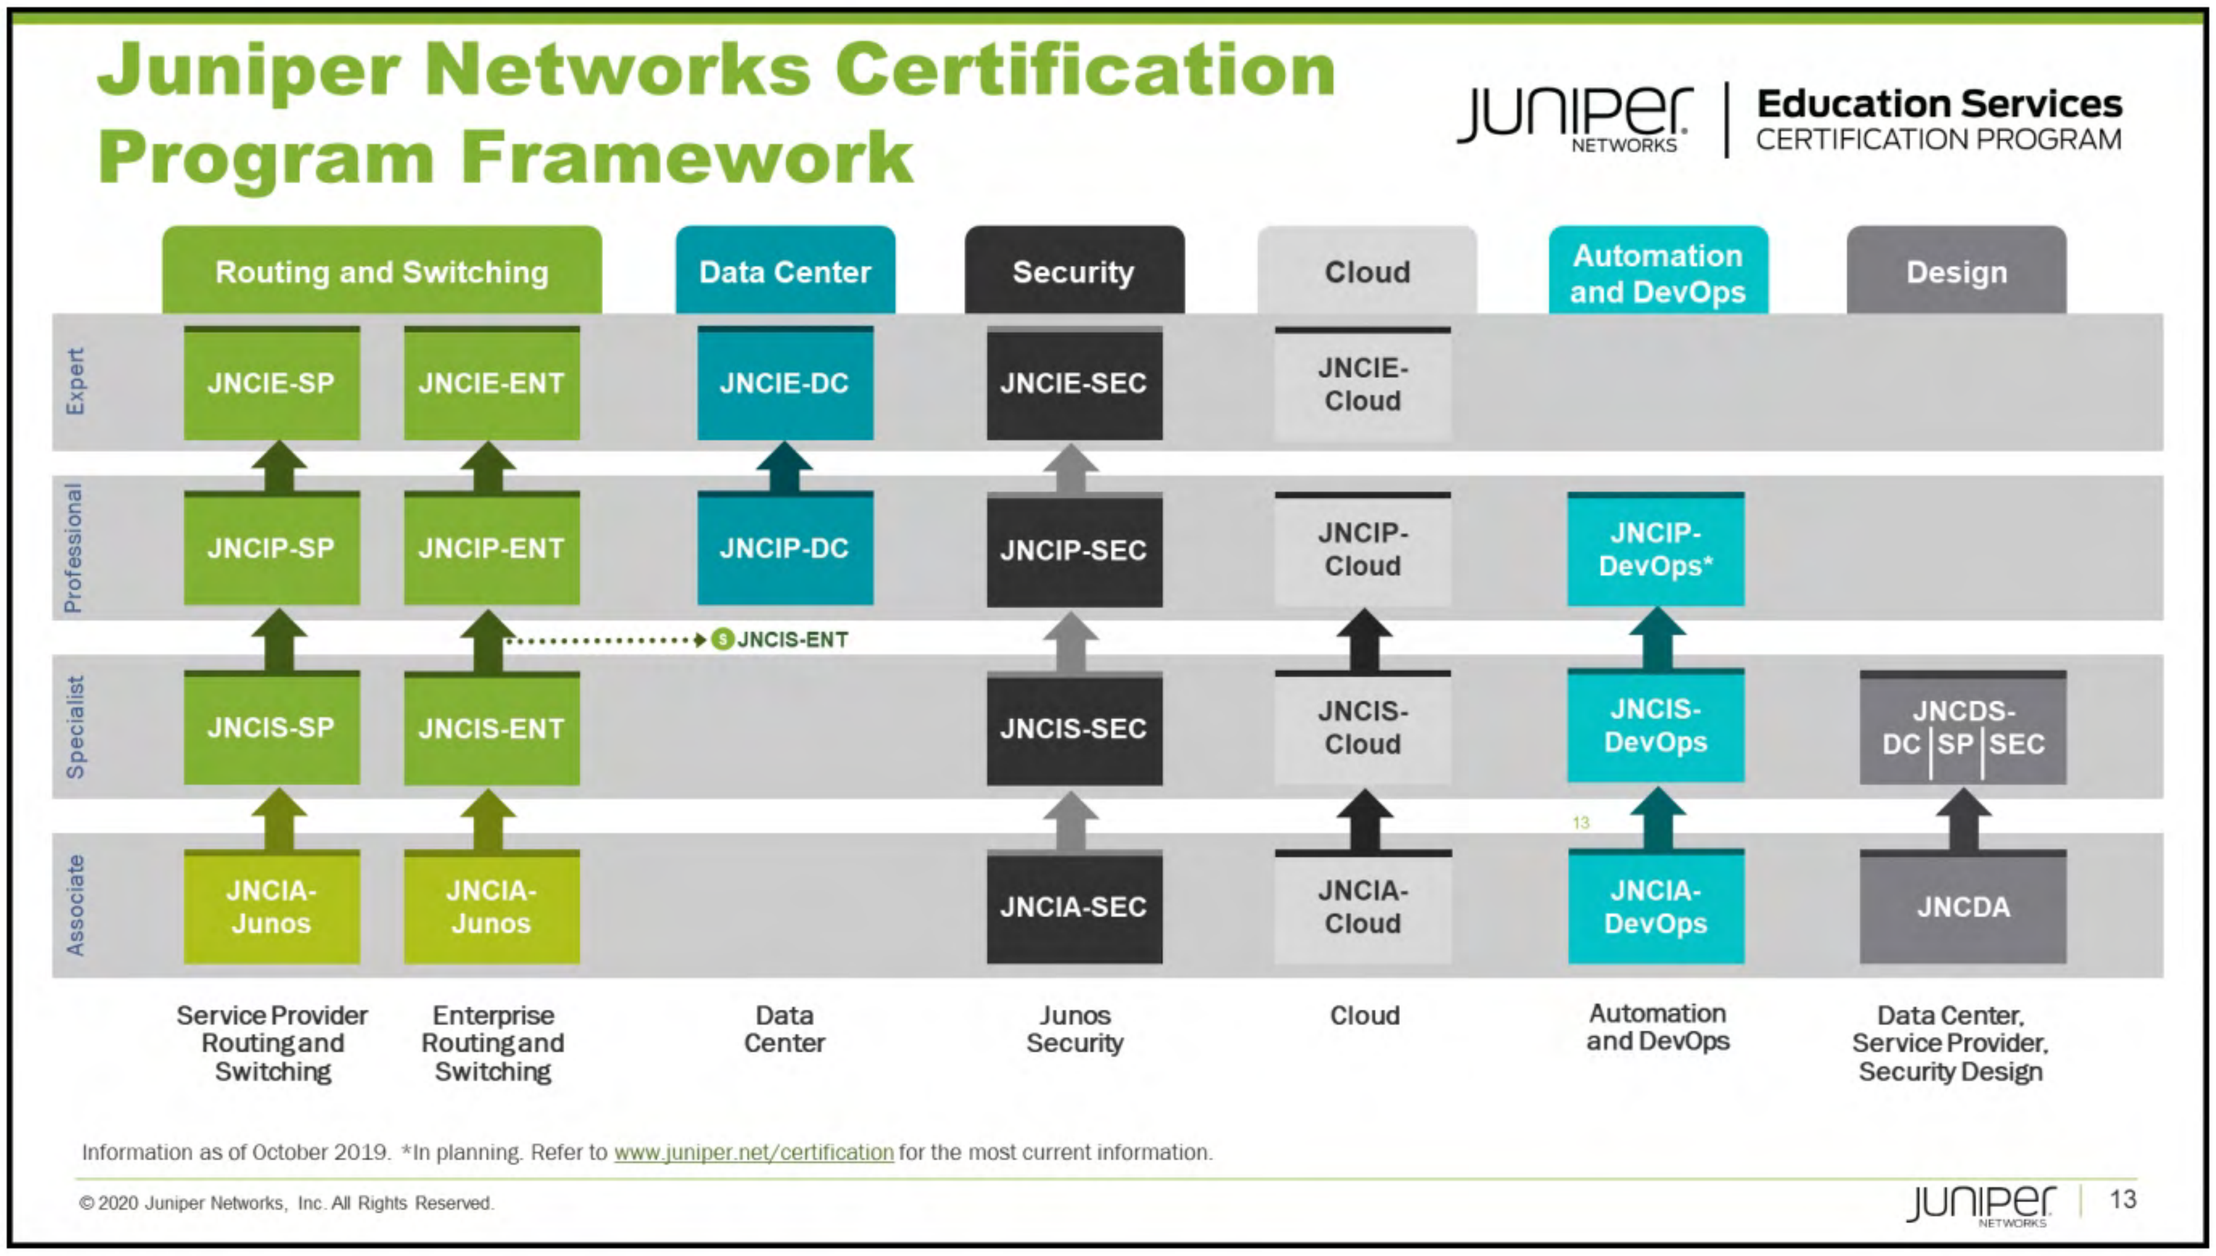

Juniper Networks Certification Program Overview

The Juniper Networks Certification Program (JNCP) consists of multitiered tracks that enable participants to demonstrate competence with Juniper Networks technology through a combination of written proficiency exams and hands-on configuration and troubleshooting exams. Successful candidates demonstrate a thorough understanding of Internet and security technologies and Juniper Networks platform configuration and troubleshooting skills.

The JNCP offers the following features:

- Multiple tracks;
- Multiple certification levels;
- Written proficiency exams; and
- Hands-on configuration and troubleshooting exams.

Each JNCP track has one to four certification levels—Associate-level, Specialist-level, Professional-level, and Expert-level. The Associate-level, Specialist-level, and Professional-level exams are computer-based exams composed of multiple choice questions administered at Pearson VUE testing centers worldwide.

Expert-level exams are composed of hands-on lab exercises administered at select Juniper Networks testing centers. Please visit the JNCP website at http://www.juniper.net/certification for detailed exam information, exam pricing, and exam registration.

Certification Preparation

- Training and study resources:
  - Juniper Networks Certification Program website: www.juniper.net/certification
  - Education Services training classes: www.juniper.net/training
  - Juniper Networks documentation and white papers: www.juniper.net/documentation
- Community:
  - J-Net: http://forums.juniper.net/t5/Training-Certification-and/bd-p/Training_and_Certification
  - Twitter: @JuniperCertify

© 2020 Juniper Networks, Inc. All Rights Reserved.

Preparing and Studying

The slide lists some options for those interested in preparing for Juniper Networks certification.

Junos Genius: Certification Preparation App

- Prepare for certification
  - Access e-learning videos, practice tests, and more
    - Any time, any location, any device
    - Ongoing training opportunities as new content is regularly added
    - Prepare for associate, specialist, and professional-level JNCP certification
    - Purchase on-demand training courses
    - View assets offline (app only; JNCP practice tests not included)

*app only – JNCP practice tests are not included.

Juniper Networks | Education Services JUNOS GENIUS

www.juniper.net/junosgenius

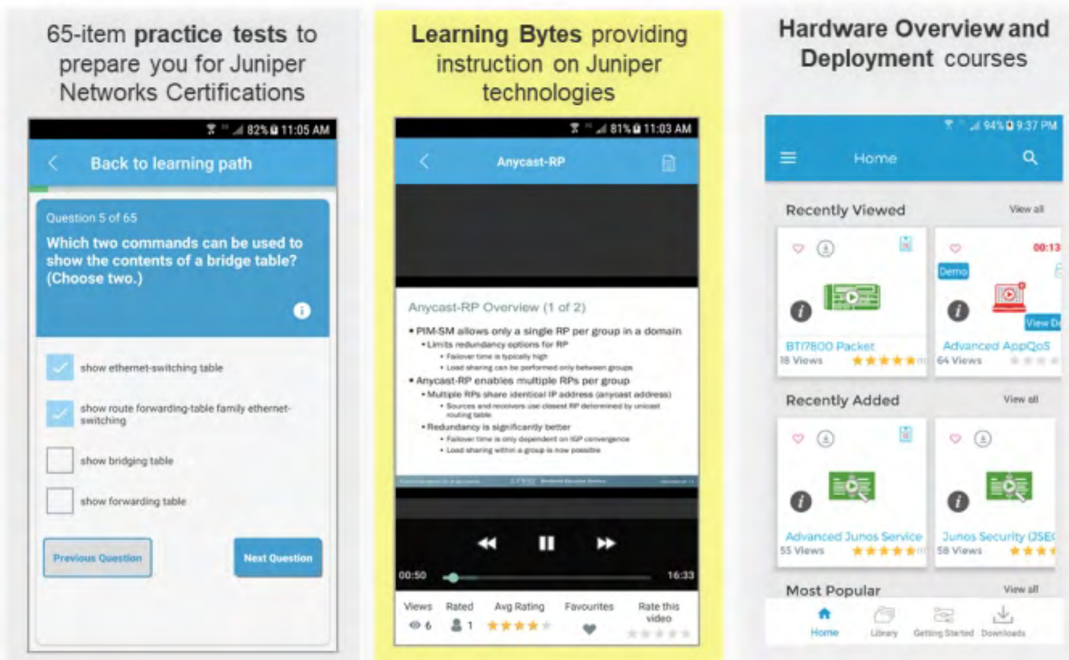

© 2020 Juniper Networks, Inc. All Rights Reserved.

Juniper Networks | 15

Junos Genius

The Junos Genius mobile learning platform (www.junosgenius.net) helps you learn Juniper technologies and prepare for Juniper certification exams on your schedule. An app for iOS and Android devices, along with laptops and desktops, Junos Genius provides certification preparation resources, practice exams, and e-learning courses developed by experts in Juniper technology. Courses cover automation, routing, switching, security, and more.

Find Us Online

http://www.juniper.net/jnet

http://www.juniper.net/facebook

http://www.juniper.net/youtube

http://www.juniper.net/twitter

© 2020 Juniper Networks, Inc. All Rights Reserved.

Find Us Online

The slide lists some online resources to learn and share information about Juniper Networks.

Questions

© 2020 Juniper Networks, Inc. All Rights Reserved.

Any Questions?

If you have any questions or concerns about the class you are attending, we suggest that you voice them now so that your instructor can best address your needs during class.

Education Services

Advanced Junos Service Provider Routing

Chapter 2: OSPF

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - Describe the various OSPF LSA types
  - Explain the flooding of LSAs in an OSPF network
  - Describe the SPF algorithm
  - Explain OSPF link metrics
  - Describe the various OSPF authentication methods
  - Explain the differences between OSPFv2 and OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Will Discuss:

- OSPF link-state advertisements (LSAs);
- How LSAs are flooded throughout the network;
- The shortest path first (SPF) algorithm;
- OSPF link metrics;
- OSPF authentication methods; and
- The differences between OSPFv2 and OSPFv3.

Agenda: OSPF

→OSPFv2 Review
- Link-State Advertisements
- Protocol Operations
- OSPF Authentication
- OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPFv2 Review

The slide lists the topics we will discuss. We discuss the highlighted topic first.

OSPFv2 Review (1 of 3)

- OSPF is a link-state IGP used within an AS
- Neighbors use hello packets to form adjacencies
- Routers on a broadcast segment elect a DR
  - Use `interface-type p2p` on a point-to-point Ethernet link to eliminate the need for this election

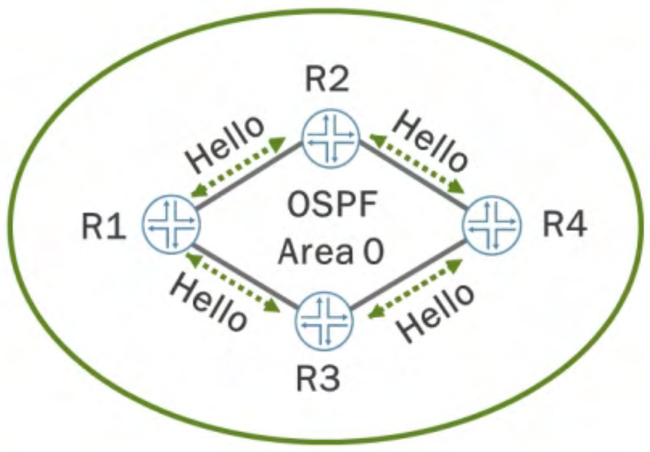

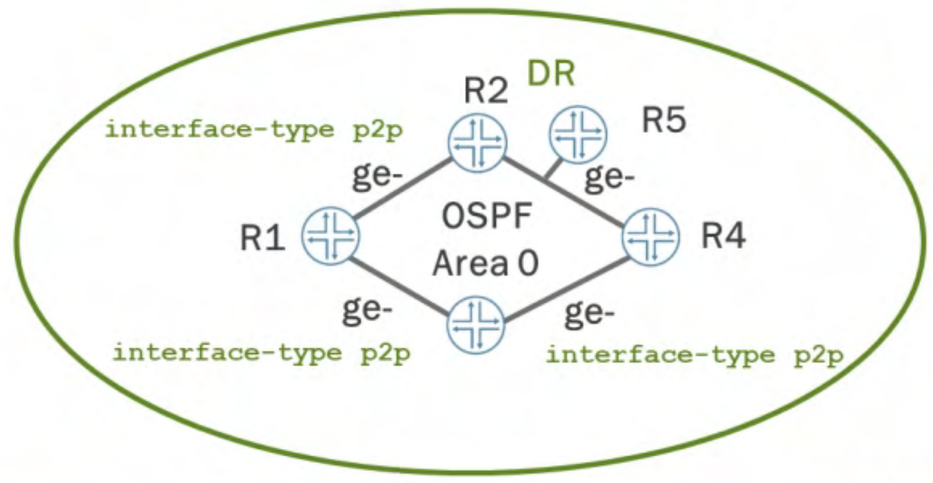

© 2020 Juniper Networks, Inc. All Rights Reserved.

Link-State Protocol

OSPF is an interior gateway protocol (IGP) based on the concept of a link-state database (LSDB). As such, each OSPF-speaking router in the network attempts to form an adjacency with each neighboring OSPF router. When these adjacencies are in place, each router generates and floods LSAs into the network in a reliable manner. The LSAs are placed into the LSDB on each router where the SPF algorithm is calculated to find the best path to each end node in the network.

Neighbors Use Hello Packets

OSPF routers send out hello packets to form and maintain adjacencies with their neighbors.

LSAs Are Flooded

OSPF uses IP protocol number 89 and the AllSPFRouters multicast address of 224.0.0.5 to flood LSAs. OSPF routers forward all LSAs through all OSPF-configured interfaces except the one on which an LSA was received. The LSAs are then installed into the OSPF database, which forms the topology map of the network. Finally, the SPF algorithm is run to determine the shortest path to each destination subnet.

Designated Router Election

On a broadcast segment, OSPF routers elect a single node to represent the segment to the network. This node is called the designated router (DR). It forms an OSPF adjacency with all routers on the segment and floods a network LSA into the appropriate area. The criteria for electing the DR is the highest configured priority on the segment, which is set to 128 by default. The second criteria for electing a DR is the highest router ID on the segment.

The election of a DR on a broadcast segment is a nondeterministic event. Thus, the router with the best criteria might not always be the DR for the segment. An operational DR maintains its status on the segment until it stops operating.

The first OSPF router on a link determines itself as the DR if it does not detect a Hello from another router within the 40 second election period.

Currently, a lot of Ethernet segments are used as point-to-point, full-duplex links. This eliminates the need for a DR election. Use the **`interface-type p2p`** command on both sides of the link to change the default interface type. Using this option can save up to forty seconds of wait time to get the OSPF adjacency to a full state.

OSPFv2 Review (2 of 3)

- All OSPF routers maintain a copy of the database
  - Database contents consist of information learned through LSAs and must match on all routers within an area
  - SPF algorithm uses the contents of the LSDB as input data to calculate the best loop free paths

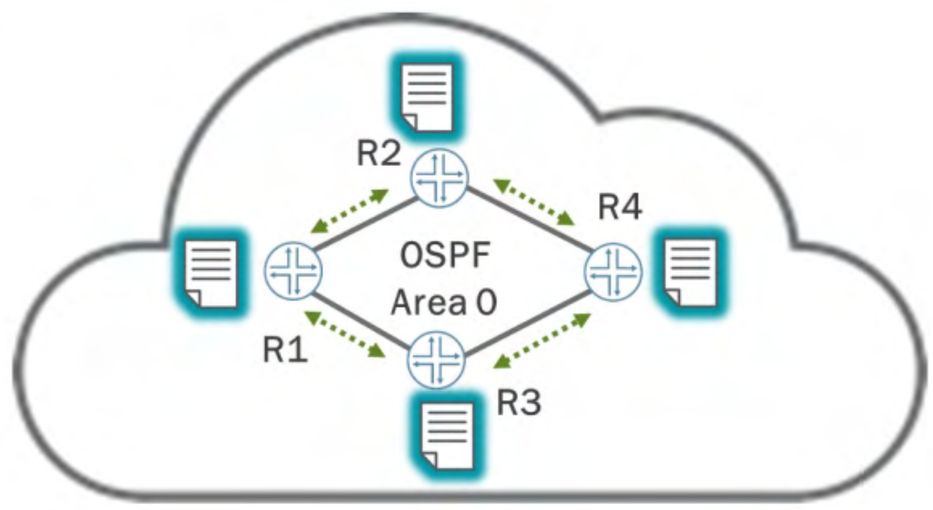

© 2020 Juniper Networks, Inc. All Rights Reserved.

The Link-State Database

In addition to discovering neighbors and flooding LSAs, a third major task of the OSPF protocol is to create and maintain the LSDB. The link-state, or topological, database stores the LSAs as a series of records. The records stored within the LSDB include details such as the advertising router's ID, its attached networks and neighboring routers, and the cost associated with those networks or neighbors. According to the OSPF RFC, each router in an OSPF area must have an identical LSDB to ensure accurate routing knowledge. We discuss OSPF areas later in this course.

The information recorded in the database is used as input data to calculate the shortest paths (that is, least-cost paths) for all destination prefixes within the network. OSPF uses the SPF algorithm (also known as the Dijkstra algorithm) to calculate, all at once, the shortest paths to all destinations. It performs this calculation by creating a tree of shortest paths incrementally and picking the best candidate from that tree. The results of this calculation are then handed to the router's routing table for the actual forwarding of data packets. For the purposes of this course, the SPF and Dijkstra algorithm are considered the same.

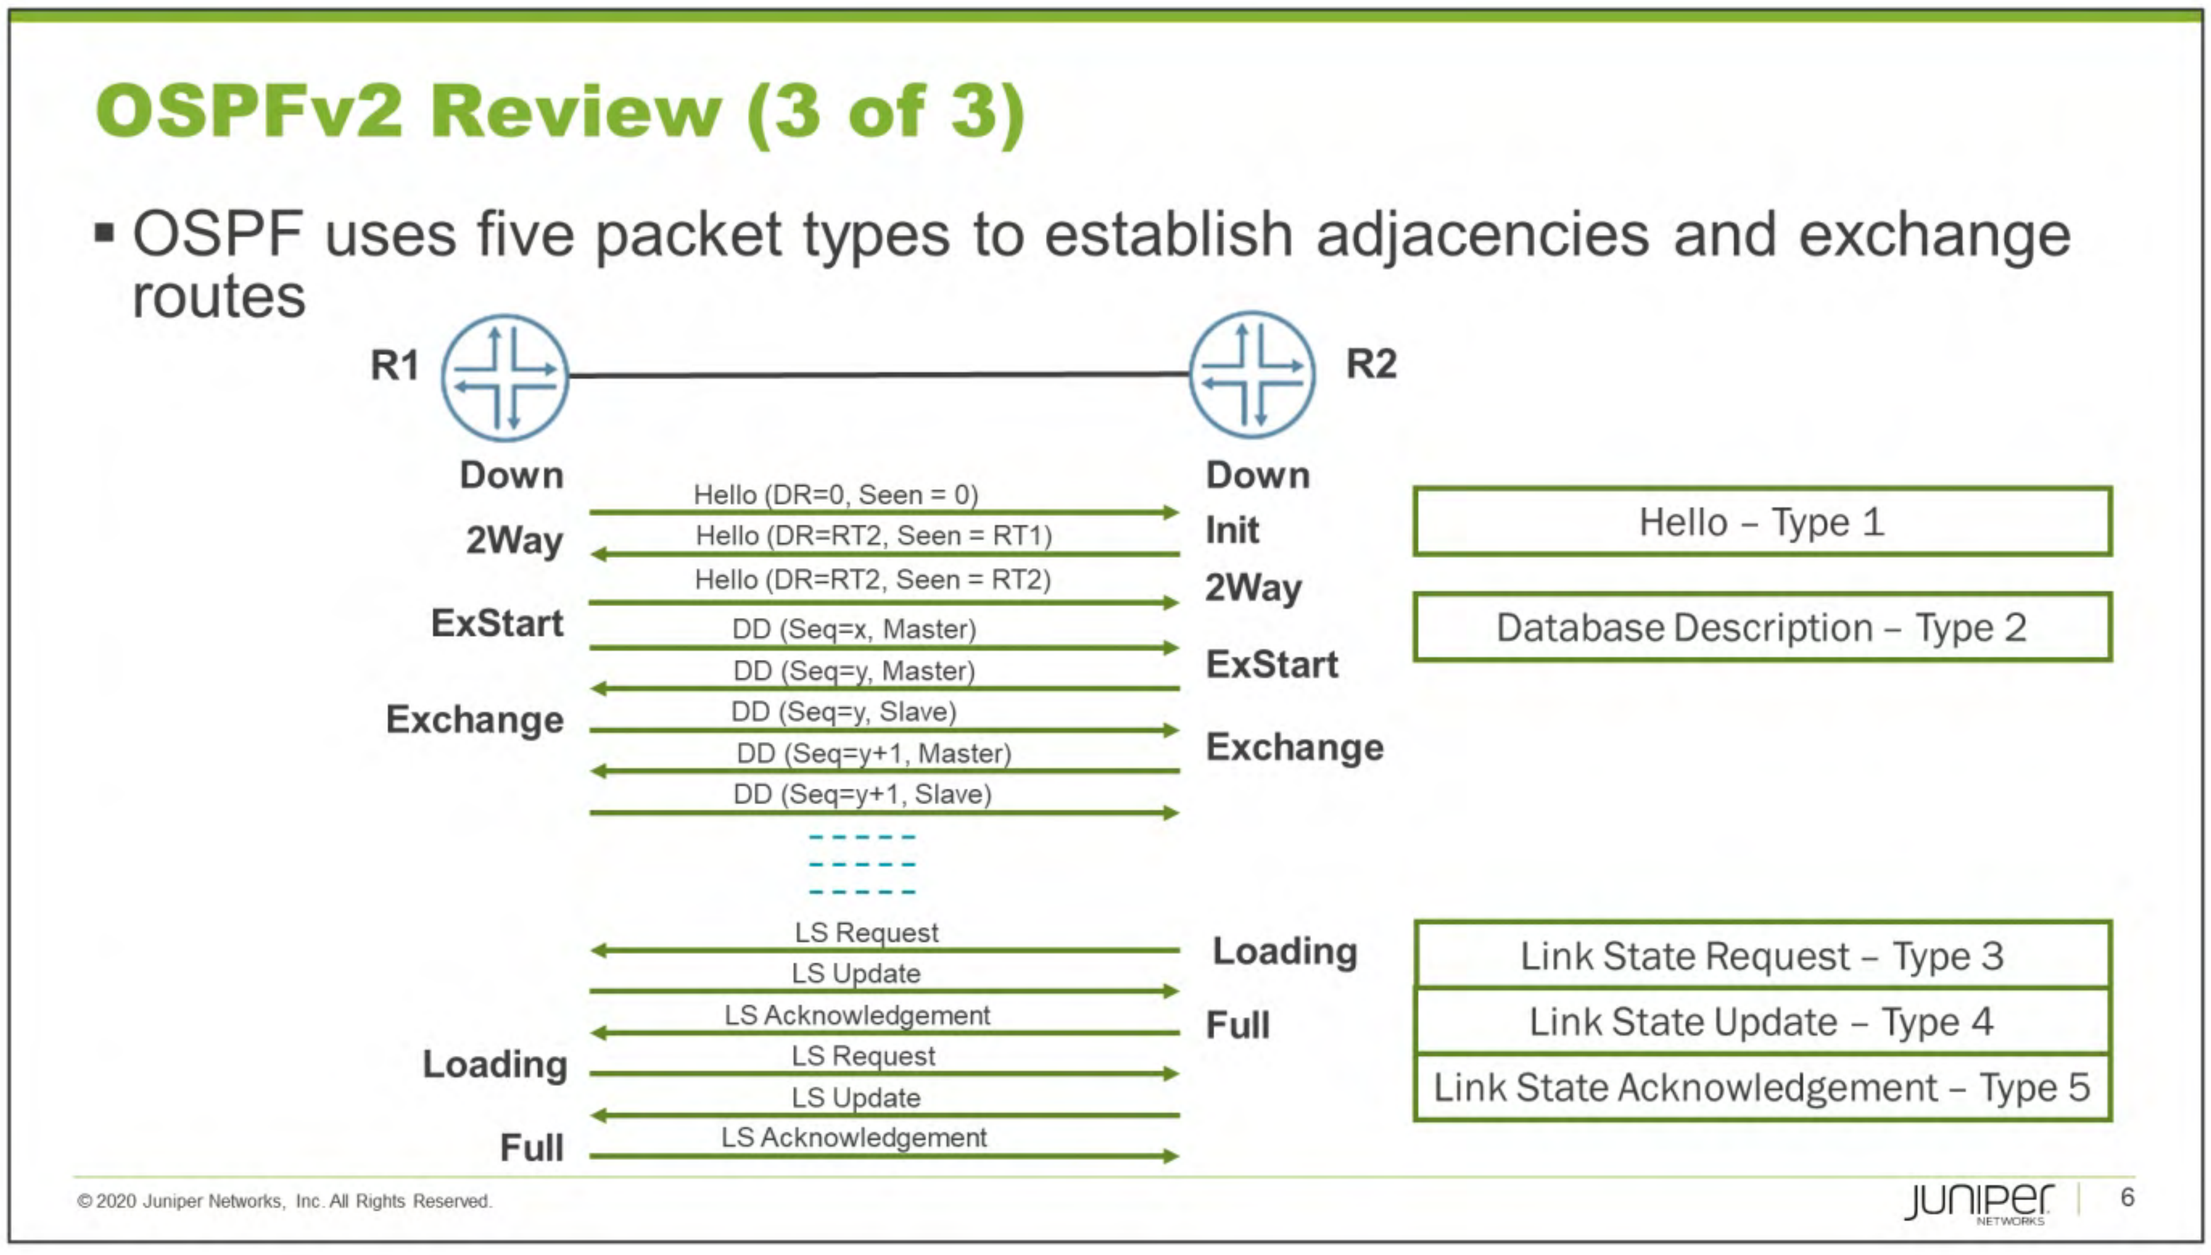

Packet Types

Every OSPF router uses a specific set of packets to perform its functions. The packet types include the following:

- *Hello*: Sent by each router to form and maintain adjacencies with its neighbors.
- *Database description*: Used by the router during the adjacency formation process. It contains the header information for the contents of the LSDB on the router.
- *Link-state request*: Used by the router to request an updated copy of a neighbor's LSA.
- *Link-state update*: Used by the router to advertise LSAs into the network.
- *Link-state acknowledgment*: Used by the router to ensure the reliable flooding of LSAs throughout the network.

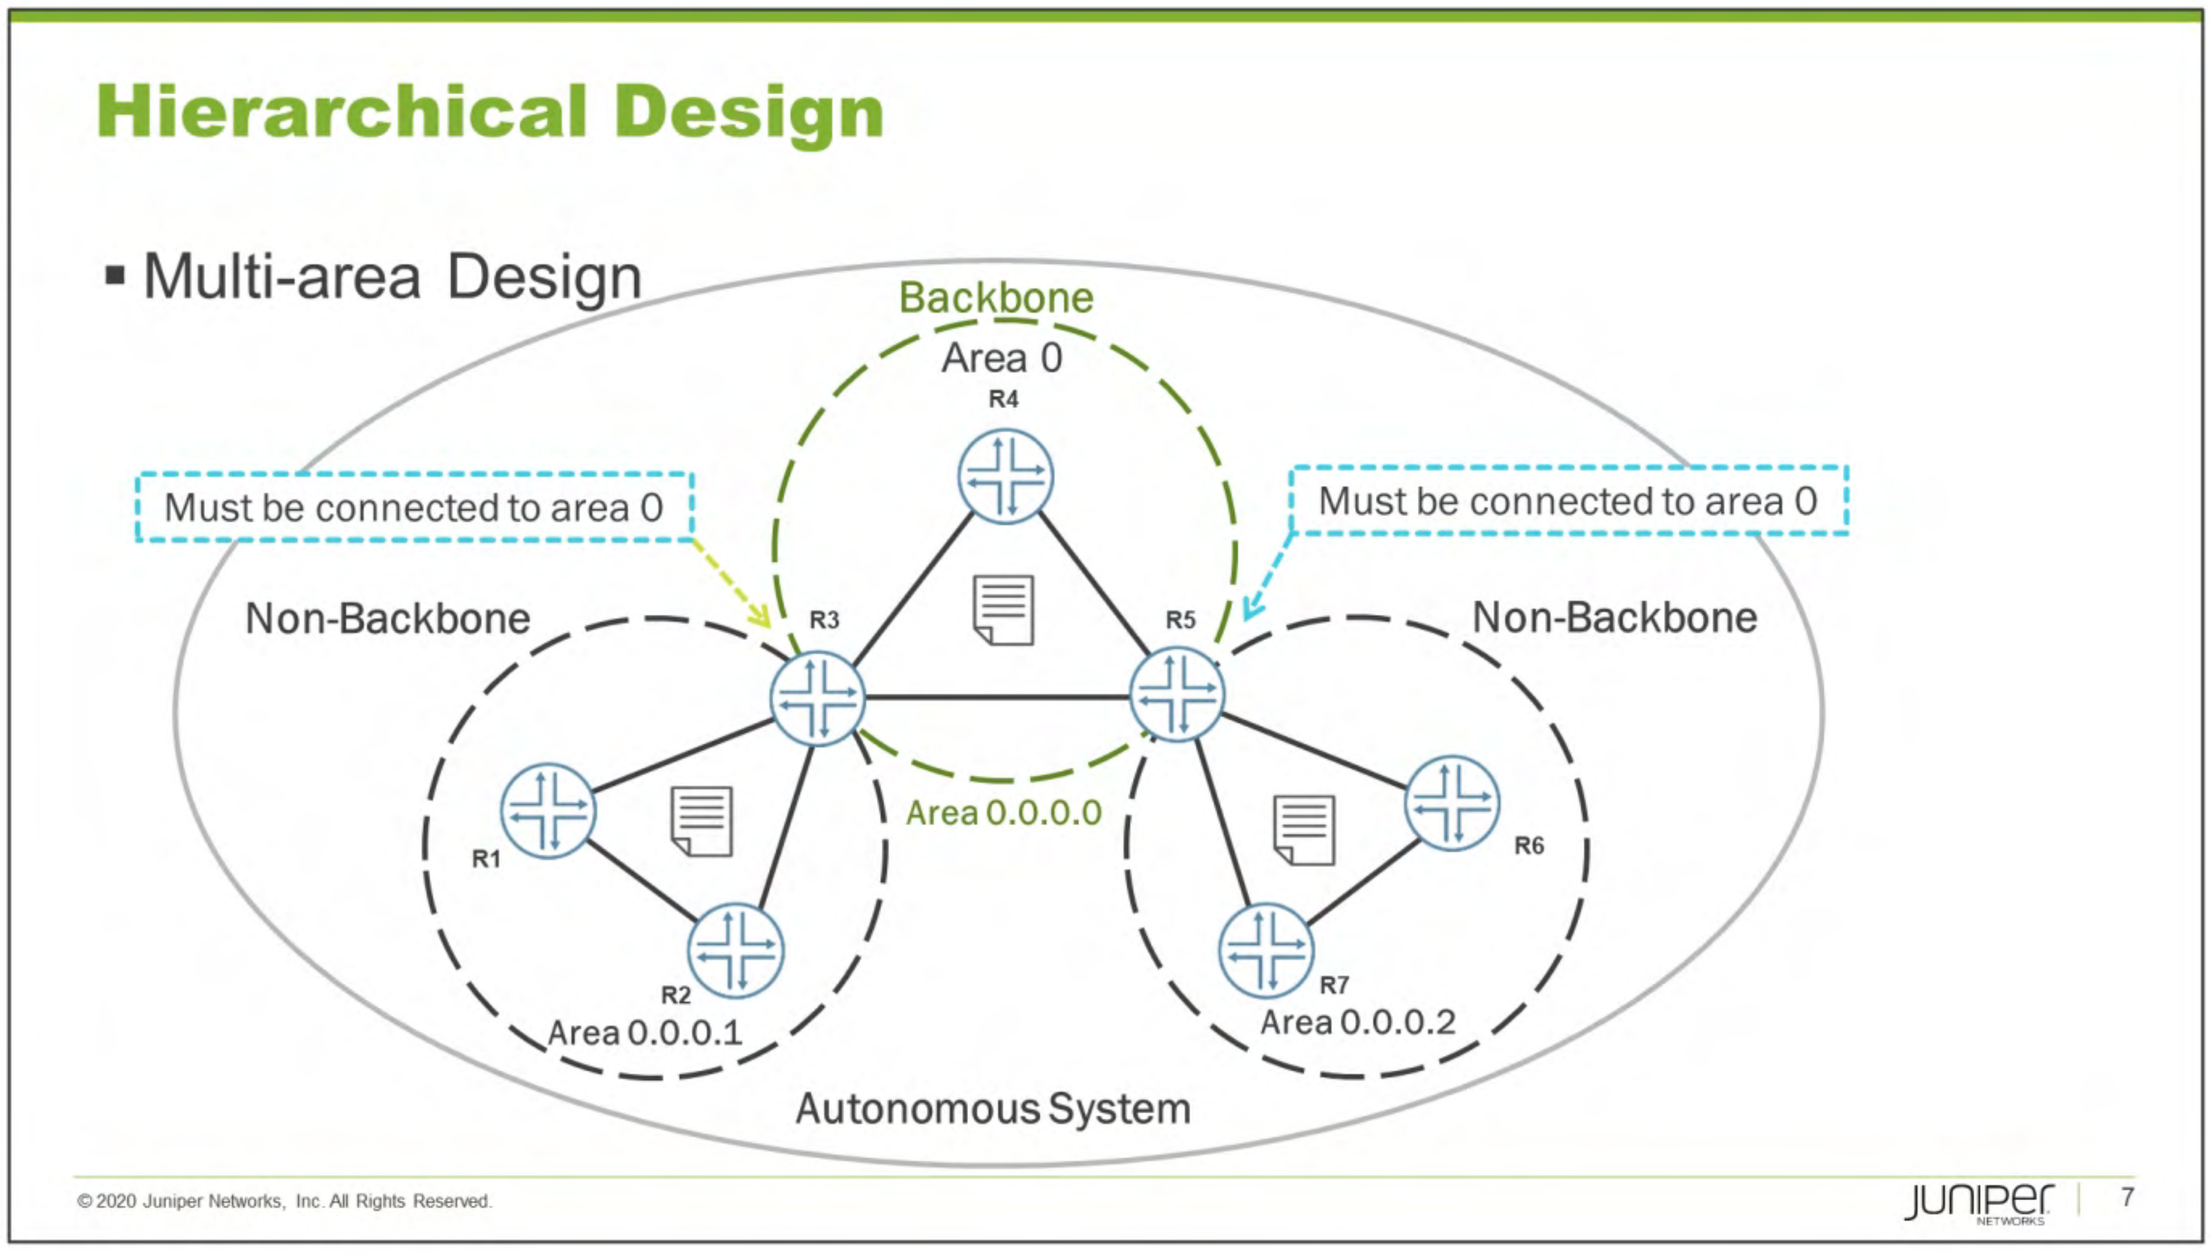

Hierarchical Design

The slide graphically displays the hierarchical nature of OSPF. A backbone area (0.0.0.0) is the connecting point for all other areas. By specification, each area must attach to the backbone in at least one location.

Area 1, Area 2, and Area 3 each contain routers internal to that area as well as a single area border router (ABR). The ABR generates summary LSAs that represent the routes within its area and floods those to the backbone. The ABR is also responsible for generating summary LSAs that represent the backbone routes and injecting those into its attached area.

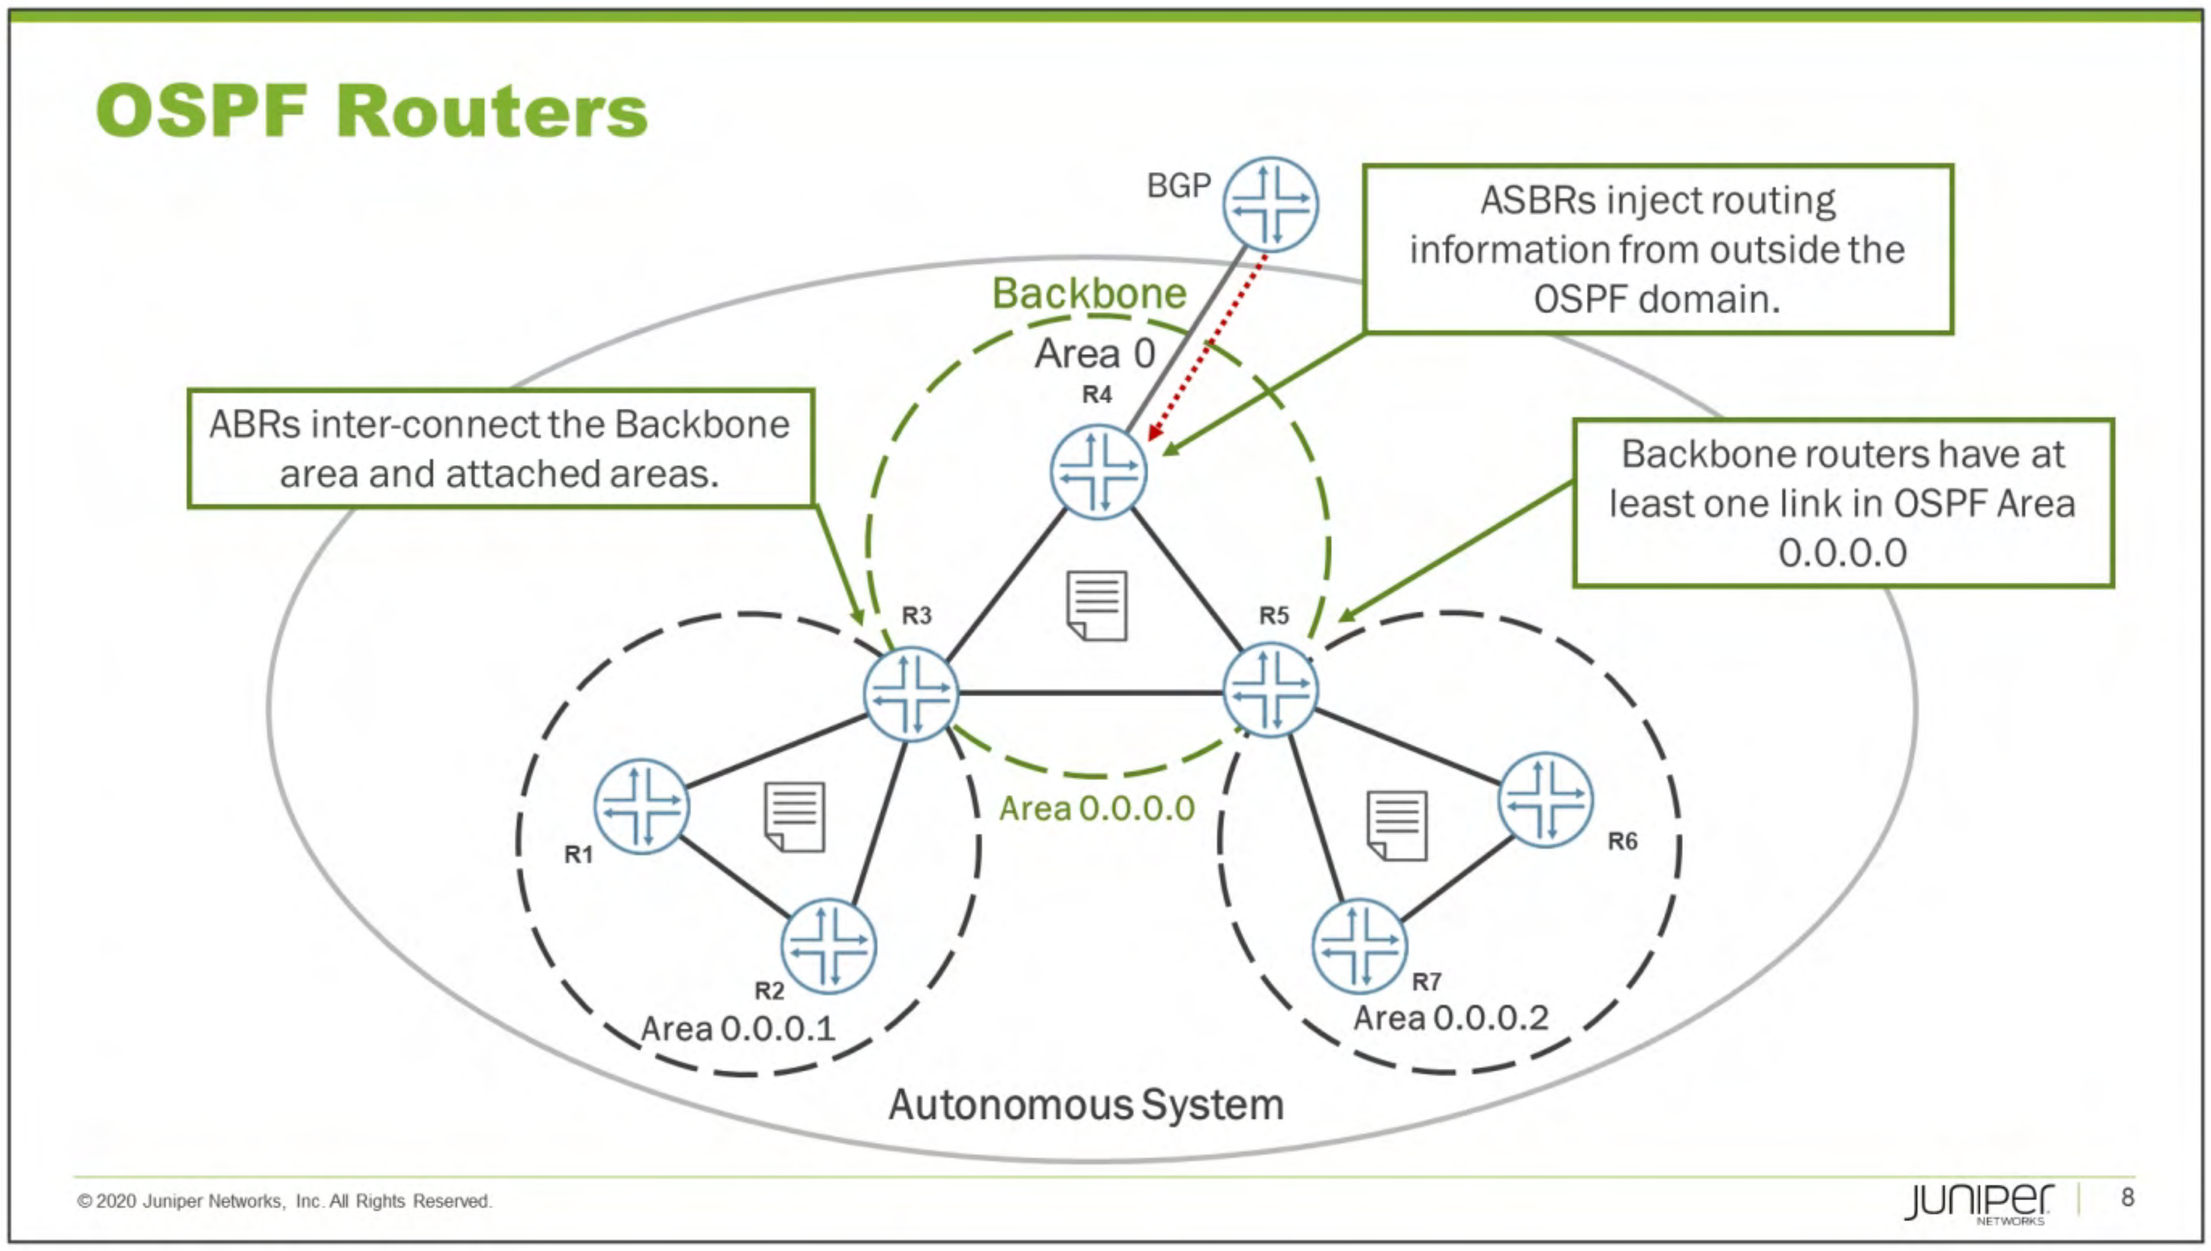

OSPF Routers

OSPF routers can take on a number of different roles within an OSPF domain. The following list shows the common types of OSPF routers along with a brief description:

- *Area border router* (ABR): An OSPF router with links in two or more areas, the ABR is responsible for connecting OSPF areas to the backbone. The ABR generates summary LSAs that represent the routes within its area and floods those to the backbone. The ABR is also responsible for generating summary LSAs that represent the backbone routes and injecting these LSAs into its attached area.
- *Autonomous system boundary router (ASBR)*: An OSPF router that injects routing information from outside the OSPF autonomous system (AS), an ASBR is typically located in the backbone. However, the OSPF specification allows an ASBR to be in other areas as well.
- *Backbone router*: Defined as any OSPF router with a link to Area 0 (the backbone). This router can be completely internal to Area 0 or an ABR depending on whether it has links to other, non-backbone areas.
- *Internal router*: An internal router is an OSPF router with all its links within the same area.

Single Area OSPF Configuration

- Configured at the `[edit protocols ospf]` hierarchy level
- List the interfaces associated with the area:

```
[edit protocols]
user@router# show
ospf {
    area 0.0.0.0 {
        <area options>;
        interface ge-0/0/0.0 {
            interface-type p2p;
            metric 100;
        }
        interface lo0.0;
    }
}
```

Area ID

Interface within Area 0

Interface-specific options

Useful for troubleshooting, IBGP, and tunnel endpoints

© 2020 Juniper Networks, Inc. All Rights Reserved.

Single Area Configuration Example

The slide illustrates the basic hierarchy and syntax used to configure an OSPFv2 single area. OSPFv2 is configured at the `[edit protocols ospf]` hierarchy as shown in the example.

OSPF Interfaces

All logical interfaces associated with the area should be listed under the area.

The loopback interface is not advertised by default on Juniper routers. Best practice recommends that the loopback interface is added to the OSPF configuration (in Area 0) for use in troubleshooting. This practice will also support the configuration of IBGP that peers to loopbacks. The loopback interface is also commonly used for MPLS Label Switch Path (LSP) tunnel endpoints.

OSPF RID

- Each OSPF router selects a 32-bit value to use as its RID
  - Uniquely identifies the router within the network
  - Populated within the LSAs sent out by each router
  - Used by the LSDB to run SPF
  - Two directly connected routers with the same RID will not form an adjacency
  - Two routers with the same RID that are not directly connected will create problems
- You can set the RID explicitly within `[edit routing-options]`

```
[edit routing-options]
user@router# set router-id 192.168.1.1
```

- If the RID is not set explicitly, the address of the primary interface of the router is chosen as the value of the RID.
  - Normally this is the loopback interface address if it has a non-Martian route IPv4 address configured
- In a multi-area network, best practice is to put the loopback in area 0.0.0.0

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Router ID

Each OSPF-speaking router in a network must select a 32-bit value to use as the router ID (RID) of the router. This RID uniquely identifies the router within the OSPF domain. It is transmitted within the LSAs used to populate the LSDB and is used to calculate the shortest-path tree using the method described on a previous slide.

It is very important in a link-state network that no two routers share the same RID value. If two routers share the same RID value, the LSDB might not be consistent on all routers. This inconsistency happens because the RID is one method to verify whether an LSA is already present in the database. Therefore, new critical information from one of the routers is never present in all the routers. In addition, because the Dijkstra calculation uses the RID, it is possible that an individual router might not generate a loop-free shortest path topology. This scenario could have a disastrous effect on IP routing.

Manual Router ID Selection

To avoid possible problems with the OSPF RID, you can set the router ID manually within the Junos configuration. Within the `[edit routing-options]` configuration hierarchy, the **`router-id`** command assigns a 32-bit value to the RID.

Router ID is the Primary Interface

If you do not configure a router ID explicitly, the IP address of the first interface to come online is used. This is usually the loopback interface, in this case the smallest non-127/8 IP address configured on the loopback interface is used as RID. Otherwise, the first hardware interface with an IP address is used.

Because the OSPF routing process uses the primary router interface as the RID, its selection can have a very important consequence. If the router addressing changes, the OSPF RID might also change, and a duplicate RID situation might arise. The hazards of this type of situation were outlined in the previous sections.

What If...?

- ...you have multiple addresses configured on an interface, but do not want to advertise all of them into OSPF?

```
[edit interfaces]
user@router# show lo0
unit 0 {
    family inet {
        address 172.16.1.1/32 {
            primary;
        }
        address 172.16.2.1/32;
        address 172.16.3.1/32;
        address 172.16.4.1/32;
    }
}
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface lo0.0;
}
```

```
[edit protocols ospf area 0.0.0.0]
user@router# run show ospf database detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq
Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000003
22  0x22 0xea97  72
  bits 0x0, link count 4
  id 172.16.1.1, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  id 172.16.4.1, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  id 172.16.3.1, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  id 172.16.2.1, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
```

Specifying the interface *name* will result in *all* of its configured addresses being advertised into OSPF.

© 2020 Juniper Networks, Inc. All Rights Reserved.

What If...?

In the slide, interface lo0.0 has four addresses configured on it. Typically, you would put interface lo0.0 into OSPF using the **`set area <area-id> interface <interface-name>`** command. However, doing this will advertise four stub networks in the router LSA: one for each of the configured addresses on interface lo0.0. You might come across a situation such as this for which you *do not want* to advertise all of the addresses into OSPF. A solution to this problem is offered on the next page.

The Solution...

- Specify the interface address instead of the interface name

```
[edit interfaces]
user@router# show lo0
unit 0 {
    family inet {
        address 172.16.1.1/32 {
            primary;
        }
        address 172.16.2.1/32;
        address 172.16.3.1/32;
        address 172.16.4.1/32;
    }
}
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface 172.16.1.1;
}
```

```
[edit]
user@router# run show ospf database detail

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq
Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000005
10  0x22 0xf2ff  36
  bits 0x0, link count 1
  id 172.16.1.1, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
```

Specifying an interface *address* will result in only that address being advertised into OSPF.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Then...

The slide shows a solution to our problem. By specifying the specific interface address to advertise, only one network from lo0.0 interface is advertised into OSPF.

Multi-Area OSPF Configuration

- Configured at the `[edit protocols ospf]` hierarchy level
- Each area is listed along with the interfaces associated with that area:

```
[edit protocols]
user@router# show
ospf {
    area 0.0.0.0 {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
        interface ge-0/0/2.0;
        interface lo0.0;
    }
    area 0.0.0.1 {
        nssa;
        interface ge-0/0/3.0;
    }
    area 0.0.0.2 {
        stub;
        interface ge-0/0/4.0;
    }
}
```

Area ID

Area-specific options

© 2020 Juniper Networks, Inc. All Rights Reserved.

Multi-Area Configuration Example

The slide illustrates the basic hierarchy and syntax used to configure the OSPFv2 multiarea. As in the single area example previously shown, the configuration is performed at the `[edit protocols ospf]` hierarchy.

OSPF Interfaces

All logical interfaces associated with the particular area should be listed under that area. Remember to include the loopback interface, if it should be advertised.

Agenda: OSPF

- OSPFv2 Review
- →Link-State Advertisements
- Protocol Operations
- OSPF Authentication
- OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

Link-State Advertisements

The slide highlights the topic we discuss next.

LSA Types

- LSA types:
  - Router LSAs (Type 1)
  - Network LSAs (Type 2)
  - Summary LSAs (Type 3)
  - ASBRSummary LSAs (Type 4)
  - AS external LSAs (Type 5)
  - Group membership LSAs (Type 6)
  - NSSA LSAs (Type 7)
  - External attributes LSAs (Type 8)
  - Opaque LSAs (Types 9, 10, and 11)
- Each LSA type describes a portion of the OSPF routing domain
- LSAs 6, 8, and 11 are not supported

© 2020 Juniper Networks, Inc. All Rights Reserved.

LSA Types

The following list provides the details of the LSA types:

- *Router LSA (Type 1)*: Sent by each router to describe its individual links and their status.
- *Network LSA (Type 2)*: Sent by the DR on the broadcast network.
- *Summary LSA (Type 3)*: Sent by an ABR to describe routes or other information from one area into another.
- *ASBRSummary LSA (Type 4)*: Sent by an ABR to describe ASBRs in other OSPF areas.
- *AS external LSA (Type 5)*: Sent by an ASBR to describe routes external to the OSPF domain.
- *Group membership LSA (Type 6)*: Used to support Multicast Open Shortest Path First (MOSPF).
- *NSSA LSA (Type 7)*: Sent by an ASBR in a not-so-stubby area (NSSA) to describe routes external to the OSPF domain.
- *External attributes LSA (Type 8)*: Used to *tunnel* attributes used by other routing protocols through OSPF.
- *Opaque LSAs (Types 9, 10, 11)*: Used to transmit information not currently supported by other LSA types. Examples include graceful restart and traffic engineering information.

LSA Functions

Each LSA type has a specific function within the OSPF domain. They describe a portion of the topology used by the SPF algorithm to supply routing information to a routing table. We discuss each LSA in more detail on future slides.

Each LSA has a maximum age of 3600 seconds (1 hour) associated with it. This maximum provides a mechanism for removing stale information from the database. To ensure that its LSAs are up to date, each OSPF router periodically refreshes them prior to reaching the maximum age limit. The Junos operating system performs this refresh every 3000 seconds (50 minutes).

LSAs Not Supported

The Junos OS does not support some of the LSA types listed here. These LSAs are the group membership (Type 6), external attributes (Type 8), and the domain-scope opaque (Type 11) advertisements.

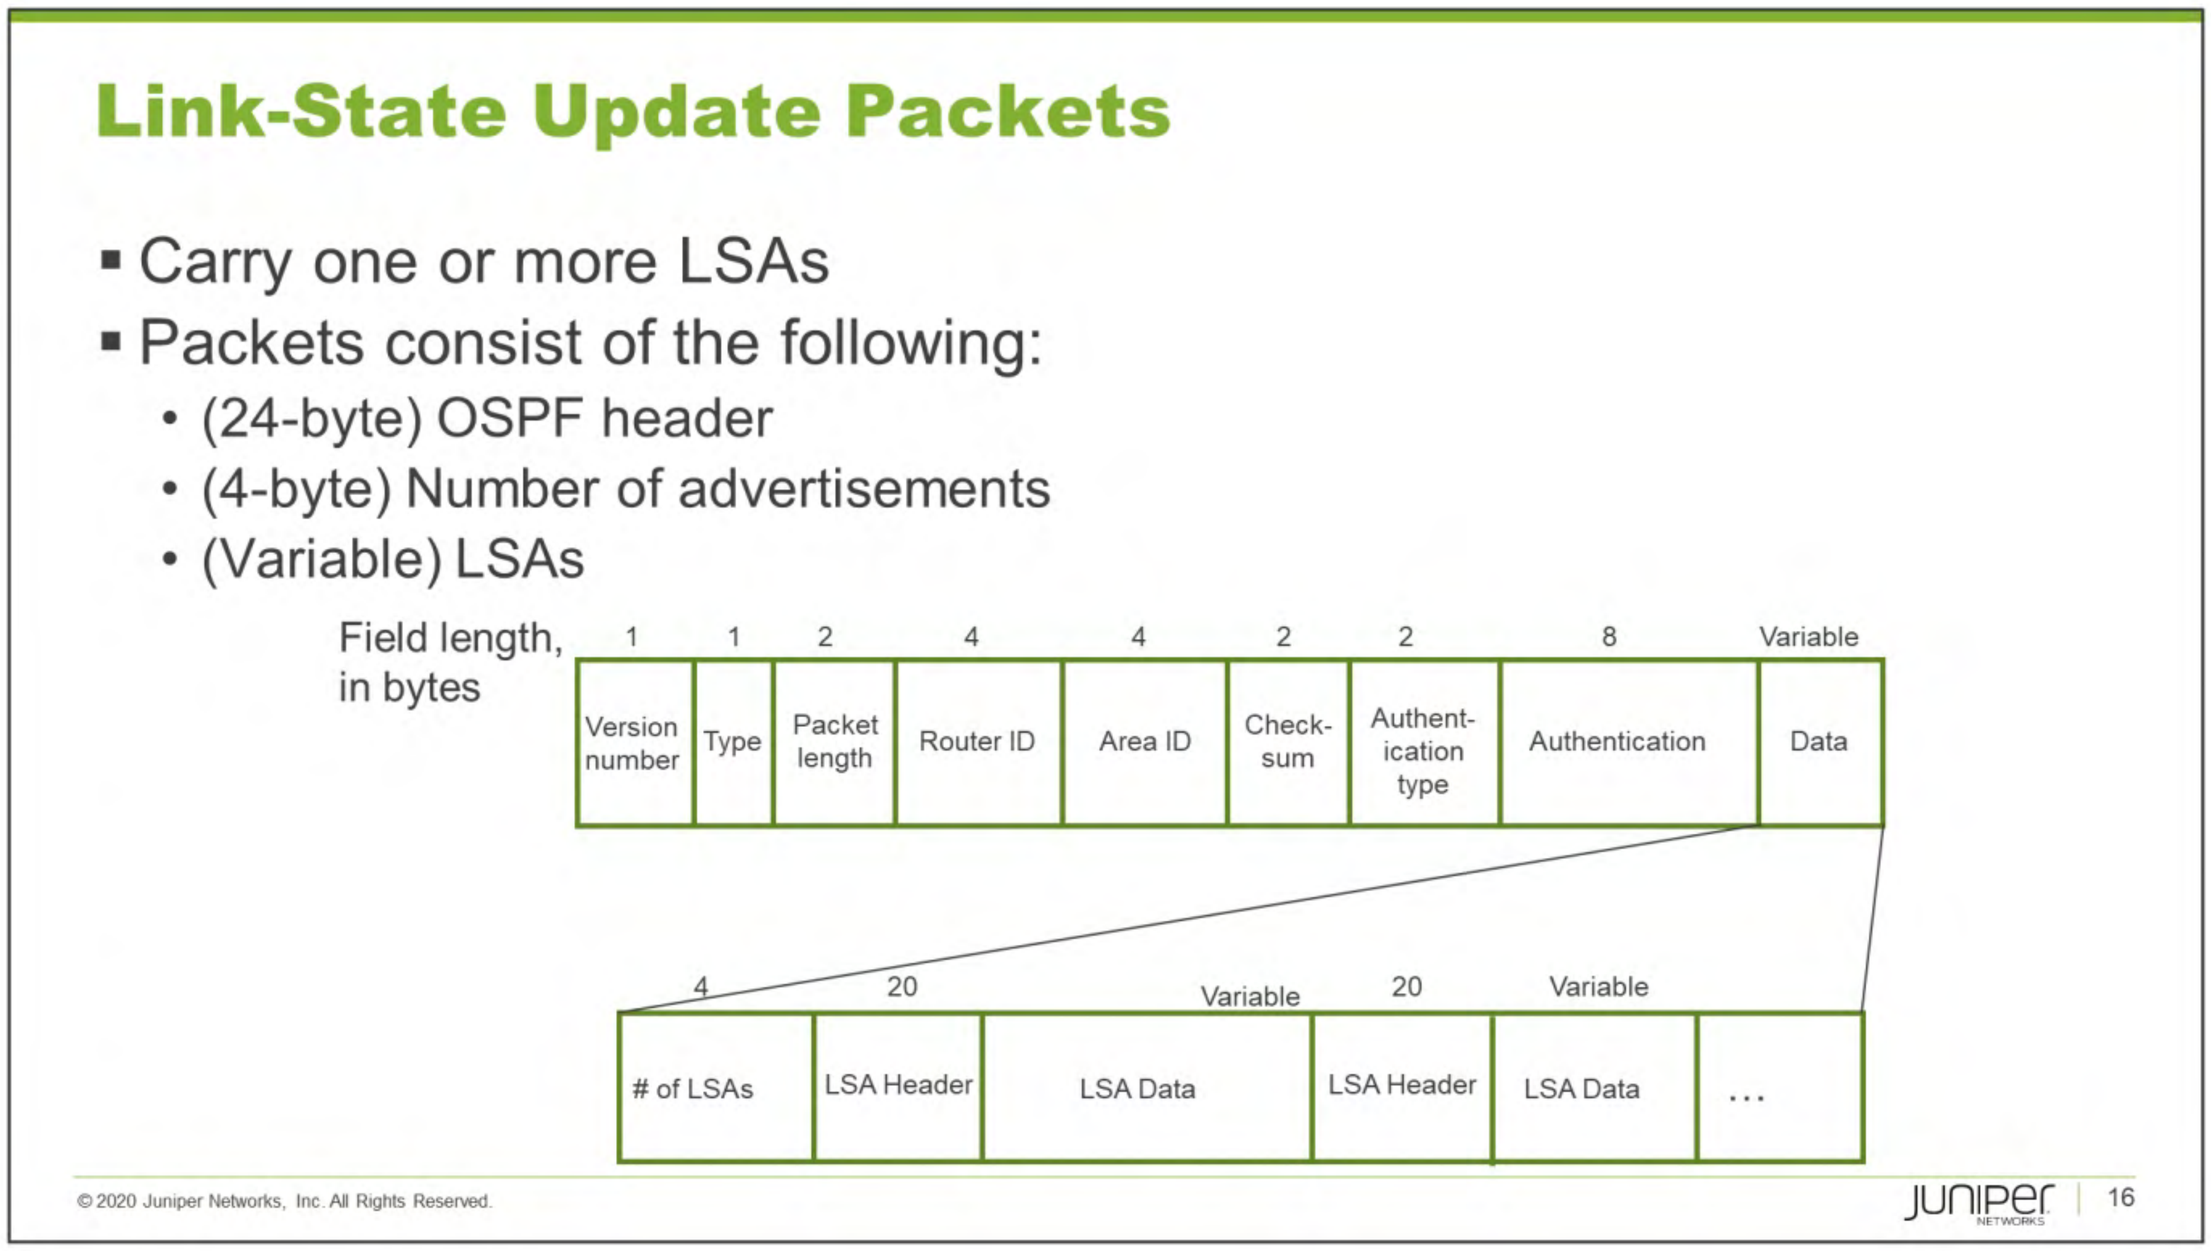

Multiple LSAs in a Single Update

Each link-state update packet sent into the network by an OSPF-speaking router can carry multiple different LSAs. This ability saves network resources by allowing routers to use transmission and routing capacity more efficiently.

Link-State Update Format

The graphic illustrates the format of the link-state update packet. Following the standard OSPF header, the update packet contains a 4-byte field used to notify other routers as to the number of advertisements stored in the update message. This field is followed by the LSAs themselves, each with its own header and format.

LSA Header

- 20 bytes of information that identify the LSA uniquely and consist of:
  - Link-state age – Count up timer
  - Options – Indicates the optional capabilities support on this router
    - P bit (position 5) set in all NSSA external LSAs
    - E bit (position 7) set in all external LSAs
  - Link-state type – LSA Type
  - Link-state ID – Varies based on LSA Type
  - Advertising router – Router ID of originating router
  - Link-state sequence number – Determines if LSA has changed
  - Link-state checksum – LSA integrity check
  - Length

© 2020 Juniper Networks, Inc. All Rights Reserved.

LSA Header

The following list details the information contained in the LSA header:

- *LS age (2 bytes)*: Measured in seconds, the link state age is the time since the LSA was first originated. Each router increments this field prior to reflooding the LSA.
- *Options (1 byte)*: Indicates support for OSPF options. Within the context of an individual LSA, the E bit (position 7) is set in all External LSAs and the P bit (position 5) is set in all NSSA external LSAs.
- *LS type (1 byte)*: Encodes the specific LSA type.
- *Link-state ID (4 bytes)*: Describes various portions of the OSPF domain. Each LSA type uses this field in a different manner. For Router LSA, it will be equal to RID of the router. For Network LSA, this field is set equal to the DR's IP address. For ASBRSummary LSA, it is equal to ASBR's RID. For Summary, External and NSSA LSAs, link-state ID is set equal to the advertised IP subnet.
- *Advertising router (4 bytes)*: The router ID of the router that first originated the LSA.
- *LS sequence number (4 bytes)*: Verifies that each router has the most recent version of an LSA. This field is incremented each time a new version is generated. Values range from `0x80000000` to `0x7FFFFFFF`.
- *LS checksum (2 bytes)*: The checksum of the entire LSA contents, minus the link state age field. This field is used to ensure data integrity in the LSDB.
- Length (2 bytes): The entire length of the individual LSA, including the header.

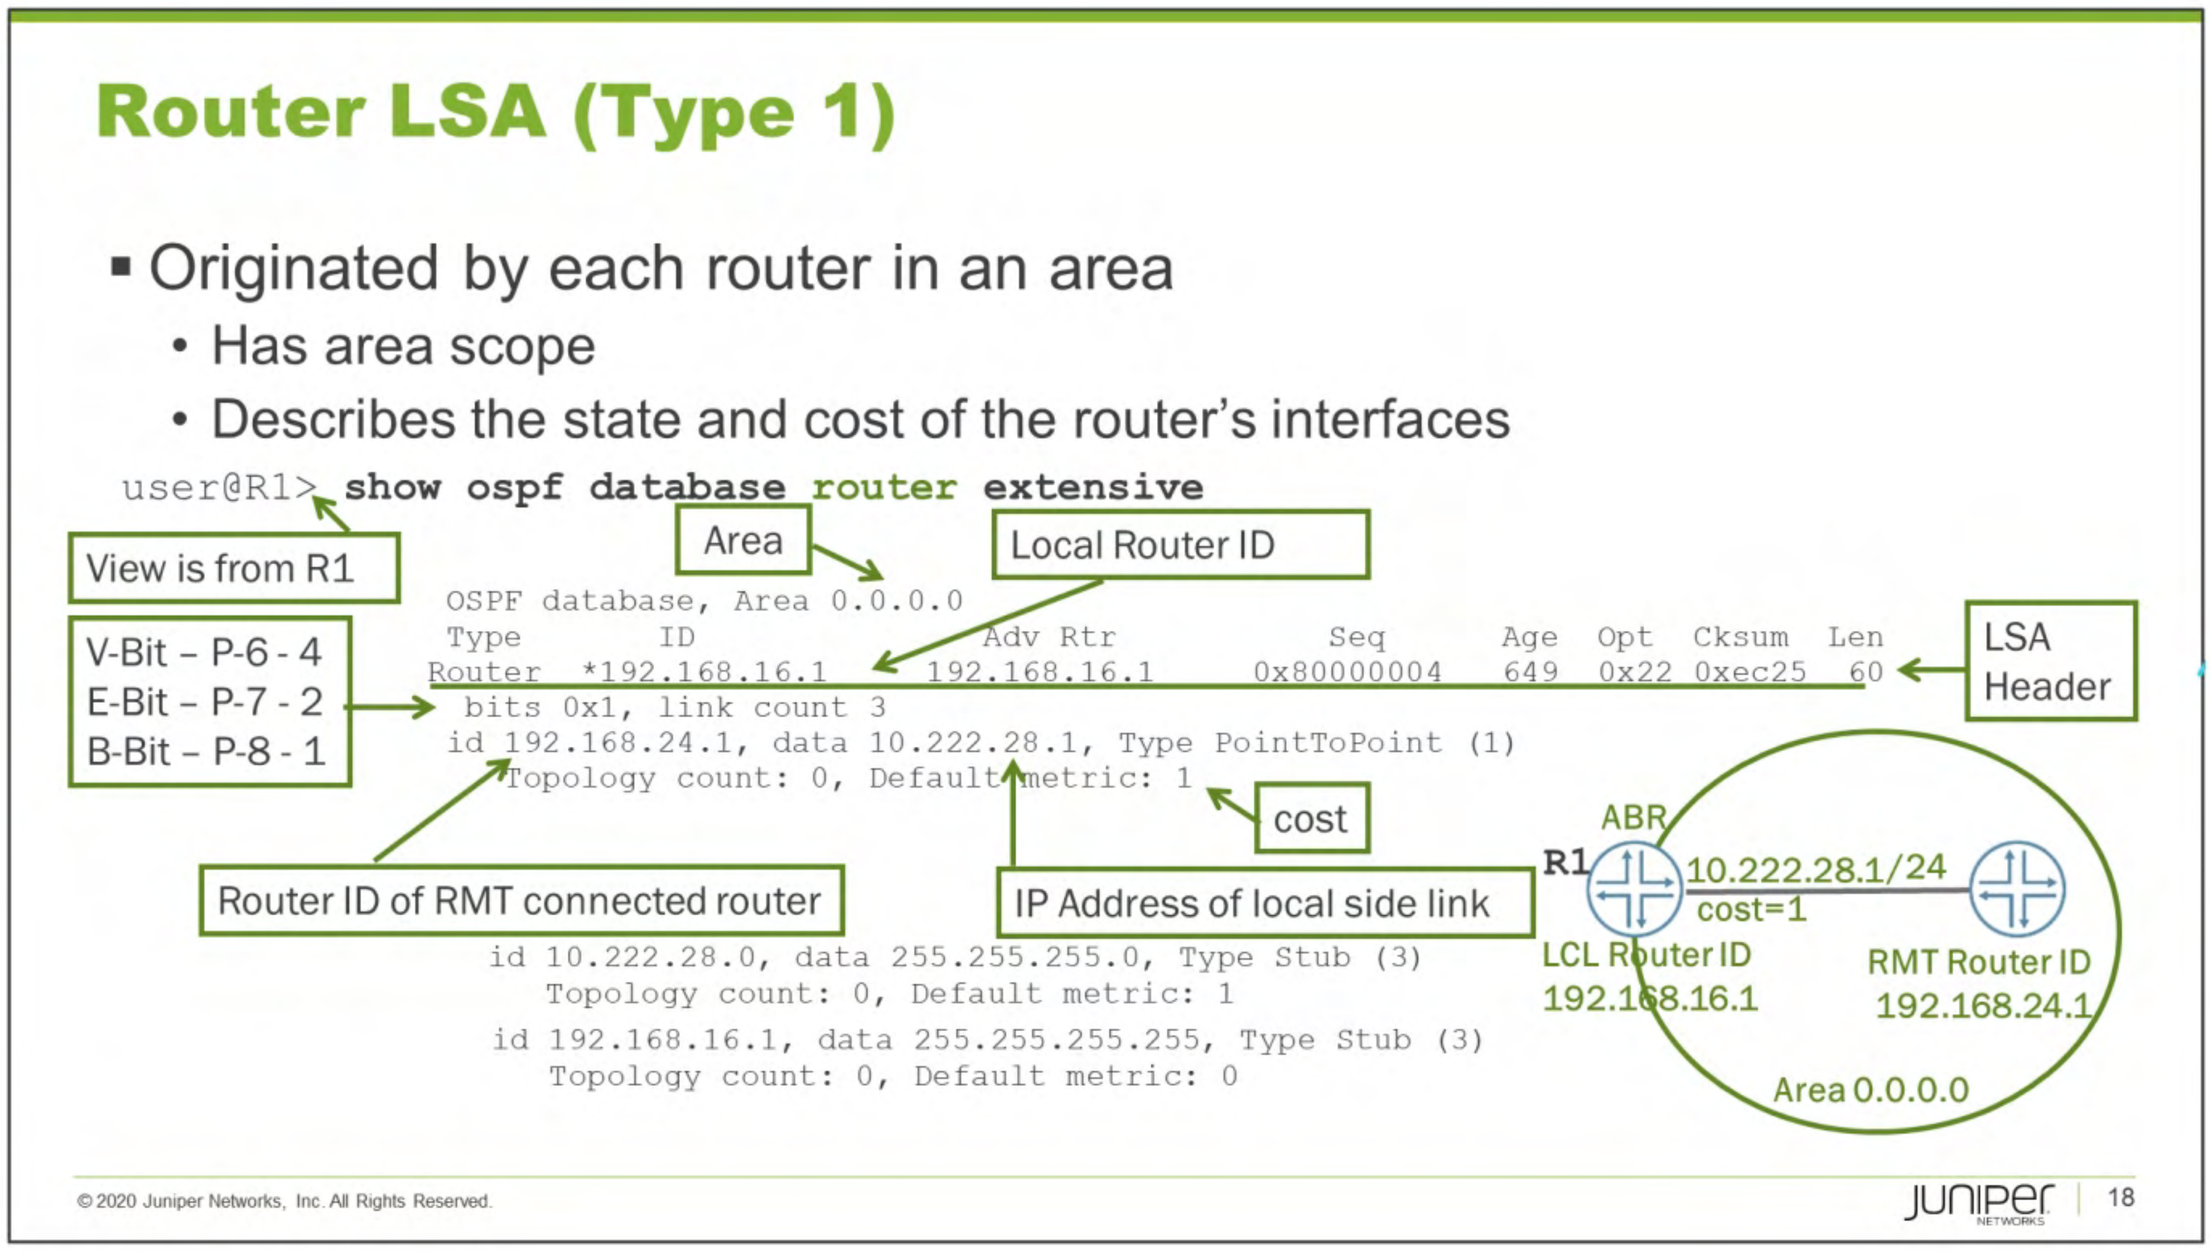

Router LSA

Each OSPF-speaking router generates a Type 1 LSA to describe the status and cost (metric) of all interfaces on the router. This LSA is flooded to each router in the OSPF area. It is defined as having an area scope; therefore, it is not flooded across an area boundary. In addition to the standard LSA header, the router LSA also contains the following fields:

- *V, E, and B bits (1 byte)*: Following five bits set to a value of 0, the V, E, and B bits represent the characteristics of the originating router. The V bit is set when a virtual link is established. An ASBR sets the E bit. An ABR sets the B bit.
- *Number of links (2 bytes)*: This value gives the total number of links represented by the following set of fields.
- *Link ID (4 bytes)*: This field represents to what the far side of the link is connected. It is used in conjunction with the link type field.
- *Link data (4 bytes)*: This field represents to what the near side of the link is connected. It is used in conjunction with the link type field.
- *Link type (1 byte)*: This field describes the type of link.
- *Number of Multi Topology IDs (MT-IDs) (1 byte)*
- *Metric (2 bytes)*: This field provides the cost to transmit data out of the interface.
- *MT-ID* (1 byte): Renamed TOS field to represent the multi topology ID.
- *MT-ID metric (2 bytes)*: Renamed TOS metric field to represent the Multi Topology (MT) metric.
- *Additional data (4 bytes)*: This field is unused.

Router LSA Example

You can see the details of the router LSA in the example on the slide. By using the keyword **`extensive`**, you can see each of the LSA fields. You can gather some important details about the local router by closely examining the LSA:

- The LSA is shown in the database within the Area 0.0.0.0 section. Recall that a router LSA has area scope, so a separate LSA is generated for each area representing the links only within that area.
- The LSA header tells us this router LSA is 649 seconds old and has a length of 60 bytes.
- This router LSA was originated by the same router from which the capture was taken. Notice the asterisk (`*`) next to the link-state ID value of `192.168.16.1`.
- This router is an ABR. We see this by the setting of bits 0x1. Recall that position 7 (0x2) is for the E bit, which is set when the originating router is an ASBR. Bit position 8 (0x1) is for the B bit, which is set when the originating router is an ABR. Since the value is 0x1, it is an ABR. If the value was 0x3, the router would be an ABR and an ASBR.
- This router has three links connected to Area 0, which can be determined in two ways. First, the link count field is set to a value of 3 and secondly, the three individual links are listed below the LSA header.
- A single point-to-point link exists, and two links are connected to stub networks.

Link ID and Link Data Fields

- Interpretation depends on value of the link type field

Link Type	Link ID	Link Data
Point-to-point (Type 1)	Neighbor's RID	Local router's interface IP address
Transit (Type 2) supports M/A network links	DR's interface IP address	Local router's interface IP address
Stub (Type 3) supports P2P, passive, and lookback Interfaces	Network number	Subnet mask
Virtual link (Type 4)	Neighbor's RID	Local router's interface IP address

© 2020 Juniper Networks, Inc. All Rights Reserved.

Link ID and Link Data Fields

The information in the router LSA's link ID and link data fields is associated with the type of link OSPF is operating. The following link types are supported:

- *Point-to-point (Type 1)*: On a point-to-point interface, an OSPF router always forms an adjacency with its peer over an unnumbered connection. As such, the link ID field contains the neighbor's router ID. The link data field contains the IP address of the interface on the local router.
- *Transit (Type 2)*: A connection to a broadcast segment is always noted as a transit link. The link ID field contains the interface IP address of the segment's DR. The link data field contains the interface IP address of the local router.
- *Stub (Type 3)*: A router advertises a stub network when a subnet does not connect to any OSPF neighbors. Advertising a stub network occurs for the loopback interface and any passive interfaces. In addition, the IP subnet for any point-to-point interface is advertised as a stub because the adjacency was formed over an unnumbered interface. The link ID field for a stub network contains the IP network number and the link data field contains the subnet mask.
- *Virtual link (Type 4)*: A virtual link operates between an ABR connected to Area 0 and an ABR that is not connected to Area 0. Once established, the virtual link appears in the Area 0 router LSA of each endpoint. The link ID field contains the neighbor's router ID, and the link data field contains the interface IP address of the local router.

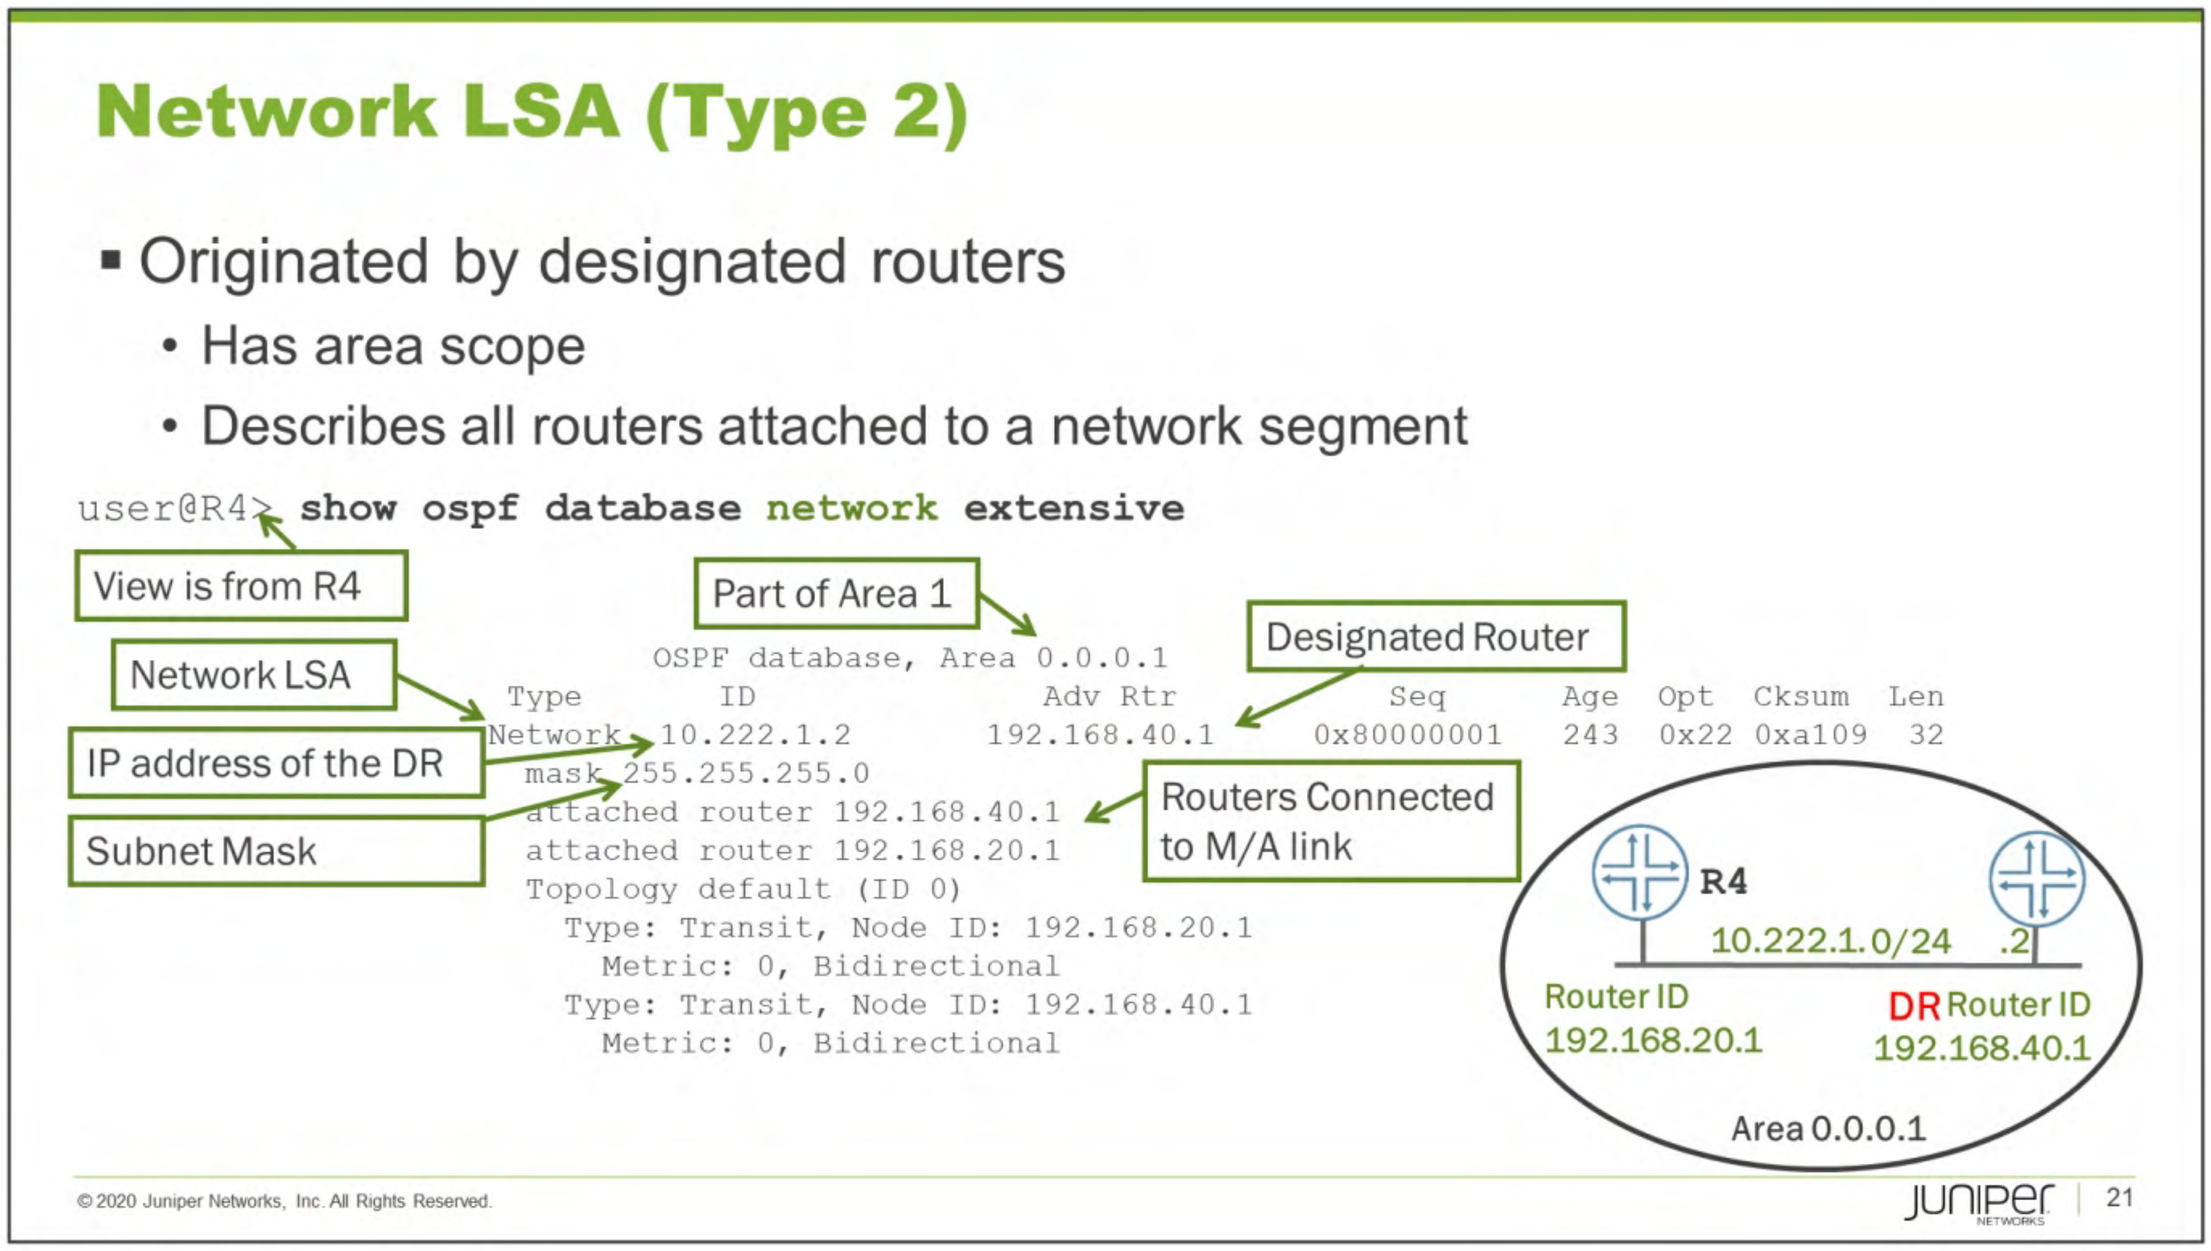

Network LSA

Each OSPF router elected to be the DR on a broadcast link generates a Type 2 LSA. This LSA lists each router connected to the broadcast link, including the DR itself. In addition to the standard LSA header, the network LSA also contains the following fields:

- *Network mask (4 bytes)*: This field denotes the IP subnet mask for the interface connected to the broadcast network.
- *Attached router (4 bytes)*: This field is repeated for each router connected to the broadcast network. The value of each instance is the router ID of the attached routers.

Network LSA Example

You can see the details of the network LSA in the example on the slide. By using the keyword **`extensive`**, you can see each of the LSA fields. You can gather some important details about this network by closely examining the LSA:

- The DR for this network is 10.222.1.2, which you can determine by examining the link-state ID field. This field lists the DR's IP address. The Router ID of that device is 192.168.40.1, which is determined from the Adv Rtr field.
- Because only two instances of the attached router field are present, you can safely deduce that only two routers are connected to the link. The router IDs of those two routers are 192.168.20.1 and 192.168.40.1.
- The router that advertised this LSA was 192.168.40.1 and that is not this router (no asterisk) and there are only two routers on the link, that must make this router with a Router ID of 192.168.20.1.
- A broadcast segment connects the two routers together whose address is 10.222.1.0/24, which is determined from the mask field.
- To determine the IP address of the local device you would need to find the router LSA on this network with a Link Type of Transit.

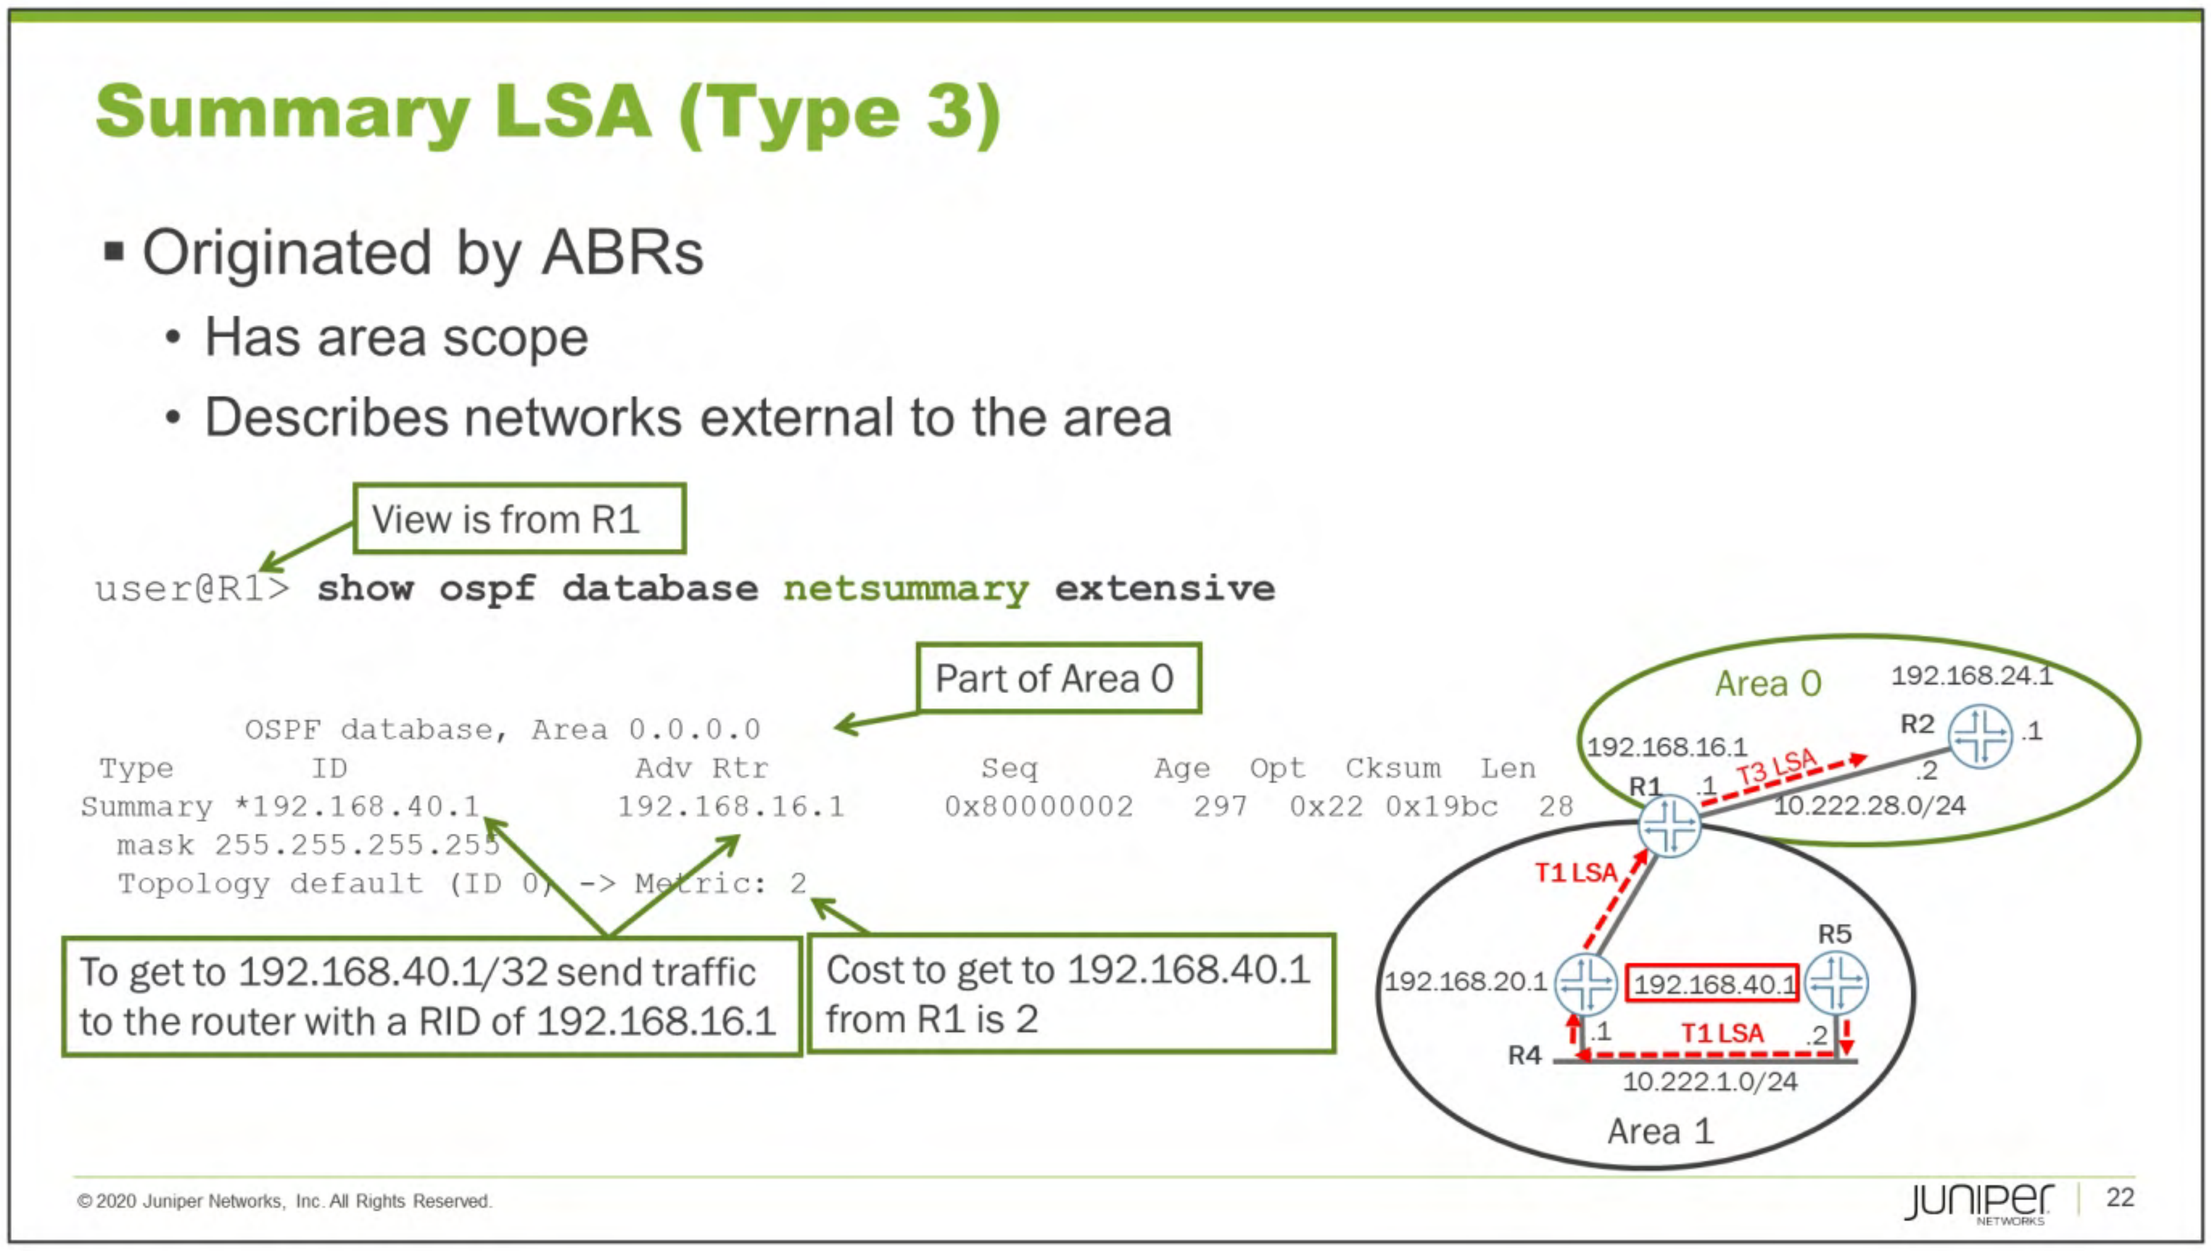

Summary LSA

Each ABR that transmits information from one OSPF area into another generates a Type 3 LSA to describe that information. This LSA is flooded to each router in the OSPF area. This LSA has an area scope; therefore, it is not reflooded across the area boundary by another ABR. Instead, the receiving ABR generates a new Type 3 LSA describing the link and floods it into the adjacent area.

In addition to the standard LSA header, the summary LSA also contains the following fields:

- *Network mask (4 bytes)*: This field represents the subnet mask associated with the network advertised. It is used in conjunction with the link-state ID field, which encapsulates the network address in a Type 3 LSA.
- *Metric (3 bytes)*: This field provides the cost of the route to the network destination. When the summary LSA is representing an aggregated route (using the **`area-range`** command), this field is set to the largest current metric of the contributing routes.
- *MT-ID (1 byte)*: This field represents the MT-ID value used in a Multi Topology configuration.
- *MT-ID metric (3 bytes)*: This field represents the MT-ID metric used in a Multi Topology configuration.

Summary LSA Example

You can see the details of the summary LSA in the example on the slide. By using the keyword **`extensive`**, you can see each of the LSA fields. You can gather some important details about the local router by closely examining the LSA:

- We are currently on R1, and this router created the summary LSA. It represents the network LSA generated from the R5 router (The 192.168.40.1 was the loopback of R5.).
- This LSA will be advertised into Area 0 and is being advertised by R1. (Adv Rtr is the Router ID of R1 which is 192.168.16.1)
- This LSA includes a metric cost of 2, which is the cost to get from R1 to R5.

External LSA (Type 5)

- Originated by ASBRs
  - Has domain scope - Forwarded unchanged by ABRs
  - Describes networks external to the OSPF domain

View is from R1

```
user@R1> show ospf database external extensive

    OSPF AS SCOPE link state database
    Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
   Extern   203.0.113.0      192.168.40.1     0x80000004    21  0x22 0x9649  36
     mask 255.255.255.0
     Topology default (ID 0)
       Type: 2, Metric: 0, Fwd addr: 0.0.0.0, Tag: 0.0.0.0
```

Not stored as part of any area

E-Bit

To get to the 203.0.113.0/24 network, send traffic to the router with a RID of 192.168.40.1

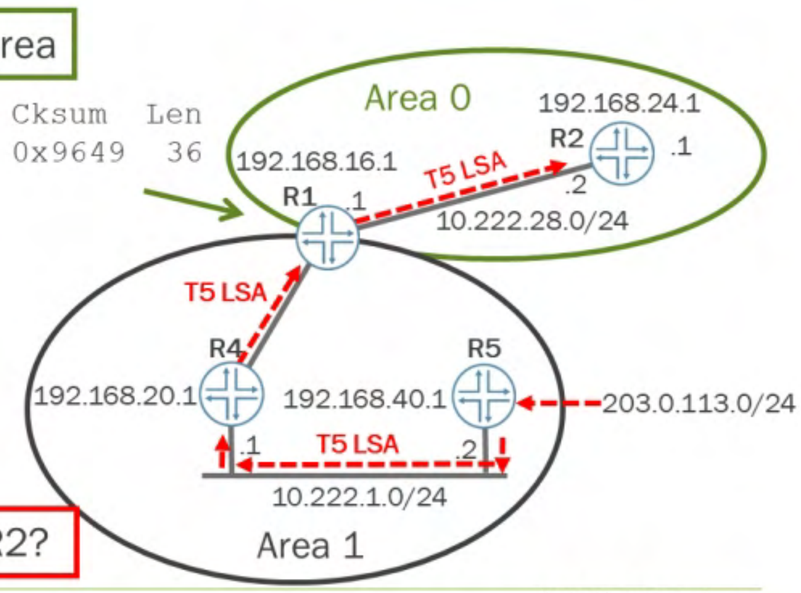

Is there a problem getting to the router with a RID of 192.168.40.1 from R2?

© 2020 Juniper Networks, Inc. All Rights Reserved.

AS External LSA

Each ASBR generates a Type 5 LSA to advertise any networks external to the OSPF domain. This LSA is flooded to each nonstub router in the entire OSPF domain. In addition to the standard LSA header, the AS external LSA also contains the following fields:

- *Network mask (4 bytes)*: This field represents the subnet mask associated with the network advertised. It is used in conjunction with the link-state ID field, which encapsulates the network address in a Type 5 LSA.
- *E bit (1 byte)*: The E bit determines the type of external metric represented by the metric field. It is followed by 7 bits, all set to 0 to make up the entire byte. A value of 0, the default value, indicates that this is a Type 2 external metric. Thus, any local router should use the encoded metric as the total cost for the route when performing an SPF calculation. A value of 1 indicates that this is a Type 1 external metric. Therefore, the encoded metric of the route should be added to the cost to reach the advertising ASBR. This additive value then represents the total cost for the route.
- *Metric (3 bytes)*: This field represents the cost of the network as set by the ASBR.
- *Forwarding address (4 bytes)*: This field provides the address toward which packets should be sent to reach the network. A value of 0.0.0.0 represents the ASBR itself.
- *External route tag (4 bytes)*: This 32-bit value field can be assigned to the external route. OSPF does not use this value, but it might be interpreted by other protocols.
- *MT-ID and MT-ID metric fields (4 bytes)*: These fields represent the MT-ID and MT-ID metric values used in a Multi Topology configuration.

You can see the details of the AS external LSA in the example on the slide. By using the keyword **`extensive`**, you can see each of the LSA fields. You can gather some important details about the local router by examining the LSA closely:

- The external LSDB lists all the LSAs. These LSAs have a domain scope do not belonging to a spcific OSPF area.
- This router is generating the LSA.
- This route is a Type 2 External LSA meaning any local router should use the encoded metric as the total cost for the route when performing an SPF calculation.

ASBR Summary LSA (Type 4)

- Originated by ASBRs
  - Has area scope - Created by ABRs
  - Describes networks external to the OSPF domain

View is from R1

```
user@R1> show ospf database asbrsummary extensive

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
ASBRSum *192.168.40.1     192.168.16.1     0x80000001   833  0x22 0x9795  28
  mask 0.0.0.0
  Topology default (ID 0) -> Metric: 2
```

Sent into adjacent Area

To get to the ASBR router with a RID of 192.168.40.1 send traffic to the router with RID of 192.168.16.1

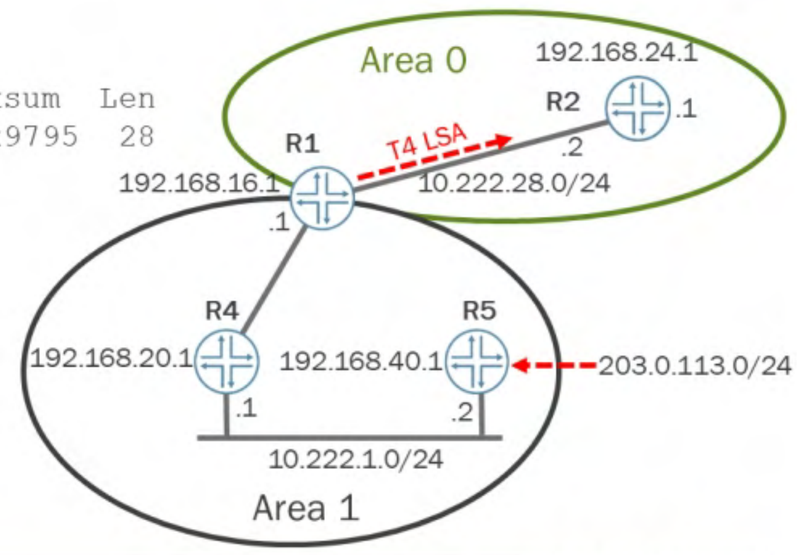

© 2020 Juniper Networks, Inc. All Rights Reserved.

ASBR Summary LSA

Each ABR that must transmit information about an ASBR from one OSPF area into another generates a Type 4 LSA to describe that information. This LSA is flooded to each router in the OSPF area. It is defined as having an area scope; therefore, another ABR does not reflood it across the area boundary. In addition to the standard LSA header, the ASBR summary LSA also contains the following fields:

- *Network mask (4 bytes)*: This field has no meaning in a Type 4 LSA and is set to 0.0.0.0. The address of the ASBR is encoded in the link-state ID field.
- *Metric (3 bytes)*: This field provides the cost of the route to the ASBR.
- *MT-ID (1 byte)*: This field represents the MT-ID value used in a Multi Topology configuration.
- *MT-ID metric (3 bytes)*: This field represents the MT-ID metric used in a Multi Topology configuration.

ASBR Summary LSA Example

You can see the details of the ASBR summary LSA in the example on the slide. By using the keyword **`extensive`**, you can see each of the LSA fields. You can gather some important details about the local router by examining the LSA closely.

- This router is an ABR because it contains a database for Area 0.0.0.0.
- The ASBR being advertised is 192.168.40.1. You can see this information coded in the link-state ID field. Because the router ID of the ASBR is a full 32-bit value, the network mask is not needed and is set to a value of 0.0.0.0.
- The metrics to each of the ASBRs are encoded in the appropriate field in the LSA. As with the summary LSA (Type 3), the local router must add the cost to reach a remote ABR to the encoded metric prior to calculation of the SPF algorithm.
- Router 192.168.40.1 has an export policy configured within OSPF, which means that it has set the E bit in its router LSAs.
- Based on the E bit setting in the router LSAs, the ABR generates an ASBR summary LSA across the area boundaries. In our case, we see the Type 4 LSA within Area 0.

NSSA LSA (Type 7) (1 of 2)

- Originated by an ASBR within the NSSA area
  - Has same format as an AS external LSA (Type 5)
  - Has area scope
  - Describes networks external to the OSPF domain

View is from R5

```
user@R5> show ospf database nssa extensive

OSPF link state database, Area 0.0.0.1
Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
NSSA     203.0.113.0     192.168.40.1     0x80000001   514  0x28 0x59b1  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 192.168.40.1, Tag: 0.0.0.0
```

Part of Area 1

Uses the Fwd addr field

To get to the 203.0.113.0/24 network send traffic to the router with Router ID of 192.168.40.1

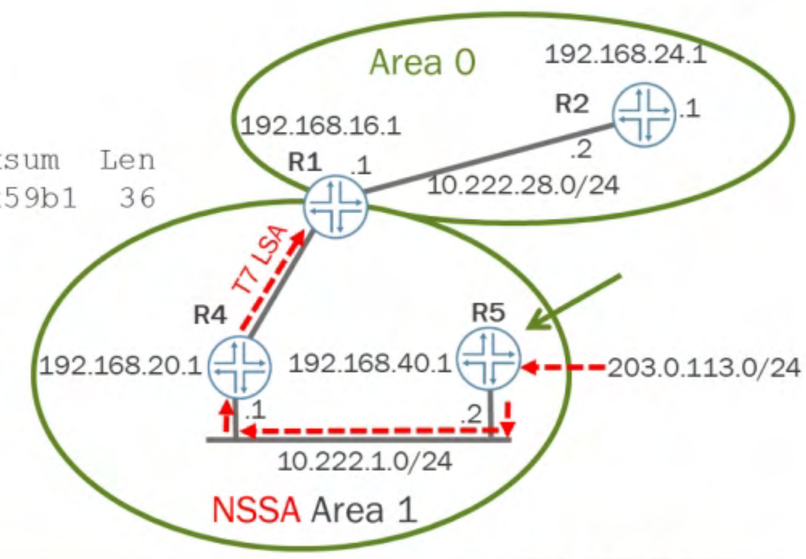

© 2020 Juniper Networks, Inc. All Rights Reserved.

NSSA External LSA Origination

Each ASBR within the NSSA generates a Type 7 LSA to advertise any networks external to the OSPF domain. This LSA is flooded to each router within the NSSA. Because the LSA has only an area flooding scope, it is not sent into other adjacent areas.

The format of the NSSA external LSA is exactly the same as the AS external LSA (Type 5) described on an earlier slide. The only difference between the two LSAs is in the use of the forwarding address field. With the Type 7 LSA, the forwarding address depends on whether the network connecting the NSSA to the adjacent system is known as an internal route to the NSSA. If the network is known, the next-hop address of the remote router is placed into the forwarding address field. If the network is not known, the forwarding address field contains the router ID of the ASBR.

By definition, only NSSA-capable routers can interpret a Type 7 LSA. Having only NSSA-capable routers able to interpret this LSA type causes a problem with flooding the routing knowledge throughout the OSPF domain because not all routers are configured to support NSSA. Furthermore, the NSSA external LSA represents the same type of information as the AS external LSA, and each OSPF router expects to receive an LSA for all external routes. This apparent contradiction is resolved by allowing the ABR to generate an AS external LSA (Type 5) on behalf of the NSSA ASBR. For each Type 7 LSA received, the ABR translates the information into a Type 5 LSA and sends that information into the backbone.

External LSA (Translated Type 7) (2 of 2)

- Translated into an AS external LSA (Type 5) by the ABR at the NSSA border. When multiple ABRs exist, the ABR with the highest RID performs the translation.
- Does not require a Type 4 ASBR Summary from the translating ABR

```
user@R1> show ospf database external extensive

OSPF link state database
Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   203.0.113.0     192.168.16.1     0x80000005   289  0x22 0x4790  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 0, Fwd addr: 192.168.40.1, Tag: 0.0.0.0
```

To get to the 203.0.113.0/24 network send traffic to the router with Router ID of 192.168.16.1

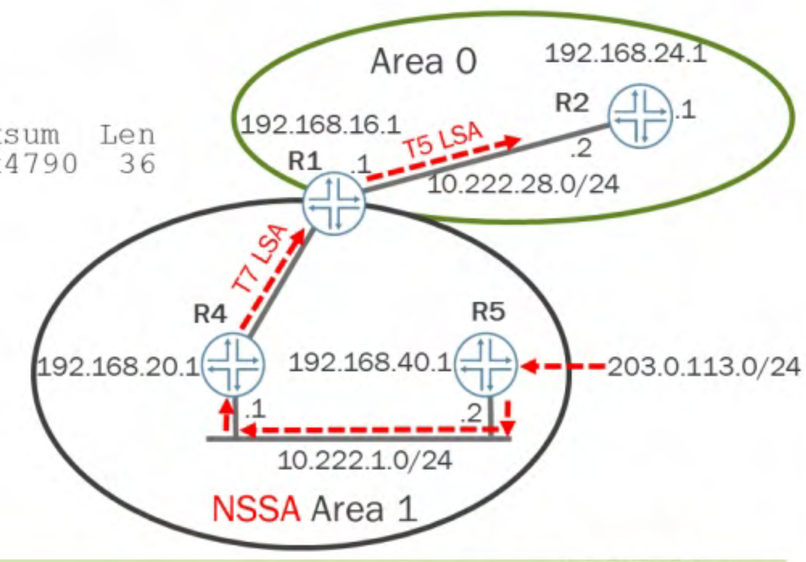

© 2020 Juniper Networks, Inc. All Rights Reserved.

NSSA External LSA Translation

The ABR with the highest Router ID connected to an NSSA area is responsible for translating the NSSA LSA into an External LSA for area 0. Because the ABR uses its own Router ID as the Advertising Router field ofthe External LSA, A Type 4 (ASBR Summary) LSA is not needed.

When an ASBR is also an ABR with an NSSA area attached to it, a Type 7 LSA is exported into the NSSA area by default. If the ABR is attached to multiple NSSAs, a separate Type 7 LSA is exported into each NSSA by default. Use the **no-nssa-abr** command to disable the export.

Opaque LSA (Types 9–11)

- Allows for the future extensibility of OSPF
  - The Junos OS uses Type 9 for graceful restart capability
  - The Junos OS uses Type 10 for MPLS traffic engineering
  - Type 11 is currently not supported
- Consist of a standard LSA header followed by application-specific information
  - OSPF or other applications can use information field directly

© 2020 Juniper Networks, Inc. All Rights Reserved.

Future Extensibility

RFC 2370 defines the OSPF opaque LSA. It was designed to extend the protocol to support future enhancements without generating new LSA types. As of this writing, production networks use both the Type 9 and Type 10 opaque LSAs. The Type 9 LSA supports graceful restart. The Type 10 LSA supports MPLS traffic engineering. The Type 11 opaque LSA is not used at this time.

Flooding Scope

The main difference between the opaque LSAs is in the flooding scope of each type: The Type 9 LSA has a link-local scope, the Type 10 LSA has an area scope, and the Type 11 LSA has domain scope.

Opaque LSA Format

The format of the opaque LSA has the standard LSA header followed by some number of octets containing application-specific information. You can interpret the total number of octets contained in the message by examining the length field in the LSA header. In addition, the link-state ID field is segmented into a 1-byte opaque type field and a 3-byte opaque ID field. The Internet Assigned Numbers Authority (IANA) has the responsibility of assigning new opaque type codes in the 0–127 range. Values between 128 and 255 are set aside for private and experimental use only.

Opaque-capable routers communicate their ability to each other by using a new bit in the options field. Bit position 2 is defined as the O bit. A value of 1 in this bit position allows a remote router to flood opaque LSAs to the local router. A value of 0 tells the remote router not to flood those LSA types.

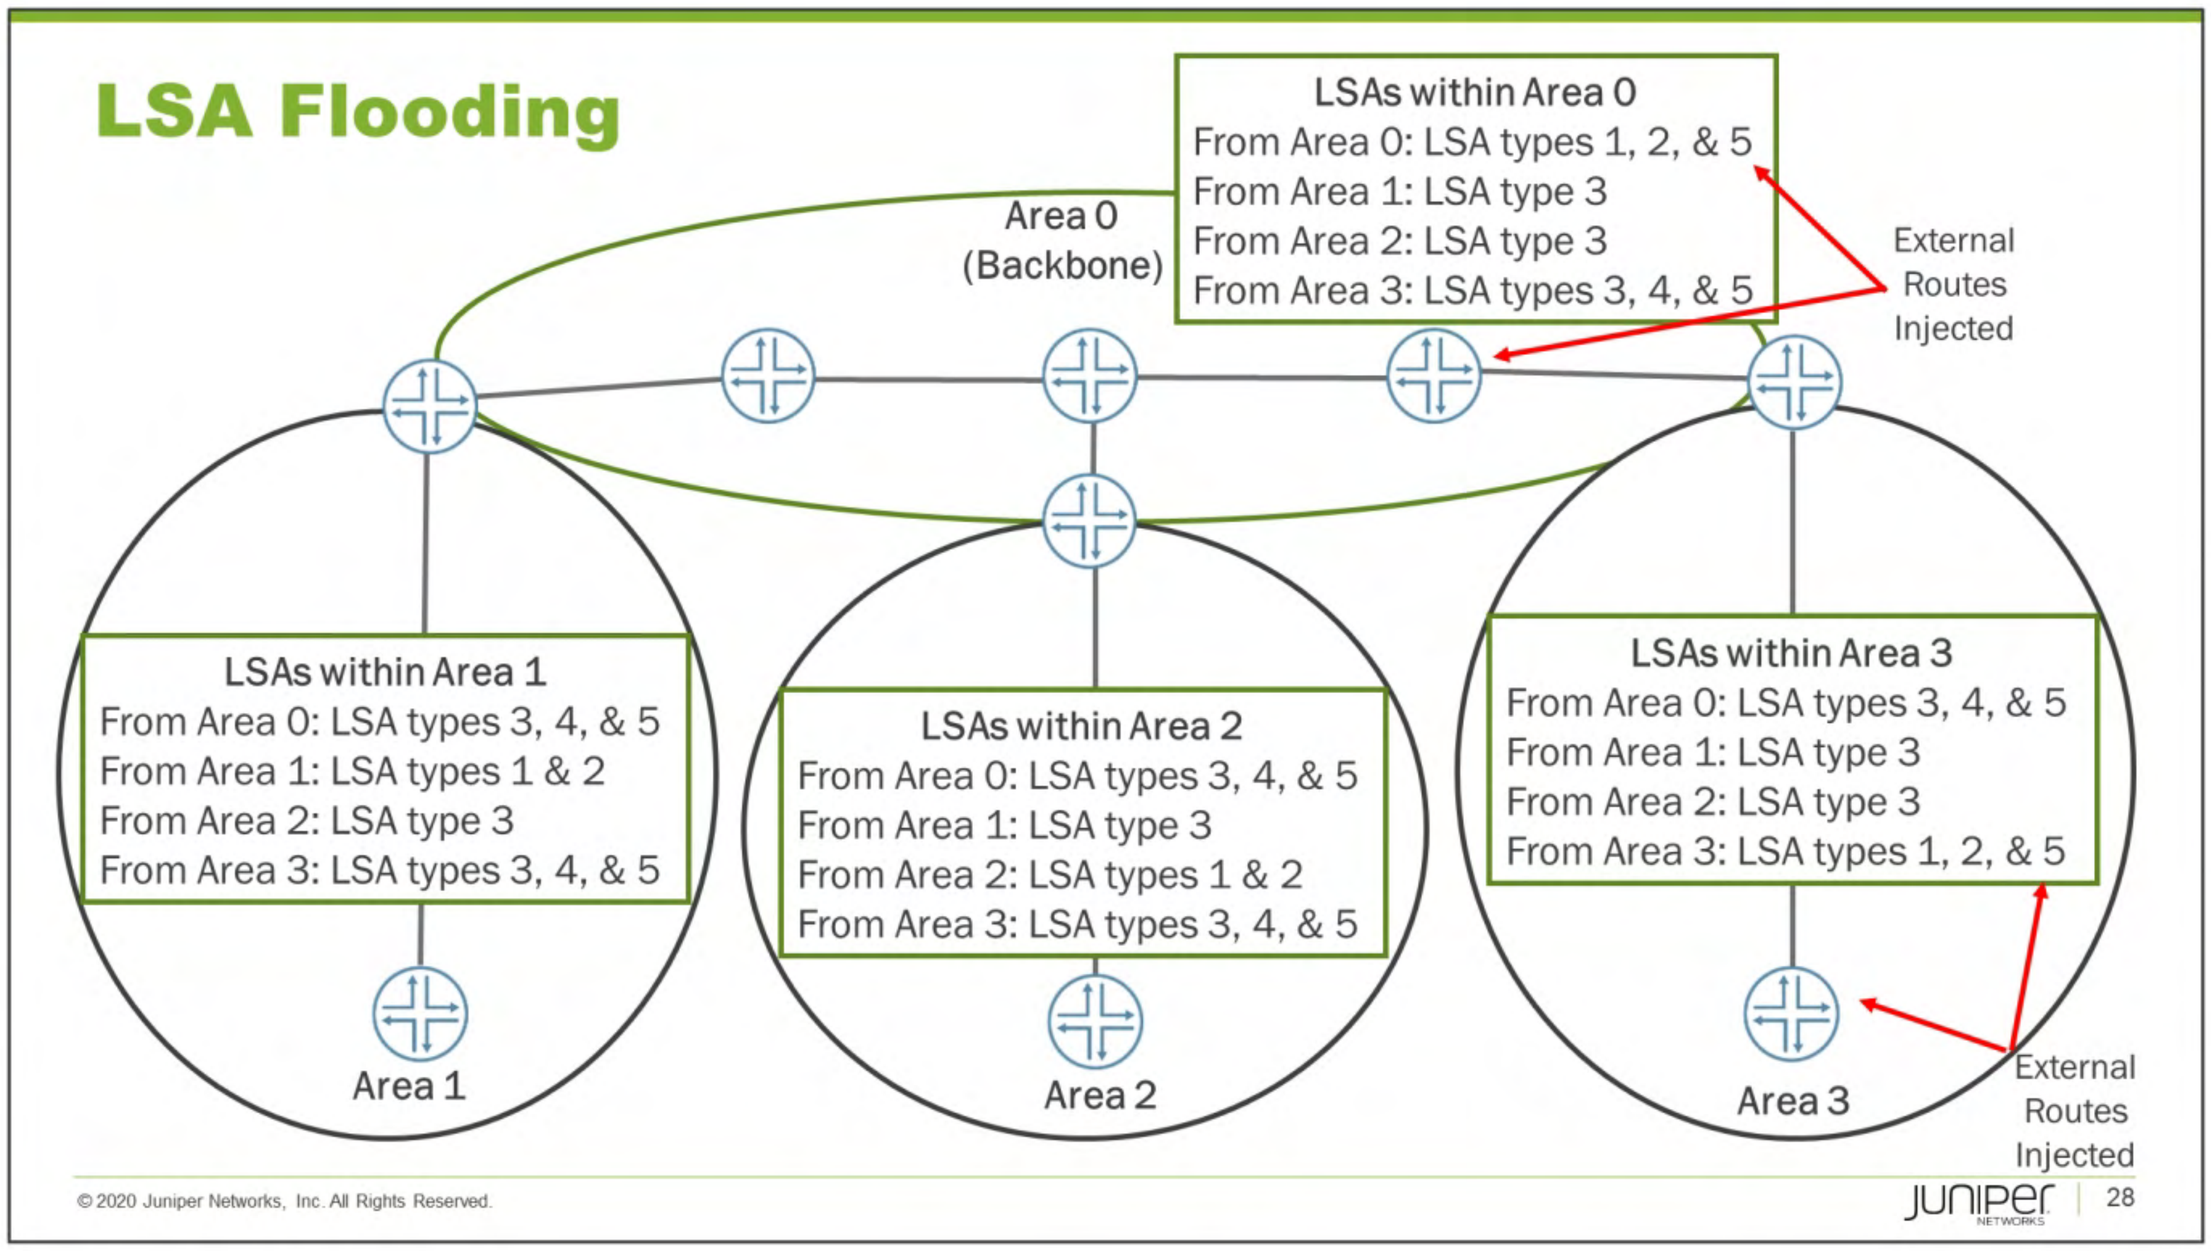

LSA Flooding Scopes

As discussed in previous slides, each LSA in the OSPF protocol has a specific scope within which it can be flooded. The slide graphically displays those flooding scopes.

LSA Types 1 and 2 are generated within each area. Because these LSAs have an area flooding scope, they remain within their own particular area and are not seen in other areas. The ABR places the routing information contained within these LSAs into a Type 3 LSA and forwards it across the area boundary. In addition, existing Type 3 LSAs within the backbone are forwarded into the nonbackbone areas to provide routing knowledge to those routers. This results in Type 3 LSAs within every area that represent all OSPF routes. Within Area 1, for example, Type 3 LSAs exist that represent Area 0, Area 2, and Area 3.

In the diagram, two ASBRs inject routes into the OSPF domain. The AS external LSAs (Type 5) that represent those routes have domain flooding scope (with the exception of stub areas). As such, they exist within each OSPF area. To reach those external routes, the OSPF routers use either router LSAs or ASBR summary LSAs (Type 4) to have knowledge of the ASBRs. Much like the Type 3 LSAs, the ASBR summary LSAs have area scope, so the area border binds them. Each ABR regenerates the Type 4 summary LSAs into their respective areas to provide the knowledge of how to route toward the ASBR that is advertising a given external prefix. This capability leads to Type 4 LSAs for Area 0 being injected into Areas 1, 2, and 3, while a Type 4 for Area 3 is injected into Areas 0, 1, and 2.

Sample OSPF Database

```
user@router> show ospf database

    OSPF link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *192.168.16.1     192.168.16.1     0x80000004   177  0x2  0xd45b  60
Router   192.168.36.1     192.168.36.1     0x80000005   305  0x2  0xda47  60
Summary *10.222.1.0       192.168.16.1     0x80000002   412  0x2  0xfafa  28
Summary *10.222.29.0      192.168.16.1     0x80000002   631  0x2  0xbb1f  28
Summary *192.168.20.1     192.168.16.1     0x80000001   412  0x2  0x87c6  28
ASBRSum  192.168.32.1     192.168.36.1     0x80000001   240  0x2  0x3b07  28

    OSPF link state database, area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *192.168.16.1     192.168.16.1     0x80000007    39  0x2  0xcc62  60
Router   192.168.20.1     192.168.20.1     0x80000002   415  0x2  0xd7d9  48
Network  10.222.1.1       192.168.20.1     0x80000001   418  0x2  0x6a75  32
Summary *192.168.32.1     192.168.16.1     0x80000001   238  0x2  0xe96b  28
Summary *192.168.36.1     192.168.16.1     0x80000002   631  0x2  0xb19f  28
ASBRSum *192.168.32.1     192.168.16.1     0x80000001   238  0x2  0xdb78  28
ASBRSum *192.168.36.1     192.168.16.1     0x80000001   574  0x2  0xa5ab  28
    OSPF external link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *192.168.17.0     192.168.16.1     0x80000001   631  0x2  0x3812  36
Extern   192.168.21.0     192.168.20.1     0x80000001   420  0x2  0x8693  36
Extern   192.168.33.0     192.168.32.1     0x80000001   590  0x2  0x1713  36
Extern   192.168.37.0     192.168.36.1     0x80000001   576  0x2  0xce53  36
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Sample OSPF Database

You can see key details of the OSPF database in the database example on the slide, which has been edited for brevity. To this point, we have examined small portions of the database through the use of the **`extensive`** keyword.

In the slide above, The OSPF link state database is shown using the **show ospf database** command. This command displays the LSA header information, which provides a quick view of the different LSAs and their types stored in the database.

Agenda: OSPF

- OSPFv2 Review
- Link-State Advertisements
- →Protocol Operations
- OSPF Authentication
- OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

Protocol Operations

The slide highlights the topic we discuss next.

OSPF Database Protection

- Limits the number of LSAs not generated by the local router in a given OSPF routing instance
- Protects the LSDB from being flooded with excessive LSAs
- Useful if VPN routing and forwarding is configured on your provider edge and customer edge routers are using OSPF as the routing protocol

```
user@router# show protocols ospf
database-protection {
    maximum-lsa 1000;
    warning-only;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Database Protection

This feature limits the number of LSAs that were not generated by the local router. This limit protects the LSDB from being flooded with excessive LSAs. Once the specified maximum LSA count is exceeded, the database typically enters the ignore state. In this state, all neighbors are brought down, and nonself-generated LSAs are destroyed. In addition, the database will send out hellos but ignore all received packets, not form any full neighbors, and therefore not learn about new LSAs. However, if you had configured the warning-only option, only a warning is issued and the database does not enter the ignore state but continues to operate as before.

Overall, this feature can be very useful if a VPN routing and forwarding (VRF) instance is configured to use OSPF between the PE and CE routers.

Shortest Path First Algorithm

- Based on the Dijkstra algorithm
  - Link-state database
  - Candidate database
  - Tree database
- Run on a per-area basis on each router
  - Independent calculation of the topology
- Result is passed to the Junos OS routing table
  - The route selection algorithm (route preference value) determines whether the route is marked active

© 2020 Juniper Networks, Inc. All Rights Reserved.

Dijkstra Algorithm

After a router receives a new LSA and places it into the LSDB, the router runs the SPF algorithm. This computation uses the database as a data source and results in a loop-free network topology using the best metric from the local router to all nodes in the network.

During the course of this calculation, the algorithm uses three databases—the LSDB, the candidate database, and the tree database. As we have explored, the LSDB is the total compilation of routing knowledge in the network. Conceptually, it consists of multiple tuples in the form of: router ID, neighbor ID, cost. These tuples describe each link in the network.

Run on a Per-Area Basis

The router calculates the Dijkstra algorithm for each LSDB present on the local router. Recall from an earlier slide that each OSPF area maintains a separate copy of the database. These copies allow each area to have a separate forwarding topology independent of another area.

Results Are Passed to the Routing Table

Once the algorithm is run, the routing table on the Routing Engine receives the results of the tree database. At this point, the Routing Engine controls the determination of whether to use any individual OSPF route to forward traffic. The preference value assigned to each route often handles this choice.

The action of calculating the best OSPF route prior to the route being placed into the routing table has a consequence with regard to the filtering of routing knowledge. An individual filter (or policy) operates on IP routes that enter the router as IP routes and are placed into the routing table. This form of data flow is not present in OSPF because the routing information enters the router as an LSA and is only placed into the routing table after the Dijkstra algorithm is performed. Hence, the best method of filtering OSPF routing knowledge is to keep that information from being placed into the database (on a per-area basis) in the first place.

OSPF does offer a limited import-policy functionality. You can use the **`import`** statement at the `[edit protocols ospf]` configuration hierarchy to block external routes—those learned from Type 5 and Type 7 LSAs—from being installed in the routing table. This import policy does not, however, prevent LSAs from being flooded. We strongly discourage the use of OSPF import policy because of the potential for unknowingly discarding traffic.

SPF Example (1 of 6)

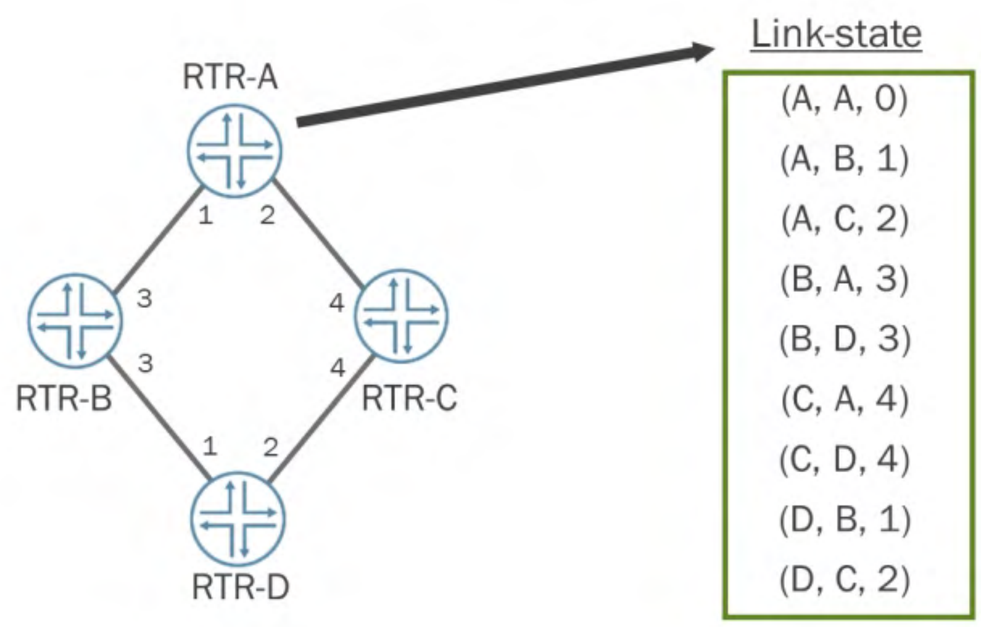

© 2020 Juniper Networks, Inc. All Rights Reserved.

SPF Calculation Example: Part 1

In the following slides, a sample SPF calculation is displayed. This graphic shows the beginning state of the network including the routers involved, the configured link metrics, and the LSDB. The network and the LSDB have recently converged and the local router, RTR-A, is running an SPF calculation to determine the shortest path to each node in the network.

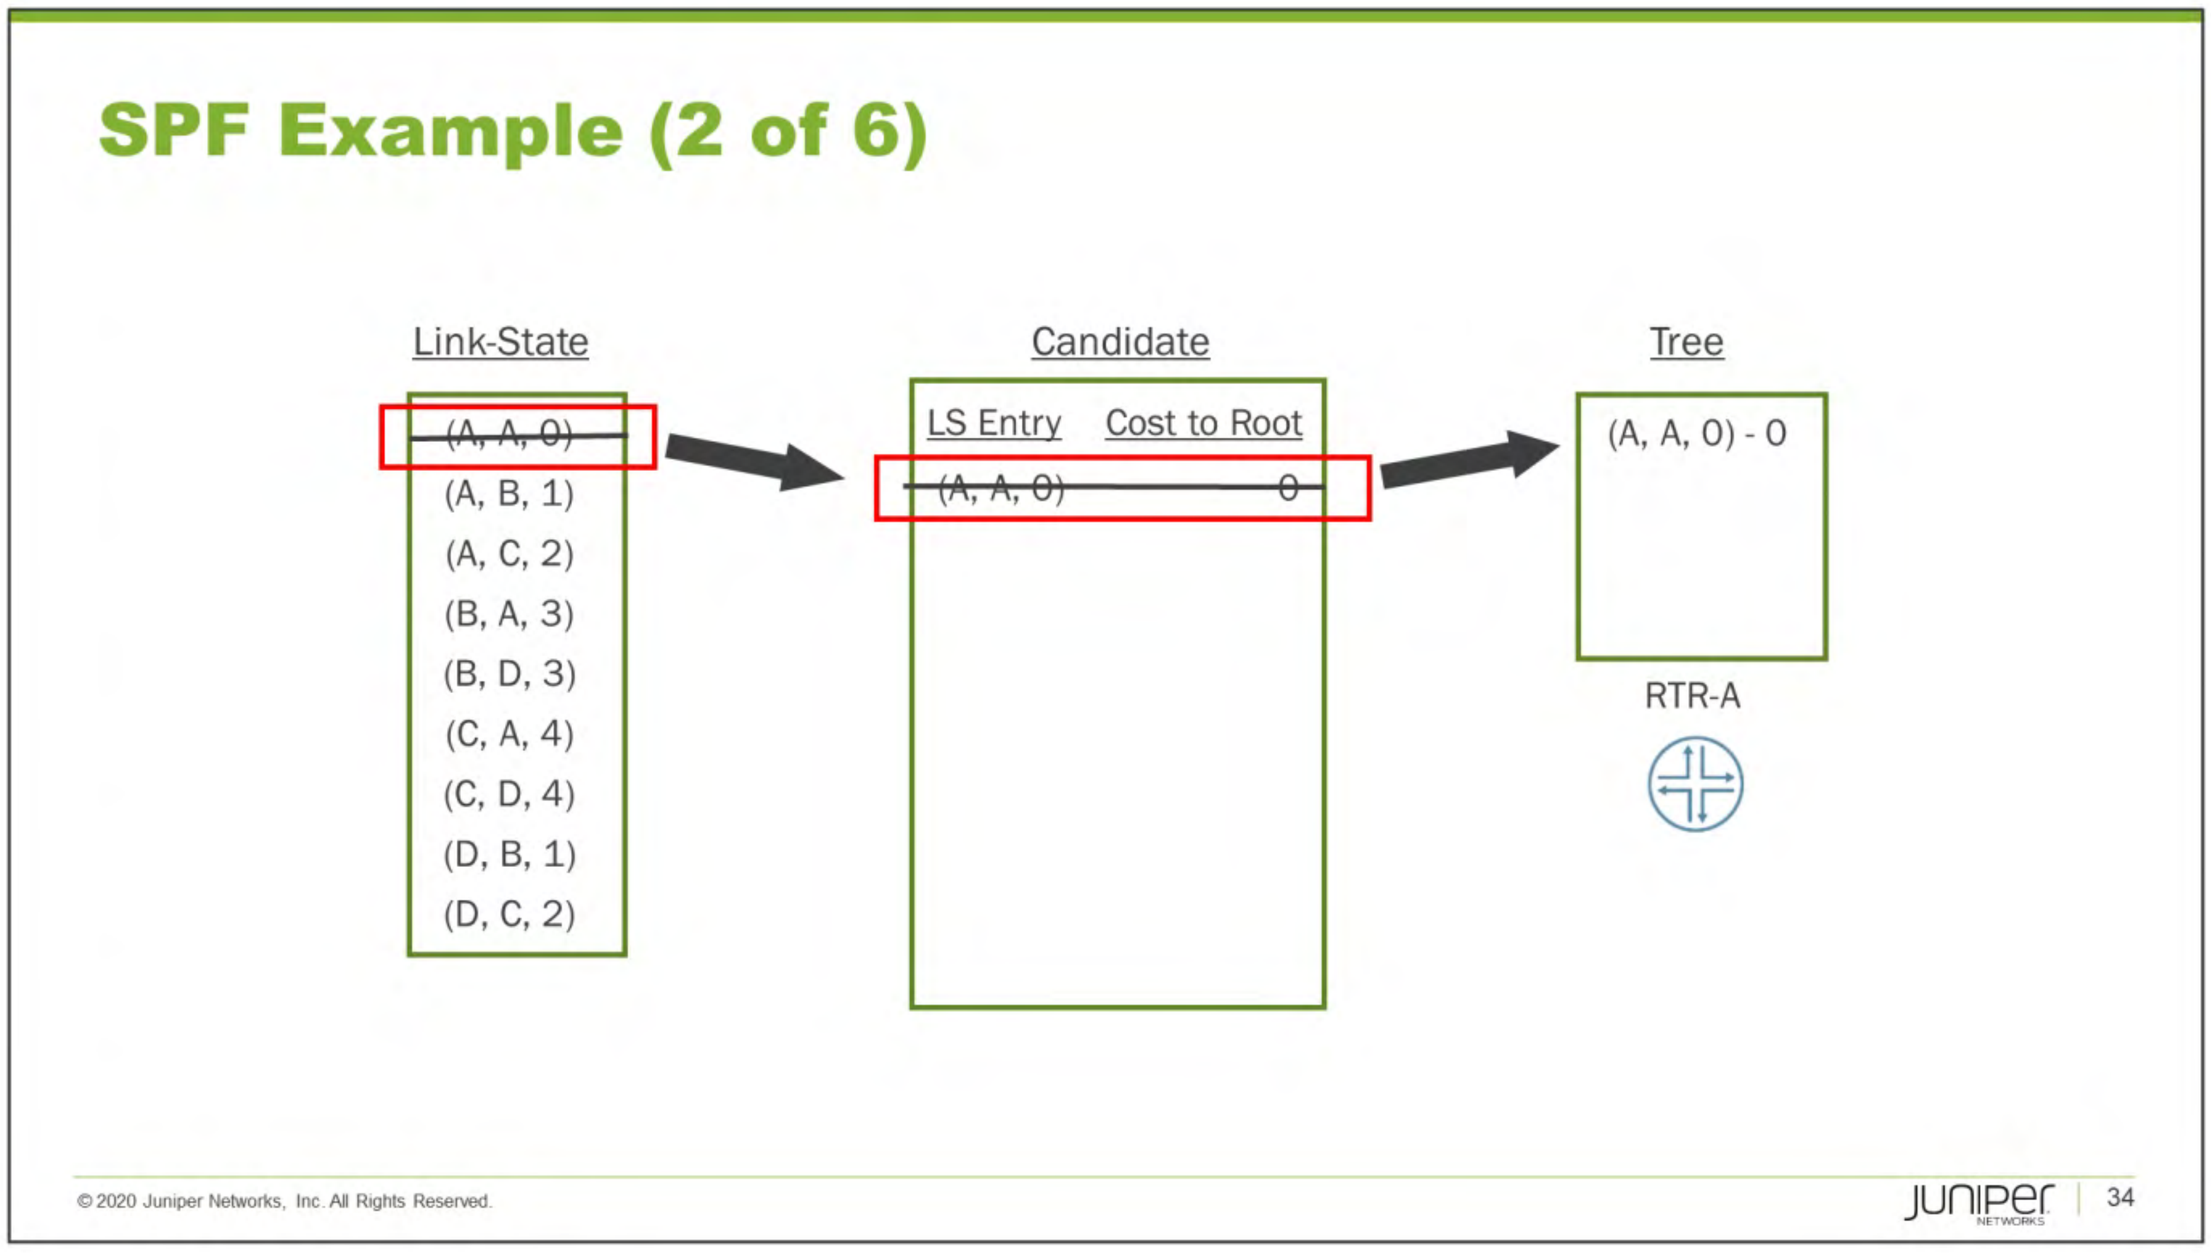

SPF Calculation Example: Part 2

RTR-A begins by moving its own local database tuple (A,A,0) into the candidate database. The total cost from the neighbor ID to the root is calculated, which results in a 0 value. In other words, RTR-A is directly connected to itself!

The lowest, and only, tuple in the candidate database is moved to the tree database and RTR-A places itself on the network map.

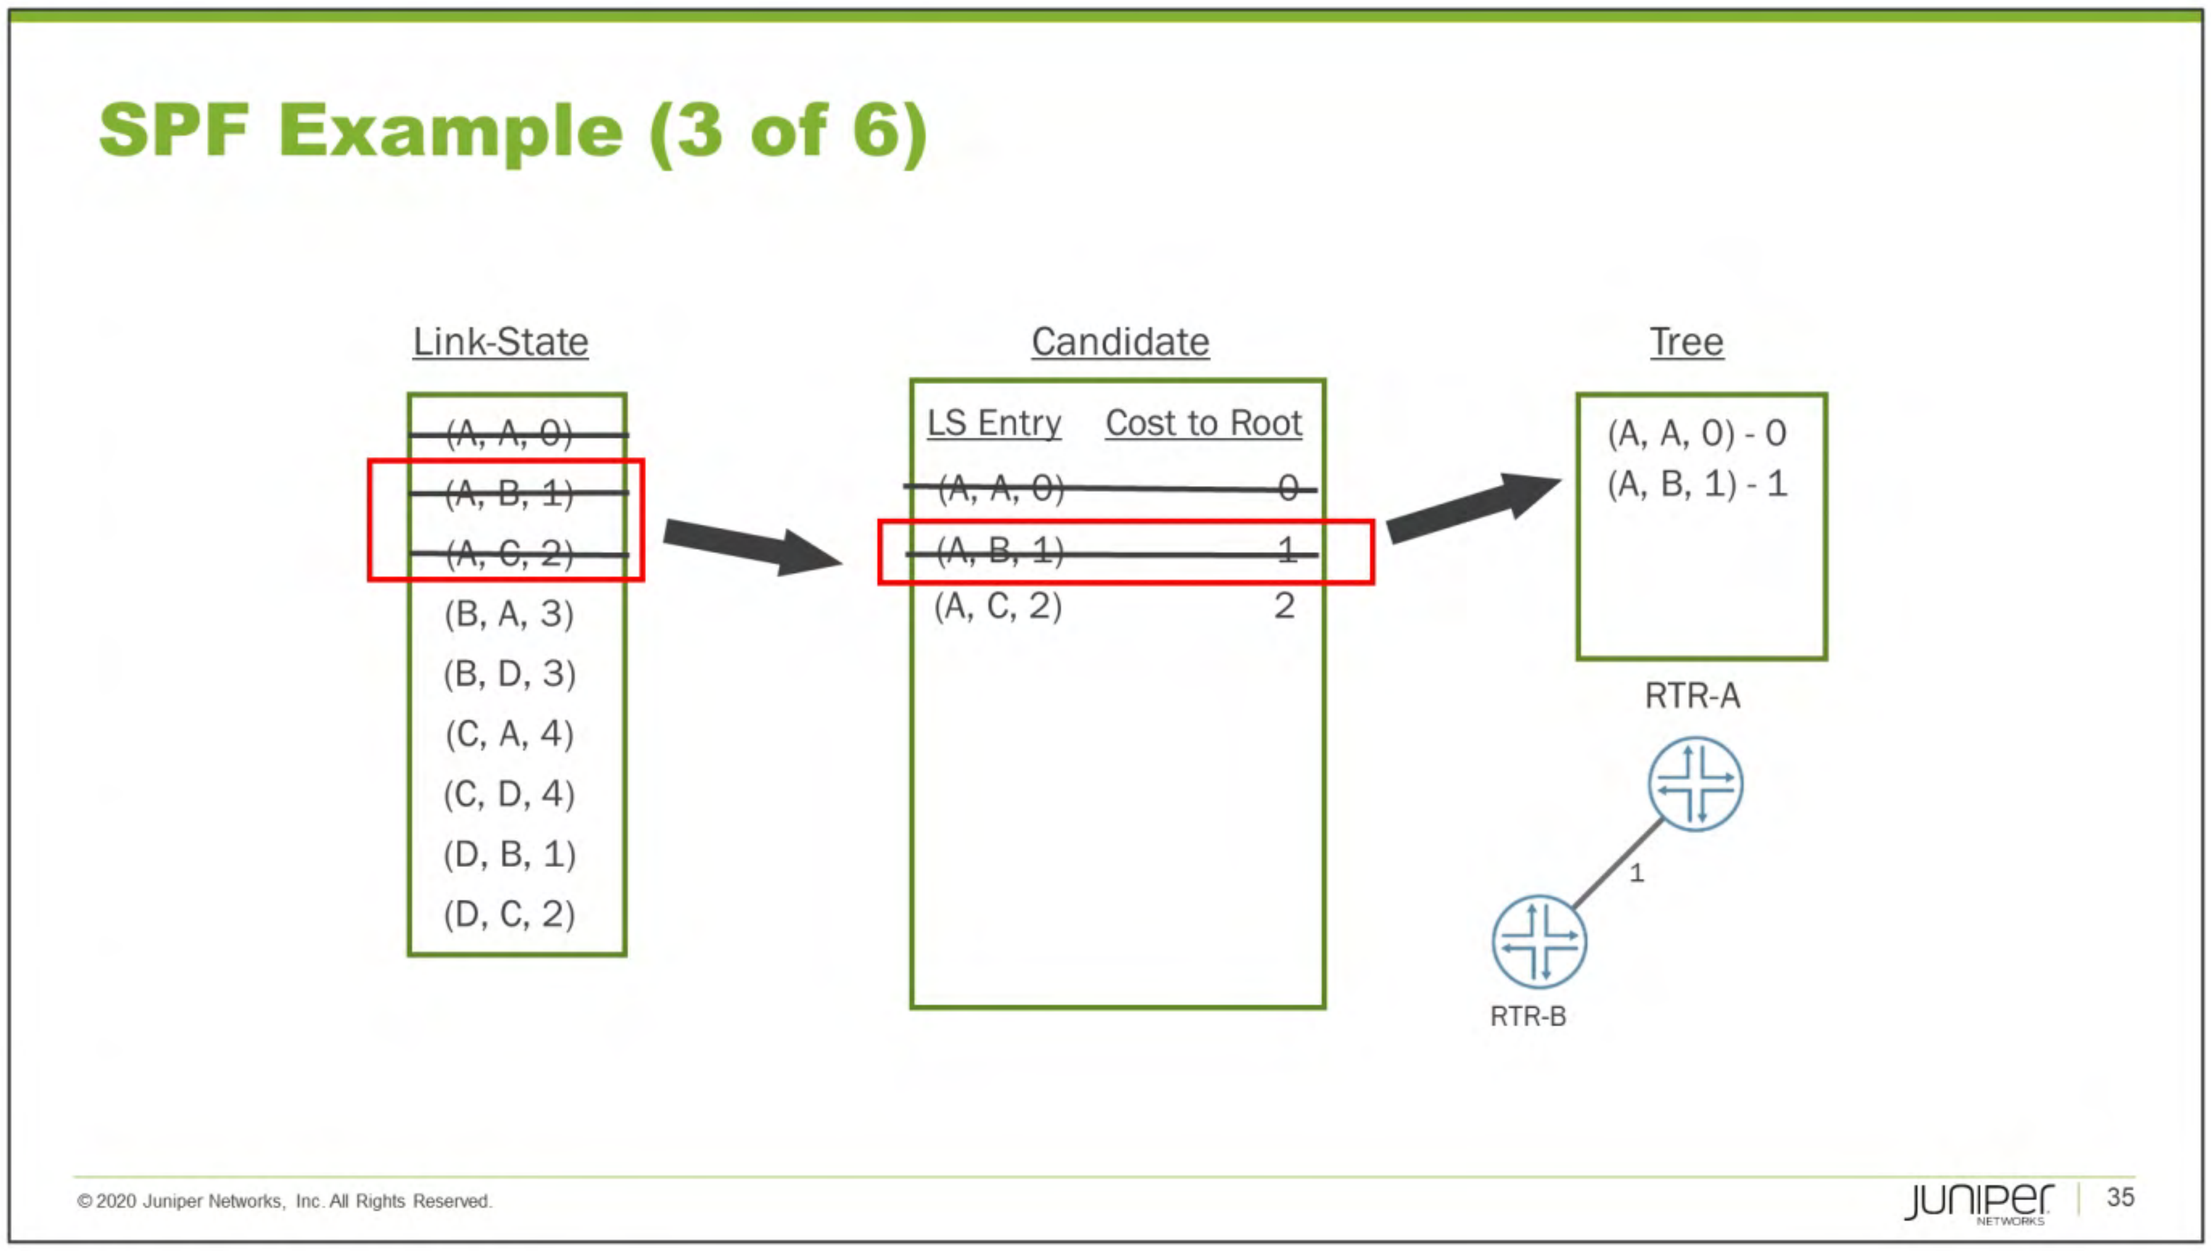

SPF Calculation Example: Part 3

All tuples from the most recent node added to the tree database are now added to the candidate database. Because RTR-A is the most recent entry to the tree database, all of RTR-A's tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are removed from the candidate, of which there are none. (For example, if B were already in the tree database, the tuple (A,B,1) would be eliminated.)

The cost to each neighbor ID from the root is then calculated. It costs RTR-A 0 to reach itself and 1 to reach RTR-B, so the total cost to RTR-B is 1. The same calculation is done for RTR-C, and the total cost of 2 is placed into the candidate database.

The lowest cost tuple in the candidate database, (A,B,1), is now moved to the tree database, and RTR-B is placed on the network map.

The candidate database is not empty, so the algorithm continues.

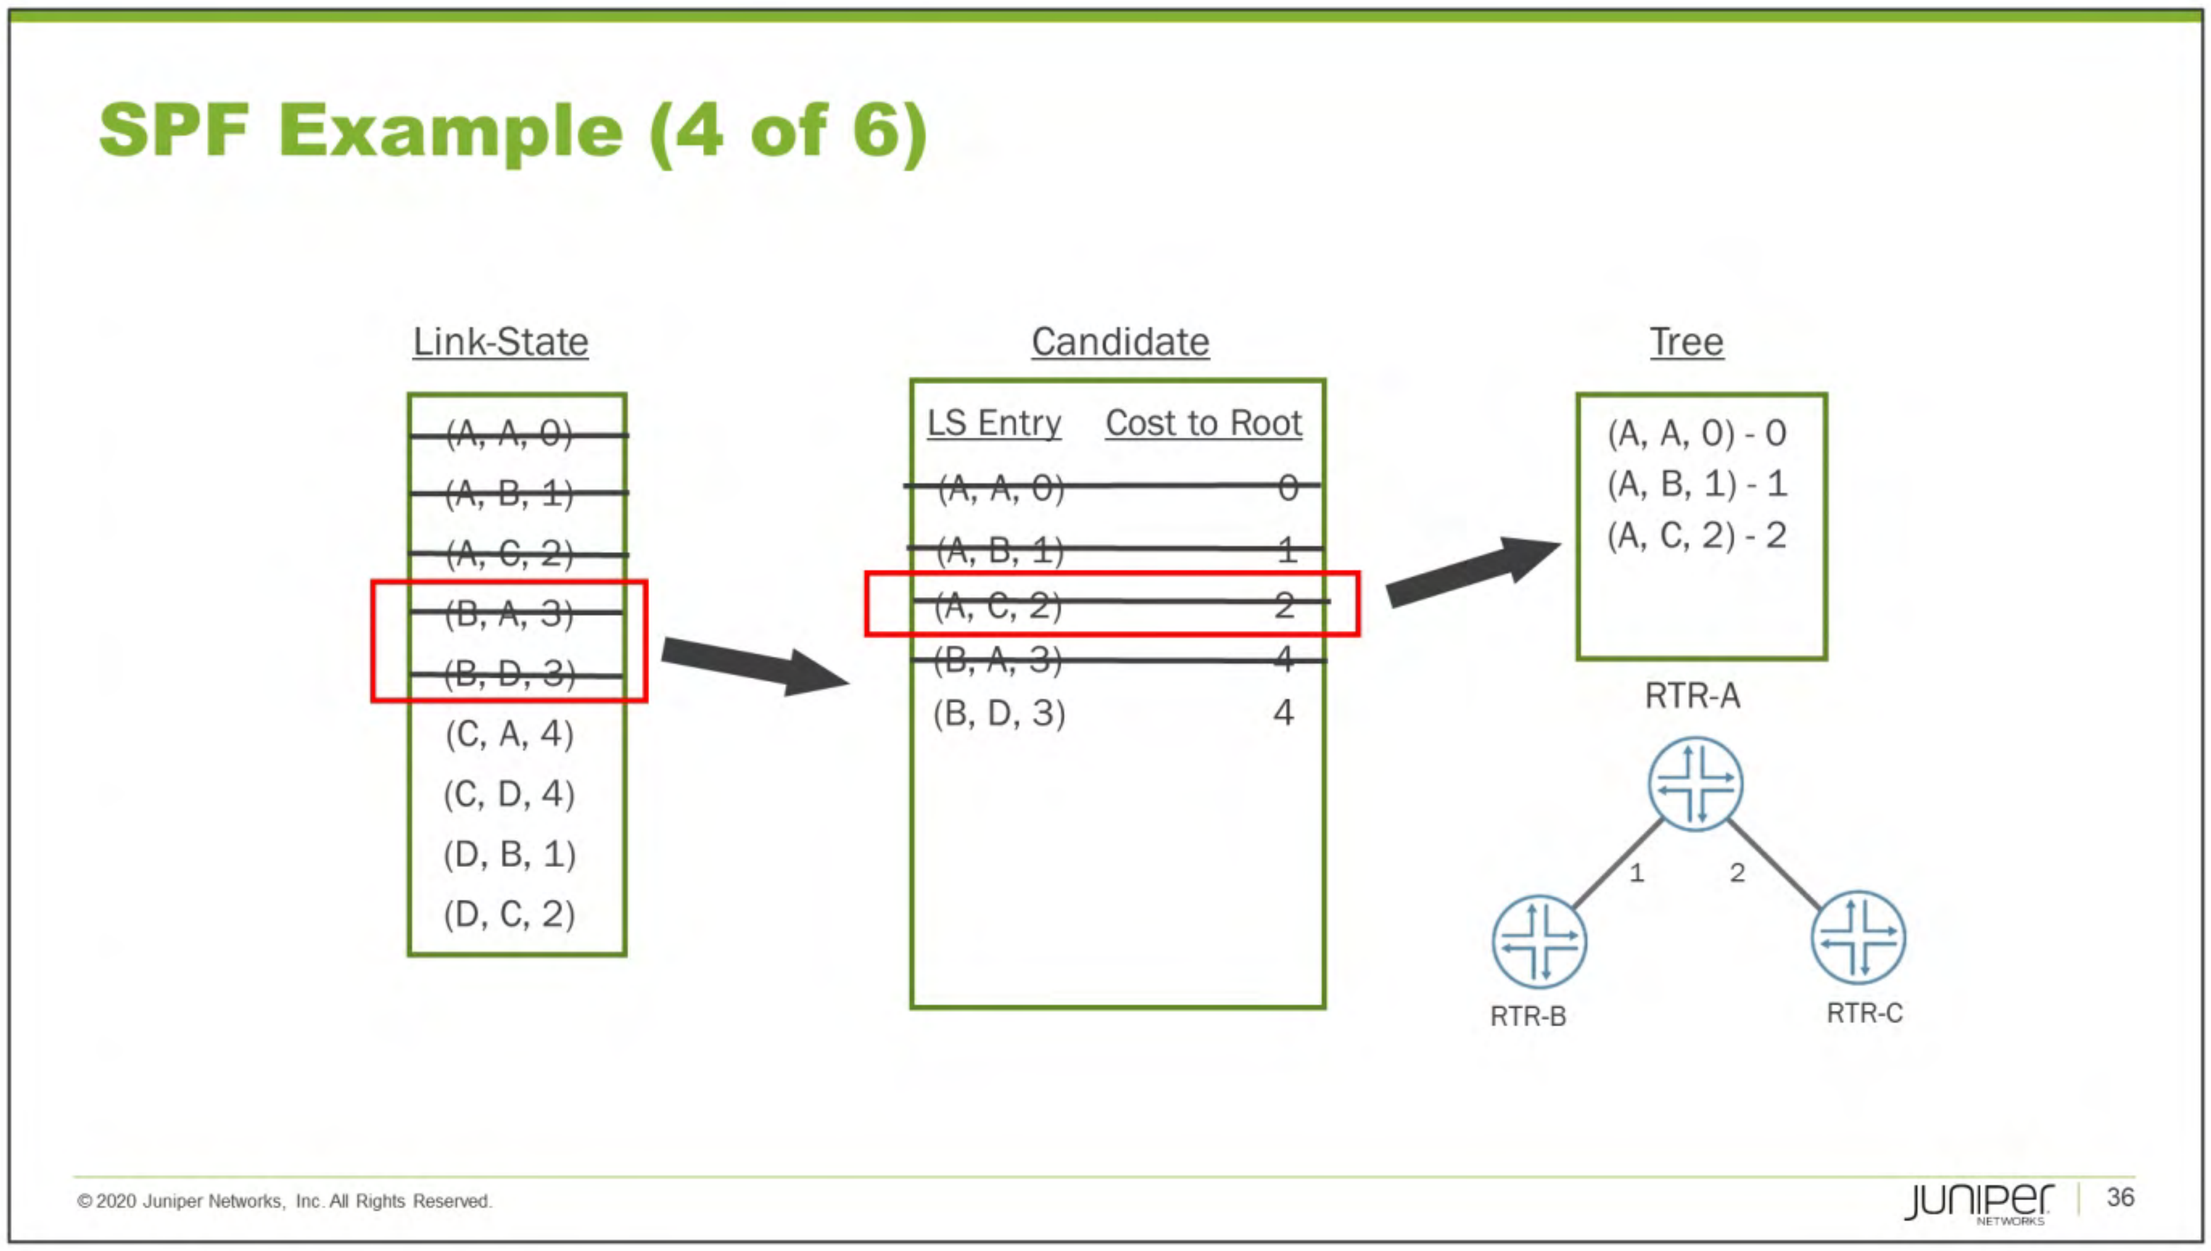

SPF Calculation Example: Part 4

RTR-B is the most recent entry to the tree database, so all RTR-B's tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are then removed from the candidate. Thus, the (B,A,3) tuple is removed because RTR-A already has the shortest path to RTR-B.

The cost to each neighbor ID from the root is then calculated. It costs RTR-B 3 to reach RTR-D, and it costs 1 to reach RTR-B from the root. Therefore, the total cost to reach RTR-D from the root through RTR-B is 4.

The lowest cost tuple in the candidate database, (A,C,2), is now moved over to the tree database, and RTR-C is placed on the network map.

The candidate database is not empty, so the algorithm continues.

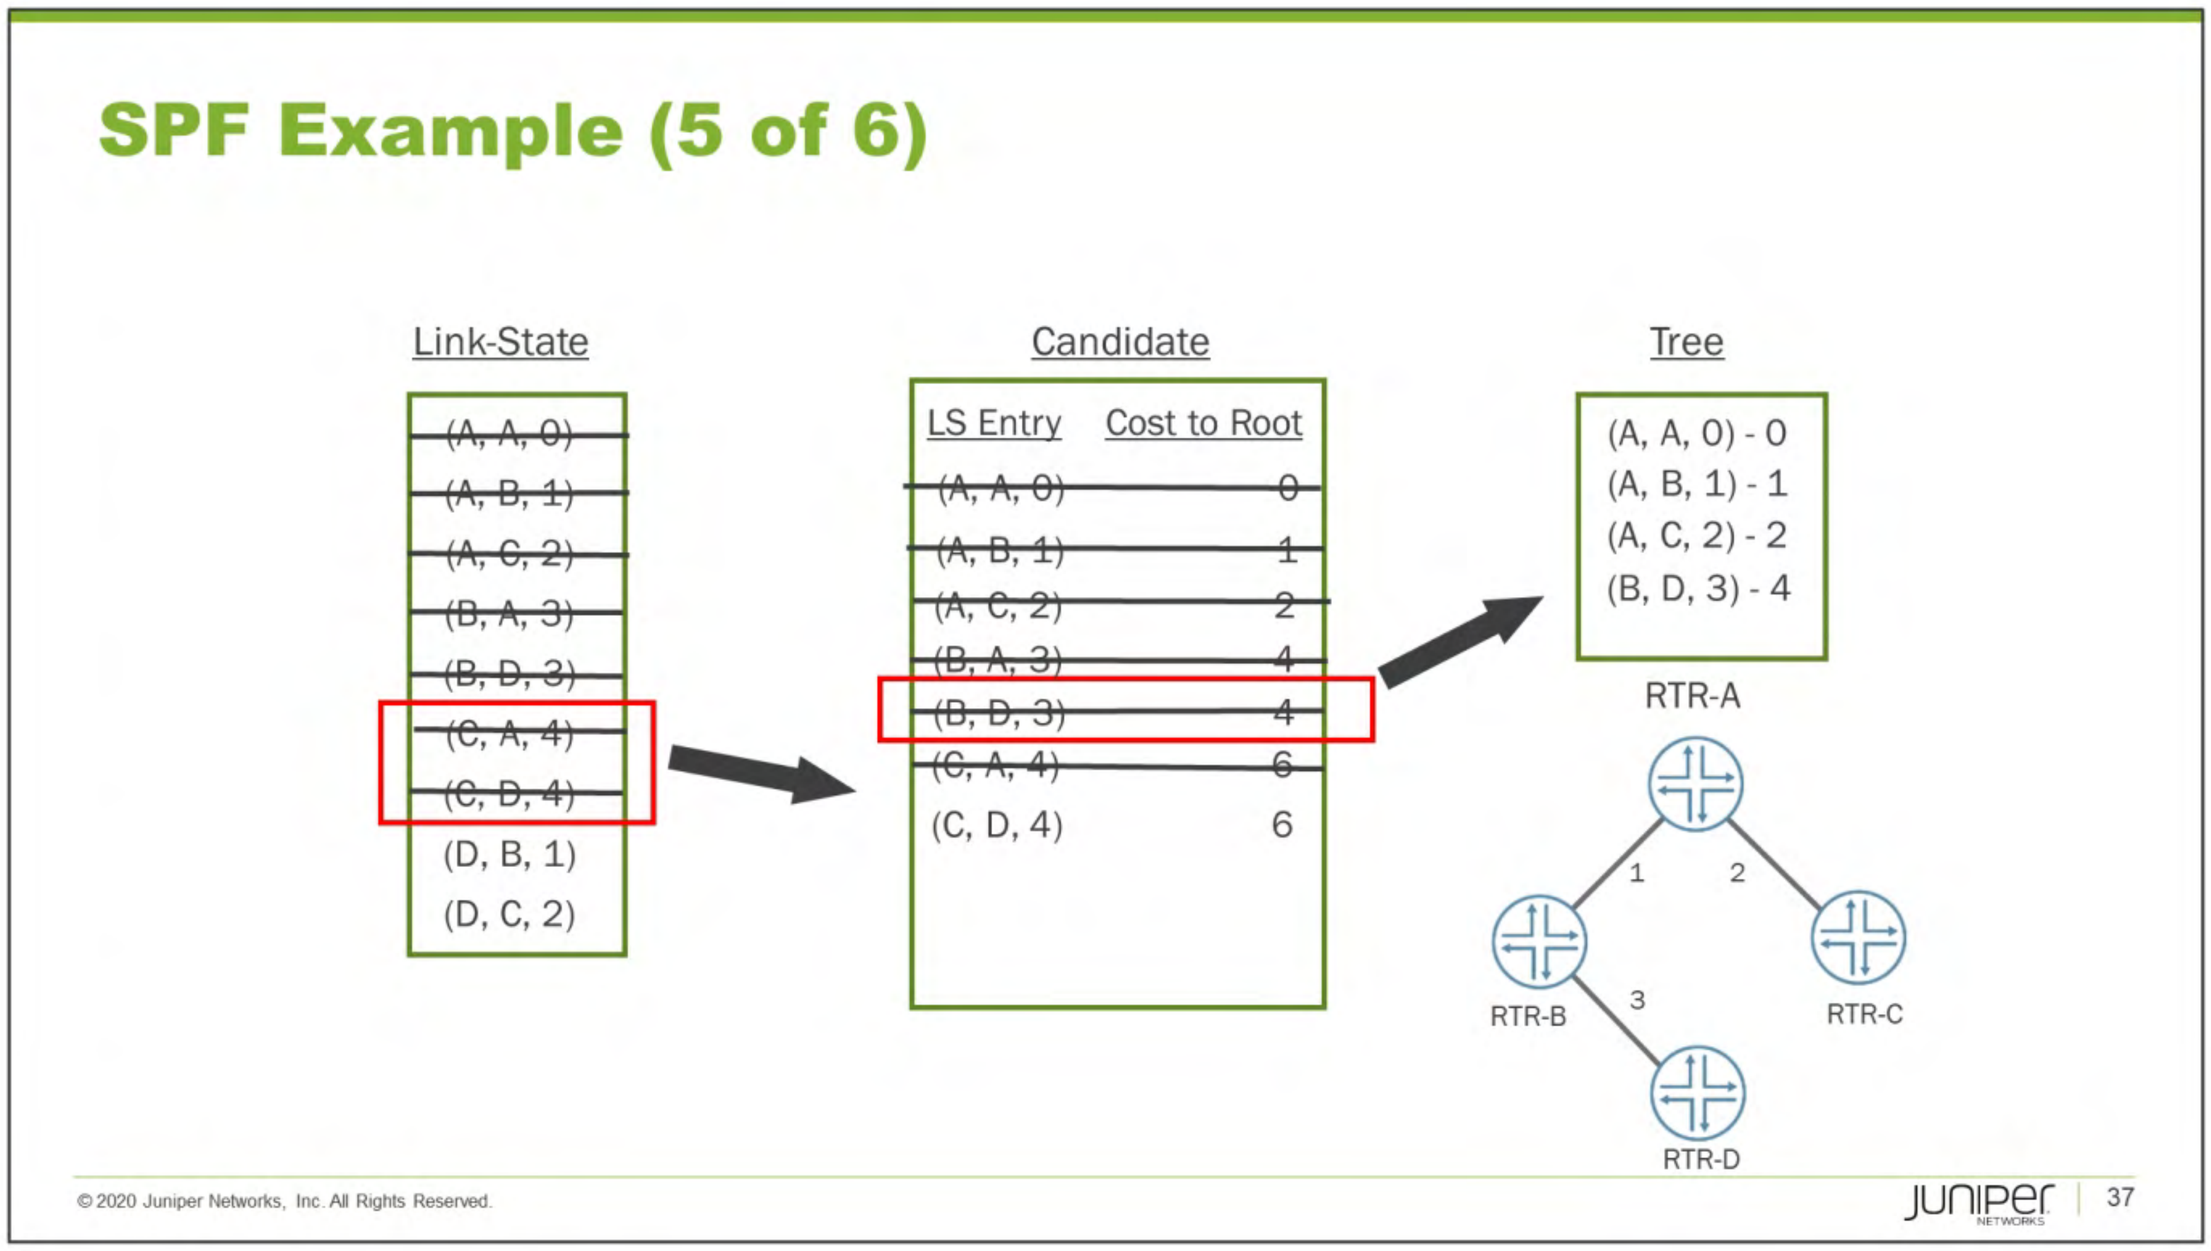

SPF Calculation Example: Part 5

As RTR-C is the most recent entry to the tree database, its tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are then removed from the candidate. Thus, the (C,A,4) tuple is removed because RTR-A already has the shortest path to RTR-C.

The cost to each neighbor ID from the root is then calculated. It costs RTR-C 4 to reach RTR-D, and it costs 2 to reach RTR-C from the root. So the total cost to reach RTR-D through RTR-C is 6.

The lowest cost tuple in the candidate database, (B,D,3), is moved to the tree database, and RTR-D is placed on the network map.

The candidate database is not empty, so the algorithm continues.

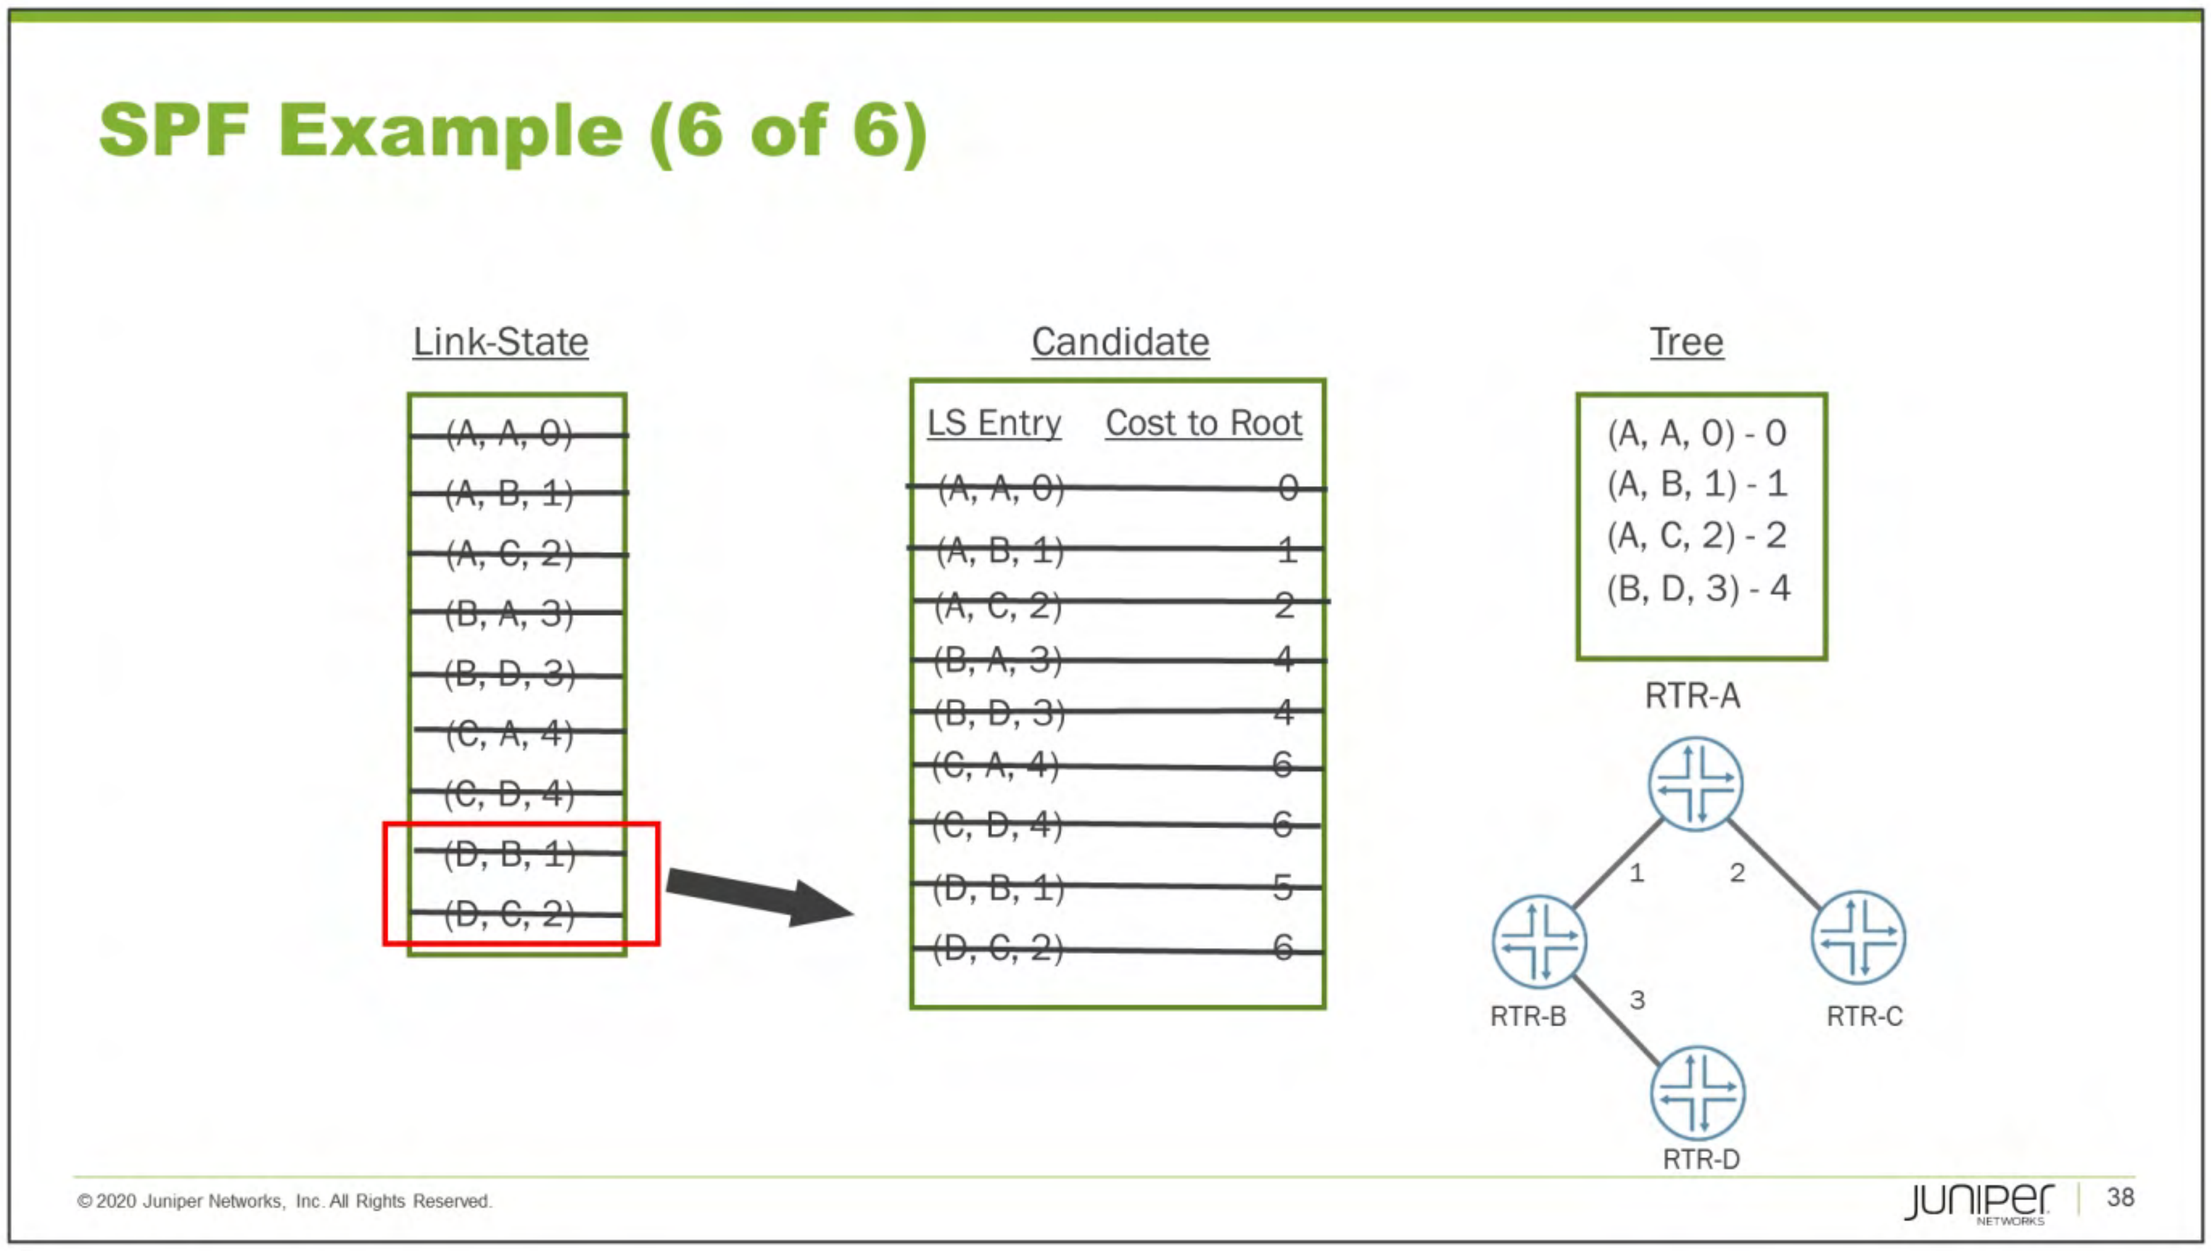

SPF Calculation Example: Part 6

RTR-D, through its link to router B, is the most recent entry to the tree database. Therefore, its tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are then removed from the candidate. Thus, the (C,D,4), (D,B,1), and (D,C,2) tuples are removed because RTR-A already has paths to RTR-B, RTR-C, and RTR-D. The candidate database is now empty of all tuples, so the algorithm stops at this point.

RTR-A now has a complete network map built with the total cost to each node calculated. This information is then passed to the Junos routing table for its use.

Controlling SPF Calculations

- Three consecutive SPF runs can occur before a mandatory hold-down occurs
  - Keeps the network stable during change
  - 5-second timer is now configurable
  - Possible values range from 2000 to 20,000 ms
- Default 200 ms delay between a topology change and running the SPF algorithm
  - Altered with the **`spf-options delay`** command
  - Possible values range from 50 to 8000 ms

```
[edit protocols ospf]
user@router# set spf-options delay 100
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Consecutive SPF Calculations

To enhance the convergence time of an OSPF network during a time of topology changes, the Junos OS by default allows for the SPF calculation to be run three times in a back-to-back fashion before encountering a mandatory hold-down period. The 5-second hold-down timer used to be hardcoded into the software but is now configurable under `[protocols ospf spf-options holddown]` configuration hierarchy. The timer ensures that the network can continue to forward packets, with potentially incorrect routing knowledge, during the convergence period. The timer also keeps the routing protocol process from hogging the CPU at the expense of other router functions.

Preconfigured Delay Between Calculations

A configurable timer slightly delays consecutive SPF calculations. The default timer value is 200 milliseconds; you can alter it with the **`spf-options delay`** statement. The **`spf-options delay`** statement is supported both at the global OSPF configuration hierarchy and within an OSPF routing instance. Delay values ranging from 50 milliseconds to 8000 milliseconds (8 seconds) are supported. Regardless of the configuration of this timer, the mandatory hold-down timer still starts after the third consecutive rapid SPF calculation.

A current best practice in today's networking environment is setting the **`spf-options delay`** value to be slightly larger than the worst possible propagation delay found in your network. For example, if the delay across the network is 200 milliseconds, you might set the **`spf-options delay`** to 250 milliseconds. This setting allows time for all of the duplicate LSAs to arrive at all routers *before* the SPF calculation is performed.

SPF Calculation Order

- The SPF calculations are performed per area in the following preference order.
  - Intra-Area Links — First
  - Inter-Area Links — Second
  - External Type E1 — Third
  - External Type E2 — Fourth

How would traffic passing through R1 get to R3?

How would traffic passing through R1 get to R3 if the link from R1 to R3 went down?

Which BGP router will R3 prefer if the same routes are advertised by R5 & R6?

What if the routes were redistributed as E2 routes?

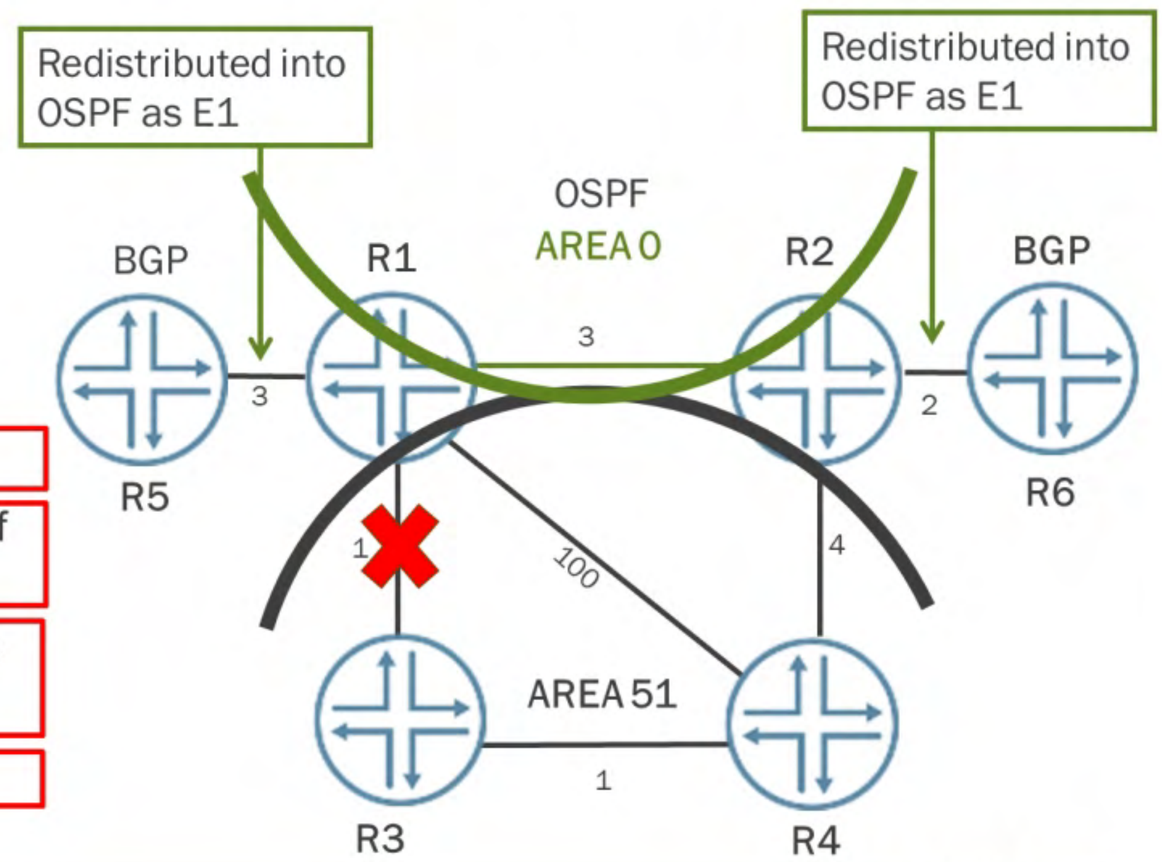

© 2020 Juniper Networks, Inc. All Rights Reserved.

SPF Calculation Order

When the SPF algorithm builds the best path to a given destination, it will do so using Intra-area links only. For example, R1 is running the SPF algorithm and needs to determine the best path. In this instance it would be to go directly from R1 to R3. What if there was no direct link between R1 and R3 what path would it use then? It would use R1 to R4 to R3. This may seem unusual since the path through R2 would be 8 versus the cost of 101 when it takes the R1 to R4 to R3 path, but OSPF first considers all paths inside an area, just like routers will consider OSPF paths before RIP even ir the RIP router has a lower metric.

Once all Intra-area paths have been considered, then Inter-area routes are then used to determine if there is a possible path. In this case if the path from R1 to R4 did not exist, then the only and best path would be from R1 to R2 to R4 to R3.

Once all Intra-area and Inter-area routes were considered, then External Type 1 LSAs are considered followed by External Type 2 LSA. In the example above, if both R5 and R6 were advertising the same route, which path would R4 take to get to an prefix if R! and R2 both imported the routes as External Type 1 routes and the R1 to R3 link was up? R4 would go through R1.

What path would R4 take if R1 and R2 advertised them as External Type 2s. R4 would go through R2.

SPF calculation order can have a big impact on the routes chosen, but so can link costs.

Reference Bandwidth

- Default OSPF cost for all links is 10^8/bandwidth (bps)
  - Links with a bandwidth ≥ 100 Mbps have a cost of 1
  - Cost calculation results in a value <1, so it is rounded up
- Use the **`reference-bandwidth`** command to change the 10^8 default within the `[edit protocols ospf]` hierarchy

```
[edit protocols ospf]
user@router# set reference-bandwidth 100g

[edit protocols ospf]
user@router# show
reference-bandwidth 100g;
area 0.0.0.0 {
    interface ge-0/0/0.0;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Default Cost

Each router uses an algorithm to determine the cost of each interface—the calculation is 100,000,000 divided by the bandwidth of the interface. If the result of this calculation is less than 1, it is rounded up to a value of 1. Because 100,000,000 is the same as the bandwidth of a Fast Ethernet interface (100 Mbps), all interfaces operating at a faster bandwidth have their calculated cost less than 1, which becomes rounded up to 1.

Change the Cost Calculation

Although each interface can have a cost assigned to it statically, assigning static costs often proves to be an administrative hassle in all but the smallest networking environments. A middle ground is available to administrators who would like an automatic calculation but have large bandwidth links in their network.

The solution is to alter the values used in the bandwidth calculation. This alteration not only allows for an automatic cost assignment but also maintains a consistent ratio across all the router interfaces.

Reference Bandwidth

To alter the calculation values, use the **`reference-bandwidth`** command within the `[edit protocols ospf]` configuration hierarchy. The value entered has a value in bits per second. You can use the standard Junos shortcut notations of *k* (kilobits per second), *m* (megabits), and *g* (gigabits). The entered value becomes the new numerator value in the bandwidth calculation. You still can assign a metric statically to an individual interface.

OSPF Cost

- Cost is used to represent a preference to use a specific interface.
- The lower the cost, the more preferred the link
- Cost can be set on a per-interface basis

```
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface ge-0/0/0.0 {
        metric 12;
    }
    interface ge-0/0/1.100 {
        metric 73;
    }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Cost of an Interface

Prior to advertising a network into the OSPF domain, each router must determine the cost associated with that network. Cost, which is the value used as the OSPF metric, is used to indicated a preference to use a specific link based on its bandwidth. By default, OSPF will automatically calculate a cost of each interface, however, this value can also be set manually and based on any criteria. When the SPF algorithm runs, OSPF will evaluate a lower cost as being more preferred than a higher cost regardless of how the value was set. Because each OSPF-speaking router advertises these cost values in its router LSAs, each router can determine the total cost to reach any destination in the network.

Set on a Per-Interface Basis

If you want to alter the automatic cost calculated by the formula above, you can manually set the cost for each interface. Within the interface portion of the **[edit protocols ospf]** configuration hierarchy, the **metric** command assigns a permanent cost to that interface. Each individual interface on the router can have a separate cost assigned to it.

Effects of Altering Metrics

- Metric values are advertised in Type 1 LSAs and populate LSDB
- As each router runs the SPF algorithm, each LSA is examined individually for the cost of the outgoing interface
  - The final metric calculation uses that cost
- Routers can disagree about the cost on a network link
  - Can result in asymmetric routing in the network
  - R1 sees a cost of 20 to reach the R4 router through R2
  - R4 sees a cost of 20 to reach the R1router through R3

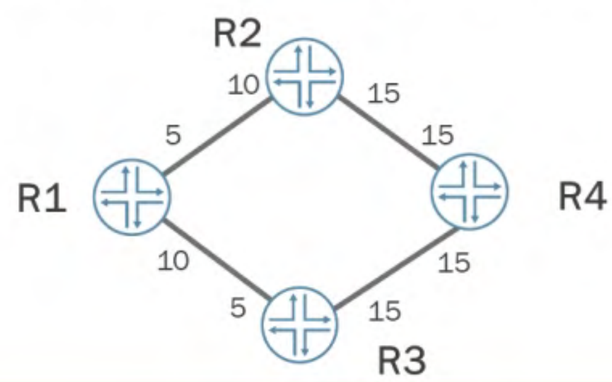

© 2020 Juniper Networks, Inc. All Rights Reserved.

Advertisement of Metric Values

After the OSPF process on a router decides what metric to assign to each interface, that information is flooded into the OSPF domain within a Type 1 LSA. Because these LSAs have an area flooding scope, each router in the OSPF area receives a copy of the metric values.

SPF Calculations

After receiving a new LSA from another router in the area, the local router performs an SPF calculation using all the values contained in the router and network LSAs. As mentioned on a previous slide, the cost is calculated from the root node to every other node in the network using the metric cost of the outgoing interfaces.

Routers Can Disagree

Because each individual router performs a separate SPF calculation, it is possible for two routers on either side of a link to disagree on the metric of that link. The example on the slide shows that each router has decided upon a different metric value for its attached links. This discrepancy results in the R1 router calculating a total cost of 45 to reach the R4 router, while the R4 router calculates a total cost of 60 to reach the R1 router.

While the configuration of different metrics for a single link does not affect the operation of OSPF, asymmetric routing might occur within the network. Asymmetric routing might cause response delays for some end users and make it challenging for network administrators to troubleshoot the network.

Overload Settings

- Used for transit traffic only if no other path is available
  - Sets metric to 65,535 in router LSA on all transit links
  - Flooding of changed LSA causes SPF calculations in network
- Can be set permanently or with a timeout value
  - Timer is between 60 and 1800 seconds
  - Timer only runs after RPD starts

```
[edit protocols ospf]
user@router# show
overload;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface ge-0/1/0.0;
}
user@router> show ospf database router extensive
    OSPF link state database, area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   192.168.56.1     192.168.56.1     0x80000005    71  0x2  0x540b  60
  id 192.168.48.1, data 10.222.61.1, type PointToPoint (1)
  TOS count 0, TOS 0 metric 65535
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Avoid Transit Traffic

OSPF borrows the *overload* concept from the IS-IS protocol. Its function is to advertise information into the OSPF area, but it is not to be used for transit traffic, if possible. This function can be useful in two situations—when a router must be taken out of the network for maintenance, and when the router has many BGP peers. In the first case, traffic should avoid the router because it can compromise its ability to forward traffic. In the second case, a network administrator might want the router to establish its BGP peering sessions fully before handling transit traffic.

Because the overload concept is not native to OSPF, the software is modified to allow for this functionality. When you configure an OSPF-speaking router for overload, the metrics for all transit interfaces are set to the maximum value of 65,535. The overload router floods these LSAs into the network, where an SPF calculation is performed in each router. The large metric values generally ensure that transit traffic through the overload router uses alternative paths through the network. Unlike IS-IS, traffic transits the overload router if no alternative path exists to a given destination.

Overload Settings

You can turn on the overload setting, turn it off as a permanent value, or have a timer associated with it. If the timer is omitted, the metric values are changed when you commit the configuration. The values will remain until you remove the overload setting from the configuration and commit it again. However, if you assign a timer value, the metrics are not changed automatically. The timer associated with the overload setting only initializes when the routing protocol process initializes. This timer can run from 60–1800 seconds. When the timer expires, the metrics return to normal in the router LSAs, but the configuration still contains the **`overload`** option.

Case Study: Overload

- Case study assumptions:
  - We have a meshed network with multiple paths
  - R2 is scheduled for maintenance
  - Overload is configured on R2 during the maintenance
  - Transit traffic is routed around R2

```
user@R1> show ospf database extensive | match "router|metric"
Router  *10.0.4.1          10.0.4.1          0x80000006  1008  0x22 0x813a  48
    Topology count: 0, Default metric: 1
    Topology count: 0, Default metric: 1
      Metric: 1, Bidirectional
      Metric: 1, Bidirectional
Router   10.0.4.2          10.0.4.2          0x80000004  1009  0x22 0xe6d9  48
    Topology count: 0, Default metric: 65535
    Topology count: 0, Default metric: 65535
      Metric: 65535, Bidirectional
      Metric: 65535, Bidirectional
user@R1> show route 10.0.5.1

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.5.0/30        *[OSPF/10] 00:36:25, metric 3
                    > to 10.0.7.1 via ge-1/1/7.0
```

```
[edit protocols ospf]
user@R2# show
overload;
area 0.0.0.0 {
    interface ge-1/1/4.0;
    interface ge-1/0/5.0;
}
```

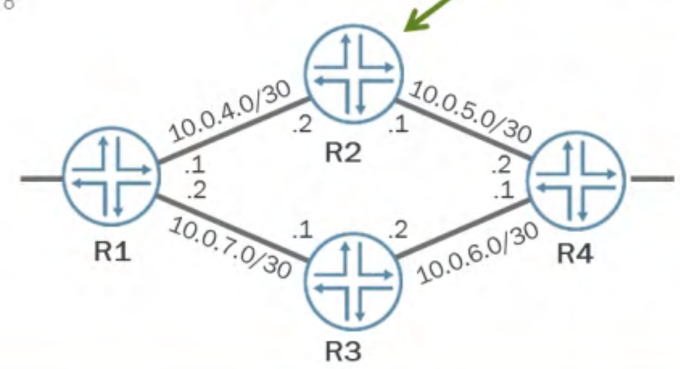

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study

Enterprise networks are typically built with multiple paths from ingress and egress points for redundancy. During maintenance operations on a router, it can be beneficial to prevent the router from receiving and forwarding transit traffic. The overload feature provides this function.

In the graphic on the slide, R2 is scheduled for maintenance. An alternate path exists through R3. Once R2 is put in overload mode the other routers will be notified and transit traffic will traverse R3. Any traffic destined for networks that terminate on R2 will still be forwarded to R2.

Agenda: OSPF

- OSPFv2 Review
- Link-State Advertisements
- Protocol Operations
- →OSPF Authentication
- OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Authentication

The slide highlights the topic we discuss next.

OSPF Authentication

- Three types of authentication are supported: `none`, `simple`, and `MD5`
- By default, the authentication type is set to `none`
  - Effectively means no authentication is performed
- Type `simple` uses a plain-text password

```
[edit protocols ospf]
user@router# show
area 0.0.0.20 {
   interface ge-0/0/2.0 {
        authentication {
             simple-password "$9$vxr8X-Djqz39s24ZDjf5"; ## SECRET-DATA
        }
    }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Authentication

The three different forms of authentication that the Junos OS supports are *none*, *simple*, and *MD5*. As of Junos Release 8.3, IP Security (IPsec) was added.

Authentication Default

The default operation of the OSPF process is the `none` mode. Thus, no authentication key is configured on any interface.

Plain-Text Passwords

With `simple` authentication type configured, each OSPF packet includes a plain-text password. This password can be captured easily through a packet analysis system. Therefore, although this password protects against an inadvertent configuration, it does not provide any security.

MD5 Authentication (1 of 2)

- Includes an encrypted checksum with all packets
  - Provides better security than `simple` type
- Each interface requires an authentication key
  - Multiple interfaces can use the same key
  - Keys are always encrypted in the configuration
- Each key requires a key ID value ranging from 0–255

```
[edit protocols ospf]
user@router# show
area 0.0.0.20 {
    interface ge-0/0/2.0 {
        authentication {
            md5 30 key "$9$wc24ZzF/O1h"; ## SECRET-DATA
        }
    }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Encrypted Checksum

To provide better security in an OSPF network, we recommend that you use an authentication type of MD5. MD5 includes an encrypted checksum in all OSPF packets instead of a simple-text password. Each OSPF-speaking router uses the same MD5 algorithm to calculate the checksum value, so interoperability and a correct result can be virtually guaranteed.

Authentication Keys

The actual password to verify and authenticate packets is contained within the authentication command. You can configure each individual interface with an authentication value.

Key ID Values

You configure each individual interface with a key value. All interfaces in the area might share the same key value, or each interface might contain a separate value.

MD5 Authentication (2 of 2)

- MD5 authentication allows for multiple key ID values
  - Highest value used by default
  - For easy transition, assign each key ID a start time

```
[edit protocols ospf area 0.0.0.0]
user@router# show
interface ge-0/0/1.0 {
    authentication {
        md5 1 key "$9$fQF/SyK7-w"; ## SECRET-DATA
        md5 2 key "$9$.fQntu1ylM/ClM8XbwmfT" start-time "2016-1-20.12:00:00 +0000";
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Multiple Key Values

When configuring multiple key values, you can also specify a start time for the router to begin using a new MD5 key. This setting aids in the graceful transition from one password key to another. The example in the slide displays the format of the **start-time** command. When the local router's time reaches the specified value, the router begins to transmit all OSPF packets using the new key ID value and password. To begin using a new MD5 key ID immediately, you can use the keyword of **now** in place of a specific time and date. The router automatically places the current system time in the configuration for you.

Verifying Authentication

- Authentication information available with the **`show ospf interface detail`** command

```
user@router> show ospf interface detail
Interface            State   Area            DR ID           BDR ID          Nbrs
ge-0/0/1.0           DR      0.0.0.0         10.1.0.111      2.2.2.2            1
  Type: LAN, Address: 10.1.0.111, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  DR addr: 10.1.0.111, BDR addr: 10.1.0.222, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 4, Start time: 2017 Jan 20 12:00:00 UTC
  Protection type: None
  Topology default (ID 0) -> Cost: 1
ge-0/0/2.0           DR      0.0.0.0         10.1.0.111      2.2.2.2            1
  Type: LAN, Address: 10.2.0.111, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  DR addr: 10.2.0.111, BDR addr: 10.2.0.222, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: Password
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

View Authentication Information

The **`show ospf interface detail`** command displays the type of authentication used on the interface with the `Auth type` output. Displaying this output occurs regardless of whether you use per-area or per-interface authentication in your network.

The possible values displayed in the output are `None`, `Password`, and `MD5`. When MD5 authentication is in use, the router also displays the current key ID value and the time at which that ID was first used. If **`start-time`** is not specified in the configuration, the value of the `Start time` field in the **`show ospf interface detail`** command output is the UNIX epoch (1970 Jan 1 00:00:00 UTC).

Agenda: OSPF

- OSPFv2 Review
- Link-State Advertisements
- Protocol Operations
- OSPF Authentication

→OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPFv3

The slide highlights the topic we discuss next.

OSPFv3

- OSPFv3
  - Fundamental mechanics of OSPF unchanged
    - Areas (Regular, Stub, Not-so-Stubby, Totally Stubby)
    - LSA flooding
    - DR elections
    - Options
    - Summarization
    - Junos includes support for Virtual links, Graceful restart, External prefix limits, and Bidirectional Forwarding Detection.
  - Significant changes to account for the difference between IPv4 and IPv6 addressing
  - Configuration mainly requires substituting `ospf3` for `ospf`

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPFv3

OSPFv3 is defined in RFC 5340, with some additional features, such as graceful restart and authentication, defined in separate documents (RFC 5781—OSPFv3 Graceful Restart and RFC 4552— Authentication/Confidentiality for OSPFv3).

OSPFv3 maintains the fundamental mechanisms of OSPF, including LSA flooding scopes, areas, DR election, stub areas, NSSAs, and so on. Though OSPFv3 is often associated with IPv6 addressing, it is also completely compatible with IP version 4 (IPv4) addressing schemes. However, some changes are necessary to account for the differences in IPv4 versus IP version 6 (IPv6) addressing. We address these differences on the following slides.

In terms of the Junos OS configuration, all that is required to convert from OSPFv2 to OSPFv3 is, in most cases, to substitute `ospf3` for `ospf` in the configuration.

```
[edit]
user@router# show protocols ospf?
Possible completions:
> ospf                    OSPF configuration
> ospf3                   OSPFv3 configuration
```

OSPF and OSPFv3 Configuration

IPv4 OSPF

```
[edit protocols ospf ]
area 0.0.0.0 {
    interface ge-0/0/4.0;*
    interface ge-0/0/5.0;*
    interface lo0.0;*
}
area 0.0.0.1 {
    nssa {
        default-lsa default-metric 10;
    }
    interface ge-0/0/6.0;*
}
* Must have a family inet address configured
```

IPv6 OSPF3

```
[edit protocols ospf3 ]
area 0.0.0.0 {
    interface ge-0/0/4.0;**
    interface ge-0/0/5.0;**
    interface lo0.0;**
}
area 0.0.0.1 {
    nssa {
        default-lsa default-metric 10;
    }
    interface ge-0/0/6.0;**
}
**Must have a family inet6 address configured
```

IPv4 OSPF3

```
[edit protocols ospf3 realm ipv4-unicast ]
area 0.0.0.0 {
    interface ge-0/0/4.0;***
    interface ge-0/0/5.0;***
    interface lo0.0;***
}
area 0.0.0.1 {
    nssa {
        default-lsa default-metric 10;
    }
    interface ge-0/0/6.0;***
}
***Must have a family inet address configured
   Must also have family inet6 configured
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF and OSPFv3 Configuration Examples

The slide illustrates the differences of configuring IPv4 in OSPF and OSPFv3. It also shows how to configure IPv6 using OSPFv3.

Note that in order to implement IPv4 only using OSPFv3, you must configure `family inet6` under the interfaces that will be part of the OSPFv3 network. This is because OSPFv3 uses IPv6 Link-Local addresses to pass messages between routers on the same network segment.

OSPFv3 Router ID

- Same as OSPF Router ID
  - OSPFv3 maintains the 32-bit RID that represents the router in the link-state database
    - This is not an IPv4 address, it just looks like one!
  - The RID can't be derived from an IPv6 address as it is possible with IPv4
  - Requires explicit configuration (assuming no IPv4 addresses are present) Set the RID manually under `routing-options`

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPFv3 Router ID

OSPFv3 requires a Router ID to provide device uniqueness within the OSPFv3 database. Like OSPF, every OSPFv3 router has a single RID; which is a 32-bit number in dotted quad notation. If the RID is not explicitly configured under `routing-options`, the Junos OS uses IPv4 addressing to derive a RID based upon the same rules as OSPFv2 (that is, it is likely the loopback address). However, it is conceivable that IPv6 devices be deployed without any IPv4 addresses with which to derive the RID, and so we recommend configuring the router ID explicitly.

Differences from OSPFv2 (1 of 2)

- Differences between OSPFv2 and OSPFv3
  - Use of link-local addresses
    - Used to originate packets
  - Authentication removed
    - Done at the IP (IPv6) layer
  - LSA format changes
    - New LSAs and renaming of old ones
  - Options field expanded
    - V6 bit – Indicate if the link should be excluded from IPv6 route calculations
    - R bit – Used like the Overload bit. Is the originator an active router?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Differences Between OSPFv2 and OSPFv3: Part 1

Though much remains the same, several differences exist between OSPFv2 and OSPFv3:

- *Use of link-local addresses*: OSPF packets are sent using the IPv6 link-local addressing. These link-local addresses are also used as next-hop information during forwarding.
- *Authentication removed*: Authentication uses the IPsec framework built into IPv6. As a result, it is not required at the application layer and was removed.
- *LSA format changes*: Summary LSAs are now referred to as inter-area-prefix LSAs, and ASBR summary LSAs are now inter-area-router LSAs to account for differences in address size. The intra-area-prefix LSA is also included, carrying prefix information internal to areas that were previously carried inside router and network LSAs.
- *The options field was expanded from 8 bits to 24 bits*: The option field is included in OSPF hello packets, database description packets, and certain LSAs (router LSAs, network LSAs, inter-area-router LSAs, and link LSAs). Previously defined option bits are present, as well as added support for the V6-bit and R bits. The V6-bit is used to indicate whether the route or link should be excluded from IPv6 routing calculations. The R bit is used like the IS-IS overload bit and indicates whether the originator is an active router. If the R bit is clear (that is, 0) in the OSPF options field, the advertising router can participate in OSPF without being used for transit traffic. This would be a useful setting for hosts that are multihomed but never used to forward traffic between interfaces.

Differences from OSPFv2 (2 of 2)

- Differences between OSPFv2 and OSPFv3 (cont):
  - Protocol processing per link, not per subnet
  - Removal of addressing semantics
    - Router and network LSAs have no addressing
    - Uses intra-area-prefix LSA to carry addressing information
  - Unknown LSA handling – Defines whether to keep unknown LSAs local or to forward them. This information is carried in the U bit of the LSA type field
  - Flooding scope is encoded in the LSA type field using the S1 and S2 bits

S2	S1	Flooding Scope
0	0	Link-Local scope
0	1	Area scope
1	0	AS scope
1	1	Reserved

© 2020 Juniper Networks, Inc. All Rights Reserved.

Differences Between OSPFv2 and OSPFv3: Part 2

Further differences between OSPFv2 and OSPFv3:

- *Protocol processing per link, not per subnet*: You need only a single adjacency per link even if multiple IPv6 subnets exist on the link. These adjacencies are formed using link-local addresses, which removes the requirement of having matching subnet and subnet mask configurations for two routers to become adjacent.
- *Removal of addressing semantics*: Addressing semantics were removed from OSPF headers, and no prefix information is carried in the router LSAs or network LSAs. This change makes OSPFv3 somewhat protocol independent. The router and network LSAs now express the topology in a protocol-independent fashion. To carry the equivalent information that was carried in these LSAs with OSPFv2, a new LSA called the intra-area-prefix LSA is introduced. The intra-area-prefix LSA carries the IPv6 addressing information.
- *Unknown LSA handling*: The old way of discarding unknown LSA types is no longer supported because a mix of capabilities in a network, especially on a single link, causes forwarding issues. OSPFv3 codes the proper handling of an unknown LSA type in the LSA handing bit. The LSA is either treated as being of local scope only or stored and forwarded as if it were understood.
- *Flooding scope is generalized*: Flooding is now generalized and is coded into the LS type field. An LSA can be either flooded on the local link, area, or throughout the AS. This flooding also assists in the handling of unknown LSA types as the flood scope is encoded.

OSPFv3 LSA Types

LS type field:

LSA Function	LS type	Description	Like OSPFv2
1	0x2001	Router LSA	Type 1 Router LSAs
2	0x2002	Network LSA	Type 2 Network LSAs
3	0x2003	Inter-Area-Prefix LSA	Type 3 Summary LSAs
4	0x2004	Inter-Area-Router LSA	Type 4 ASBR-Summary
5	0x4005	AS-External-LSA	Type 5 AS-External LSA
6	0x2006	Group Membership LSA	Type 6 Multicast
7	0x2007	Type-7 LSA	Type 7 NSSA External
8	0x0008	Link LSA	None
9	0x2009	Intra-Area-Prefix LSA	Types 1 and 2

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPFv3 LSA Types

The link state (LS) type field is encoded in the LSA header and indicates the flooding behavior and the information contained in the LS update. LSA types are now encoded with LSA properties and an LSA function code.

LSA properties include:

- *U bit*: Used to indicate how a router that does not understand the LS function code should handle the LSA. When the code is 0, the router should treat the LSA as if it has link-local scope only. When the code is 1, the router should store and flood it as if it were understood.
- *S-bits and flooding scope*: Used to indicate the flooding scope for the LSA:

S-bits and flooding scope

S2	S1	Flooding scope
0	0	Link-local scope.
0	1	Area scope.
1	0	AS scope.
1	1	Reserved.

The function codes are very much like the OSPFv2 type codes. Some exceptions, however, do exist. Most notably, all addressing semantics are removed from the router and network LSAs. The addresses that were advertised in the OSPFv2 versions of the LSAs are now carried in a new intra-area-prefix LSA.

Another addition is the presence of the link LSA. The link LSA is used to advertise to directly attached OSPF neighbors the link-local IPv6 address assigned to the interface. As indicated by the LS type, the link LSA is not flooded beyond the physical broadcast domain.

Summary

- In this content, we:
  - Described the various OSPF LSA types
  - Explained the flooding of LSAs in an OSPF network
  - Described the SPF algorithm
  - Explained OSPF link metrics
  - Described the various OSPF authentication methods
  - Explained the differences between OSPFv2 and OSPFv3

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Discussed:

- OSPF LSAs;
- How LSAs are flooded throughout the network;
- The SPF algorithm;
- OSPF link metrics;
- OSPF authentication methods; and
- The differences between OSPFv2 and OSPFv3.

Review Questions

1. Which LSA is used to support graceful restart?
2. How can you automatically alter the metric of the router's links?
3. What are the different forms of OSPF authentication?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review Questions

1.

2.

3.

Lab: OSPF Multiarea Networks

- Configure, monitor, and troubleshoot the operation of an OSPF network using multiple areas.
- Alter metrics, configure overload, and configure authentication.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab: Configuring and Monitoring OSPF

The slide lists the objectives for this lab.

Answers to Review Questions

1.

LSA Type 9 supports graceful restart.

2.

The metric or cost of a router's links can be automatically altered with the **`reference-bandwidth`** command.

3.

The different forms of OSPF authentication include none, simple, and MD5.

Education Services

Advanced Junos Service Provider Routing

Chapter 3: Advanced OSPF

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - Describe OSPF area types and operations
  - Configure various OSPF area types
  - Summarize and restrict routes

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Will Discuss:

- OSPF area types and operations;
- How to configure various OSPF area types; and
- How to summarize and restrict routes.

Agenda: OSPF Areas

→Review of OSPF Areas
- Stub Area Operation
- Stub Area Configuration
- NSSA Operation
- NSSA Configuration
- Route Summarization

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review of OSPF Areas

The slide lists the topics we will discuss. We discuss the highlighted topic first.

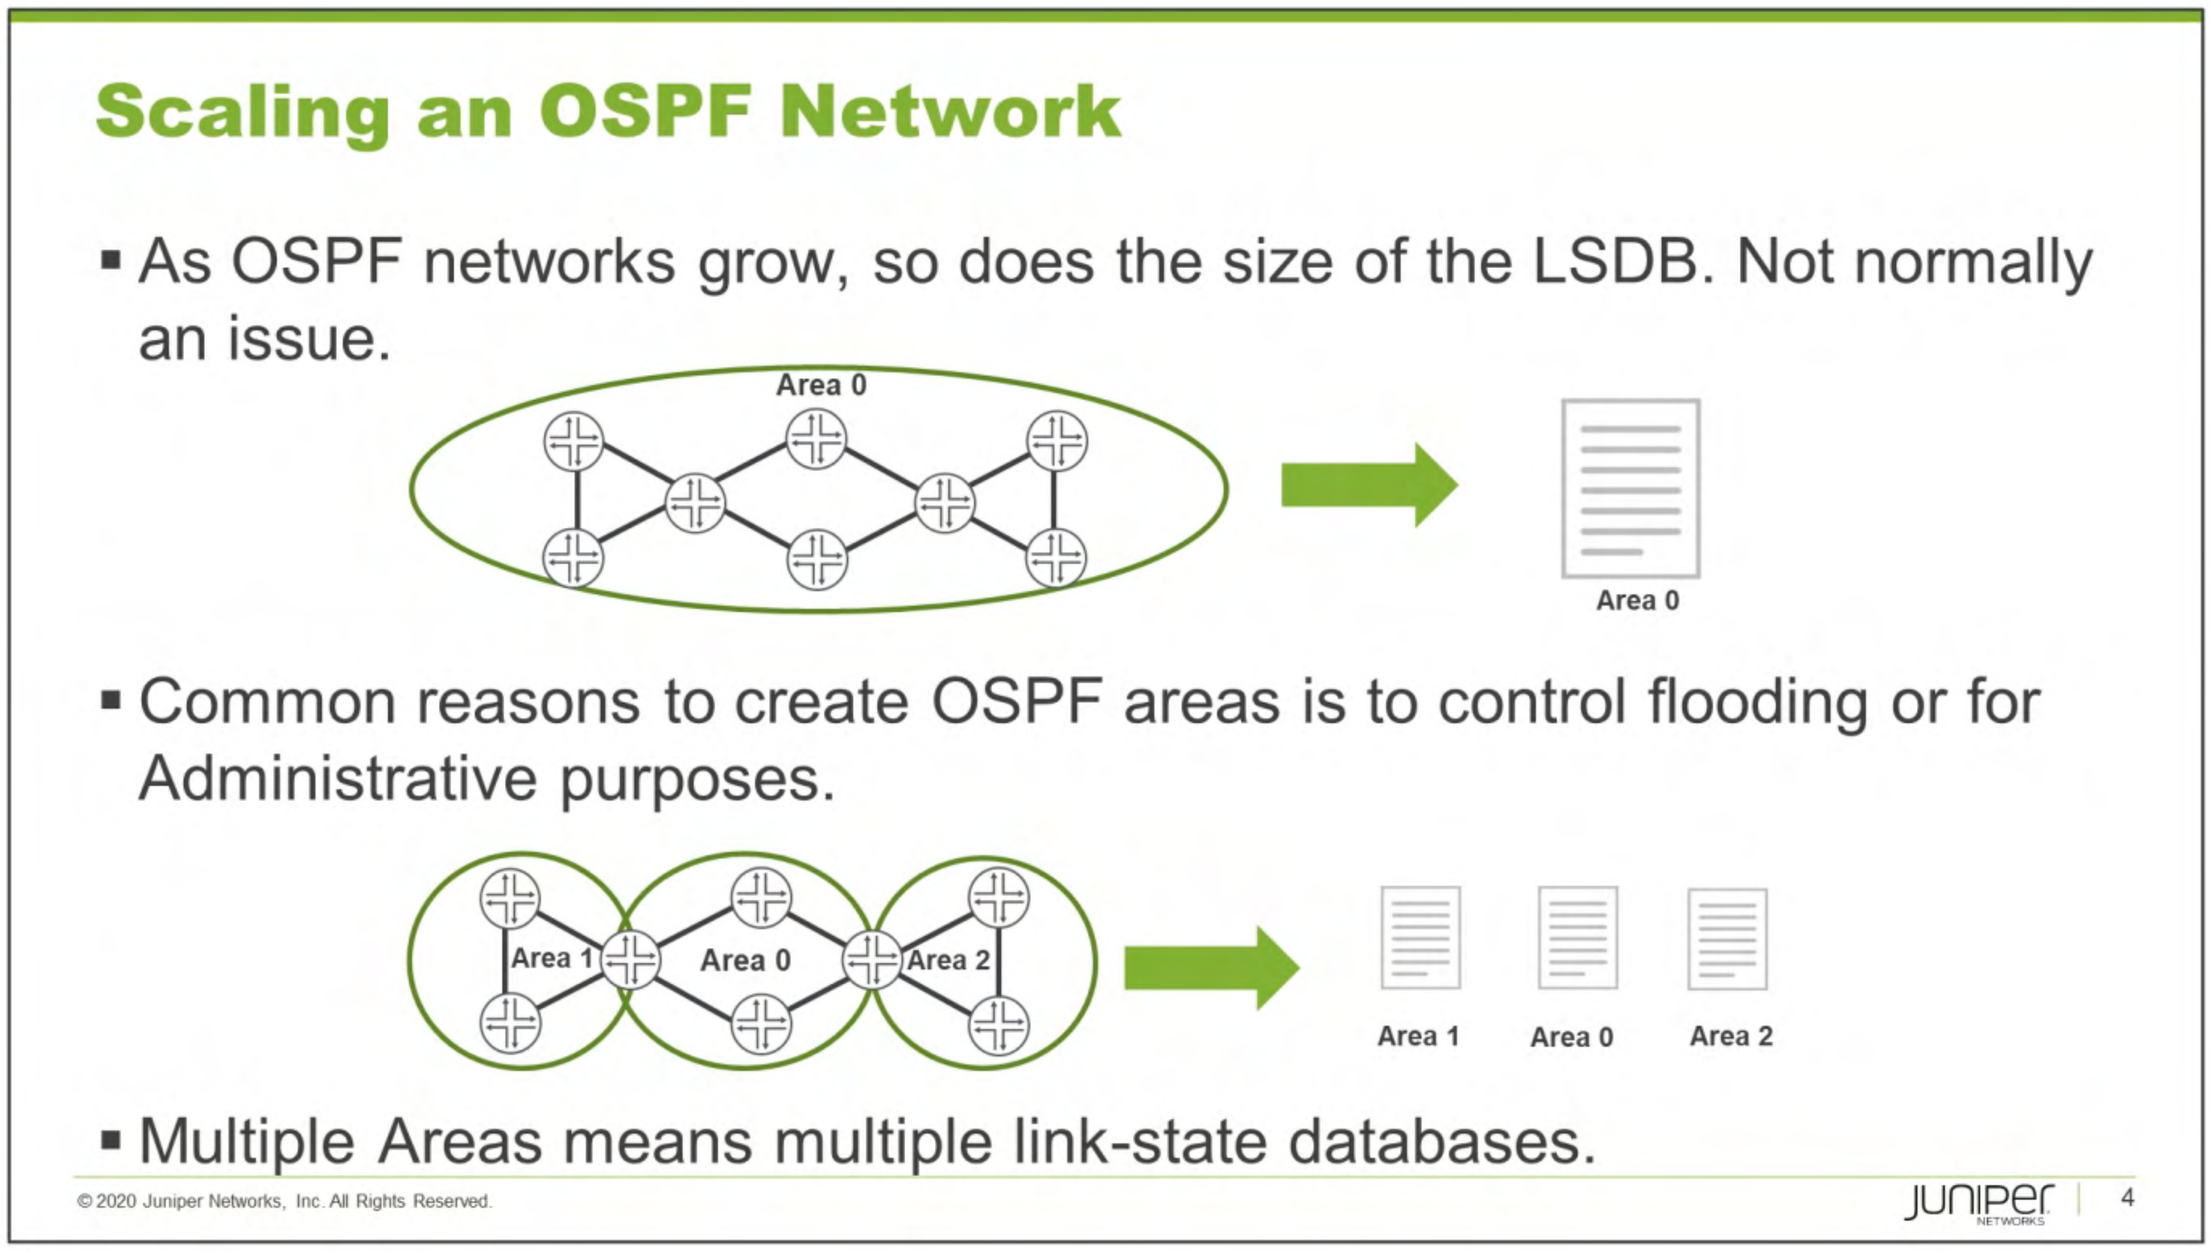

OSPF Scalability

With a link-state protocol, flooding link-state update packets and running the shortest-path-first (SPF) algorithm consumes router resources. In the early years of OSPF this was a large problem as the routers had limited processing power and memory to handle the computations required to calculate the best path while also forwarding packets. By creating separate areas and applying route summarization, companies could reduce the size of the Link State Database (LSDB) easing the demand on router resources. Today routers have been optimized with the separation of the control and forwarding planes. They also have more powerful processors and support more memory, so networks can grow much larger and remain as a single area. While it is possible for networks to still grow to a size where supporting a single OSPF area is an issue, in todays networks, areas are more likely implemented to limit flooding or for administrative reasons.

Multiarea Requirements

If creating a single or dual area OSPF network, there is no requirement for an area to be designated as Area 0. However best practice recommends the primary area should be designated as Area 0 as this area is referred to as the backbone area and has special rules when used in multiarea designs.

In multiarea designs the rule is that all other OSPF areas must connect themselves to the backbone area. This requirement provides the loop prevention mechanism of a multiarea design. By default, all data traffic between OSPF areas transits the backbone. While there are exceptions to this direct connect rule, such as implementing virtual links, this solution is meant to be temporary. Virtual links are not covered as part of this module.

Shrinking the Link-State Database

Creating multiple areas does provide an opportunity to shrink the size of the OSPF LSDB for each area. There are various OSPF area types that help in limiting LSA flooding. We discuss the various OSPF area types on a subsequent slide.

In addition to adding new OSPF areas that restrict link-state advertisement (LSA) flooding, you can also perform route summarization on the borders between OSPF areas. Route summarization has two key benefits:

- It can further reduce the size of the LSDB; and
- It can hide some instabilities in one OSPF area from all other OSPF areas.

For route summarization to be effective, you must carefully plan the addressing within your OSPF network so that subnets can be more easily summarized. Complete coverage of route summarization is beyond the scope of this course.

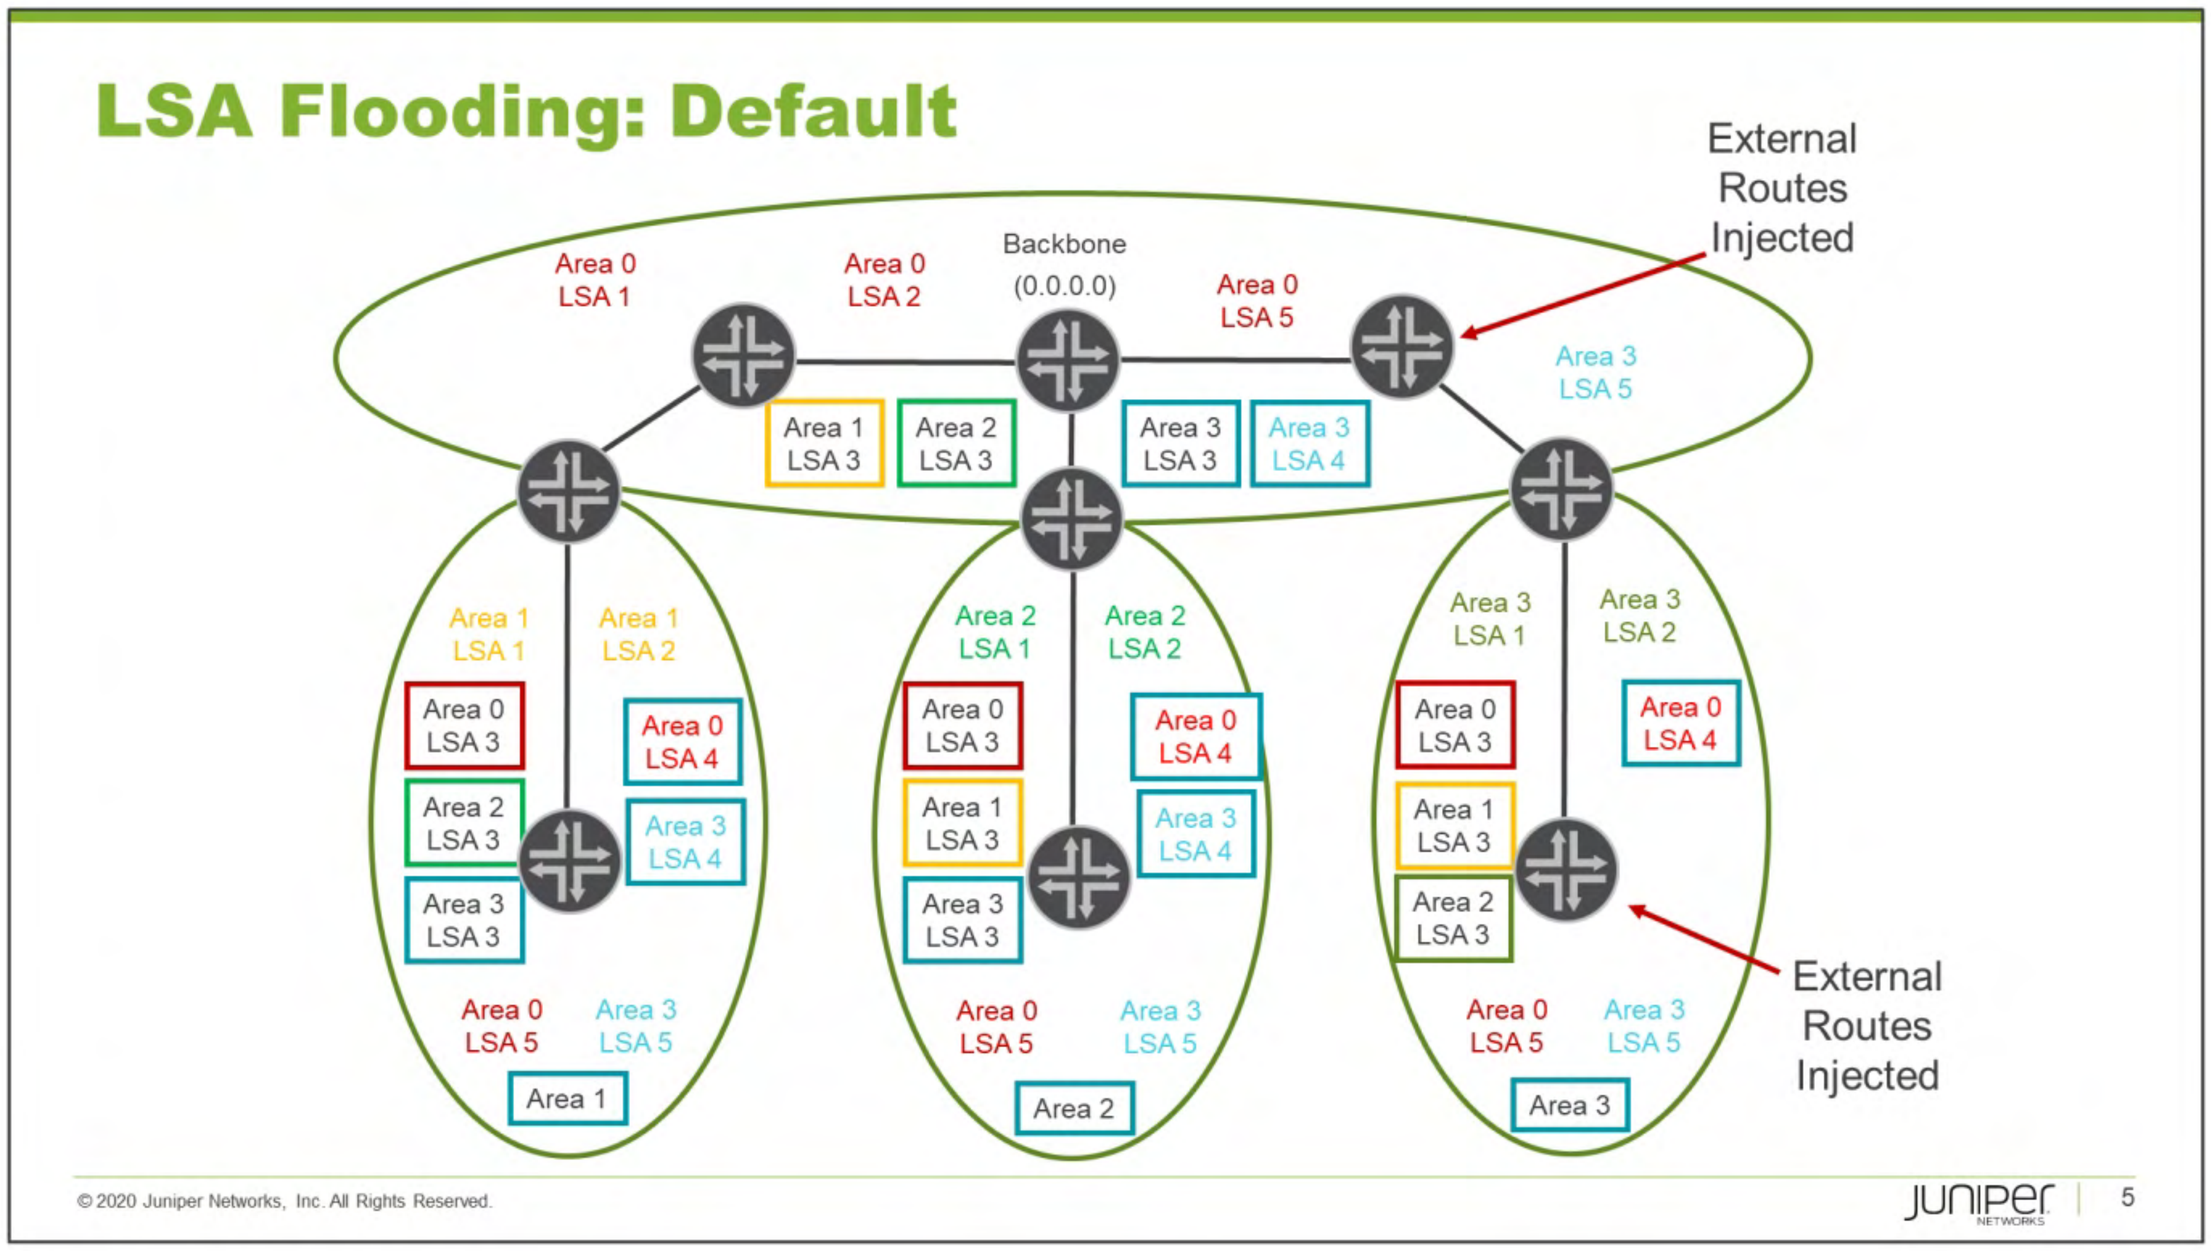

OSPF Default Flooding Rules

As OSPF interfaces are added into each router's configuration, T1 and T2 LSAs are created. In this environment, routes from one area are unable to reach another area. To resolve this issue the ABR creates summary LSAs (T3) to represent the prefixes from one area and inject those T3 LSAs into the adjoining area.

In the example above, the T1 and T2 LSA from Area 0 are injected into Areas 1, 2 and 3. In this way, devices in Area 1, 2, and 3 will be able to reach prefixes in Area 0, but not to one another. To overcome this issue, the ABR connected to each area will create a T3 LSA for each T1, and T2 in the non-backbone area and forward them into the backbone area as T3 LSAs. Now this allows devices from Area 0 to access devices in the non-backbone areas, but still has not solved the problem of the non-backbone areas communicating with each other.

In order to achieve this inter-area communications the ABR connected to Area 0 has a special capability. An ABR connected to Area 0 is the only ABR that can take a T3 LSA from one area and create a new T3 LSA and forward it into another non-backbone area. In this way all areas are able to reach the networks in every other area.

External routes injected into Area 0 become T5 LSAs. T5 LSAs have a domain scope so ABRs forward T5 LSAs into all areas. However the other areas will not know how to reach the advertising router. To overcome this issue, ABRs will create a supplemental T4 LSA to provide the SPF algorithm with information on how to forward the traffic towards the advertising router.

To manage the default flooding of LSA, special area type can be used.

OSPF Areas Types

- Stub
- Totally Stubby
- Not-so-Stubby (NSSA)
- Totally Not-so-Stubby

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Area Types

Using OSPF, you can segment an AS into smaller groups of routers known as *areas*. Using areas achieves the OSPF hierarchy that can facilitate growth and scalability. In order to simplify the segmentation process, OSPF has special designations that can be applied to areas that will automatically limit flooding into those areas. These special areas area classified as follows:

- Stub Areas
- Totally Stubby Areas
- Not-so-Stubby-Areas (NSSA)
- Totally Not-so-Stubby Areas

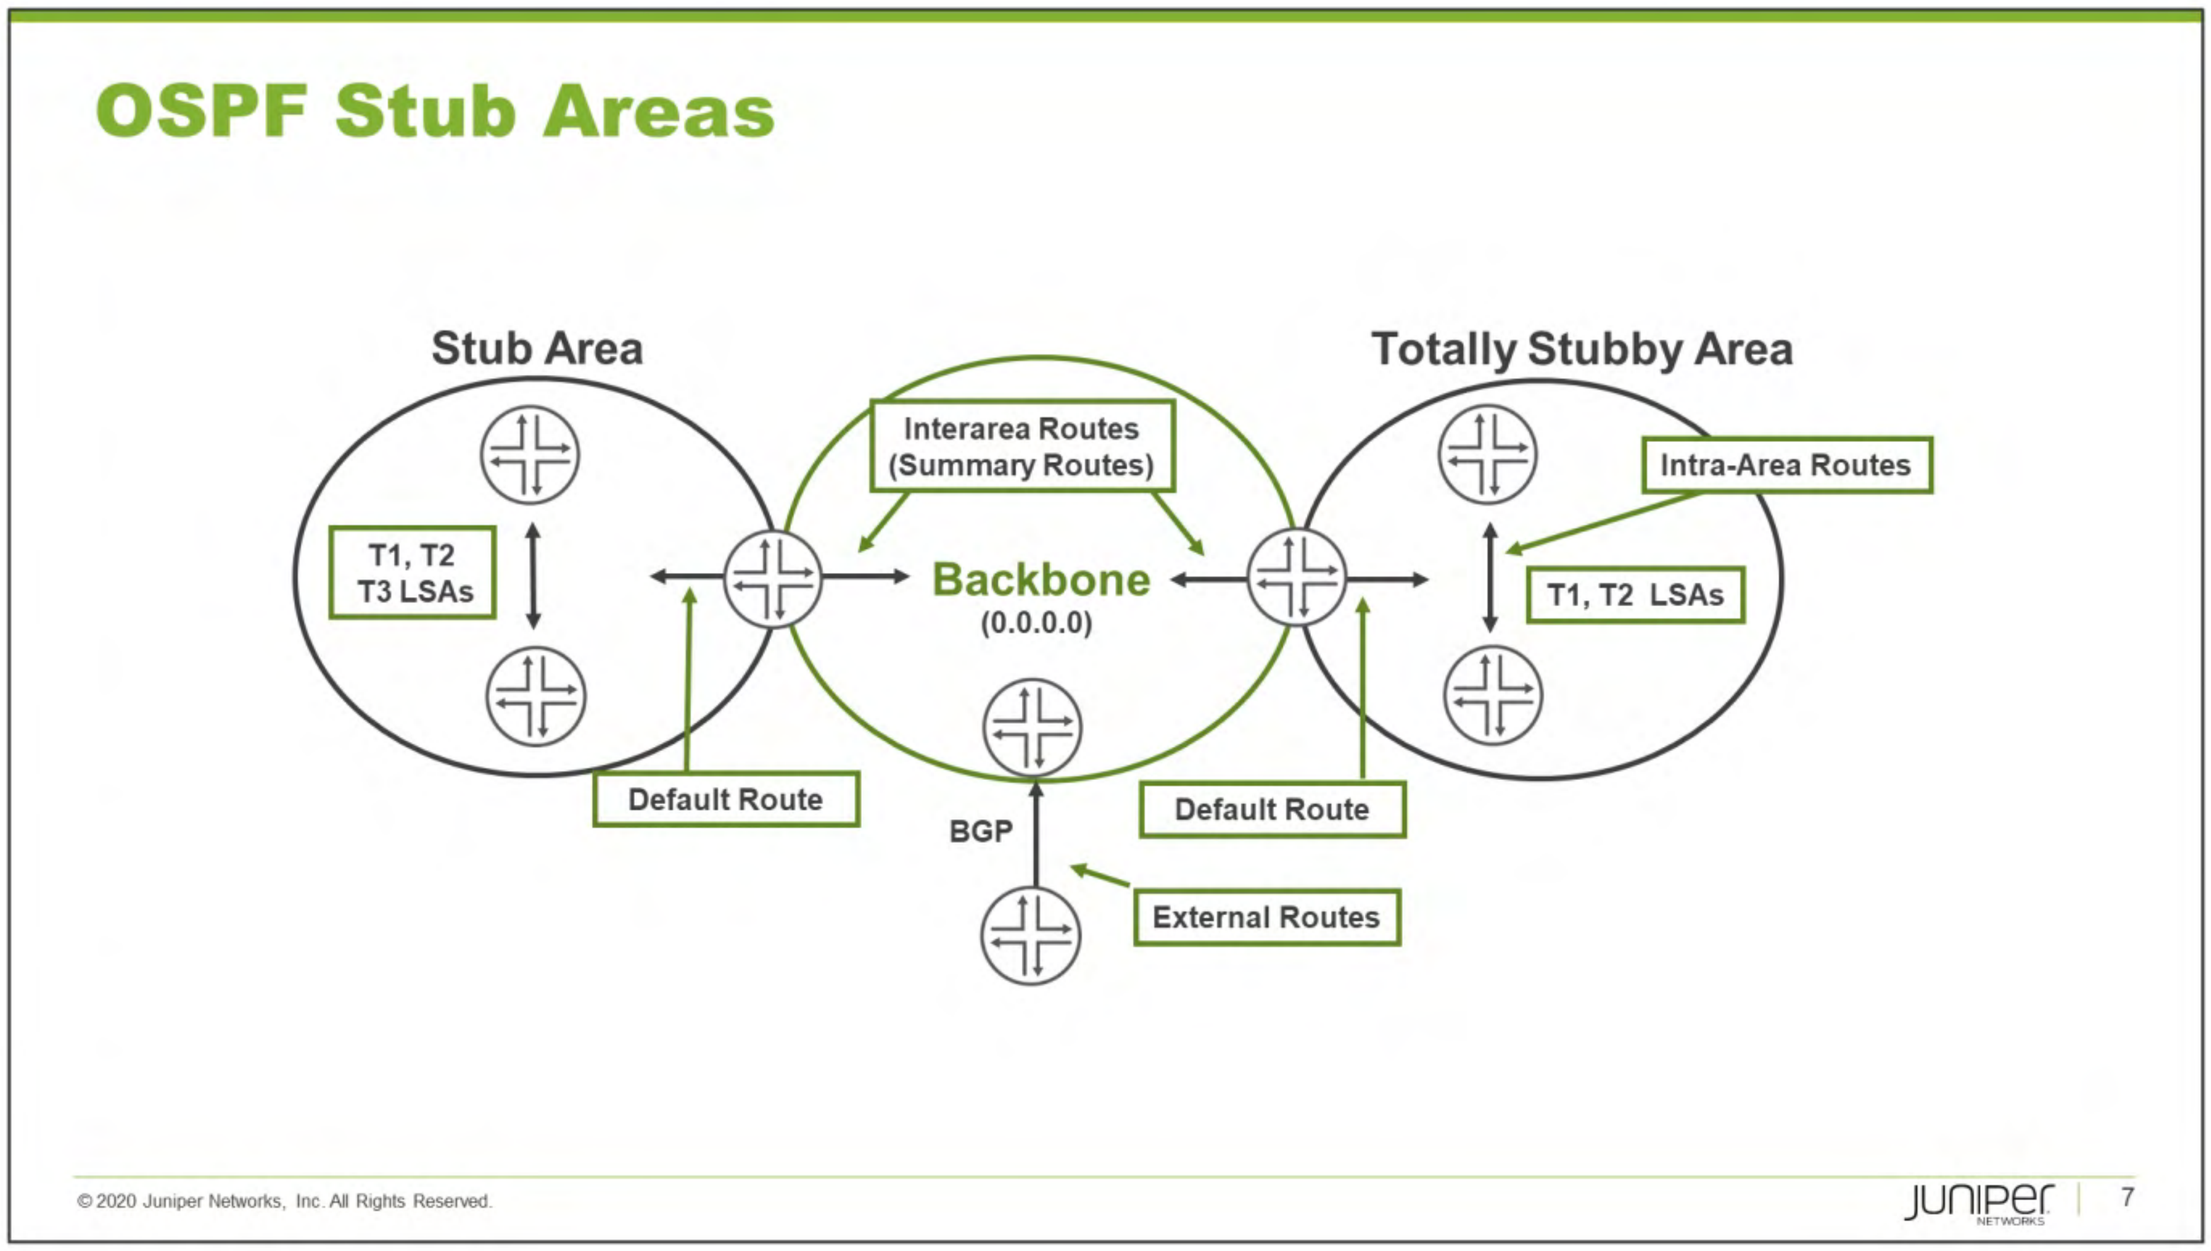

OSPF Area Types

This slide introduces some OSPF area types and illustrates the relationship between these areas, including the types of routes exchanged between them. On the slide, all areas are connected directly to the backbone. In the rare situations where a new area is introduced that cannot have direct physical access to the backbone, you can configure a virtual link. Virtual links are covered in a different chapter.

OSPF classifies different types of routing information as follows:

- *Intra-area* or *internal* routes: Routes that are generated from within an area, where the destination belongs to the area;
- *Inter-area* or *summary* routes: Routes that originate from other areas; and
- *External* routes: Routes that originate from other routing protocols, or different OSPF processes, and that are injected into OSPF through redistribution.

OSPF Stub Areas

A Stub areas can contain intra-area route information if the form of T1, T2 LSAs and Inter-area routing information in the form of T3 LSAs.

Totally Stubby Area

A Totally Stubby area only contains intra-area route information in the form of T1, T2 LSAs. A single T3 LSA is possible if a default route has been configured to be injected by the ABR.

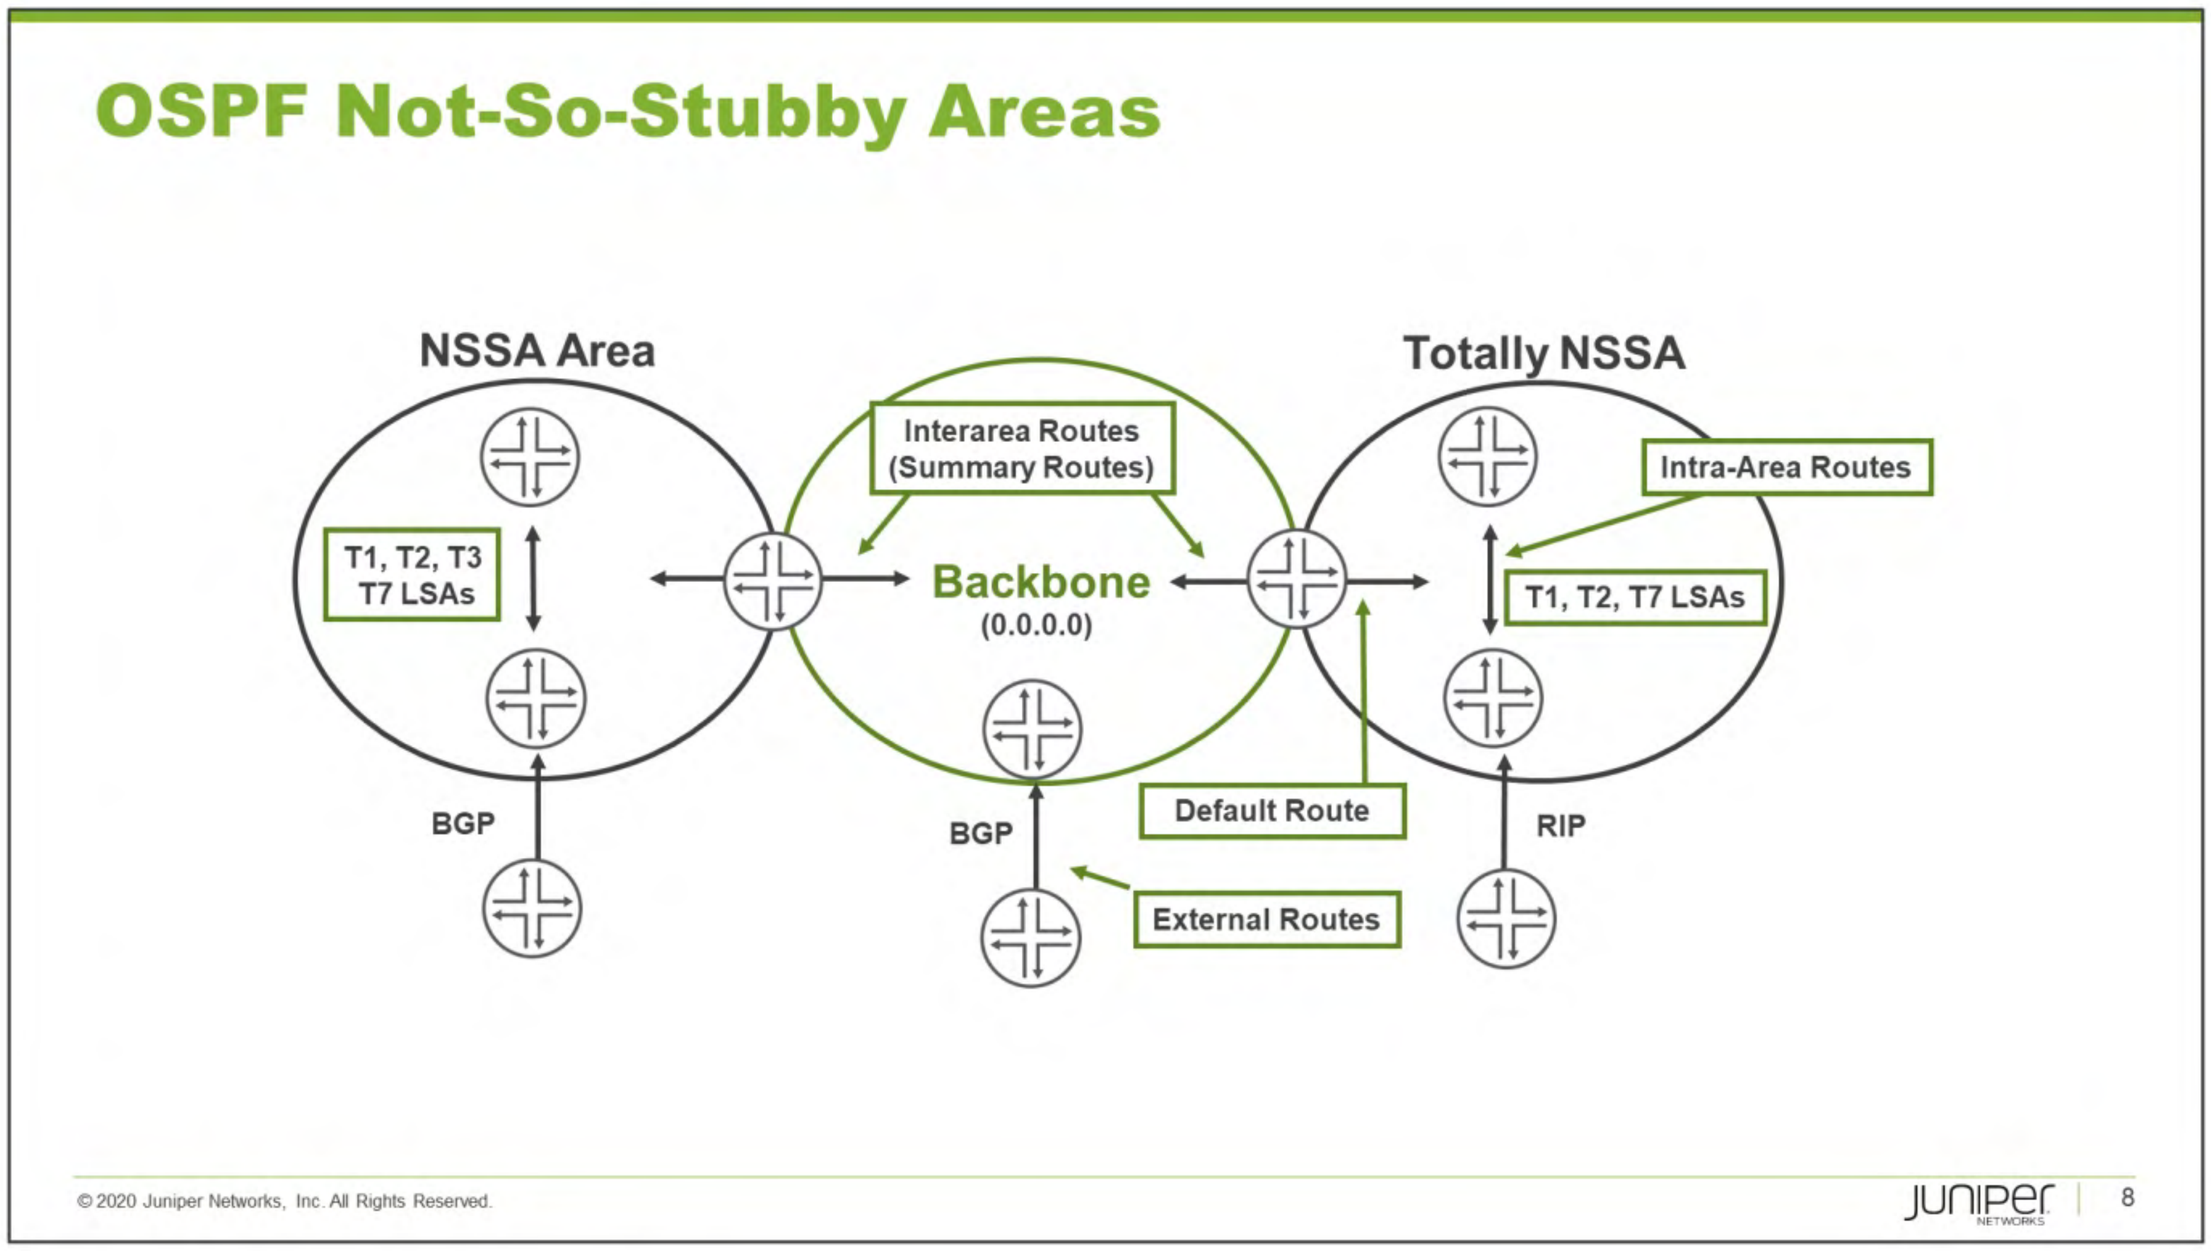

OSPF Not-so-Stubby Areas (NSSA)

An NSSA will have intra-area route information in the form of T1 and T2 LSAs. Inter-area route information is carried in T3 LSAs, A NSSA can also support external route information in the form of T7 LSAs. A default route injected into the NSSA will be a T7 LSA.

Totally Not-so-Stubby Area

A Totally NSSA will have intra-area route information in the form of T1 and T2 LSAs. A Totally NSSA can also support external route information in the form of a T7 LSA. A default route injected into the Totally NSSA as a T3 LSA, but can be changed back into a T7, if desired.

Agenda: OSPF Areas

- Review of OSPF Areas
- →Stub Area Operation
- Stub Area Configuration
- NSSA Operation
- NSSA Configuration
- Route Summarization

© 2020 Juniper Networks, Inc. All Rights Reserved.

Stub Area Operation

The slide highlights the topic we discuss next.

Stub Areas

- Reduces the size of the LSDB
  - ABR does not inject Type 4 LSA into area
  - ABR does not flood Type 5 LSA into area
- Reachability for routes external to OSPF is achieved using a 0/0 default route injected by the ABR
  - Manual configuration step for added administrator control
- ASBR in a stub area cannot flood LSAs for external routes
- Virtual links cannot transit a stub area

© 2020 Juniper Networks, Inc. All Rights Reserved.

Reduce LSDB Size

The operation of an OSPF stub area reduces the size of the LSDB within that area by eliminating AS external routing information. When all of an area's routers agree to operate in stub mode, the ABR stops forwarding Type 5 external LSAs into the area. The ABR also stops generating Type 4 summary LSAs, which are no longer needed because of the suppression of Type 5 external LSAs.

ABR Provides Reachability

The routers within the stub area still might require IP reachability to the external routes they no longer have in their databases. A 0.0.0.0/0 default route generated by the ABR in the area accomplishes this reachability. Within the Junos OS, the advertisement of a default route does not occur automatically; this allows for better control of which ABR, if any, should be used for outbound traffic flows.

No ASBRs in a Stub Area

Because the definition of a stub area does not allow the use of external LSA information within the area, no functional AS boundary routers can exist within a stub area. If any ASBR configuration exists, the router generates one or more Type 5 LSAs and places them into its local database. These external LSAs will not be flooded out any interfaces configured as part of the stub area.

No Virtual Links in a Stub Area

For similar reasons, a stub area cannot support a virtual link because the area attached to the backbone through the stub area might require external LSA information. Because the *transit* routers cannot forward the Type 5 LSAs, the far-end area would not receive the correct information.

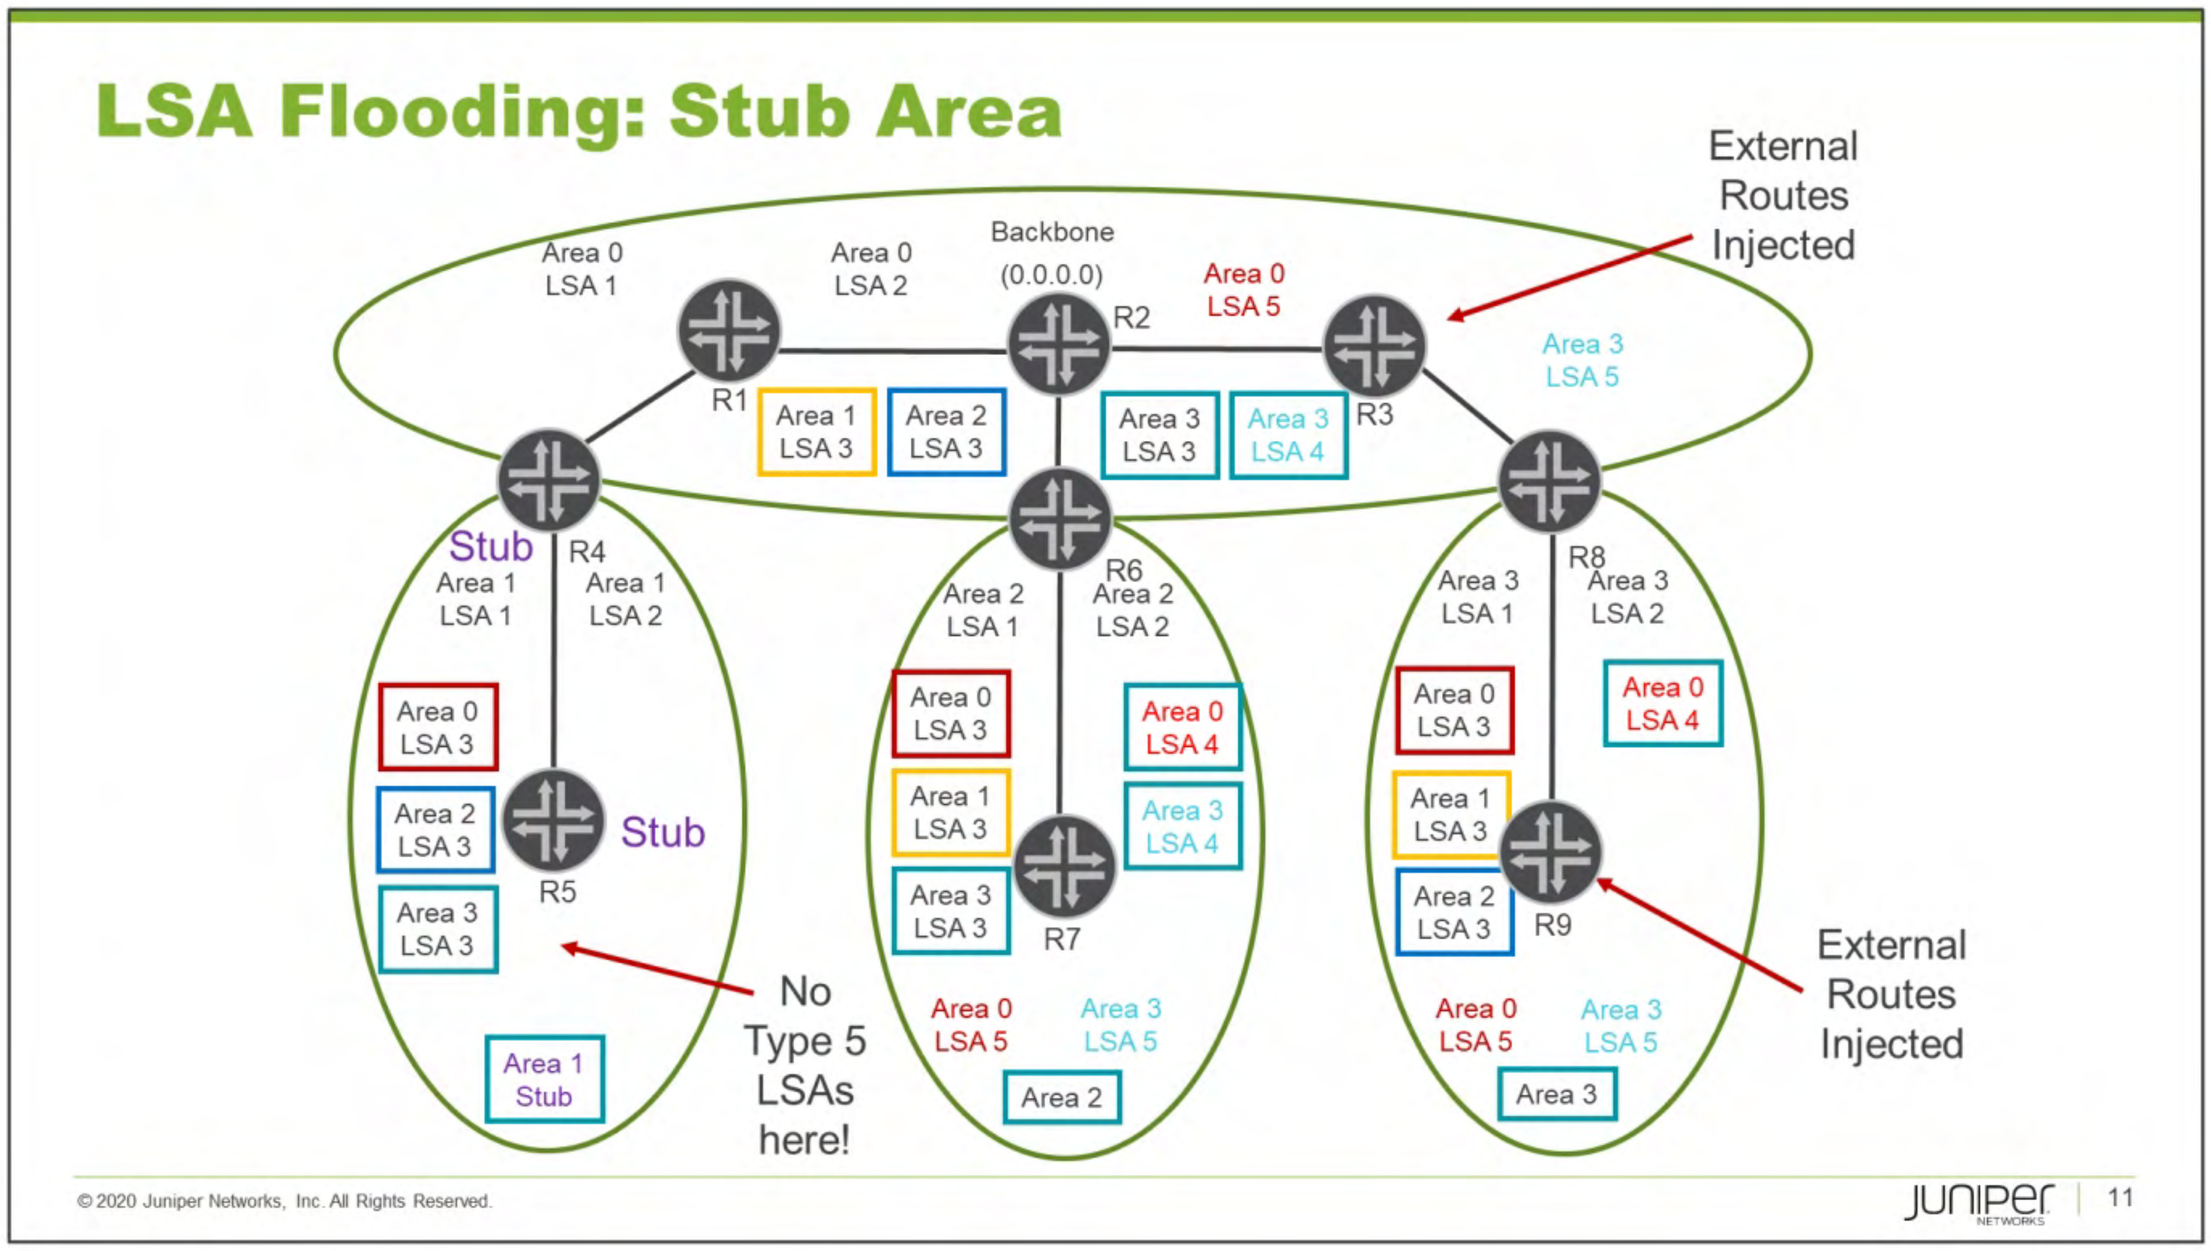

Stub Area Flooding Scopes

In the example on the slide, Area 1 is now configured as a stub area. This configuration causes the ABR of Area 1 to stop forwarding the Type 5 LSAs from Area 0 and Area 3. Because the routers within Area 1 no longer have the external LSA information, they also no longer require the knowledge of any ASBRs in the network. As such, the Area 1 ABR stops generating Type 4 LSAs for reachability of the ASBRs in Area 0 and Area 3.

Totally Stubby Area

- Stub area configured with the **`no-summaries`** option
- Further reduces the size of the LSDB
  - ABR does not inject Type 3 LSA into area
  - ABR does not inject Type 4 LSA into area
  - ABR does not flood Type 5 LSA into area
- Reachability for external routes is available using a 0/0 default route injected by the ABR
  - Again, a manual configuration step exists for administrator control
- ASBR in a **`no-summaries`** area cannot flood LSAs for external routes
- Virtual links cannot transit a **`no-summaries`** area

© 2020 Juniper Networks, Inc. All Rights Reserved.

Further Reduces LSDB Size

The operation of an OSPF stub area with no summaries further reduces the size of the LSDB within that area by also omitting Type 3 summary LSAs. This behavior has the effect of the routers within the area only having Type 1 router and Type 2 network LSAs in their databases. As noted previously, this type of area is known as a totally stubby area because a default route is now needed to reach AS external and interarea prefixes.

ABR Provides Reachability

The routers within a stub area with no summaries still might require IP reachability to interarea and external routes that are no longer represented in their databases. A 0.0.0.0/0 default route generated by the ABR in the area accomplishes this reachability. Within the Junos OS, this advertisement is a manual step for the network administrator that allows for better control of which ABR, if any, should be used for traffic flowing to interarea or external destinations.

No ASBRs in a Stub Area

As before, no functional AS boundary routers can exist within a stub area, regardless of whether summaries are permitted.

No Virtual Links in a Stub Area

Like a stub area, a stub area with no summaries cannot support virtual links.

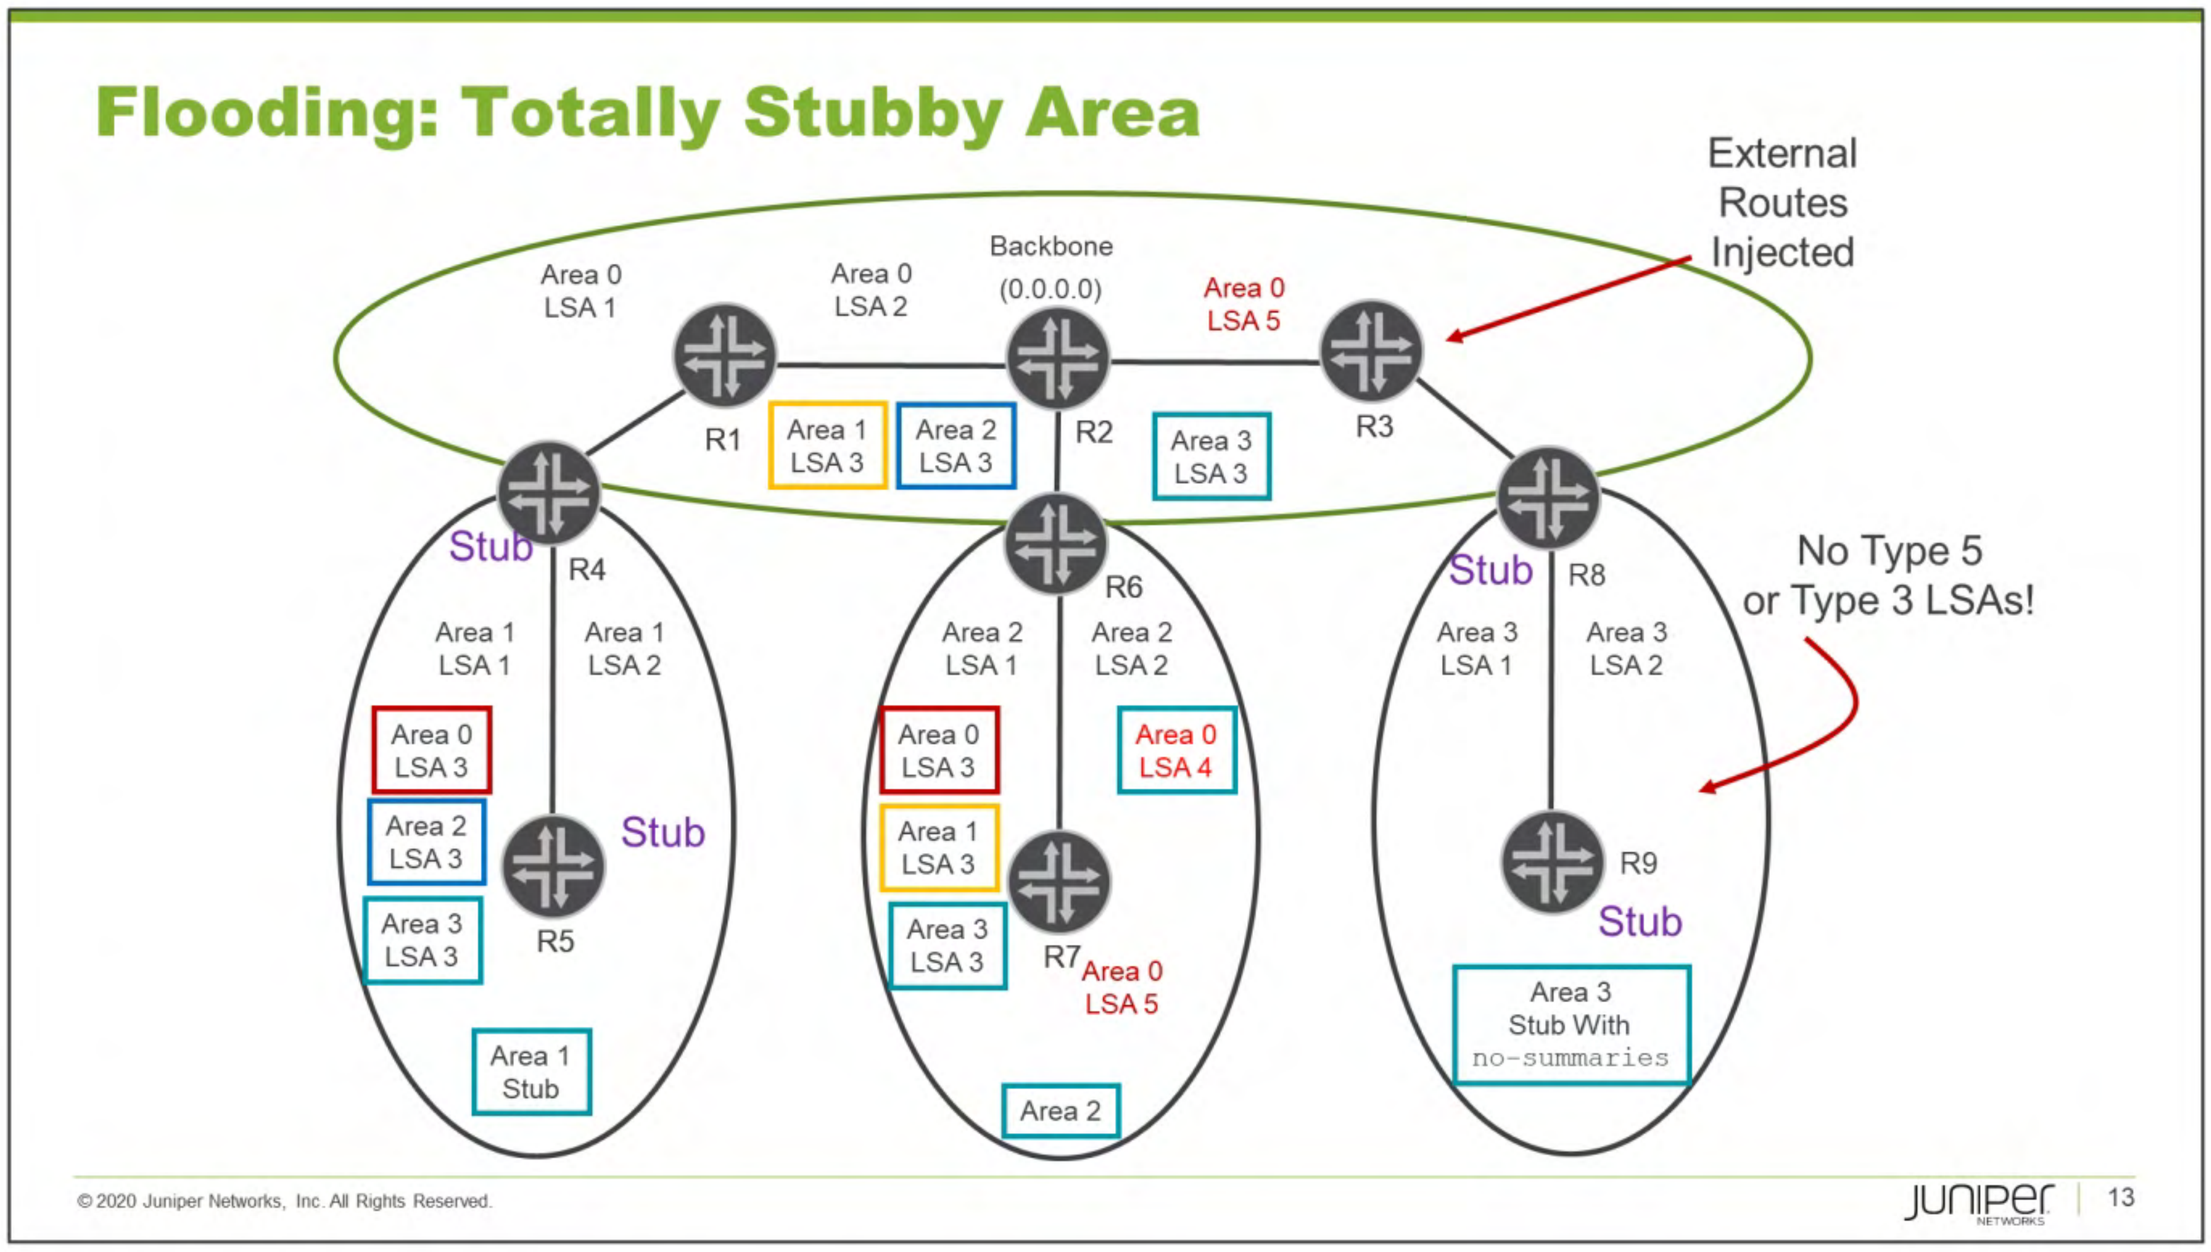

Flooding Scopes for Totally Stubby Areas

The slide shows that Area 1 is still configured as a stub area as shown on a previous slide. In addition, Area 3 has now been configured as a stub area with no summaries.

As discussed previously, the ABR of Area 3 has stopped forwarding the Type 5 LSAs from Area 0 and has also stopped generating the Type 4 LSAs from Area 0. In addition, Area 3's ABR no longer generates Type 3 LSAs for the other OSPF areas. This fact has left the LSDB for Area 3 quite small.

Another interesting phenomenon has occurred as well. Recall that one of the Area 3 routers previously acted as an ASBR. Although the ASBR portion of that router's configuration did not change, its operation changed into a stub network. This change causes the router to generate Type 5 LSAs and place them into its own database. However, because all interfaces are now operating in stub mode, the router has no ability to forward them. As such, notice that the entire network has lost knowledge of the Type 5 LSAs from Area 3.

Agenda: OSPF Areas

- Review of OSPF Areas
- Stub Area Operation
- →Stub Area Configuration
- NSSA Operation
- NSSA Configuration
- Route Summarization

© 2020 Juniper Networks, Inc. All Rights Reserved.

Stub Area Configuration

The slide highlights the topic we discuss next.

Stub Area Configuration

- All routers in a stub area must be configured as stub routers
- The ABR injects a default route when the **`default-metric`** statement is added
  - A default route is not automatically generated!

Configured on all routers in Stub area

```
[edit protocols ospf]
user@R5# set area 1 stub

[edit protocols ospf]
user@R5# show
area 0.0.0.1 {
    stub;
    interface ge-0/0/0.0;
}
```

Configured on ABR Only

```
[edit protocols ospf]
user@R4# set area 1 stub default-metric 10

[edit protocols ospf]
user@R4# show
area 0.0.0.1 {
    stub default-metric 10;
    interface ge-0/1/1.0;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configuration on All Routers

For an OSPF stub area to function successfully, you must configure each router in that area to treat the area as a stub. This configuration turns on the E-bit (sets it to one) in the options field of the OSPF hello packet. This notifies all neighbors that the local router does not support external LSAs. An OSPF router only forms an adjacency with another router when the E bit values match—hence, the requirement that all routers agree on whether the area is a stub area.

ABR Injects a Default Route

The ABR of the stub area can optionally inject a 0.0.0.0/0 default route into the stub area. Using the **`default-metric`** command accomplishes this task, which causes the ABR to generate a Type 3 summary LSA advertising the 0.0.0.0/0 route with the associated metric attached.

The need to manually assign the metric before the default route is generated provides additional control. The administrator can control how, or if, traffic will leave the stub area by controlling which (if any) ABRs advertise a default route as well as the metric for that route.

Stub Area with `no-summaries` Configuration

- Only ABRs are configured to support stub areas with no summaries
- The ABR can inject a default route when the **`default-metric`** statement is included
  - Like a conventional stub area, the default route is not generated automatically

Configured on ABR Only

```
[edit protocols ospf]
user@R2# set area 1 stub no-summaries default-metric 10

[edit protocols ospf]
user@R8# show
area 0.0.0.1 {
    stub default-metric 10 no-summaries;
    interface ge-0/0/1.0;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configuration on the ABR Only

Because the ABR generates new Type 3 summary LSAs for injection into the areas it serves, you configure the **`no-summaries`** option only on ABR routers by adding the **`no-summaries`** command to the stub area's stanza. Because non-ABR routers in the area do not require the suppression of Type 3 LSAs, you do not have to perform this configuration on those routers.

ABR Injects a Default Route

The ABR of an area configured as a stub with no summaries can optionally inject a 0.0.0.0/0 default route into the stub area. Use the **`default-metric`** command to accomplish this task, as described on previous pages.

Agenda: OSPF Areas

- Review of OSPF Areas
- Stub Area Operation
- Stub Area Configuration
- →NSSA Operation
- NSSA Configuration
- Route Summarization

© 2020 Juniper Networks, Inc. All Rights Reserved.

NSSA Operation

The slide highlights the topic we discuss next.

Not-So-Stubby Areas

- Allows a stub area to contain external routing information from a local ASBR
  - ASBR injects Type 7 LSAs into NSSA
  - ABR converts Type 7 LSAs into Type 5 LSAs and forwards them into the backbone
- Reachability for other external routes is available through a 0/0 default route injected by the ABR
  - Manual configuration step for administrator control
  - Advertised in a Type 7 LSA or a Type 3 if **`no-summaries`** is used
- Virtual links cannot transit an NSSA

© 2020 Juniper Networks, Inc. All Rights Reserved.

Altering the Stub Area Behavior

Much like a stub area, the operation of an OSPF NSSA assists the routers within that area by reducing the size of the LSDB. The main difference between a stub area and an NSSA is that external routing information, in the form of a Type 7 LSA, is allowed within the NSSA area. The ABR with the highest router ID converts the Type 7 LSAs into Type 5 LSAs and forwards them toward the backbone as if they had originated from an ASBR in a non-stub area. This process keeps knowledge of the NSSA contained within the area.

ABR Provides Reachability

The routers within the NSSA might require IP reachability to the external routes from the backbone they no longer have in their databases. A 0.0.0.0/0 default route generated by the ABR in the area can accomplish this reachability. Within the Junos OS, this advertisement is a manual step for the network administrator, which allows for better control over which ABR should be used for outbound traffic flows. For an NSSA, the default route is carried in a Type 7 LSA using the **`default-lsa`** option by itself..

No Virtual Links in an NSSA

In a similar fashion to a stub area, an NSSA cannot support a virtual link. Again, the area attached to the backbone through the NSSA area might require external LSA information. Because the *transit* routers cannot forward the Type 5 LSAs, the far-end area does not receive the correct information.

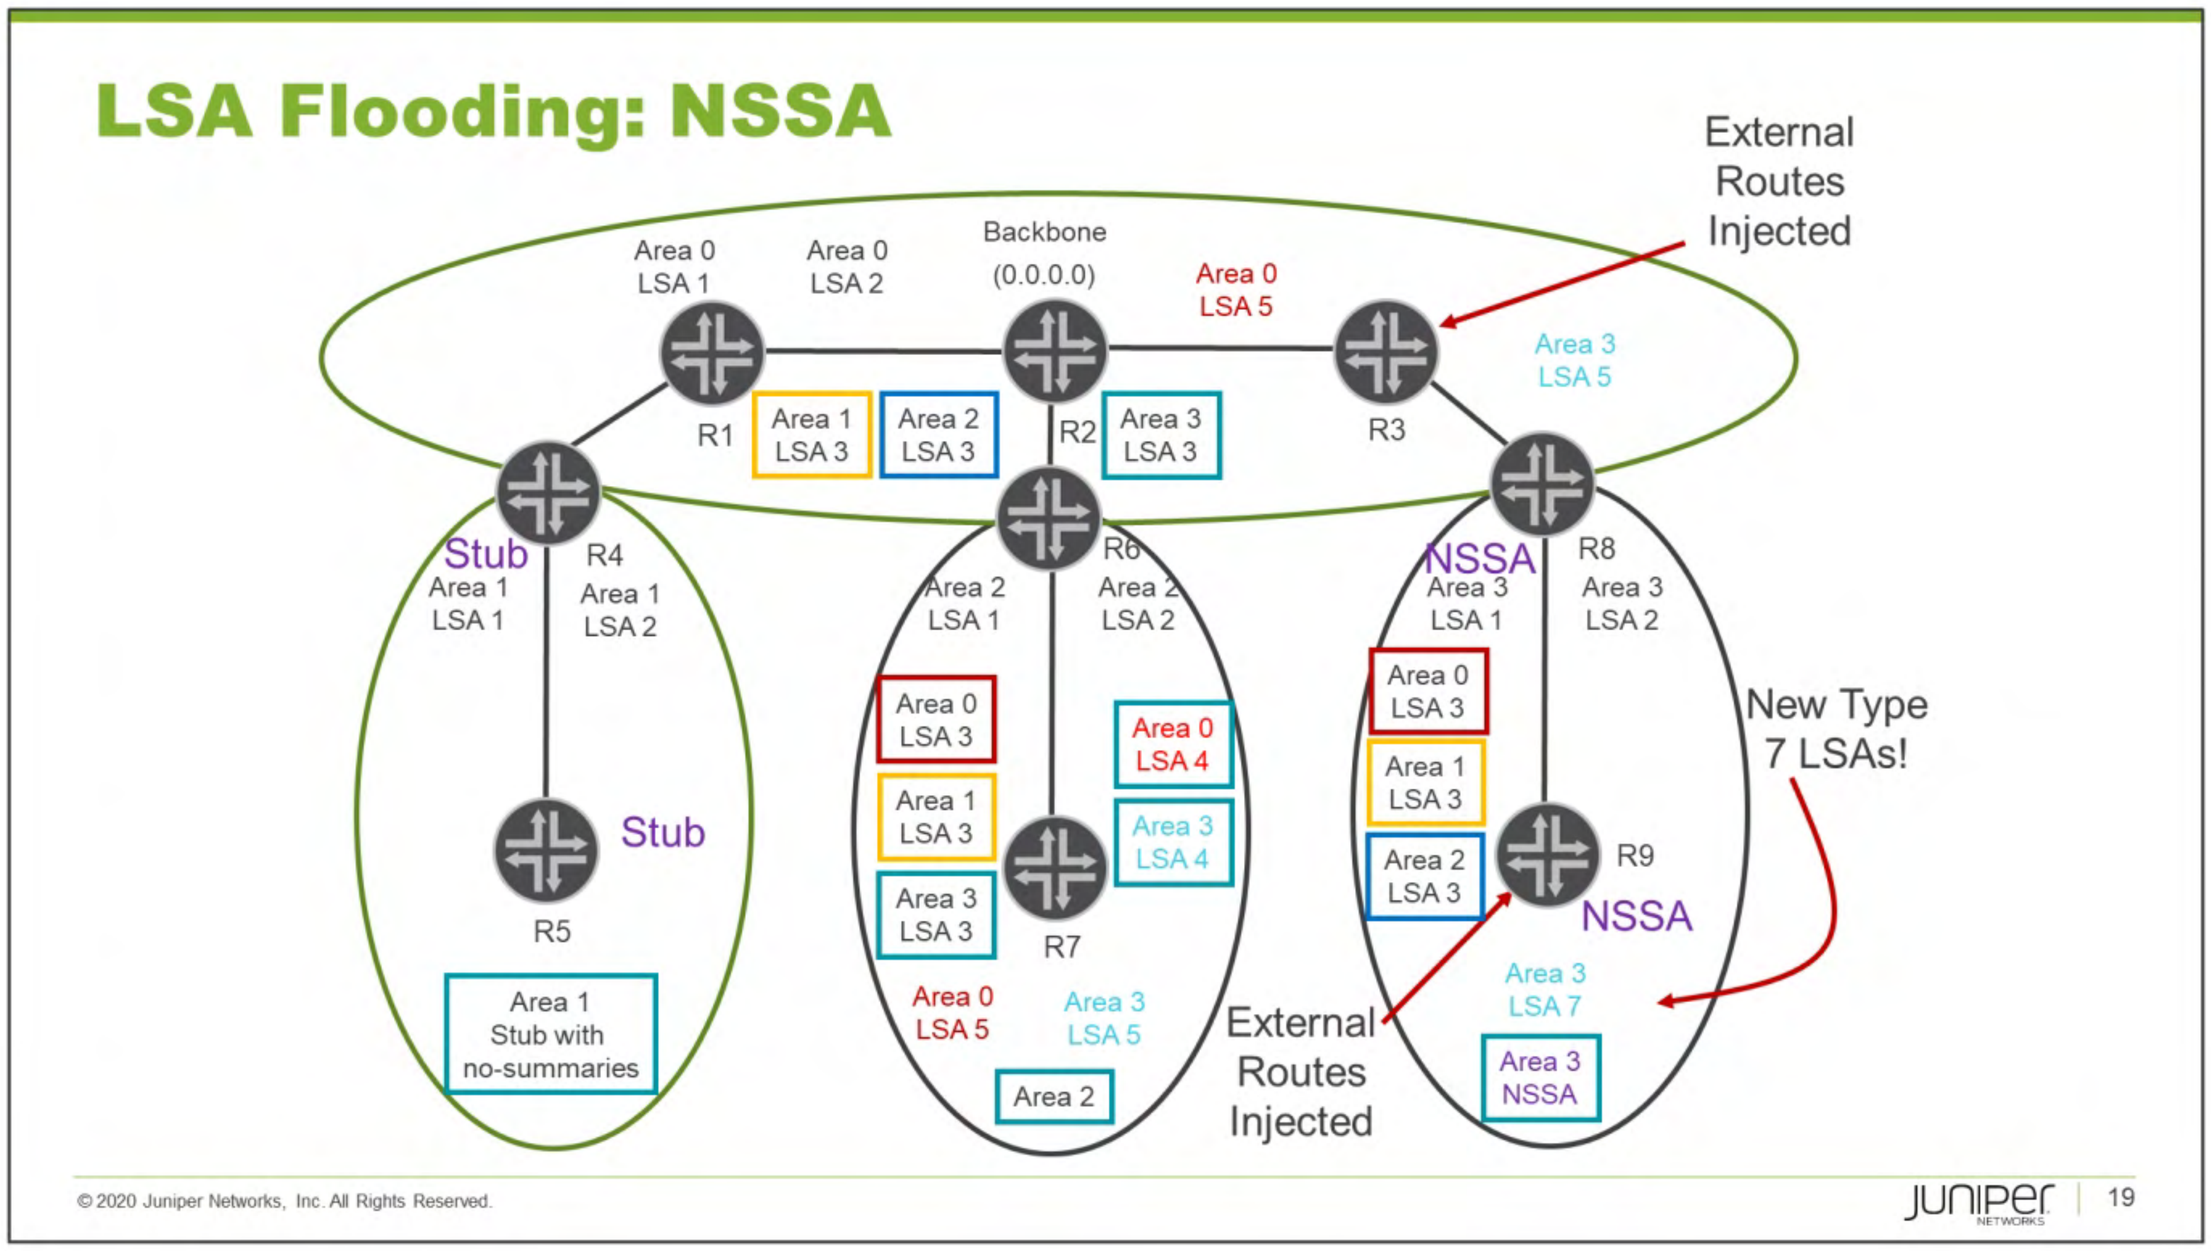

NSSA Flooding Scopes

Area 3 now is configured as an NSSA. This configuration causes the ABR of Area 3 to stop forwarding the Type 5 LSAs from Area 0. Because the routers within Area 3 cannot use the external LSA information, they also do not require the knowledge of any ASBRs in the rest of the network. As such, the Area 3 ABR stops generating Type 4 LSAs from Area 0.

As shown in the slide above, an ASBR is configured within Area 3. To allow that routing information to be propagated to the rest of the OSPF domain, the ASBR generates a Type 7 LSA for each external route. Each of the routers in Area 3 can now forward these LSAs because they are each configured as an NSSA router. When the information reaches the Area 3 ABR, it performs a Type 7 LSA to Type 5 LSA conversion and injects the Type 5 LSA into the backbone. All other OSPF routers in the domain use the Type 5 LSA as normal and have no knowledge that the routes originated within an NSSA.

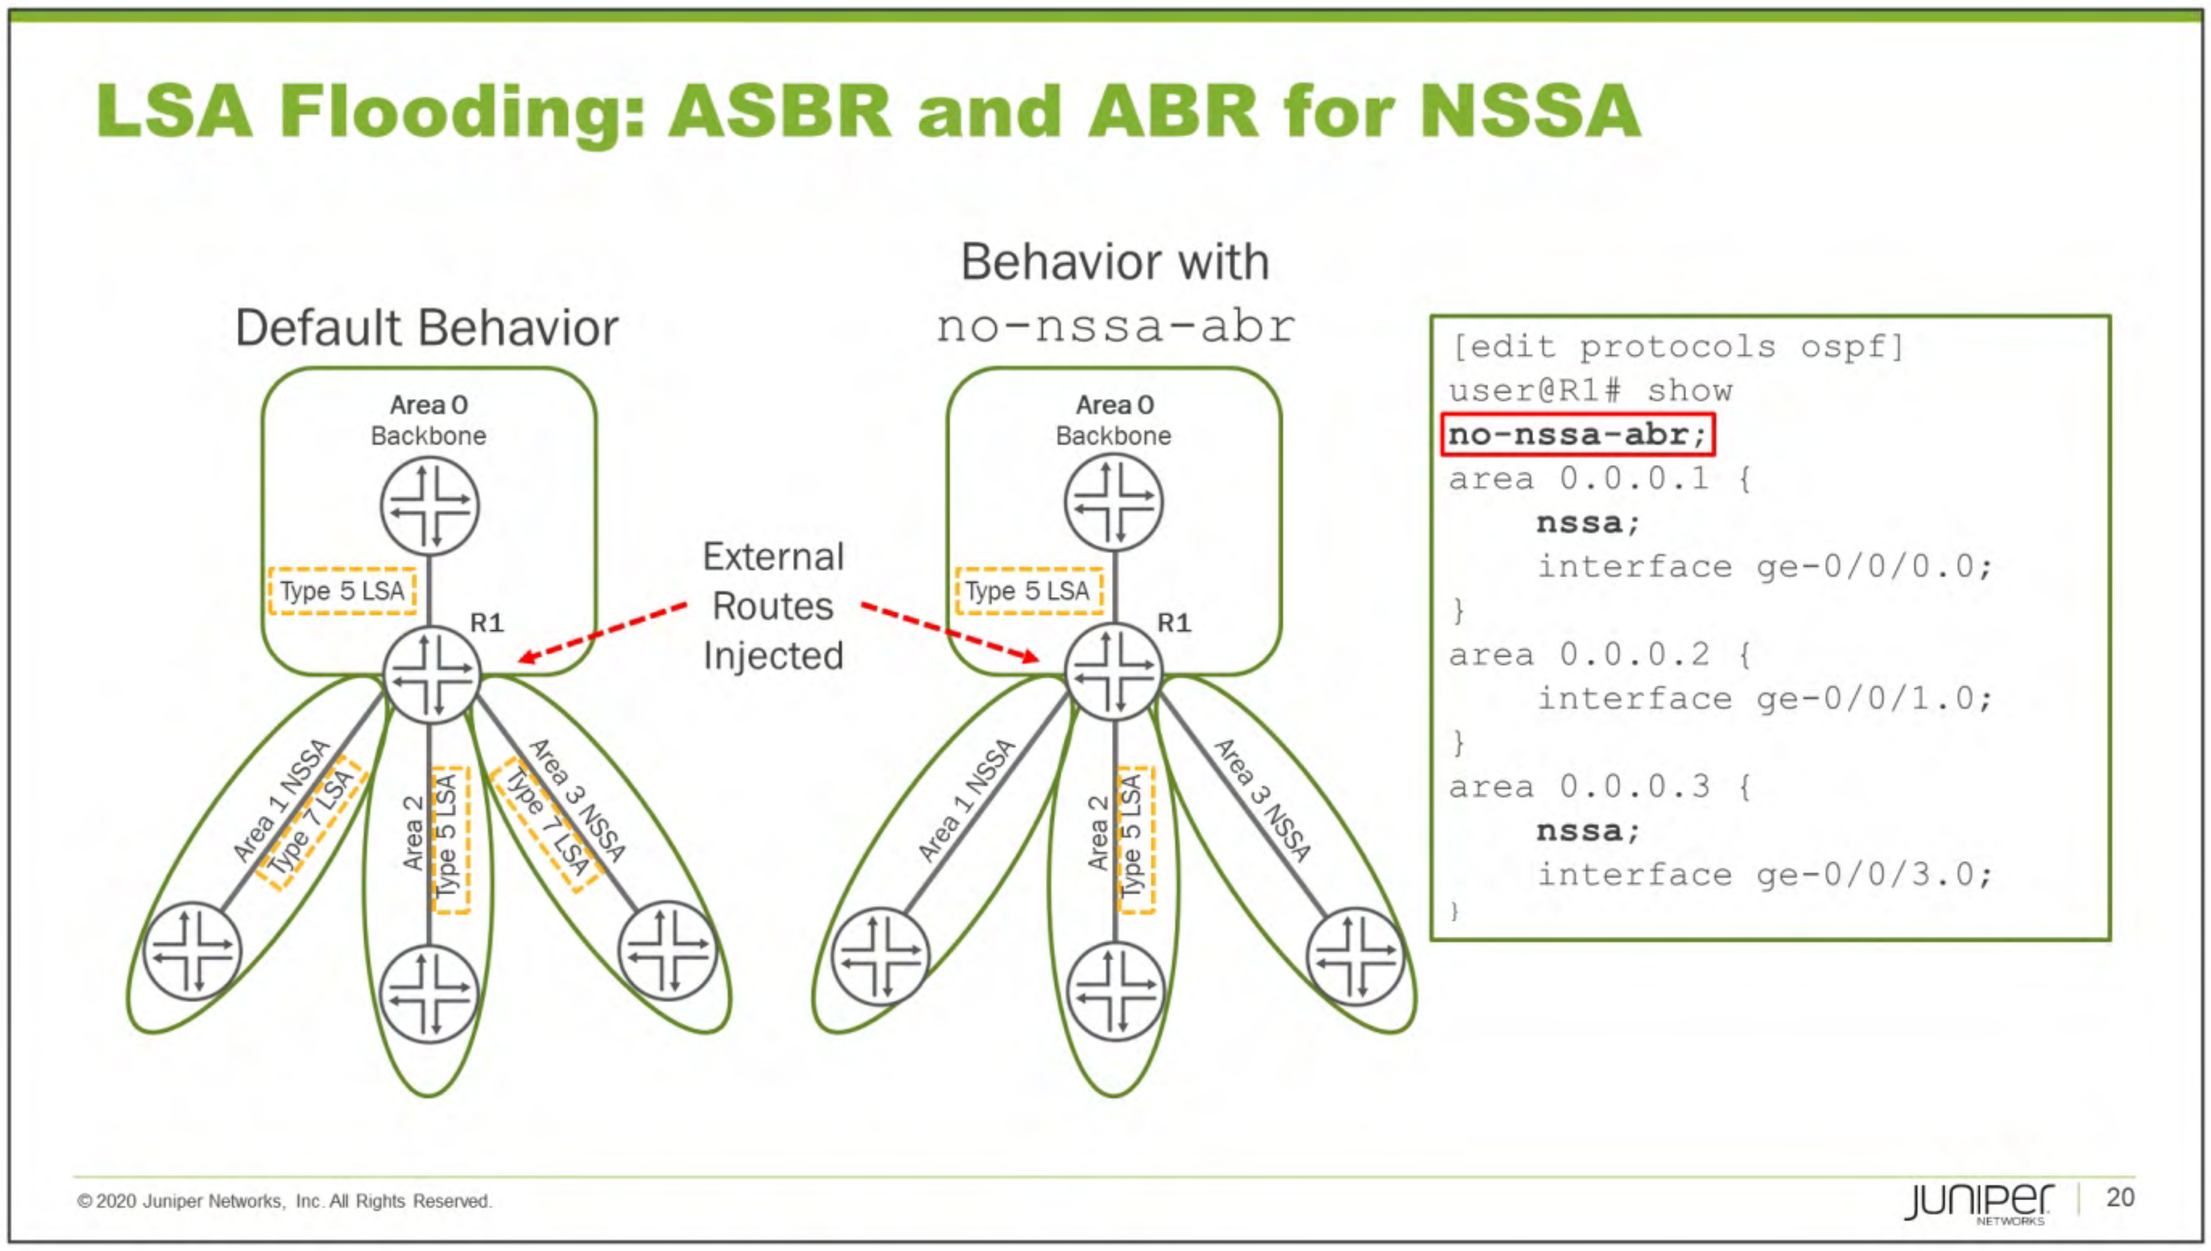

Disabling Export of LSAs into NSSAs Attached to ASBR ABRs

When an ASBR is also an ABR with an NSSA area attached to it, a Type 7 LSA is exported into the NSSA area by default. If the ABR is attached to multiple NSSA areas, a separate Type 7 LSA is exported into each NSSA area by default.

To disable exporting Type 7 LSAs into NSSAs, include the **`no-nssa-abr`** statement. This statement is needed on the ABR only.

Totally NSSA

- Behaves like a stub area with **`no-summaries`**
  - ABR does not inject Type 3 LSA from backbone into area
  - ASBR injects Type 7 LSAs into NSSA
  - ABR converts Type 7 to Type 5 and forwards into backbone
- Reachability for other external routes is available through a 0/0 default route injected by the ABR
  - Again, a manual configuration step for administrator control
    - A Type 3 summary LSA is the default
- Virtual links cannot transit an NSSA with **`no-summaries`**

© 2020 Juniper Networks, Inc. All Rights Reserved.

Behaves Like Stub Area with No Summaries

The operation of an OSPF NSSA with no summaries assists the routers within that area by further reducing the size of the LSDB. In addition to the ABR not forwarding Type 5 external LSAs, and not generating Type 4 summary LSAs, the ABR also stops flooding Type 3 LSAs into the NSSA. This process has the effect of the routers within the area having only Type 1 router, Type 2 network, and Type 7 (NSSA) LSAs in their databases.

ABR Provides Reachability

The routers within the NSSA with no summaries might require IP reachability to routes they no longer have in their databases. A 0.0.0.0/0 Type 3 LSA is generated by the ABRs to provide this reachability, when so configured. The advertisement of a default route is a manual step for the network administrator, which allows for better control over which ABR (if any) should be used for outbound traffic flows. A Type 7 LSA can be used for the advertisement of the default route for compatibility with earlier Junos OS releases.

No Virtual Links in an NSSA with No Summaries

Like a *regular* NSSA, an NSSA with no summaries cannot support a virtual link.

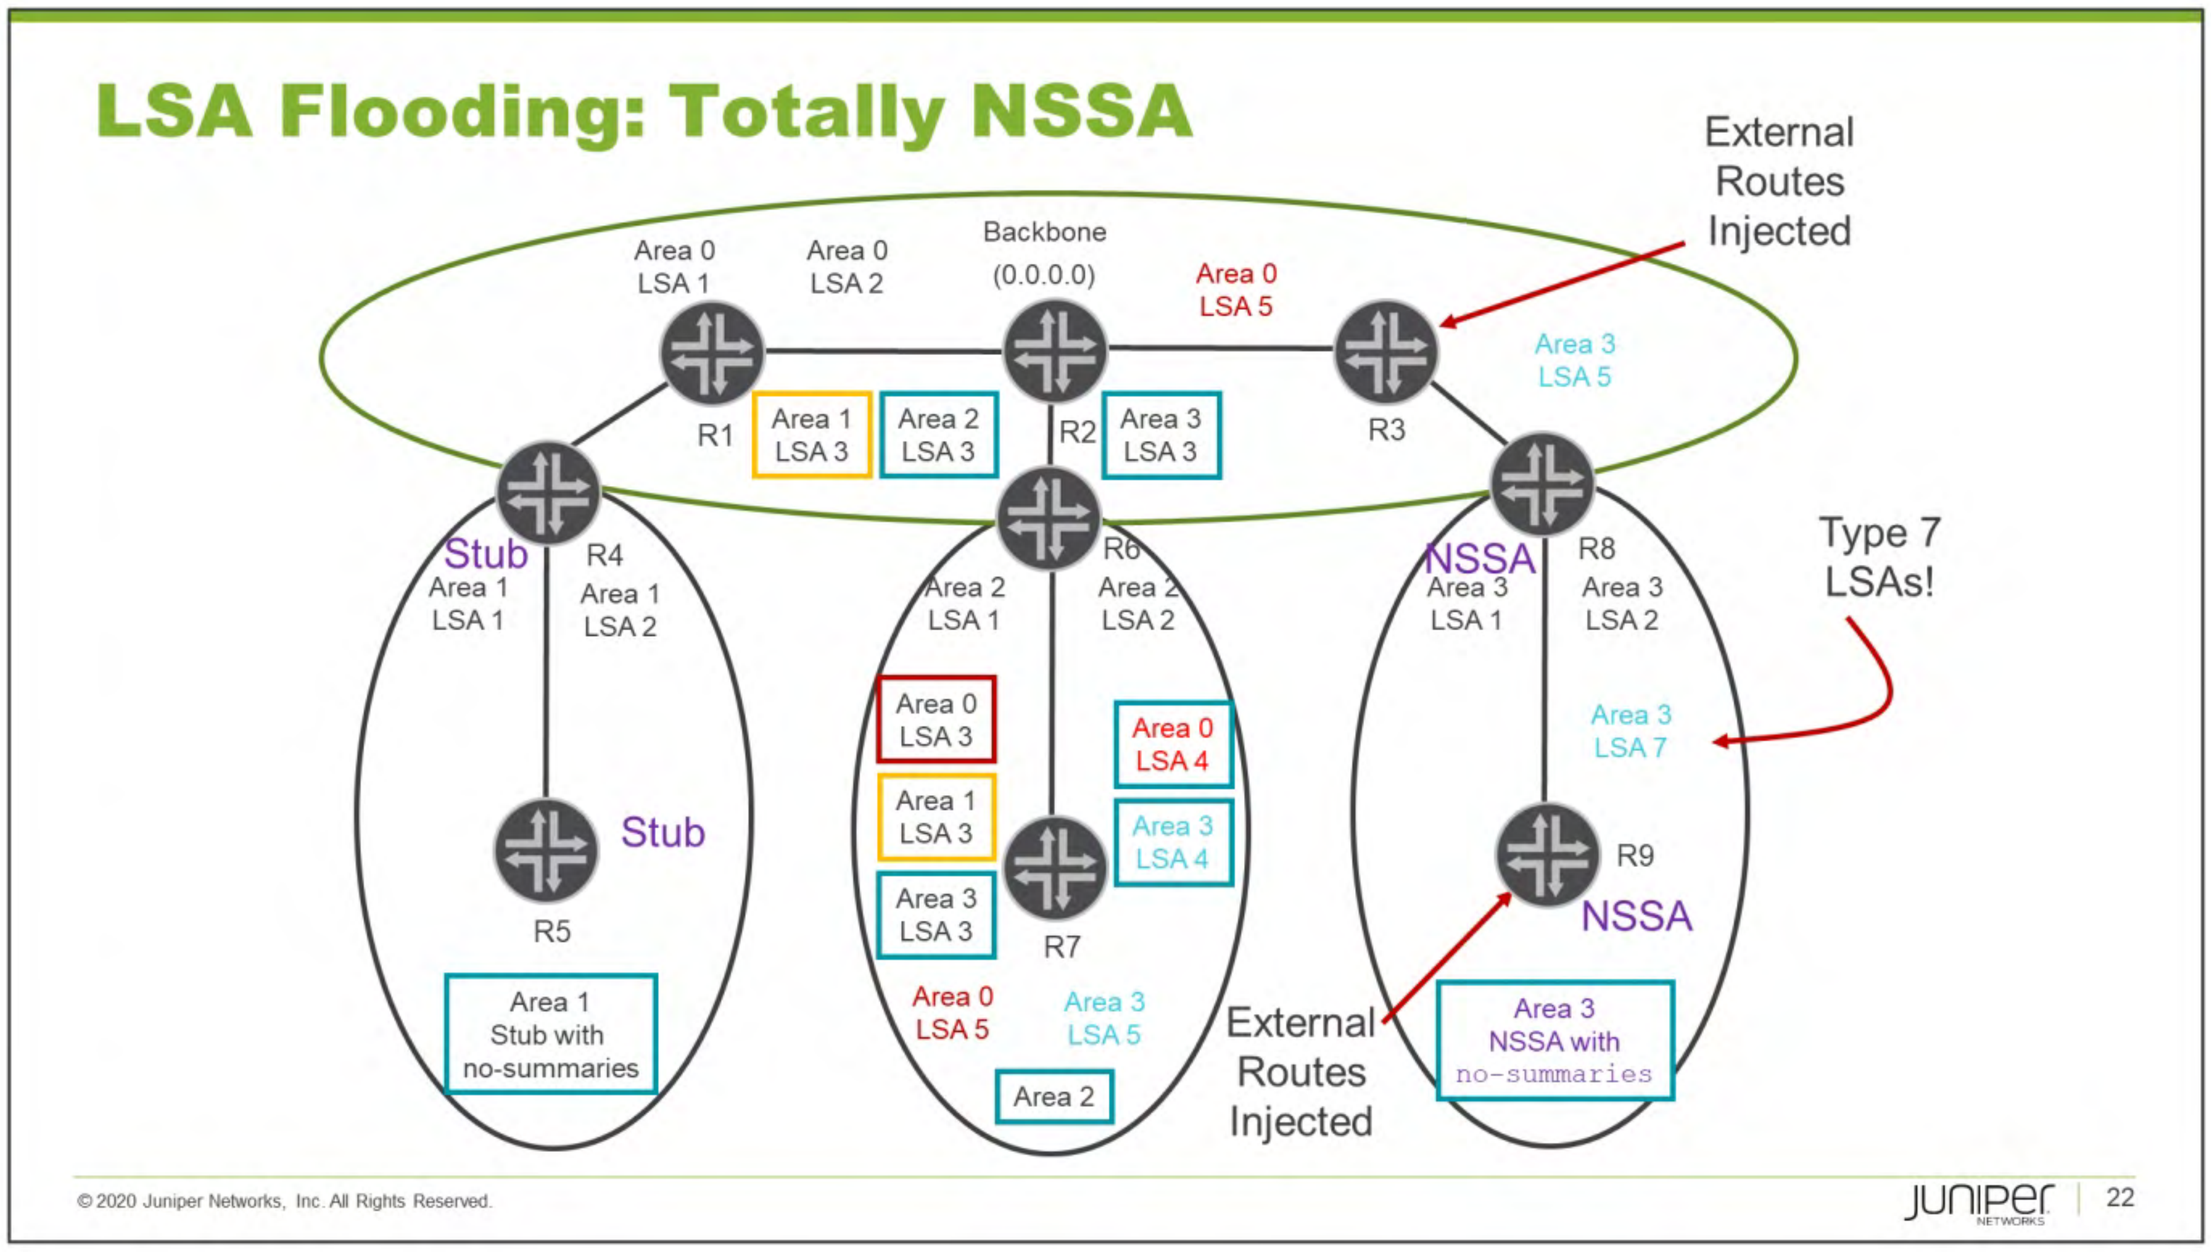

Totally NSSA Flooding Scopes

Area 3 is now configured as an NSSA with no summaries. In addition to the functionality described on a previous slide, the ABR of Area 3 stops generating Type 3 summary LSAs from Area 0.

Agenda: OSPF Areas

- Review of OSPF Areas
- Stub Area Operation
- Stub Area Configuration
- NSSA Operation
- →NSSA Configuration
- Route Summarization

© 2020 Juniper Networks, Inc. All Rights Reserved.

NSSA Configuration

The slide highlights the topic we discuss next.

NSSA Configuration

- You must configure each router in the area as an NSSA router
- The ABR can inject a default route when the **`default-metric`** statement is added
  - Specified within the **`default-lsa`** configuration hierarchy

Configured on all routers in Stub area

```
[edit protocols ospf area 0.0.0.3]
user@R8# set nssa

[edit protocols ospf area 0.0.0.3]
user@R8# show
nssa;
interface ge-0/0/0.0;
```

Configured on ABR Only

```
[edit protocols ospf area 0.0.0.3]
user@R8# set nssa default-lsa default-metric 10

[edit protocols ospf area 0.0.0.3]
user@R8# show
nssa {
    default-lsa default-metric 10;
}
interface ge-0/1/1.0;
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configuration on All Routers

To operate an OSPF NSSA successfully, you must configure each router to participate. This configuration alters the N/P bit setting in the Options field of the OSPF hello packet to notify all neighbors that the local router does not support Type 5 external LSAs, but does support Type 7 external LSAs. An OSPF router only forms an adjacency with another router when the N/P bit values match—hence, the requirement that all routers perform the configuration.

ABR Injects a Default Route

The ABR of the NSSA can optionally inject a 0.0.0.0/0 default route into the area. Within the `default-lsa` configuration hierarchy, the **`default-metric`** command accomplishes this. The command causes the ABR to generate a Type 3 summary LSA to advertise the 0.0.0.0/0 default route into the area with the configured metric.

The area-range command cannot be used to summarize Type 5 LSAs. External routes need to be summarized on the ASBR by creating an aggregate route and redistributing that route into OSPF rather than the individual routes.

NSSA Route Configuration Options

- Can import routing information from a local ASBR
  - ASBR injects Type 7 LSAs into NSSA
  - ABR with the highest RID converts Type 7 LSAs into Type 5 LSAs and forwards them into the backbone
- A default route injected by the ABR which must be configured

```
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface ge-1/0/0.0;
}
area 0.0.0.3 {
    nssa {
        default-lsa {
            default-metric 10;
        }
    }
    interface ge-1/1/0.0;
}
```

Injected as a Type 7

```
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface ge-1/0/0.0;
}
area 0.0.0.3 {
    nssa {
        default-lsa {
            default-metric 10;
        }
        no-summaries;
    }
    interface ge-1/1/0.0;
}
```

Injected as a Type 3

```
[edit protocols ospf]
user@router# show
area 0.0.0.0 {
    interface ge-1/0/0.0;
}
area 0.0.0.3 {
    nssa {
        default-lsa {
            default-metric 10;
            type-7;
        }
        no-summaries;
    }
    interface ge-1/1/0.0;
}
```

Injected as a Type 7

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configured Only on the ABR

Because it is the job of the ABR to generate new Type 3 summary LSAs into the area, configuration takes place only on this router to create a no-summaries area. To configure an ABR to support a no-summaries configuration, apply the **`no-summaries`** command to the NSSA setting within the ABR. Because all other routers in the area do not require the suppression of a Type 3 LSA, you do not need configuration on those routers.

ABR Injects a Default Route

The ABR of the Totally NSSA can inject a 0.0.0.0/0 default route into the NSSA when the **`default-metric`** command is issued. This causes the ABR to advertise a default route with the configured metric. In a regular NSSA area the default route is injected as a Type 7 LSA. Using both the **`default-lsa`** and **`no-summaries`** options together causes a Type 3 LSA to be used. A NSSA no-summaries area then has an option to change this Type 3 back to a Type 7 again by using the **`type-7`** option. All three configuration options are shown in the slide above.

Type 3 LSA Generated by Default

When the NSSA is configured with the **`no-summaries`** command, the 0.0.0.0/0 default route is advertised using a Type 3 summary LSA. This behavior reflects a change in the default operation of the Junos OS starting with Release 5.0. To interoperate with older Junos OS versions (earlier than 5.0), a Type 7 LSA might be required. The **`type-7`** command within the **`default-lsa`** hierarchy accomplishes this task, which causes the ABR to generate a Type 7 LSA advertising the 0.0.0.0/0 route with the associated metric.

This Type 7 LSA is advertised using a Type 1 metric, which means that all routers calculate the cost to the ABR and add to it the advertised default metric. You can alter this behavior with the **`metric-type`** command within the **`default-lsa`** hierarchy. The Type 7 LSA is then advertised using a Type 2 metric, which means that each area router uses only the advertised metric of the route.

Agenda: OSPF Areas

- Review of OSPF Areas
- Stub Area Operation
- Stub Area Configuration
- NSSA Operation
- NSSA Configuration

→Route Summarization

© 2020 Juniper Networks, Inc. All Rights Reserved.

Route Summarization

The slide highlights the topic we discuss next.

Summarizing Routes (Type 1 and Type 2 LSAs)

- By default, all local area routes are forwarded to the backbone
  - Stub setting changes only what enters an area, not what leaves that area
- Use the **area-range** command to summarize routing information
  - Can result in a single Type 3 LSA injected into the backbone
  - Configured on the ABRs only

Configured on ABR immediately after the Area number. Summarizes only T1 and T2 LSAs with a single T3 LSA.

```
[edit protocols ospf]
user@router# set area 1 area-range 192.168.16/20

[edit protocols ospf]
user@router# show
area 0.0.0.1 {
    area-range 192.168.16.0/20;
    interface ge-1/0/1.100;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Summarizing Routes (Type 1 and Type 2 LSAs)

The default operation of an ABR is to generate a Type 3 summary LSA into the backbone area for each Type 1 and Type 2 LSA in the LSDB. The configuration of a stub area or stub area with no summaries does not affect this operation because a stub configuration only alters information that enters the nonbackbone area.

Route Summarization

To reduce the size of the LSDB in Area 0.0.0.0 (and other remote OSPF areas), you can configure the ABR *not* to perform its default one-for-one mapping function. This is accomplished using an address range statement on the ABR with the **area-range** command. All Type 1 and Type 2 LSAs that fall within that address range will no longer be advertised individually into the backbone. Instead, a single Type 3 summary LSA is advertised. The metric associated with this summary route will be equal to the highest metric associated with the individual contributing routes.

Because only the ABR performs this mapping function, you configure the **area-range** command on ABR routers only.

Summarizing Routes (NSSA External Type 7 LSAs)

- The Junos OS forwards all ASBR Type 7 routes to the backbone by default
- Use the **area-range** command within the NSSA configuration to summarize routing information
  - Injects a Type 5 LSA into the backbone
  - Configured on ABRs only

Configured on ABR immediately after the NSSA option. Summarizes T7 LSAs in a single T5 LSA..

```
[edit protocols ospf]
user@router# set area 1 nssa area-range 192.168.16/20

[edit protocols ospf]
user@router# show
area 0.0.0.1 {
    nssa {
        default-lsa default-metric 10;
        area-range 192.168.16.0/20;
        }
    interface ge-0/1/1.0;
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Summarizing Routes (NSSA External Type 7 LSAs)

The default operation of the ABR for Type 7 LSA routes local to a nonbackbone area is to generate a single Type 5 external LSA for each Type 7 LSA in the LSDB. The configuration of an NSSA or an NSSA area with no summaries does not affect this operation because an NSSA configuration only alters information that enters the nonbackbone area.

Route Summarization

To reduce the size of the LSDB in Area 0.0.0.0 (and other remote OSPF areas), you can configure the ABR *not* to perform its default one-for-one mapping function. You can configure an address range on the ABR by using the **area-range** command within the NSSA configuration hierarchy. All Type 7 LSAs that fall within that address range are *not* advertised individually into the backbone. Instead, a single Type 5 external LSA is advertised.

Because only the ABR performs this mapping function, only the ABR is configured with the **area-range** command.

Note that the **area-range** command referenced here is specific to the NSSA configuration hierarchy and only affects Type 7 LSA routes. The **area-range** command discussed in the previous slide was within the area hierarchy itself and affected Type 1 and Type 2 LSA routes. The configuration can have these two commands in place at the same time, and they will summarize different aspects of the local area routing domain.

Suppressing Routes (Type 1, 2, and 3 LSAs)

- Adding the **`restrict`** keyword to the **`area-range`** command stops routes in that range from entering the backbone
  - No Type 3 LSA will be injected into the backbone
  - Configured on the ABRs only
  - Allows greater control over which routes are advertised to other areas

Restrict option will block T1, T2 LSAs and will not create a T3 summary.

```
[edit protocols ospf]
user@router# set area 1 area-range 192.168.16/20 restrict

[edit protocols ospf]
user@router# show
area 0.0.0.1 {
    area-range 192.168.16.0/20 restrict;
    interface ge-0/0/1.0;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Suppressing Routes (Type 1, 2, and 3 LSAs)

The default action of the **`area-range`** command causes the generation of a single Type 3 summary LSA into the backbone for all prefixes that fall within the defined range. You can configure the ABR with the **`restrict`** keyword to block one or more prefixes from advertisement into the OSPF backbone. Such a configuration prevents routing information from each Type 1 and Type 2 LSA that falls within the address range from being advertised to the backbone, which in turn can block connectivity to those prefixes for routers in other areas.

Use the **`restrict`** function when you want to prevent inter-area routing, or when you want a default route to be used instead of the more preferred summary information that would otherwise be generated.

Because only the ABR is responsible for this mapping function, you only configure ABR routers with the **`area-range restrict`** command.

Suppressing Routes (Type 7 LSAs)

- Adding the **`restrict`** keyword to the **`area-range`** command stops routes in that range from entering the backbone
  - No Type 5 LSA is injected into the backbone
  - Configured on ABRs only

Restrict option will block T7LSAs and will not create a T5 summary.

```
[edit protocols ospf]
user@router# set area 1 nssa area-range 192.168.16/20 restrict

[edit protocols ospf]
user@router# show
area 0.0.0.1 {
    nssa {
        default-lsa default-metric 10;
        area-range 192.168.16.0/20 restrict;
        }
    interface ge-1/0/4.0;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Suppressing Routes (Type 7 LSAs)

The default action of the **`area-range`** command is to send a single Type 5 external LSA into the backbone. This functionality already alters the default OSPF protocol operation. You then can configure the ABR additionally with the keyword **`restrict`**. This keyword prevents routing information from each Type 7 LSA that falls within the address range from being advertised to the backbone at all.

Because only the ABR is responsible for this mapping function, you only configure the NSSA ABR with the **nssa `area-range restrict`** command.

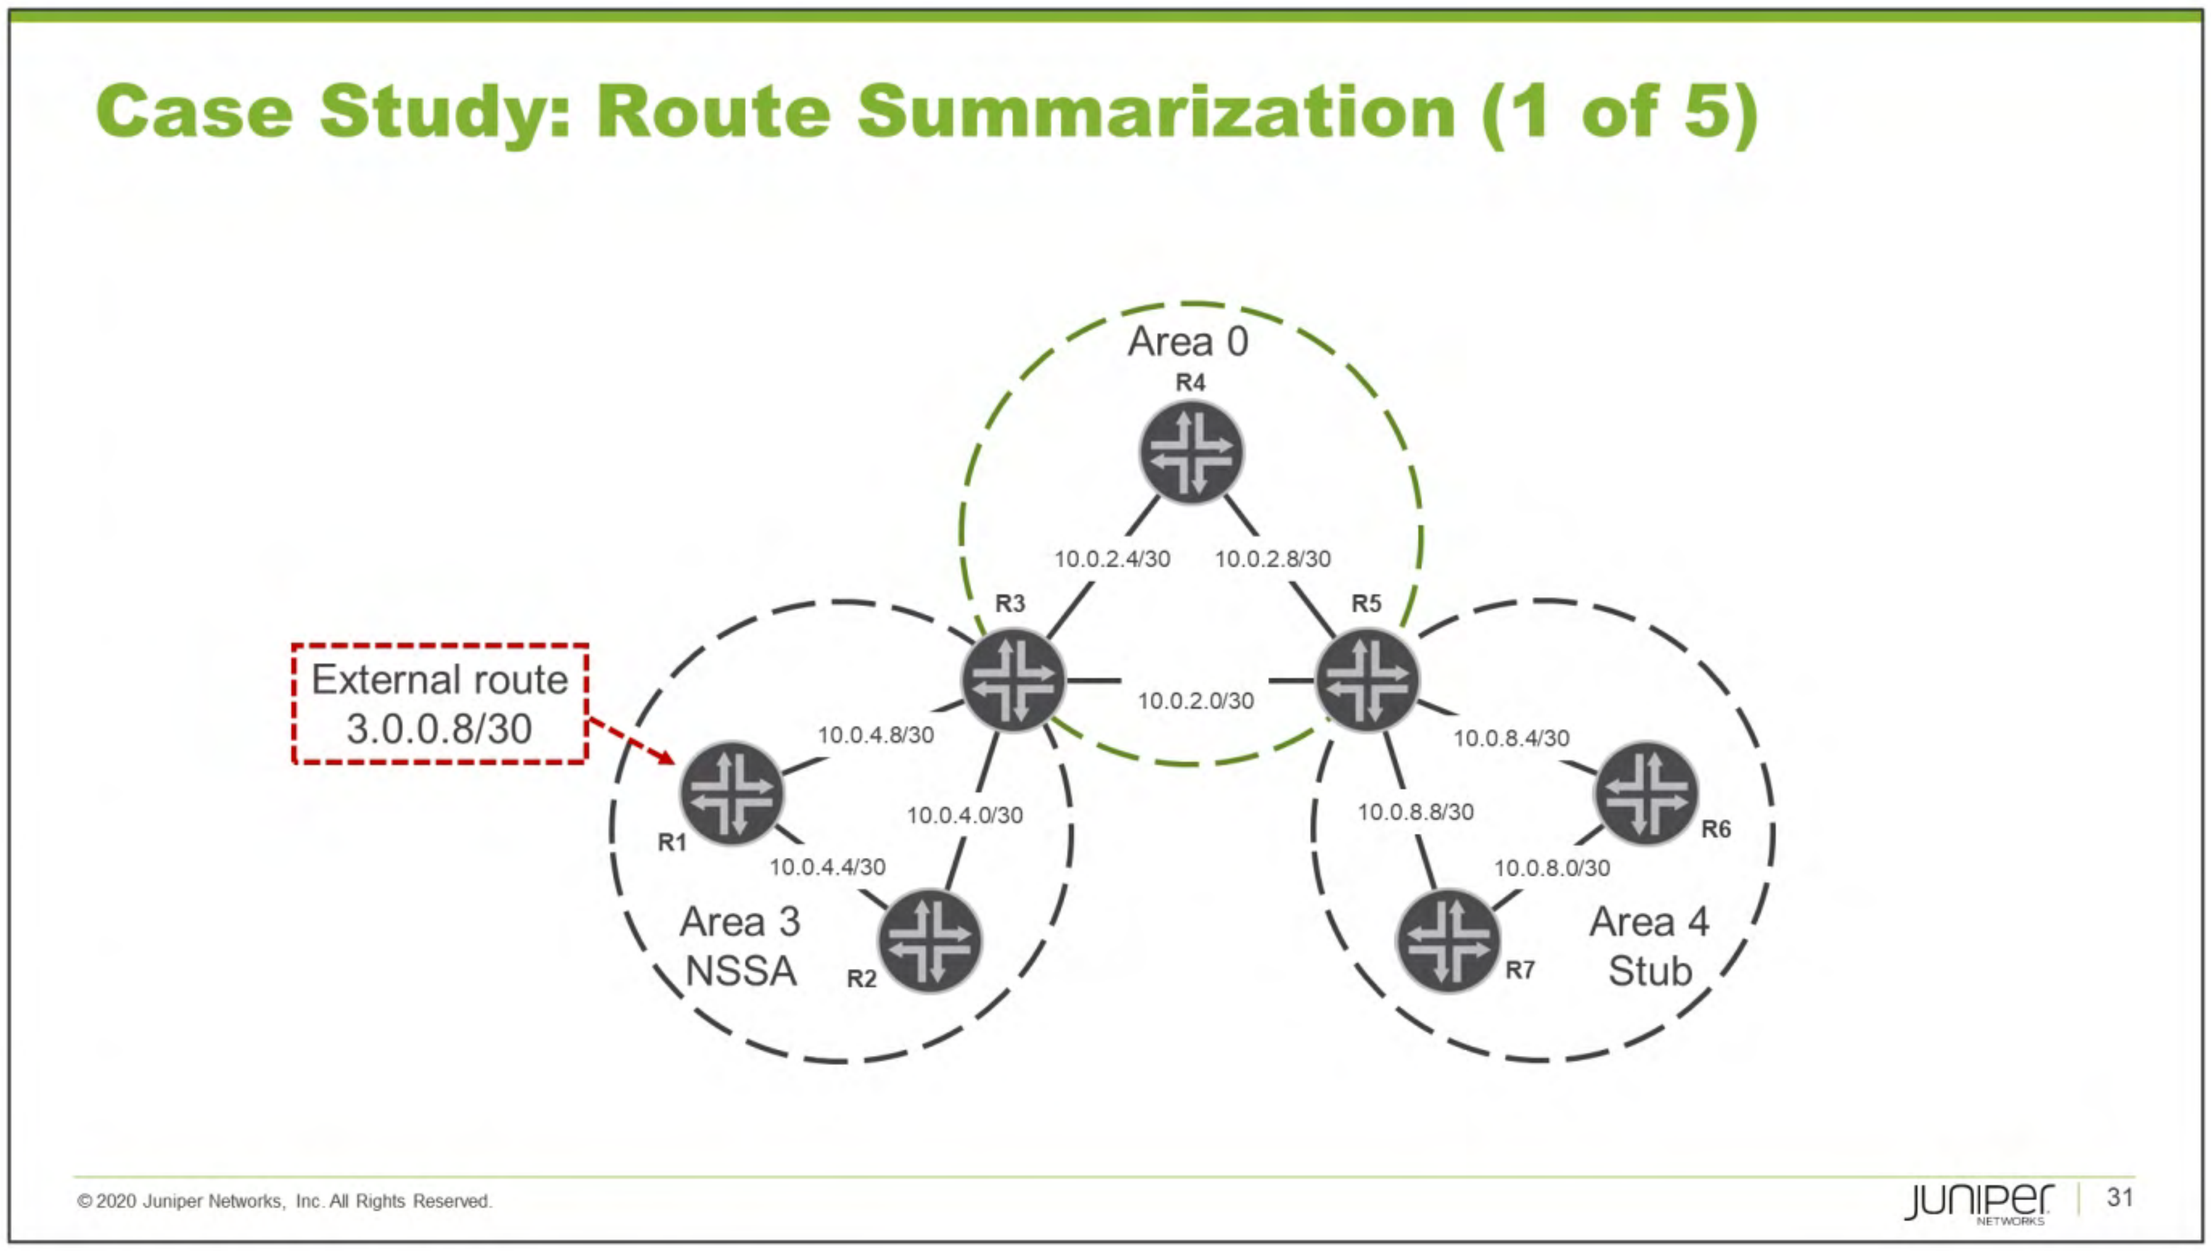

Case Study: Route Summarization—Part 1

This slide introduces the simple topology for this case study. Current highlights include the following:

- Router R3 is an ABR between NSSA Area 3 and the Area 0 backbone. Its networks are 10.0.4.0/30, 10.0.4.4/30, and 10.0.4.8/30. These can be summarized into 10.0.4.0/28. Furthermore, the 203.0.113.8/30 route exists in this area which you prevent from flooding throughout the OSPF domain.
- Router R5 is an ABR between stub Area 4 and the Area 0 backbone. Its networks are 10.0.8.0/30, 10.0.8.4/30, and 10.0.8.8/30. These can be summarized into 10.0.8.0/28.

Case Study: Route Summarization (2 of 5)

```
user@R4> show route

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3.0.0.8/30          *[OSPF/150] 00:00:29, metric 0, tag 0
                     > to 10.0.2.6 via ge-1/0/4.0
10.0.2.0/30         *[OSPF/10] 00:50:38, metric 2
                       to 10.0.2.6 via ge-1/0/4.0
                     > to 10.0.2.10 via ge-1/1/2.0
10.0.2.4/30         *[Direct/0] 00:57:53
                     > via ge-1/0/4.0
10.0.2.5/32         *[Local/0] 00:57:53
                       Local via ge-1/0/4.0
10.0.2.8/30         *[Direct/0] 00:57:53
                     > via ge-1/1/2.0
10.0.2.9/32         *[Local/0] 00:57:53
                       Local via ge-1/1/2.0
10.0.4.0/30         *[OSPF/10] 00:00:29, metric 2
                     > to 10.0.2.6 via ge-1/0/4.0
10.0.4.4/30         *[OSPF/10] 00:00:29, metric 3
                     > to 10.0.2.6 via ge-1/0/4.0
10.0.4.8/30         *[OSPF/10] 00:00:29, metric 2
                     > to 10.0.2.6 via ge-1/0/4.0
10.0.8.0/30         *[OSPF/10] 00:01:12, metric 3
                     > to 10.0.2.10 via ge-1/1/2.0
10.0.8.4/30         *[OSPF/10] 00:01:12, metric 2
                     > to 10.0.2.10 via ge-1/1/2.0
10.0.8.8/30         *[OSPF/10] 00:01:12, metric 2
                     > to 10.0.2.10 via ge-1/1/2.0
[...]
```

Summarize into 10.0.4.0/28

Summarize into 10.0.8.0/28

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study: Route Summarization—Part 2

First off, let's take a look at what the R4 router currently has in its routing table. From the output, you can see the networks, in Area 3 and Area 4, that we intend to summarize.

Case Study: Route Summarization (3 of 5)

```
[edit]
user@R3# set protocols ospf area 3 area-range 10.0.4.0/28

[edit]
user@R5# set protocols ospf area 4 area-range 10.0.8.0/28
```

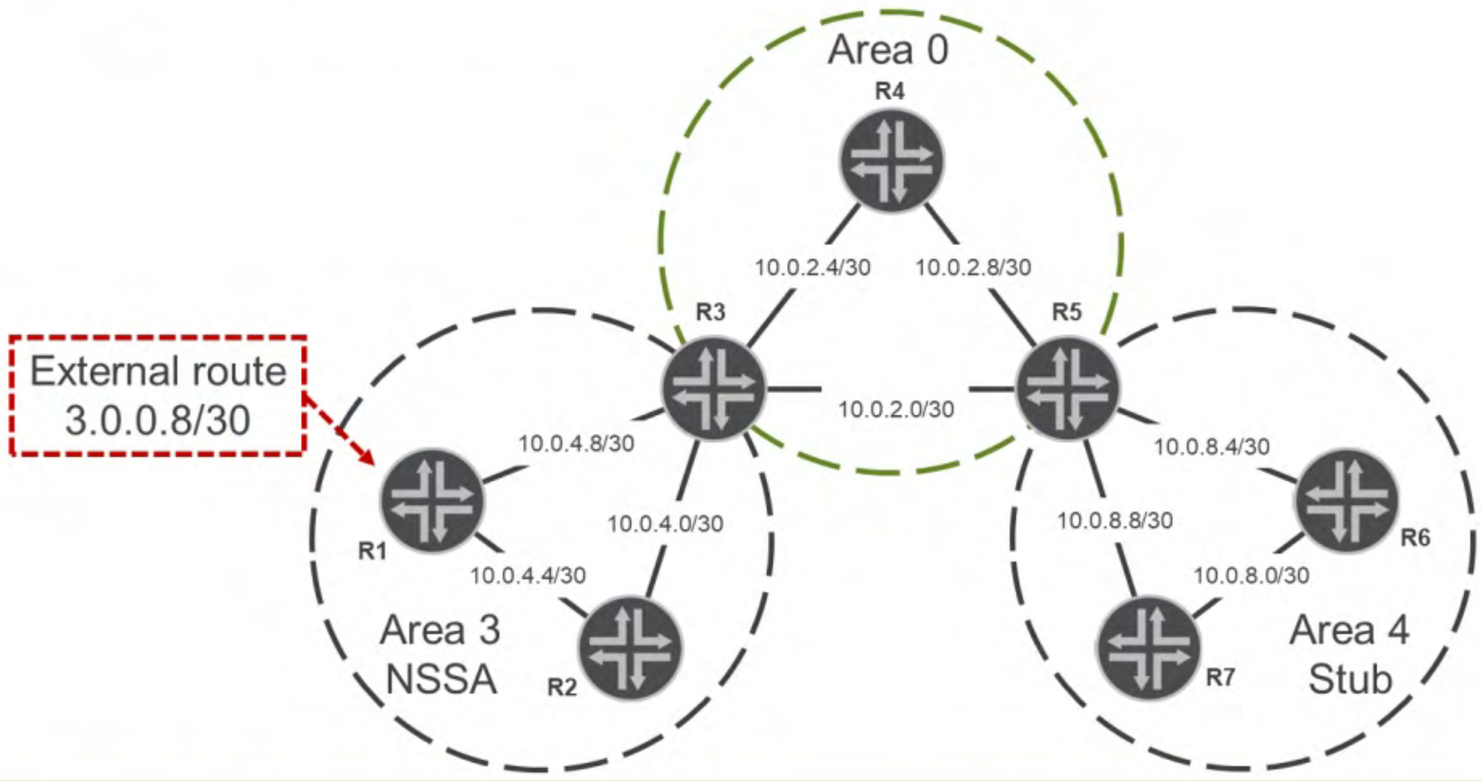

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study: Route Summarization—Part 3

The **`area-range`** option is configured as shown on the slide on both R3 and R5 to summarize the routes from each router's respective area.

Case Study: Route Summarization (4 of 5)

```
user@R4> show route

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3.0.0.8/30         *[OSPF/150] 00:32:30, metric 0, tag 0
                    > to 10.0.2.6 via ge-1/0/4.0
10.0.2.0/30        *[OSPF/10] 01:22:39, metric 2
                       to 10.0.2.6 via ge-1/0/4.0
                    > to 10.0.2.10 via ge-1/1/2.0
10.0.2.4/30        *[Direct/0] 01:29:54
                    > via ge-1/0/4.0
10.0.2.5/32        *[Local/0] 01:29:54
                       Local via ge-1/0/4.0
10.0.2.8/30        *[Direct/0] 01:29:54
                    > via ge-1/1/2.0
10.0.2.9/32        *[Local/0] 01:29:54
                       Local via ge-1/1/2.0
10.0.4.0/28        *[OSPF/10] 00:03:02, metric 3
                    > to 10.0.2.6 via ge-1/0/4.0
10.0.8.0/28        *[OSPF/10] 00:03:01, metric 3
                    > to 10.0.2.10 via ge-1/1/2.0
[...]
```

Six routes are now two

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study: Route Summarization—Part 4

This slide shows the results of our recent changes. Previously, six routes were present for the Area 3 and Area 4 networks. Now, there are only the two summarized routes.

Case Study: Route Summarization (5 of 5)

```
[edit]
user@R3# set protocols ospf area 3 nssa area-range 3.0.0.0/8 restrict
```

The 3.0.0.8/30 route is no longer present

```
user@R4> show route

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.2.0/30        *[OSPF/10] 01:31:28, metric 2
                      to 10.0.2.6 via ge-1/0/4.0
                    > to 10.0.2.10 via ge-1/1/2.0
10.0.2.4/30        *[Direct/0] 01:38:43
                    > via ge-1/0/4.0
10.0.2.5/32        *[Local/0] 01:38:43
                      Local via ge-1/0/4.0
[...]
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study: Route Summarization—Part 5

Recall that we also wanted to suppress the 203.0.113.8/30 route from flooding outside Area 3. This slide shows how to use the **`restrict`** modifier of the **`area-range`** command to accomplish our objective. As shown on the slide, after committing the configuration on R3, the 203.0.113.8/30 route is no longer present in R4's routing table.

Summary

- In this content, we:
  - Described OSPF area types and operations
  - Configured various OSPF area types
  - Summarized and restricted routes

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Discussed:

- OSPF area types and operations;
- How to configure various OSPF area types; and
- How to summarize and restrict routes.

Review Questions

1. Which LSA types are not forwarded into an OSPF stub area or NSSA? Which router accomplishes this action? Which LSA types are not forwarded into an OSPF NSSA with no summaries?
2. Which routers must you configure to support a stub area or NSSA?
3. How do routers in a stub area or NSSA reach addresses outside of the OSPF domain?
4. Which OSPF areas see the effect of the **`area-range`** command?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review Questions

1.

2.

3.

4.

Lab: OSPF Route Summarization

- Configure and monitor the operation of OSPF stub areas.
- Configure and monitor the operation of an OSPF not-so-stubby area.
- Configure and monitor address summarization with and without restrictions.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab: Configuring and Monitoring OSPF Areas and Route Summarization

The slide provides the objectives for this lab.

Answers to Review Questions

1.

The ABR does not forward Type 4 and Type 5 LSAs into a stub area or NSSA. Type 3 LSAs are also not forwarded in an OSPF NSSA with no summaries.

2.

You must configure all routers that are in the stub area or NSSA.

3.

The ABR can optionally inject a 0.0.0.0/0 default route into the stub area or NSSA.

4.

The backbone area is directly affected by the **`area-range`** command. However, any area that is able to receive summary LSAs also benefits from this command.

Education Services

Advanced Junos Service Provider Routing

Chapter 4: Advanced OSPF Options

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - Configure OSPF multi-area adjacency
  - Configure OSPF virtual links
  - Explain OSPF external reachability

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Will Discuss:

- Configuring OSPF multi-area adjacencies;
- Configuring OSPF virtual links; and
- OSPF external reachability.

Agenda: Advanced OSPF Options

→OSPF Multi-Area Adjacencies

- Virtual Links
- External Reachability

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Multi-Area Adjacencies

The slide lists the topics we will discuss. We discuss the highlighted topic first.

Objectives

- After successfully completing this content, you will be able to:
  - Configure OSPF multi-area adjacency
  - Configure OSPF virtual links
  - Explain OSPF external reachability

© 2020 Juniper Networks, Inc. All Rights Reserved.

Multi-area Adjacencies

By default, a single interface can belong to only one OSPF area. However, in some situations, you might want to configure an interface to belong to more than one area. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and used in the intra-area best path calculation.

Beginning with Junos Release 9.2, you can configure a logical interface to belong to more than one OSPF area. As defined in RFC 5185, OSPF Multi-area Adjacency, the area border routers establish multiple adjacencies belonging to different areas over the same logical interface. Each multi-area adjacency is announced as a point-to-point link in the configured area by the routers connected to the link.

For a given logical interface, it is considered as primary for one single area. To configure that same logical interface in additional areas, you must use the **`secondary`** option. In the slide, you can see that `interface ge-0/0/1.0` has been configured in two areas, Area 0 and Area 1. Because the Area 1 interface has the **`secondary`** option set, the Area 0 occurrence is considered the primary.

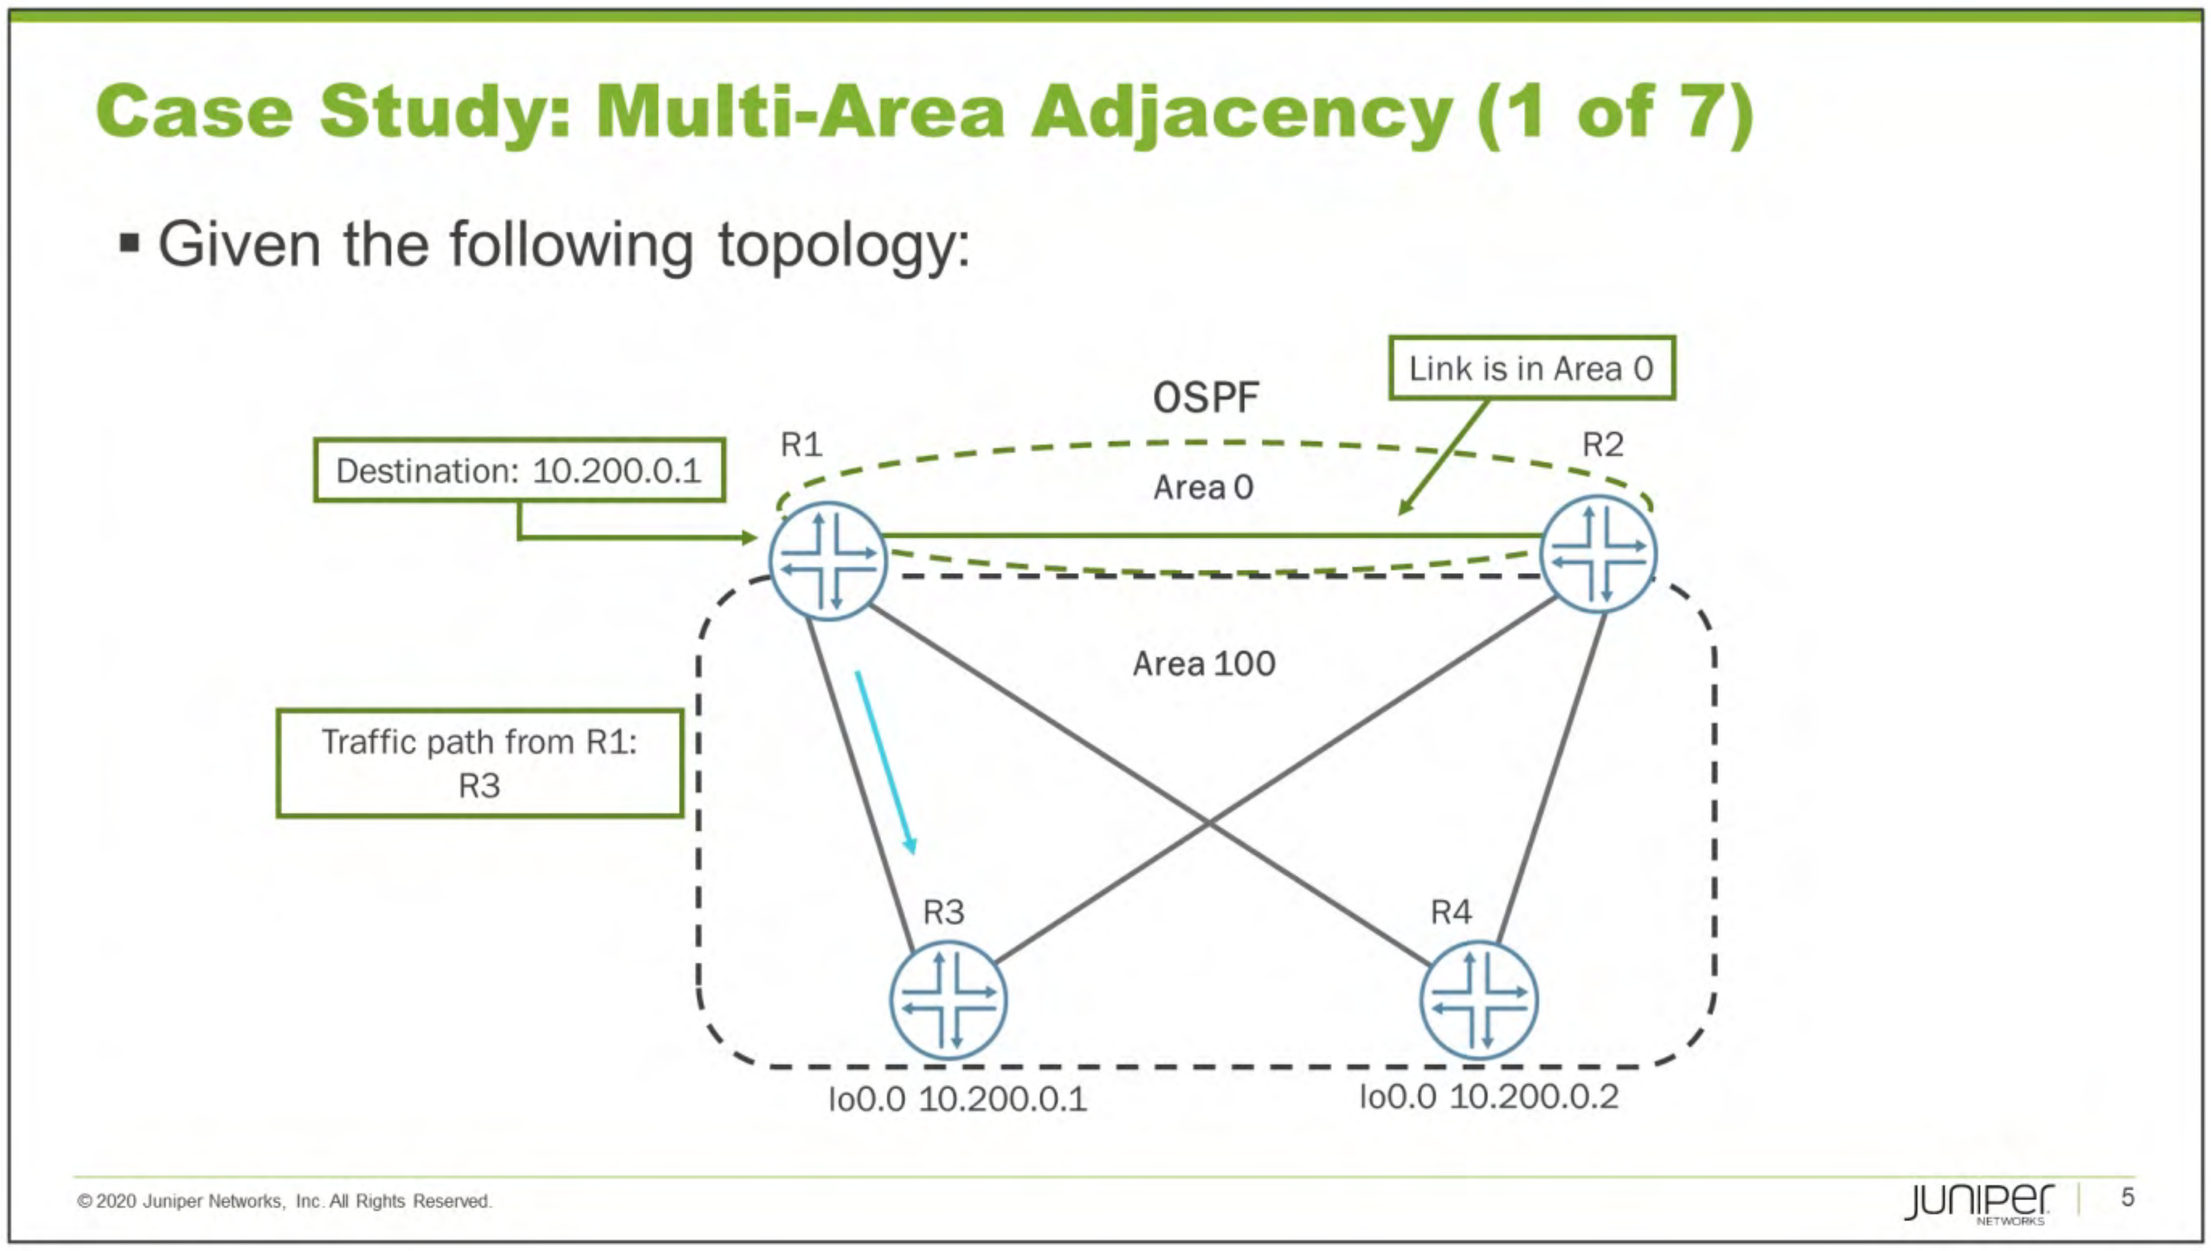

Case Study Topology

The slide displays the case study topology.

In this Case Study, traffic currently flows from R1 to R3 across the directly connected link.

Case Study: Multi-Area Adjacency (2 of 7)

- Verify that the OSPF adjacencies are in full state

```
user@R1> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.200.1.2       ge-0/0/1.0             Full      192.168.10.2     128    33
10.200.2.2       ge-0/0/2.0             Full      10.200.0.1       128    36
10.200.3.2       ge-0/0/3.0             Full      10.200.0.2       128    33

user@R2> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.200.1.1       ge-0/0/1.0             Full      192.168.10.1     128    32
10.200.4.2       ge-0/0/4.0             Full      10.200.0.1       128    35
10.200.5.2       ge-0/0/5.0             Full      10.200.0.2       128    35
```

- Trace from the R1 router to the R3 router

```
user@R1> traceroute 10.200.0.1
traceroute to 10.200.0.1 (10.200.0.1), 30 hops max, 40 byte packets
 1  10.200.0.1 (10.200.0.1)  0.533 ms  0.376 ms  0.346 ms
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF network without multi-area adjacencies

In this case study, a standard OSPF network is configured. Both R1 and R2 have 3 adjacencies. The links between R1 and R2 are configured to be part of Area 0, the remaining links are in area 100. All adjacencies are in the Full state.

The **`show ospf neighbor`** command is used to verify these adjacencies.

Perform a traceroute to verify traffic path

At this point, the default traffic path can be verified using the **traceroute** command. There is a single hop between R1 and R3.

Case Study: Multi-Area Adjacency (3 of 7)

- What happens if a link fails between R1 and R3?

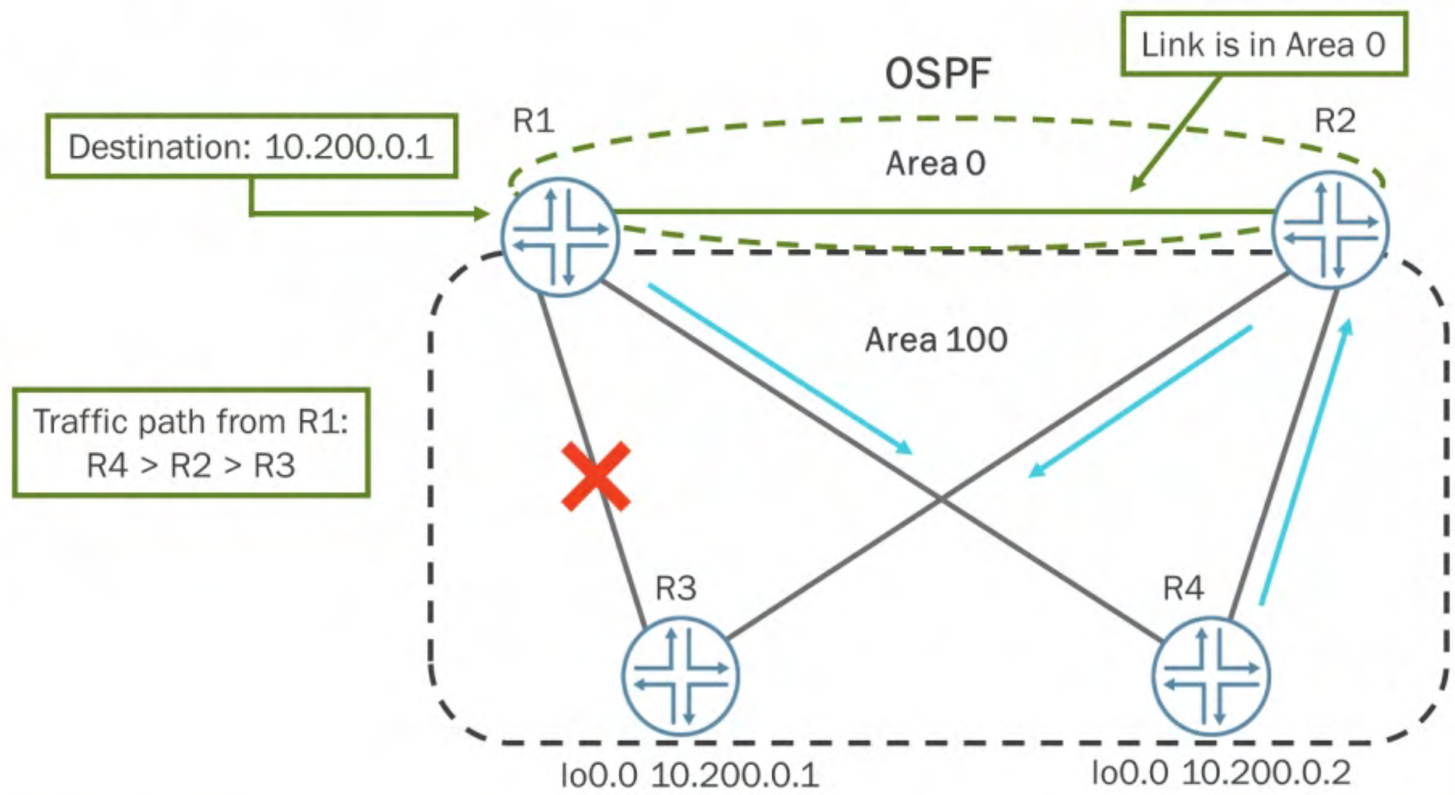

© 2020 Juniper Networks, Inc. All Rights Reserved.

Link Failure

In normal operation, if a link failure occurred between R1 and R3, traffic from R1 to R3 would flow from R4 to R2 and then to R3. Routing will happen this way regardless of metric values because OSPF always prefers intra-area paths over inter-area paths. This creates three hops to reach a router that was previously one hop away.

Case Study: Multi-Area Adjacency (4 of 7)

- The interface between R1 and R3 fails

```
user@R1> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.200.1.2       ge-0/0/1.0             Full      192.168.10.2     128    35
10.200.3.2       ge-0/0/3.0             Full      10.200.0.2       128    33
```

- New trace from the R1 router to the R3 router

```
user@R1> traceroute 10.200.0.1
traceroute to 10.200.0.1 (10.200.0.1), 30 hops max, 40 byte packets
 1  10.200.3.2 (10.200.3.2)  0.462 ms  0.293 ms  0.288 ms
 2  10.200.5.1 (10.200.5.1)  0.297 ms  0.295 ms  0.295 ms
 3  10.200.0.1 (10.200.0.1)  0.385 ms  0.376 ms  0.368 ms
```

- The trace takes the R4 – R2 – R3 path

How can you overcome this inefficient routing issue?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Impact of Interface failure.

In this case study, if the interface between R1 and R3 fails, what will be the impact?

First the adjacency between R1 and R3 will no longer exist. This can be verified by running the **`show ospf neighbor`** command.

Traffic will be rerouted to an alternate path.

Using the **traceroute** command, the new path is shown to be R1-R4-R2-R3 as shown in the output on the slide.

Case Study: Multi-Area Adjacency (5 of 7)

- Overcome the routing issue with a multi-area adjacency

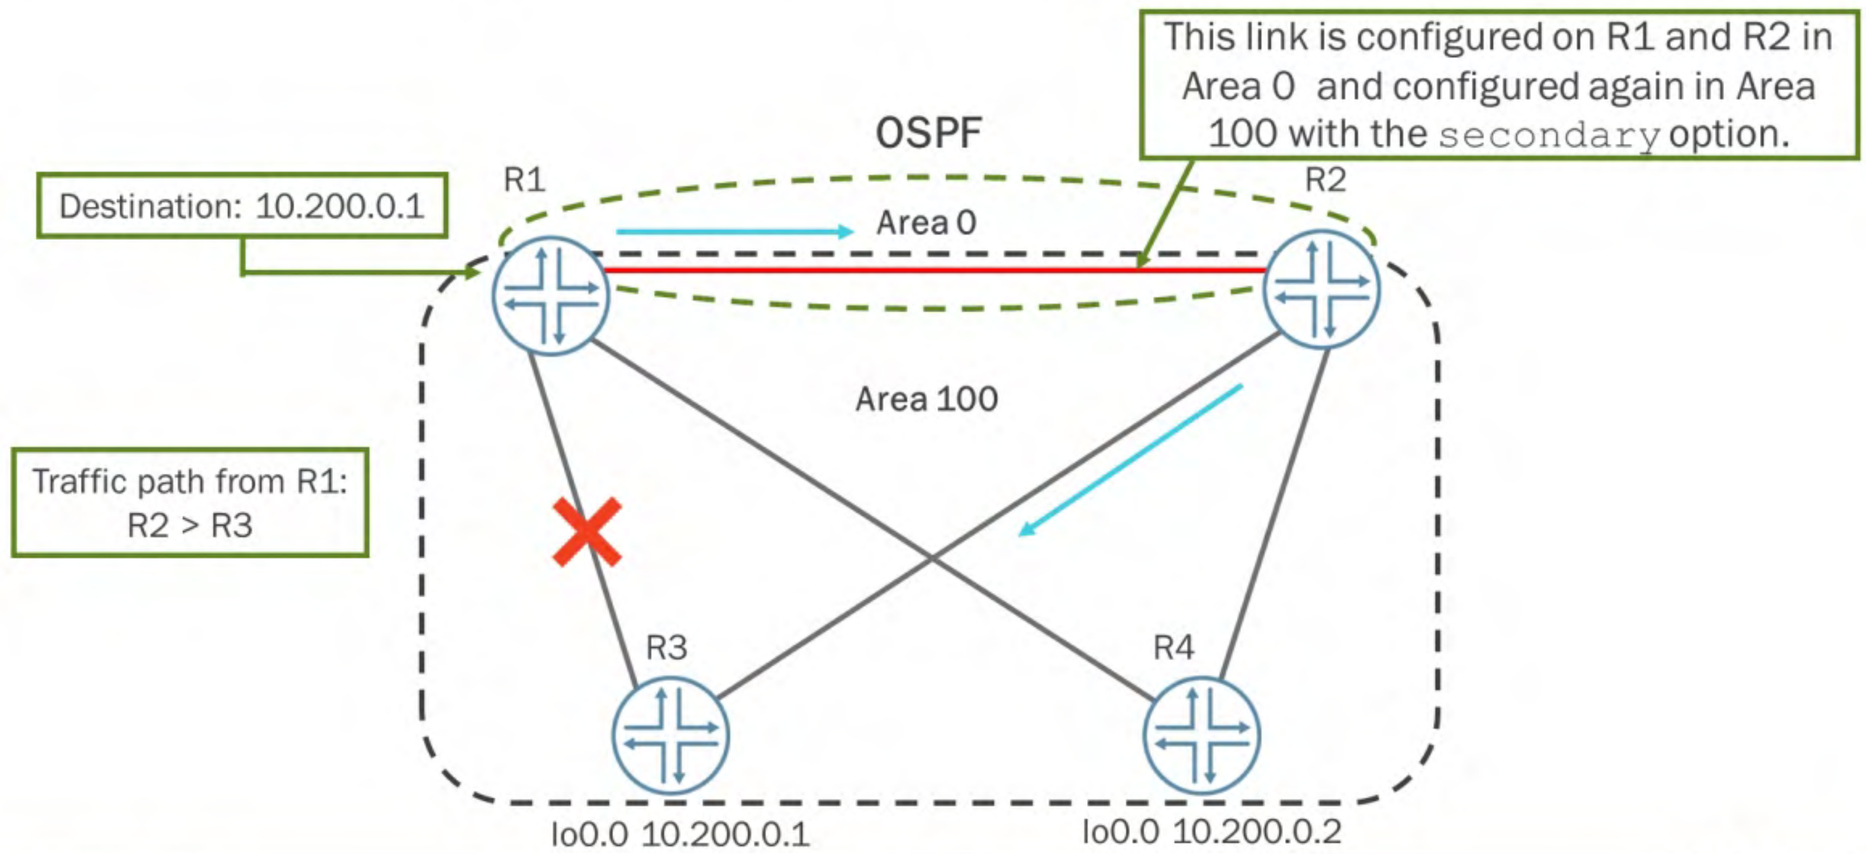

© 2020 Juniper Networks, Inc. All Rights Reserved.

Overcome the inefficient routing issues.

To overcome this issue, the link between R1 and R2 can be configured as a multi-area adjacency. This is possible by configuring the R1, R2 link that exists in area 0, a second time in Area 100 but marked with the `secondary` option.

Case Study: Multi-Area Adjacency (6 of 7)

- Configure the interface between R1 and R2 in Area 100 as a secondary interface

```
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/1.0;
        }
        area 0.0.0.100 {
            interface ge-0/0/3.0;
            interface ge-0/0/2.0;
            interface ge-0/0/1.0 {
                secondary;
```

- A point-to-point interface is established for Area 100

```
user@R1> show ospf interface
Interface          State   Area        DR ID          BDR ID         Nbrs
ge-0/0/1.0         BDR     0.0.0.0     192.168.10.2   192.168.10.1      1
lo0.0              DR      0.0.0.0     192.168.10.1   0.0.0.0           0
ge-0/0/1.0         PtToPt  0.0.0.100   0.0.0.0        0.0.0.0           1
ge-0/0/2.0         DR      0.0.0.100   192.168.10.1   10.200.0.2        1
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Multi-Area Adjacency Configuration

When configuring multi-area adjacencies, keep in mind that you are assigning a single interface to at least two distinct areas: one as the primary, and the others as secondary using the **secondary** option. As stated previously, a given logical interface can be considered as primary for only one area. For any other areas in which you configure that interface, it must be set as a **secondary** interface.

Point-to-Point Interface

The output shows that the ge-0/0/1.0 interface now appears in two distinct OSPF areas, Area 0 and Area 100. However, note the secondary link in Area 100 shows up as a point-to-point interface.

Case Study: Multi-Area Adjacency (7 of 7)

- An adjacency is also formed over a point-to-point interface for Area 100

```
user@R1> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.200.1.2       ge-0/0/1.0             Full      192.168.10.2     128    37
  Area 0.0.0.0
10.200.1.2       ge-0/0/1.0             Full      192.168.10.2     128    37
  Area 0.0.0.100
10.200.3.2       ge-0/0/2.0             Full      10.200.0.2       128    37
  Area 0.0.0.100
```

Two Full adjacencies for the same interface but in different areas

- Another traceroute to R3

```
user@R1> traceroute 10.200.0.1
traceroute to 10.200.0.1 (10.200.0.1), 30 hops max, 40 byte packets
 1  10.200.1.2 (10.200.1.2)  0.417 ms  0.295 ms  0.285 ms
 2  10.200.0.1 (10.200.0.1)  0.372 ms  0.363 ms  0.350 ms
```

  - The trace now takes the R2 – R3 path

© 2020 Juniper Networks, Inc. All Rights Reserved.

Multi-Area Adjacency Configuration

When configuring multi-area adjacencies, keep in mind that you are assigning a single interface to at least two distinct areas: one as the primary, and the others as secondary using the **`secondary`** option. As stated previously, a given logical interface can be considered as primary for only one area. For any other areas in which you configure that interface, it must be set as a **`secondary`** interface.

Point-to-Point Interface

The output shows that the ge-0/0/1.0 interface now appears in two distinct OSPF areas, Area 0 and Area 100. However, note the secondary link in Area 100 shows up as a point-to-point interface.

Agenda: Advanced OSPF Options

- OSPF Multi-Area Adjacencies
- → Virtual Links
- External Reachability

© 2020 Juniper Networks, Inc. All Rights Reserved.

Virtual Links

The slide highlights the topic we discuss next.

Network Integration

- Company acquisitions and mergers—integrating multiple OSPF networks
  - OSPF requires a stable contiguous backbone
    - Unpredictable results
    - Loops
    - Missing routes
  - Can require temporary solutions
    - Virtual links
  - Permanent solutions
    - Physical connectivity

© 2020 Juniper Networks, Inc. All Rights Reserved.

Acquisitions and Mergers

Companies are acquired or merged with other companies every day. These mergers present many interesting challenges, including how to combine the IP networks into one network. For example, imagine two companies running OSPF that must merge networks. For OSPF to work correctly, each company must connect their respective Area 0s together to form a single contiguous backbone. At first glance, you might think the easiest solution would be to create new physical connections between the routers in each company. However, this is often easier said than done. In addition, time can be a deciding factor. For these cases, a temporary solution such as a virtual link can be deployed.

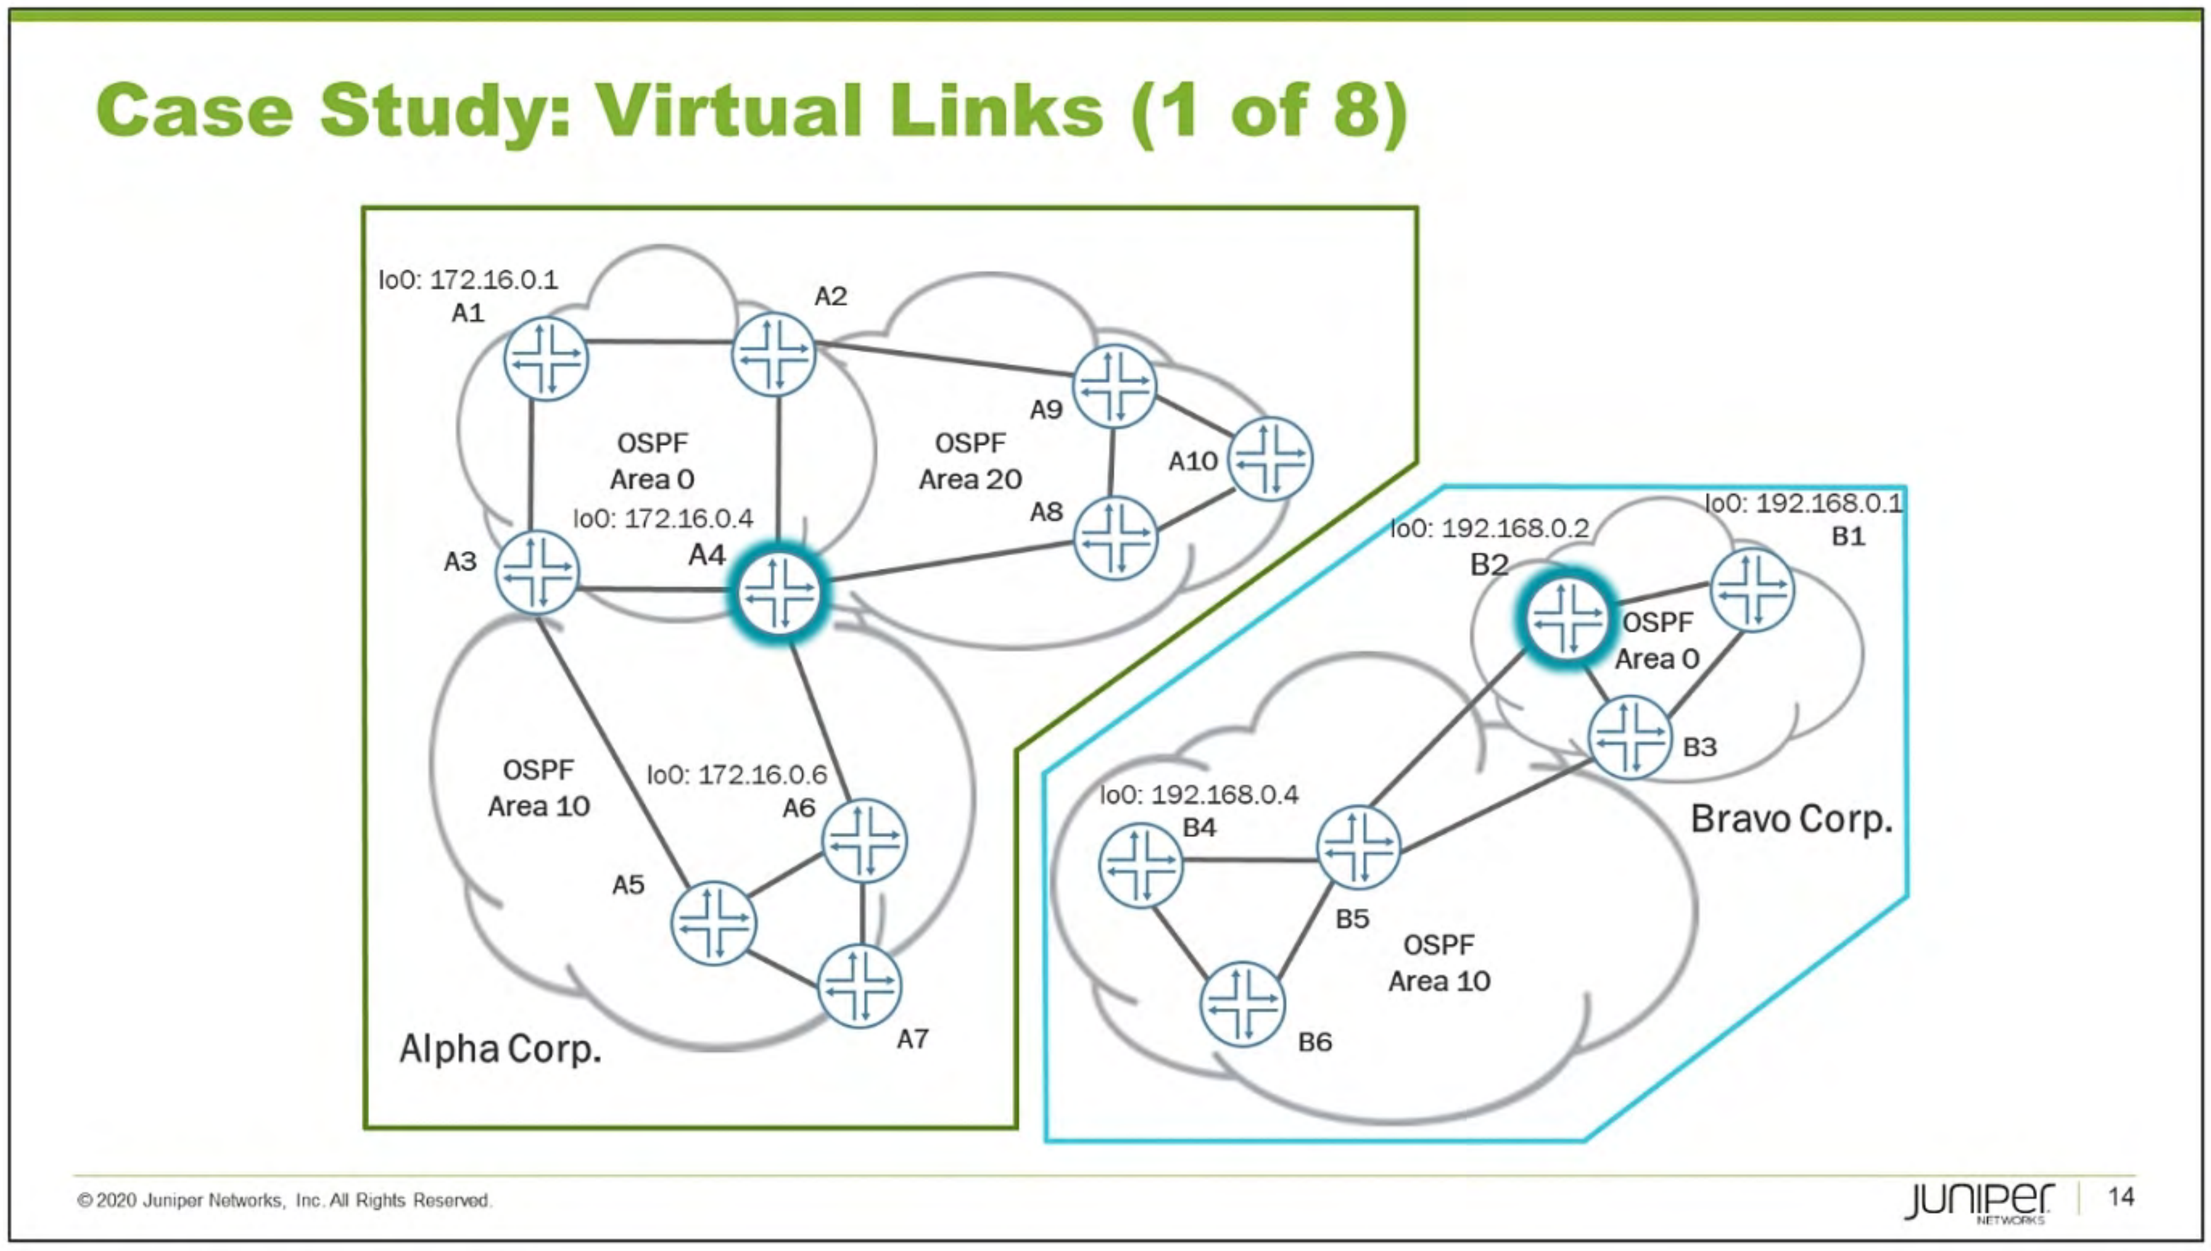

Case Study

In this case study, Alpha Corp. has acquired Bravo Corp. Both networks are running multi-area OSPF and the integration team must get both networks communicating with each other.

Case Study: Virtual Links (2 of 8)

- Alpha Corp. has acquired Bravo Corp.
  - Initial focus is on connecting both corporation's OSPF Area 0 networks.
    - Physical connectivity is not an option
    - Alternative solution has to be used

© 2020 Juniper Networks, Inc. All Rights Reserved.

Integration Case Study

During the acquisition phase, an integration team is formed to look at all facets of combining the two corporations, including their OSPF networks. The determination is that connecting both Area 0 networks together with physical connections is not a viable short term option. An alternative solution must be used.

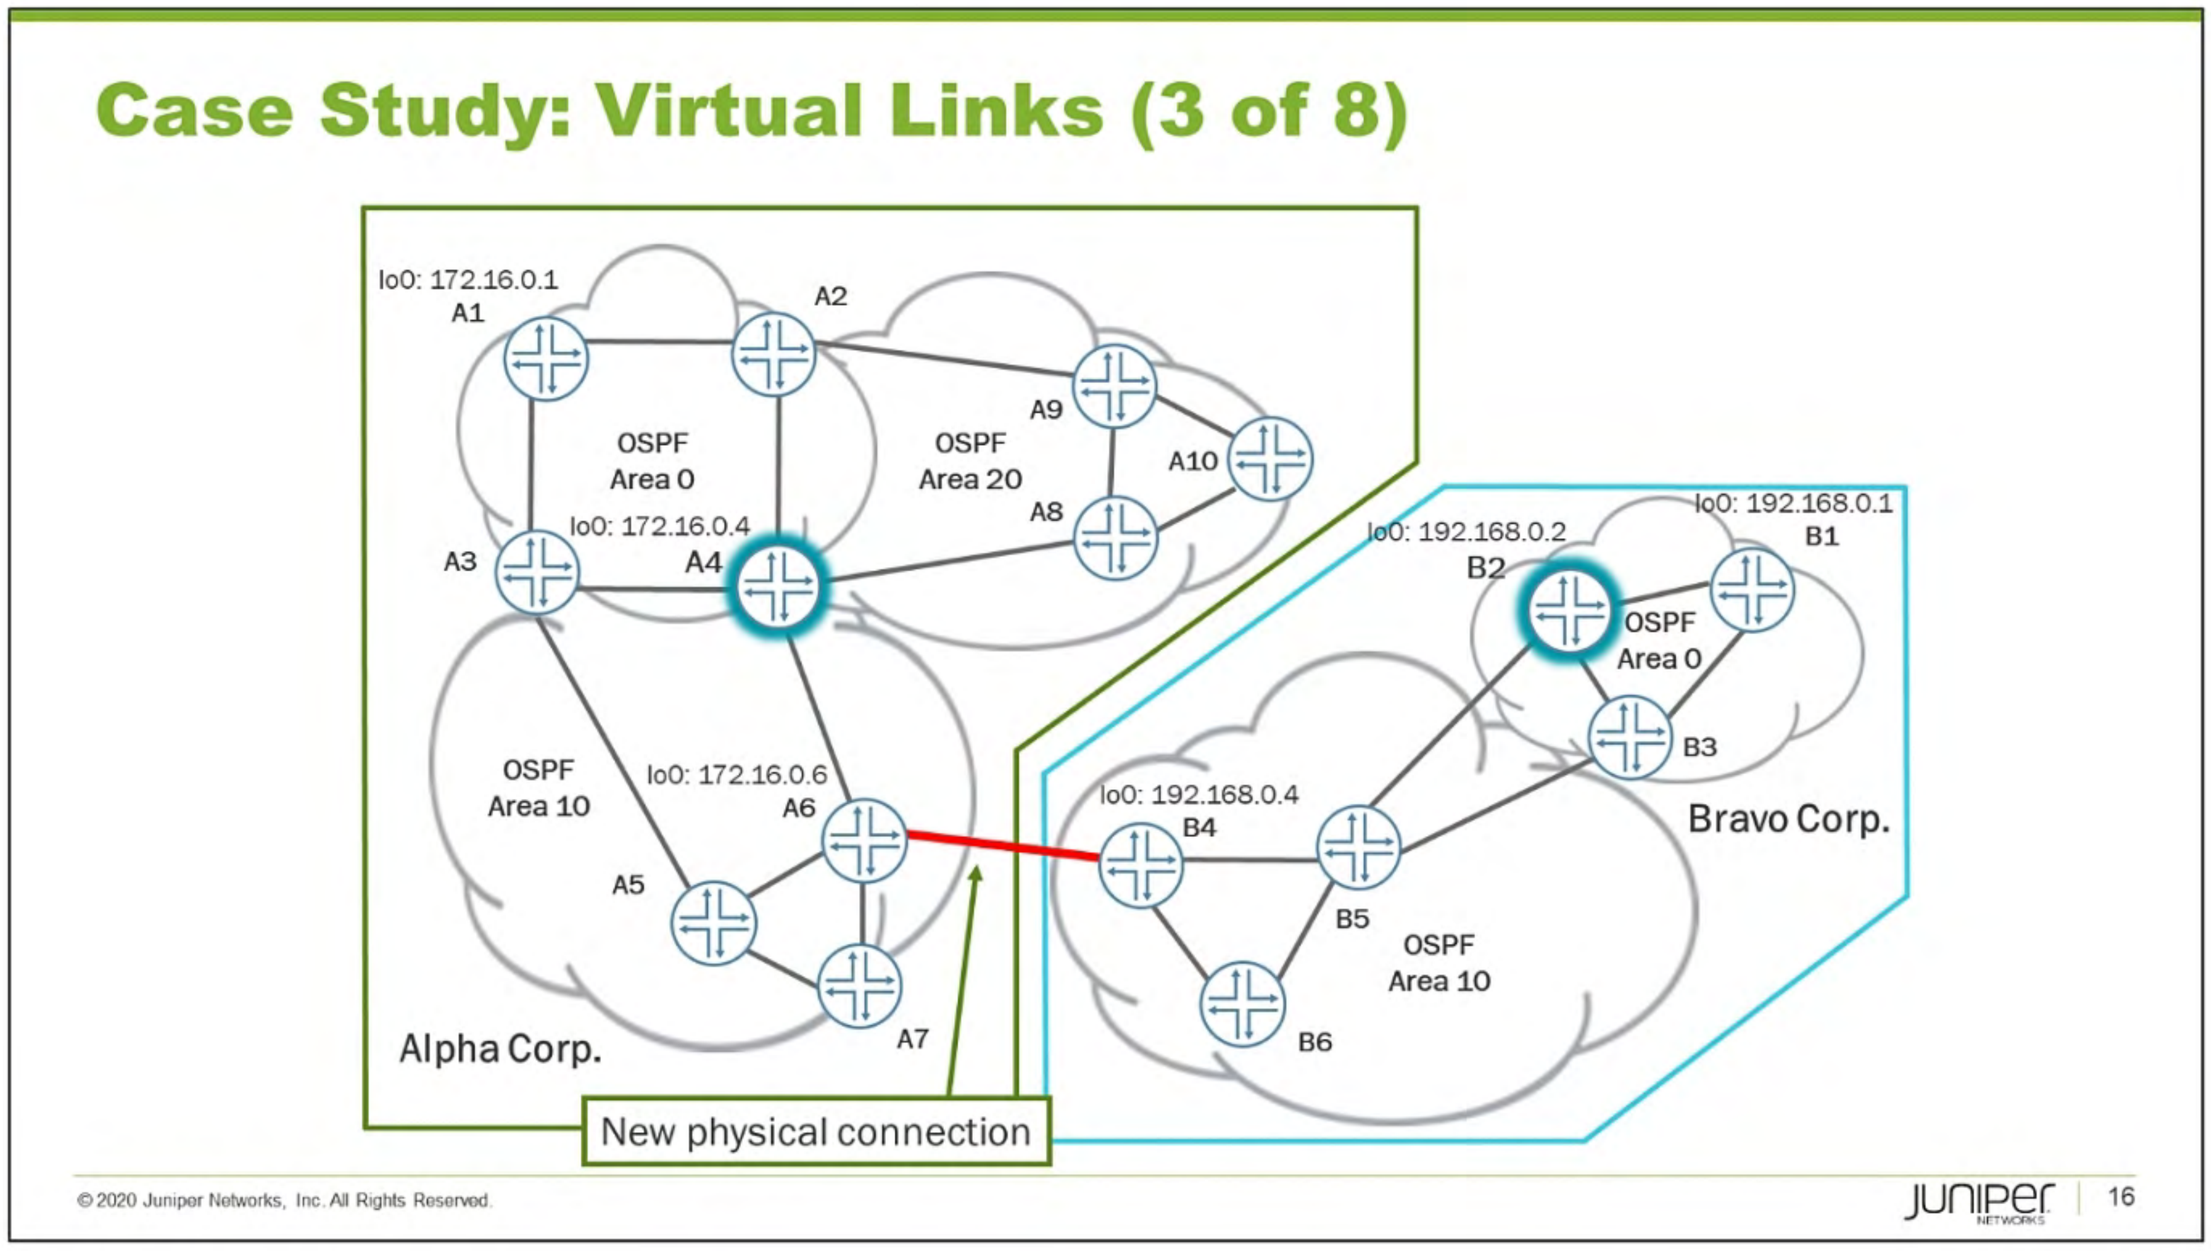

Initial Physical Connection

The first step is to establish *some* physical connectivity between the two corporations. In this case, the integration team chose to connect Alpha Corp's A6 router and Bravo Corp's B4 router. For now, the new interface will be configured in Area 10 on the A6 and B4 routers.

Case Study: Virtual Links (4 of 8)

- Limited connectivity
  - Area 10 routers can reach both Area 0 networks.

```
user@B4> ping count 5 rapid 172.16.0.1
PING 172.16.0.1 (172.16.0.1): 56 data bytes
!!!!!
--- 172.16.0.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.523/0.796/1.586/0.409 ms
```

B4 -> A1

```
user@A6> ping count 5 rapid 192.168.0.1
PING 192.168.0.1 (192.168.0.1): 56 data bytes
!!!!!
--- 192.168.0.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.519/0.581/0.788/0.104 ms
```

A6 -> B1

  - Routers in the separate Area 0 networks, however, cannot talk to each other

```
user@B1> ping count 3 172.16.0.1
PING 172.16.0.1 (172.16.0.1): 56 data bytes
ping: sendto: No route to host
ping: sendto: No route to host
ping: sendto: No route to host

--- 172.16.0.1 ping statistics ---
3 packets transmitted, 0 packets received, 100%
packet loss
```

B1 -> A1

© 2020 Juniper Networks, Inc. All Rights Reserved.

Connectivity Issues

As soon as the physical connection is created, limited connectivity is achieved. For example, the B4 router can now reach the A1 router in Alpha Corp's Area 0. However, Alpha Corp's Area 0 routers cannot reach Bravo Corp's Area 0 routers. The cause of this limited connectivity is the lack of a contiguous Area 0 backbone.

Case Study: Virtual Links (5 of 8)

- Virtual link
  - Provides a logical connection
    - Used for areas not physically connected to Area 0
    - Used to connect a discontiguous Area 0
  - Tunnels OSPF packets through a transit area
    - Creates a virtual ABR to remote routers
    - Configuration on both ends of the tunnel
    - Does NOT tunnel data packets
  - Control plane only
    - SPF will calculate the shortest path to all routes

© 2020 Juniper Networks, Inc. All Rights Reserved.

Virtual Link

One solution to the connectivity problem is to create a virtual tunnel between the two backbone areas of the corporations. This feature, known as a virtual link, provides a logical connection between areas. Essentially, OSPF packets are tunneled through a transit area to establish an OSPF adjacency and logically connect the two areas together. This establishes full connectivity between the two corporations.

Remember that a virtual link is a control plane feature only. SPF will still calculate the shortest physical path between two points, which might not be the same path as the virtual tunnel. This calculation could create some confusion when troubleshooting, which is one of the primary reasons virtual tunnels are not considered long term solutions.

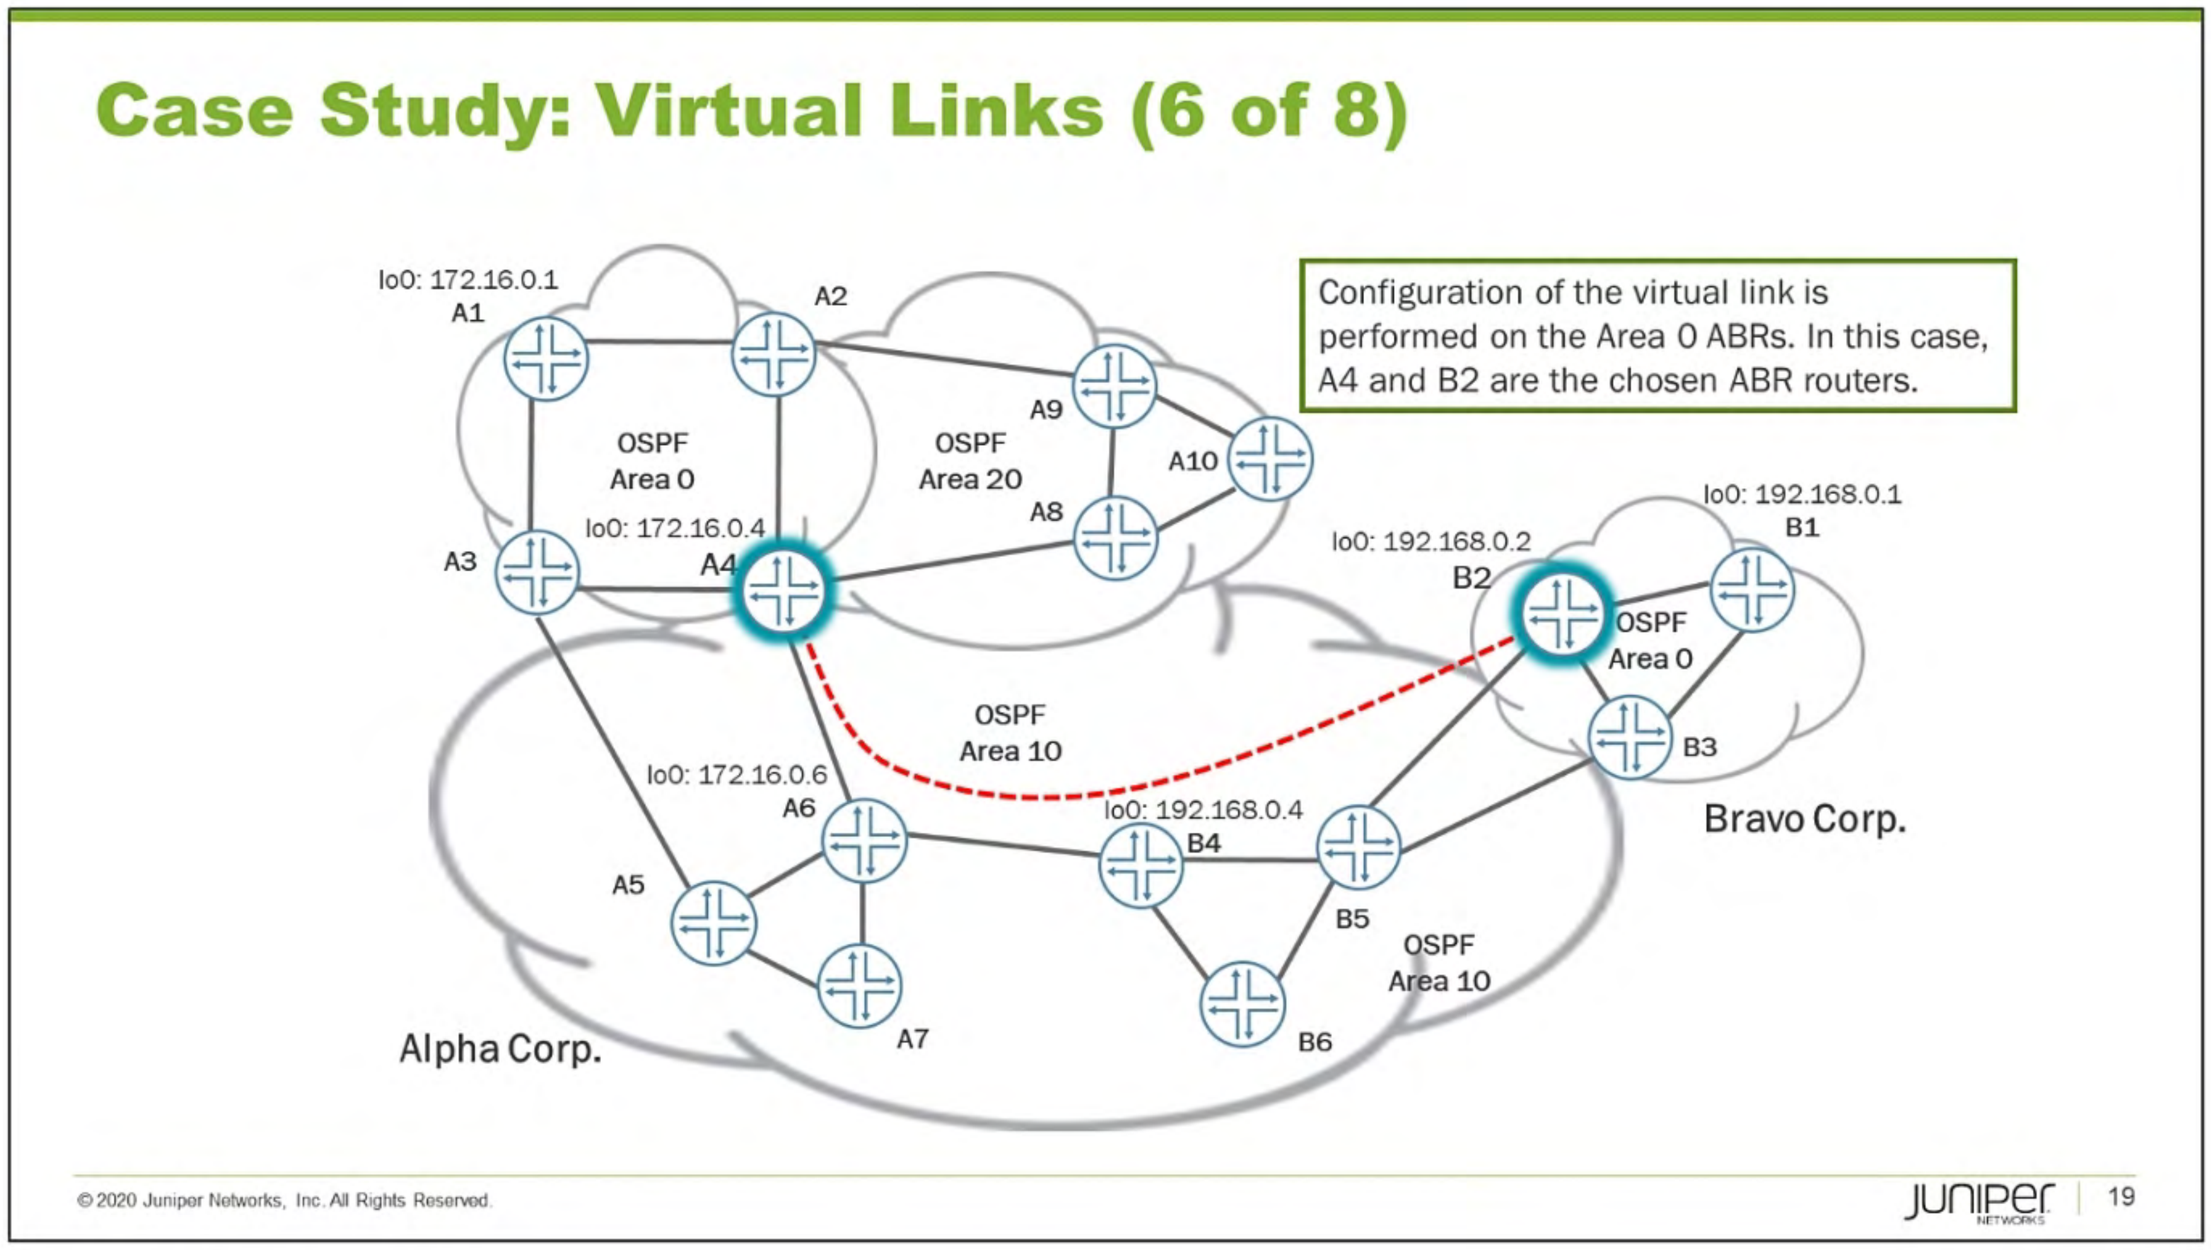

Virtual Link Established

In this case, a virtual link is established between ABRs in each corporation. These ABRs must be attached to Area 0. In the slide, A4 and B2 were chosen as the ABRs on which to establish the virtual link.

Case Study: Virtual Links (7 of 8)

- Configure a virtual link in Area 0

```
user@A4# show protocols ospf area 0
virtual-link neighbor-id 192.168.0.2 transit-area 0.0.0.10;
interface lo0.4;
interface ge-1/1/5.0 {
    interface-type p2p;
}

user@B2# show protocols ospf area 0
virtual-link neighbor-id 172.16.0.4 transit-area 0.0.0.10;
interface lo0.2;
interface ge-1/1/4.0 {
    interface-type p2p;
}
```

Virtual links are one-way and must be configured on each side of the virtual link.

- Verify the virtual link

```
user@B2> show ospf interface
ge-1/1/4.0          PtToPt  0.0.0.0        0.0.0.0        0.0.0.0        1
lo0.2               DR      0.0.0.0        192.168.0.2    0.0.0.0        0
vl-172.16.0.4       PtToPt  0.0.0.0        0.0.0.0        0.0.0.0        1
ge-1/0/5.0          PtToPt  0.0.0.10       0.0.0.0        0.0.0.0        1

user@B2> show ospf neighbor
Address          Interface              State     ID              Pri  Dead
10.8.0.1         ge-1/1/4.0             Full      192.168.0.1     128    38
10.0.0.9         vl-172.16.0.4          Full      172.16.0.4        0    34
10.8.0.6         ge-1/0/5.0             Full      192.168.0.5     128    39
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Virtual Link Configuration

The configuration of a virtual link takes place within the Area 0.0.0.0 portion of the OSPF hierarchy. The **`virtual-link`** command itself requires both a transit area and a neighbor ID to be configured. The transit area is the OSPF area through which you configure the virtual link. The neighbor ID is the 32-bit router ID (RID) of the router on the far end of the virtual link. Once each side completes this configuration, each router begins to send unicast OSPF traffic towards the far-end router to complete the link setup and form an adjacency.

Virtual Link as an Interface

Once the two ends of the link can communicate, the virtual link becomes an operational OSPF interface. It appears in **`show`** commands and within the OSPF LSDB. It is always noted in a format of `vl-`*`neighbor-id`*, where `vl` denotes it as a virtual link, and the *`neighbor-id`* is the RID of the far-end router.

Case Study: Virtual Links (8 of 8)

- Contiguous Area 0
  - Type 1 LSAs are exchanged over the tunnel

```
user@B1> show ospf database
 OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.16.0.1       172.16.0.1       0x8000000b  1716  0x22 0x9557  84
Router   172.16.0.2       172.16.0.2       0x80000002  1717  0x22 0xb217  84
Router   172.16.0.3       172.16.0.3       0x8000000a  1714  0x22 0x45b9  84
[...]
```

  - Routes are installed in the routing table

```
user@B1> show route protocol ospf 172.16/16
inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32      *[OSPF/10] 00:27:26, metric 7
                    > to 10.8.0.2 via ge-1/0/4.0
172.16.0.2/32      *[OSPF/10] 00:02:05, metric 6
                    > to 10.8.0.2 via ge-1/0/4.0
172.16.0.3/32      *[OSPF/10] 00:27:26, metric 6
                    > to 10.8.0.2 via ge-1/0/4.0
[...]
```

  - Router B1 can now reach router A1

```
user@B1> ping count 5 rapid 172.16.0.1
PING 172.16.0.1 (172.16.0.1): 56 data bytes
!!!!!
--- 172.16.0.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.582/0.681/0.879/0.140 ms
```

B1 -> A1

© 2020 Juniper Networks, Inc. All Rights Reserved.

Contiguous Area 0

Once the neighbor is established over the virtual link, connectivity is restored, all LSAs are processed, and routes to each company are installed into the routing table. As shown in the slide, router B1 can now successfully ping router A1.

Agenda: Advanced OSPF Options

- OSPF Multi-Area Adjacencies
- Virtual Links
- →External Reachability

© 2020 Juniper Networks, Inc. All Rights Reserved.

External Reachability

The slide highlights the topic we discuss next.

OSPF External Reachability (1 of 2)

- OSPF default export policy
  - Routes are not redistributed between protocols by default
  - Default export policy rejects all routes

```
policy-options {
    policy-statement ospf-default-policy {
        term reject-routes {
            then reject;
        }
    }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Default Export Policy

By default, Junos rejects all routes from other protocols. The only way to redistribute routes into OSPF is to create and apply an export policy.

The slide depicts what the hidden default policy applied to OSPF looks like. This policy will take effect unless some other policy is written and applied to OSPF.

Some administrators will manually create and apply a default policy to document the default action.

OSPF External Reachability (2 of 2)

- Route redistribution requires an export policy
  - Applied at the OSPF global level
    - Creates Type 5 External LSAs
      - Type 1 or Type 2 (default)

```
policy-options {
    policy-statement redistribute-statics {
        term static-routes {
            from protocol static;
            then {
                external type 1;
                accept;
            }
        }
    }
}

protocols ospf {
    export redistribute-statics;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Route Redistribution

For route distribution to occur, an export policy must be written and applied. Because external routes in OSPF have an inter-area flooding scope, the policies are applied globally. This feature allows external routes to be sent into all areas that allow it. When an external route is redistributed into OSPF, it will be added to the LSDB as an external Type 5 LSA (or Type 7 LSA in a NSSA area) with a Type 2 external metric. If an external LSA must have a Type 1 external metric then that must be configured in the policy.

For review:

- A Type 2 external metric uses the encoded metric as the total cost for the route when performing an SPF calculation.
- A Type 1 external metric uses the encoded metric of the route plus the cost to reach the advertising ASBR as the total cost for the route.

Case Study: OSPF Prefix Limits (1 of 3)

- Prefix limits are used as protection against redistributing too many external routes into OSPF
  - Redistribution requires an export policy
    - An incorrect OSPF export policy can inadvertently inject an excessive number of external routes
    - An external peer router can suddenly start advertising more routes than expected
  - One Type 5 LSA is created per external route
    - Each Type 5 is flooded across the OSPF domain

© 2020 Juniper Networks, Inc. All Rights Reserved.

Route Protection

Some OSPF implementations encounter problems when large numbers of external routes are injected into the LSDB. While not normally an issue for Junos devices, due to its architecture, it can become an issue depending on how many routes are being received and accepted. As an OSPF network administrator you may not need to worry about how Junos devices will perform with a large number of routes, but if you have a multi-vendor environment other routers could be adversely affected by receiving an excessive number of routes.

To protect the network from potential configuration issues that may accept a large number of routes, intentionally or unintentionally, the Junos OS allows you to limit the number of prefixes that can be accepted. The **`prefix-export-limit`** command informs the router how many routes to accept from a routing policy configuration. The command accepts a 32-bit value, which provides a range of routes from 1 to 4,294,967,295. Once the route limit is reached, the router transitions into an overload state where the local links are set to a metric of 65,535 in the router LSA. Additionally, all Type 5 LSAs from the router are purged from the database and the network in general. The local router remains in this state until the number of external routes returns to a level below the configured limit. This situation requires the administrator to manually change the existing configuration; either the configured **`prefix-export-limit`** must be increased or the routing policy must be changed to redistribute fewer routes.

Case Study: OSPF Prefix Limits (2 of 3)

- Create an OSPF export policy

```
[edit]
user@R2# show policy-options
    policy-statement redistribute-ISIS {
        term 1 {
            from protocol ISIS;
            then accept;
```

- Too many ISIS routes are being redistributed

Eight prefixes are now being exported into OSPF.

```
user@R2> show ospf database area 0.0.0.100
 OSPF link state database, Area 0.0.0.100
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
...
NSSA     0.0.0.0          10.200.0.3       0x80000003  1402  0x20 0x2d24  36
NSSA     0.0.0.0          10.200.0.4       0x80000003  1453  0x20 0x2729  36
NSSA    *10.22.0.0        10.200.0.1       0x80000001    10  0x28 0xe7f0  36
NSSA    *203.0.113.0/24   10.200.0.1       0x80000001  1324  0x28 0xdcfa  36
NSSA    *10.22.2.0        10.200.0.1       0x80000001    10  0x28 0xd105  36
NSSA    *10.22.3.0        10.200.0.1       0x80000001    10  0x28 0xc60f  36
NSSA    *10.22.4.0        10.200.0.1       0x80000001    10  0x28 0xbb19  36
NSSA    *10.22.5.0        10.200.0.1       0x80000001    10  0x28 0xb023  36
NSSA    *10.22.6.0        10.200.0.1       0x80000001    10  0x28 0xa52d  36
NSSA    *10.22.7.0        10.200.0.1       0x80000001    10  0x28 0x9a37  36
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study: A Generic Policy is created

The issue can manifest itself when a user creates a generic policy that does not restrict the number of routes that can possibly be selected. In the example all ISIS routes will be included in this policy. If ISIS had a large number of route, this policy could result in the OSPF LSDB being flooded with hundreds or thousands of routes.

The Result

For demonstration purposes, the number of routes being imported in 8. The corporate rules dictate that no more than 6 routes can be redistributed into this OSPF instance on the OSPF router. A policy may limit route redistribution by prefix but it cannot determine the number of routes that might be included.

Case Study: OSPF Prefix Limits (3 of 3)

- An OSPF prefix-limit is configured

```
[edit]
user@R2# show protocols
    ospf {
        prefix-export-limit 6;
        external-preference 90;
        export redistribute-isis;
```

The number of exported prefixes (8) exceeds the set limit of 6. Therefore, none of the prefixes are exported.

- No ISIS routes are redistributed

```
user@R2> show ospf database area 0.0.0.100

 OSPF link state database, Area 0.0.0.100
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
NSSA     0.0.0.0          10.200.0.3       0x80000003  1593  0x20 0x2d24  36
NSSA     0.0.0.0          10.200.0.4       0x80000003  1644  0x20 0x2729  36
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Prefix Limit

To stop the large amount of LSAs that could enter the router, a prefix-limit of six is configured.

The Result

The result is that no ISIS routes are distributed. This prefix-limit setting ensures that a configuration error does not affect the network.

This option can be critical to a Service Provider network in the event that an administrator writes and applies a policy that could allow thousands of routes from a protocol like BGP to be redistributed into OSPF.

Another key use in a Service Provider network is to limit the number of routes received by individual customers based on contractual agreements, This can be used to prevent inproper customer configurations from consuming capacity of the Service Providers Edge routers in excess of their agreement.

OSPF Mutual Redistribution

- Mutual redistribution
  - Redistribution from multiple routers
  - Redistribution from multiple protocols
  - Common reason for routing loops
  - Redistribution rule
    - Source route has a lower preference (smaller number)
      - No problem
    - Source route has a higher preference (larger number)
      - Sub-optimal routing
      - Routing loops

© 2020 Juniper Networks, Inc. All Rights Reserved.

Mutual Redistribution

Special care must be taken when redistribution is configured in a network. When multiple redistribution points are present, sub-optimal routing and loops could occur. Generally, if the source route has a lower preference (smaller number) than the protocol into which it is being redistributed, no issues occur. However, when the source route has a higher preference (larger number), issues can occur. Several methods exist to resolve these issues, but the easiest method usually involves modifying route preference values.

Case Study: External Routes (1 of 6)

- Case study topology:

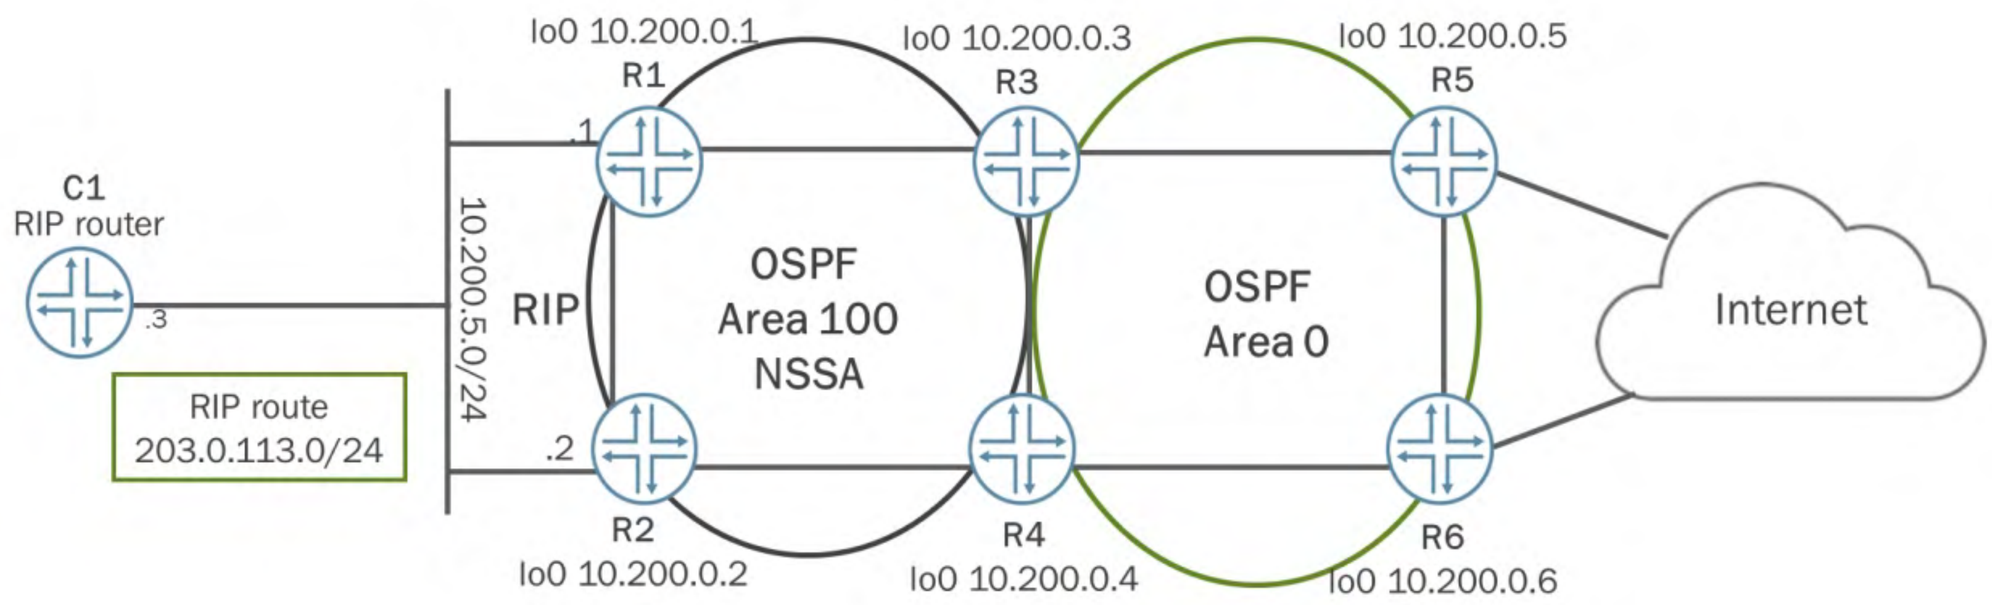

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study Topology

The slide displays the case study topology that will be used in subsequent slides. R1, R2, and C1 are all running RIP.

Case Study: External Routes (2 of 6)

- Case study background
  - The RIP router is advertising routes
    - Redistribute a single RIP route into OSPF—advertise the 203.0.113.0/24 RIP route
  - The R1 and R2 routers have an OSPF default route
    - Default route is generated by the ABR—a Type 7 LSA is created
    - Redistribute the default route into RIP

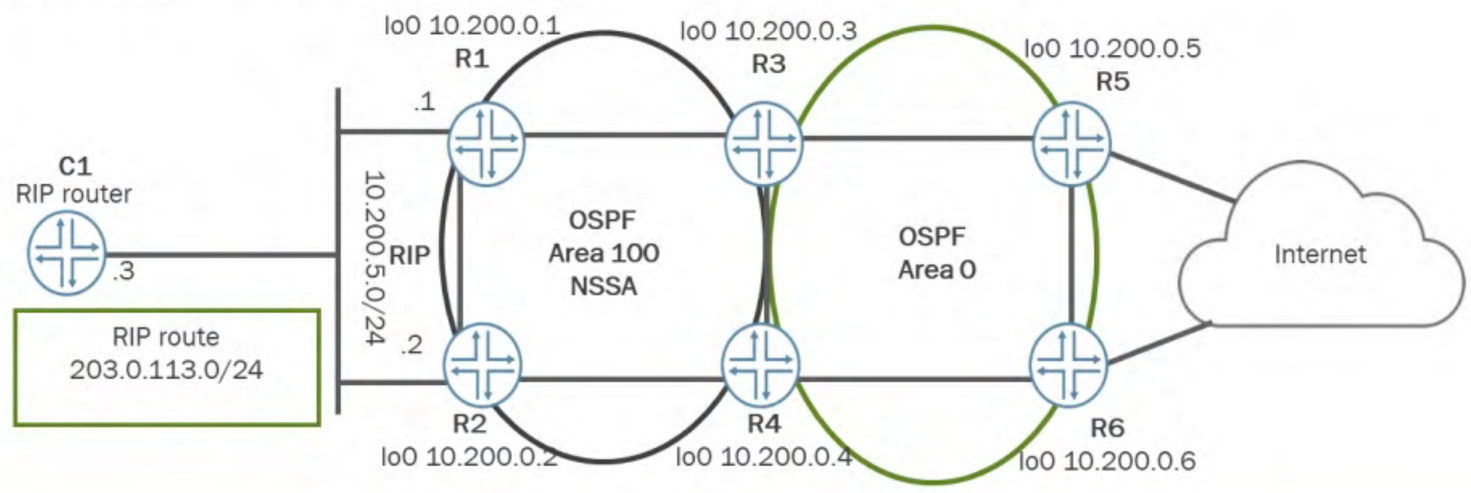

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study Background

The slide describes the goal of the case study. The goal is to advertise a single RIP route into OSPF as well as send a default route to RIP.

Case Study: External Routes (3 of 6)

- Redistribute the RIP route into OSPF

```
[edit]
user@R1# show policy-options policy-statement redistribute-rip
    term 1 {
        from {
            protocol rip;
            route-filter 203.0.113.0/24 exact;
        }
        then accept;
    }

user@R1# show protocols ospf
export redistribute-rip;
[...]
```

Same policy configured and applied on R2

- A Type 7 LSA is created for the RIP route

```
user@R1> show ospf database nssa

    OSPF database, Area 0.0.0.100
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
NSSA    *203.0.113.0       10.200.0.1      0x80000001    10  0x28 0x9632  36
NSSA     203.0.113.0       10.200.0.2      0x80000001    11  0x28 0x9e28  36
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Create a Policy for the RIP Route

The first step is to create a policy and apply the policy to OSPF. In this case, two match conditions are used, creating a logical AND. This policy will be applied to both R1 and R2 under the `[edit protocols ospf]` hierarchy by performing a **`set protocols ospf export redistribute-rip`** command.

Verify Policy Operation

Verify that the policy is working by examining the database and finding the Type 7 associated with the RIP route.

Case Study: External Routes (4 of 6)

- The R4 ABR converts the Type 7 LSA to a Type 5 LSA

```
user@R4> show ospf database external detail
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *203.0.113.0       10.200.0.4      0x80000005   289  0x22 0x4790  36
  mask 255.255.255.0
  Topology default (ID 0)
    Type: 2, Metric: 2, Fwd addr: 10.200.0.2, Tag: 0.0.0.0
```

- The ASBR address is advertised in the forwarding address of the Type 5 LSA.

© 2020 Juniper Networks, Inc. All Rights Reserved.

ABR Translation

Because the route was originated from the not-so-stubby area (NSSA), the ABR must convert the Type 7 LSA to a Type 5 for inter-area advertisement. Only ABR with the highest RID performs this translation.

Forwarding Address

When the ABR translates a Type 7 into a Type 5, it places the address of the autonomous system boundary router (ASBR) into the forwarding address field. This behavior provides for optimal routing because only one ABR translates the Type 7 LSAs to Type 5 LSAs in the presence of multiple ABRs.

Case Study: External Routes (5 of 6)

- Redistribute the default route from OSPF into RIP
  - High preference of 150 to a low preference of 100

```
user@R2> show route 0/0

inet.0: 20 destinations, 22 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[RIP/100] 05:11:22, metric 2, tag 0
                    > to 10.200.5.1 via ge-1/0/6.110
                    [OSPF/150] 05:11:19, metric 11, tag 0
                    > to 10.203.5.2 via ge-1/0/8.2
```

Toward the C1 RIP router

Toward the R4 OSPF router

  - The RIP route is the active route
  - Modify the OSPF external preference to 90

```
[edit]
user@R2# set protocols ospf external-preference 90

[edit]
user@R2# commit
```

Same command is applied to the R1 router as well

© 2020 Juniper Networks, Inc. All Rights Reserved.

Redistribution of Default Route

The next step is to redistribute the default route from OSPF into RIP using an export policy under RIP. By default, RIP has a lower preference than OSPF external routes.

Because RIP has a better preference, the default route for RIP is preferred. In the sample network, this preference creates a loop, because the OSPF routers point to the RIP router as their gateway, and the RIP router points to the OSPF ASBRs. To fix the loop, modify the OSPF external route preference to a lower value than the RIP route.

Case Study: External Routes (6 of 6)

- Mutual redistribution
  - The default route has been fixed
  - OSPF now has a better preference than RIP and is the preferred route

```
user@R2> show route 0/0

inet.0: 20 destinations, 22 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/90] 00:12:59, metric 11, tag 0
                    > to 10.203.5.2 via ge-1/0/8.2
                    [RIP/100] 05:31:17, metric 2, tag 0
                    > to 10.200.5.1 via ge-1/0/6.110
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

The Result

The result of the preference change is now a default that points properly to the ABRs in the NSSA.

Case Study: OSPF Import Policy (1 of 3)

- You need to prevent the RIP routes from appearing on R4

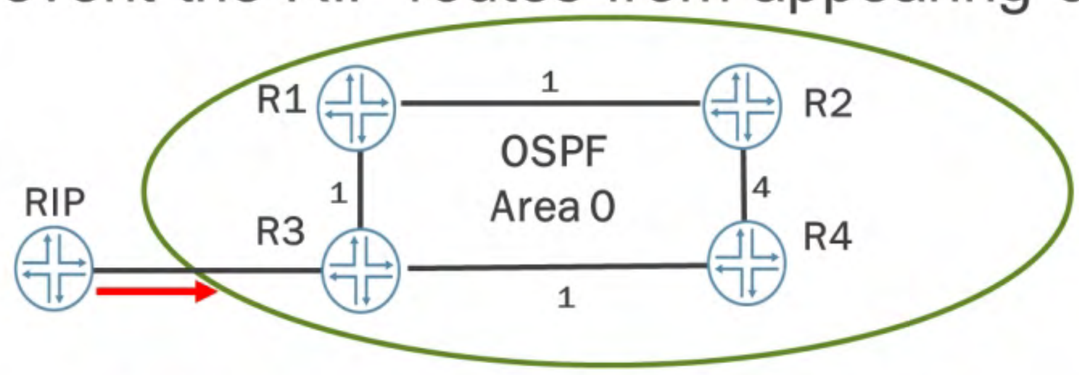

- An import policy can be used to accomplish this task
- Import policies only affect which external routes are copied from the tree database into the inet.0 table

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study: OSPF Import Policy

In this case study, you are importing RIP routes into your OSPF single area network. You have been told that these routes cannot appear in the inet.0 route table on the R4 router. Since all learned links are flooded to all routers how would this be possible.

A quick review of OSPF can reveal a solution. OSPF LSAs are received and stored in the link state database. This link state data can be viewed using the **show ospf database** command. The SPF algorithm is executed against this database and the results are stored in the tree database. This file can be viewed using the **show ospf route** command. The resulting prefix and path mappings are copied into the inet.0 table. These OSPF routes can be viewed using the **show route protocol ospf** command.

An import policy does not impact the LSDB nor the SPF calculation. It is only able to block OSPF external routes that are in the Tree database from being copied into the inet.0 table..

Even though an import policy does not impact LSA flooding nor the SPF algorithm, are there issues that an import policy can create? This is discussed on the next slide.

Case Study: OSPF Import Policy (2 of 3)

- Import policies must be used with care!

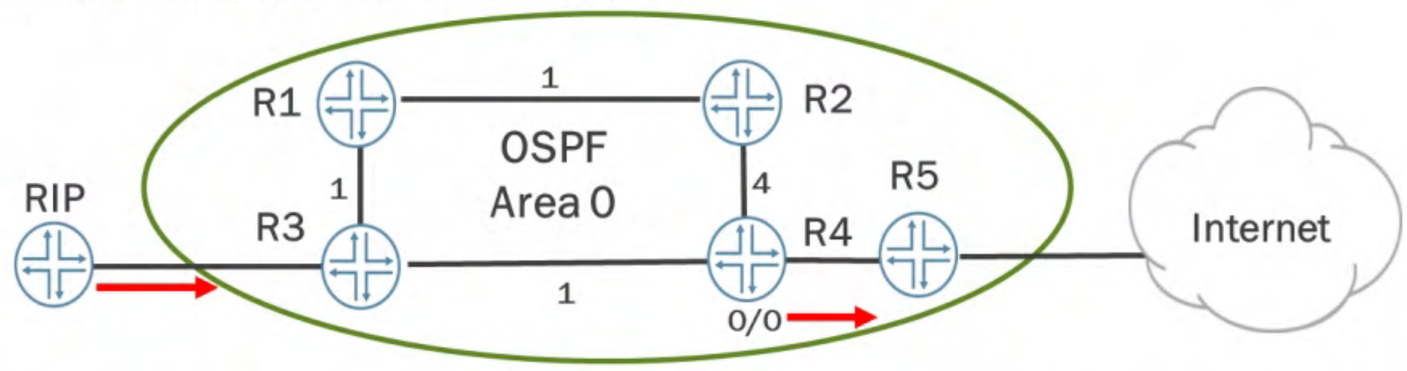

- R5 will know about the RIP routes because the import policy does not affect LSA distribution
- R5 sending traffic towards RIP router could be black holed at R4.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Import Policy Issues

Import policies must be used with care! Let.s evaluate the following possible network design. You have restricted the RIP routes from appearing in the inet.0 table. However, since the routes are still in the link state database, these external LSA will be flooded to R5. R5 does not have a similar import policy so R5 knows about the RIP routes. If traffic gets forwarded to towards the RIP router, those routes will be forwarded to R4 and be black holed.

Now what happens if R5 is also connected to the Internet and R4 has a default route pointing to R5? When R5 forwards traffic to the RIP router via R4, R4 does not have more specific routes so it uses the default route and sends it back to R5. Now a loop has been formed. The results are still the same. The traffic cannot be forwarded to its destination.

This is an important take away from this. Import policies may seem an effective way to solve a specific problem, but it needs to be used with care!

Case Study: OSPF Import Policy (3 of 3)

- Import policy is applied globally in the OSPF configuration

```
[edit]
user@R4# show policy-options
    policy-statement ospf-import-policy {
        term 1 {
            from {
                route-filter 203.0.113.0/30 exact;
            }
            then reject;
        }
    }
}
```

```
[edit]
user@R4# show protocols
    ospf {
        import ospf-import-policy;
    }
```

Due to the potential for black holing traffic, import policies should be as specific as possible and router location carefully considered

```
[edit]
user@R4# run show route 203.0.113.0/30
Run
[edit]
user@R4#
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Import Policy

An import policy is used to identify the routes to block as Junos copies the files from the Tree database to the inet.0 table. Import policies are applied globally within the OSPF hierarchy using the key word **import**.

To verify the routes are not in the table use the **show route** command.

Summary

- In this content, we:
  - Configured OSPF multi-area adjacency
  - Configured OSPF virtual links
  - Explained OSPF external reachability

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Discussed:

- Configuring OSPF multi-area adjacencies;
- Configuring OSPF virtual links; and
- OSPF external reachability.

Review Questions

1. What are the dangers of mutually exporting routes between IGPs?
2. Why must you have a backbone OSPF area?
3. Under which circumstances might you need to deploy a virtual link?
4. What issues should you be aware of when redistributing routes using two ASBRs?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review Questions

1.

2.

3.

4.

Lab: Configuring Advanced OSPF Options

- Configure a virtual link.
- Configure multi-area adjacencies.
- Examine prefix limits.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab: Configuring Advanced OSPF Options

The slide provides the objectives for this lab.

Answers to Review Questions

1.

You risk double-exporting routes, which can lead to routing loops.

2.

Although technically you do not need a backbone area, functionally you need one once three or more areas exist. The rule that all non-backbone areas must directly connect to the backbone area is the loop prevention mechanism of a multi-area OSPF network.

3.

Virtual links would be deployed if two companies were merging their networks together and physical link connectivity was not an option. This could be because of cost or time constraints.

4.

If the source route has a lower preference (more preferred), there usually are no issues. If a source route has a higher preference (less preferred), suboptimal routing or loops can occur.

Education Services

Advanced Junos Service Provider Routing

Chapter 5: Troubleshooting OSPF

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - List useful commands that are used to troubleshoot and verify OSPF
  - Isolate different OSPF issues

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Will Discuss:

- Useful commands that can be used to troubleshoot and verify correct operation of OSPF; and
- Isolating different OSPF issues.

Agenda: OSPF Troubleshooting

→Troubleshooting OSPF Adjacency Issues
- Troubleshooting LSDB Consistency Issues
- Troubleshooting OSPF Routing Issues
- Case Study – Adjacency Issues

© 2020 Juniper Networks, Inc. All Rights Reserved.

Troubleshooting OSPF Adjacency Issues

The slide lists the topics we will discuss. We discuss the highlighted topic first.

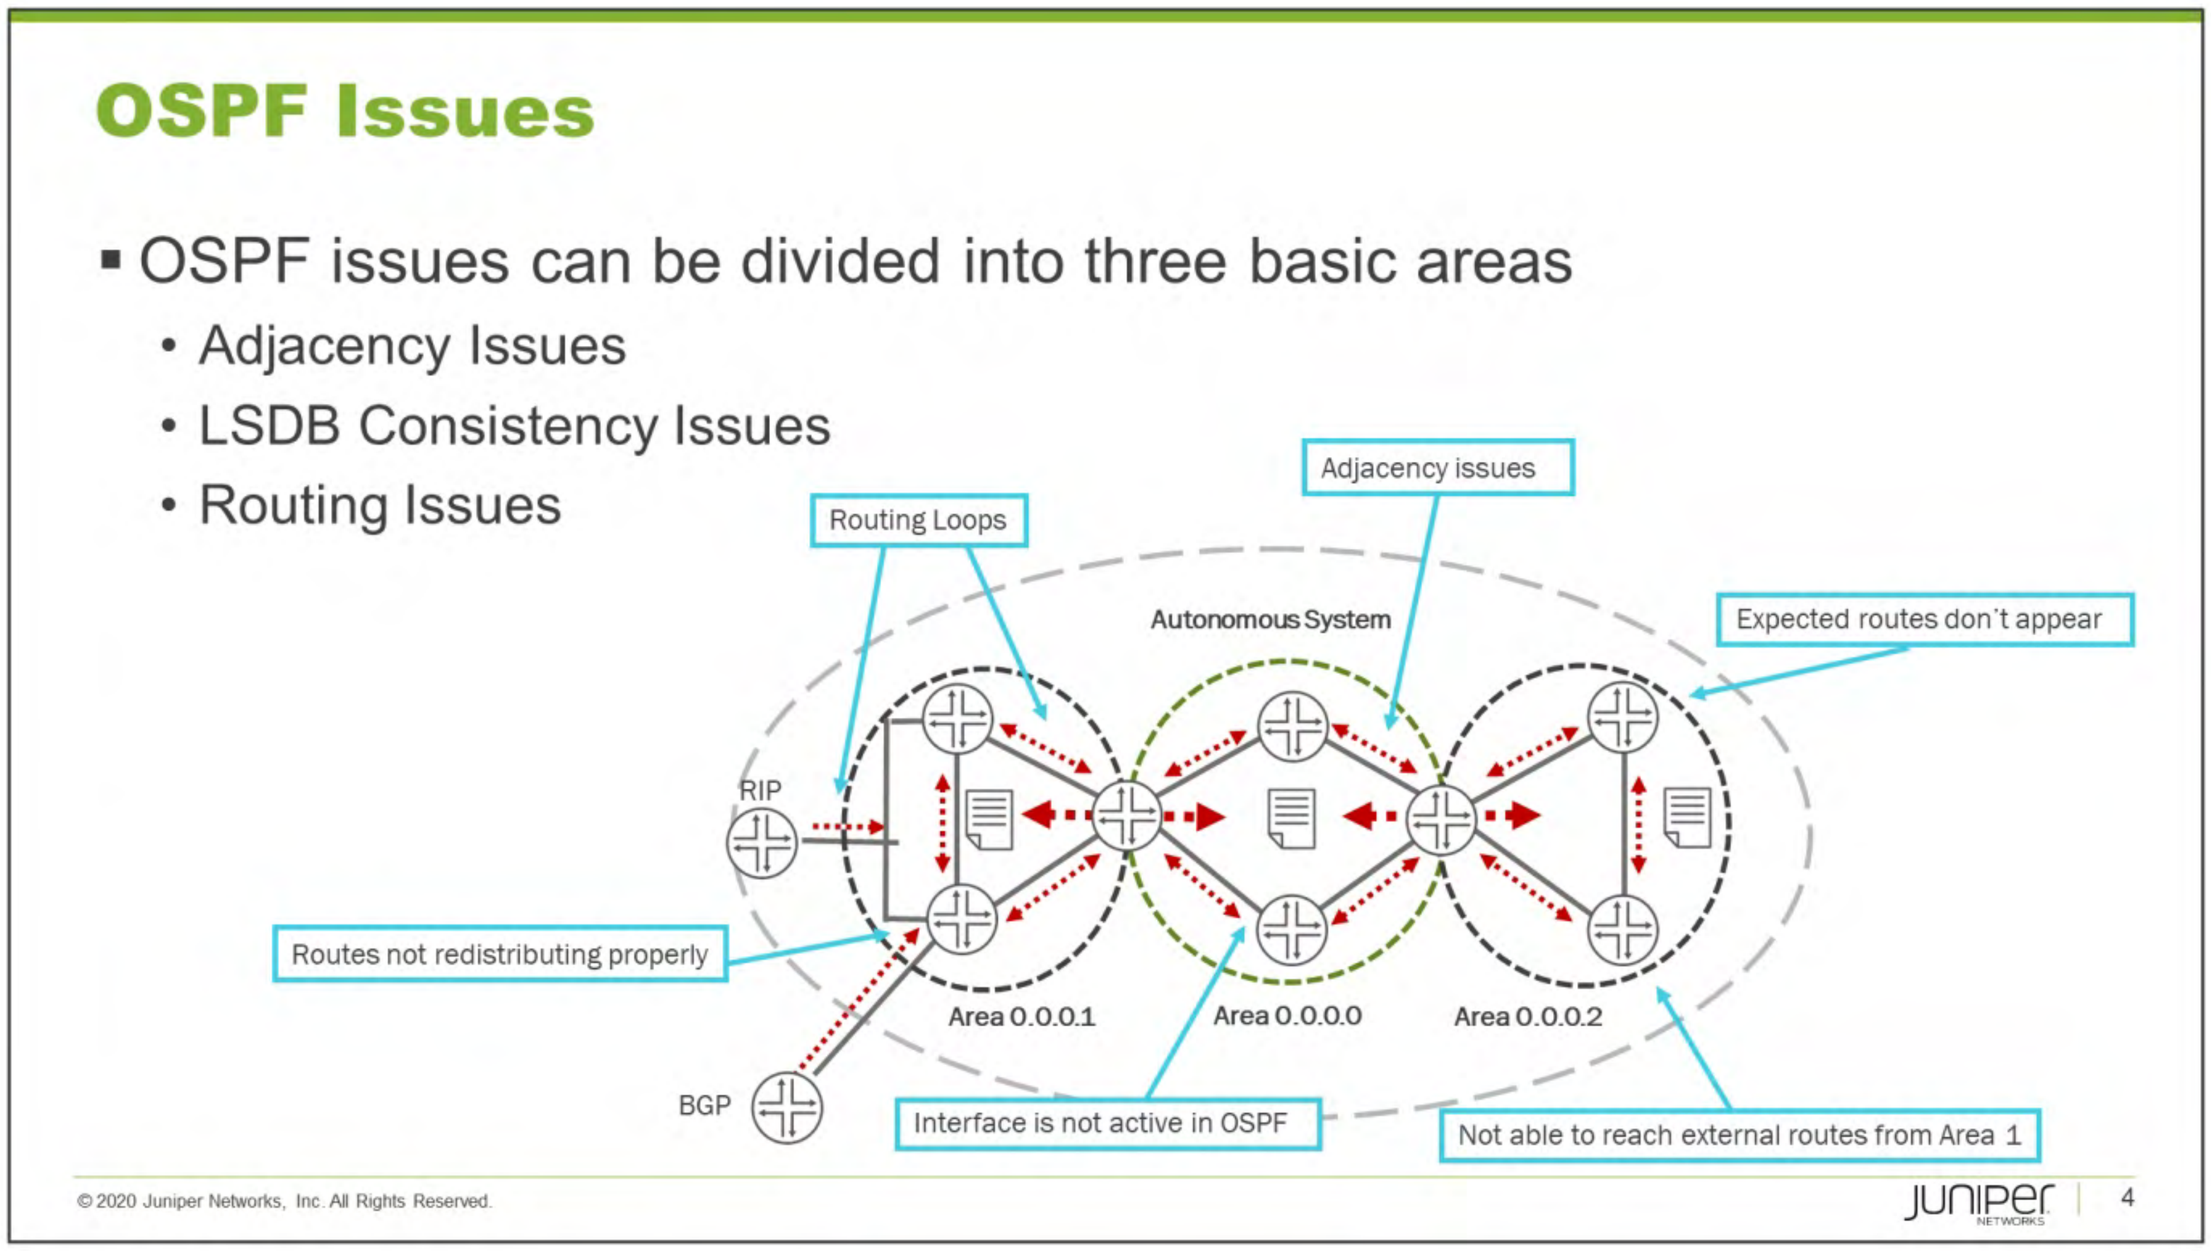

OSPF Issues

Within any OSPF network there are numerous issues that can occur. Being able to recognize the problems and provide a resolution is a critical skill for network administrators.

OSPF issues can be divided into three main types:

- Adjacency issues;
- LSDB consistency issues; and
- Routing issues.

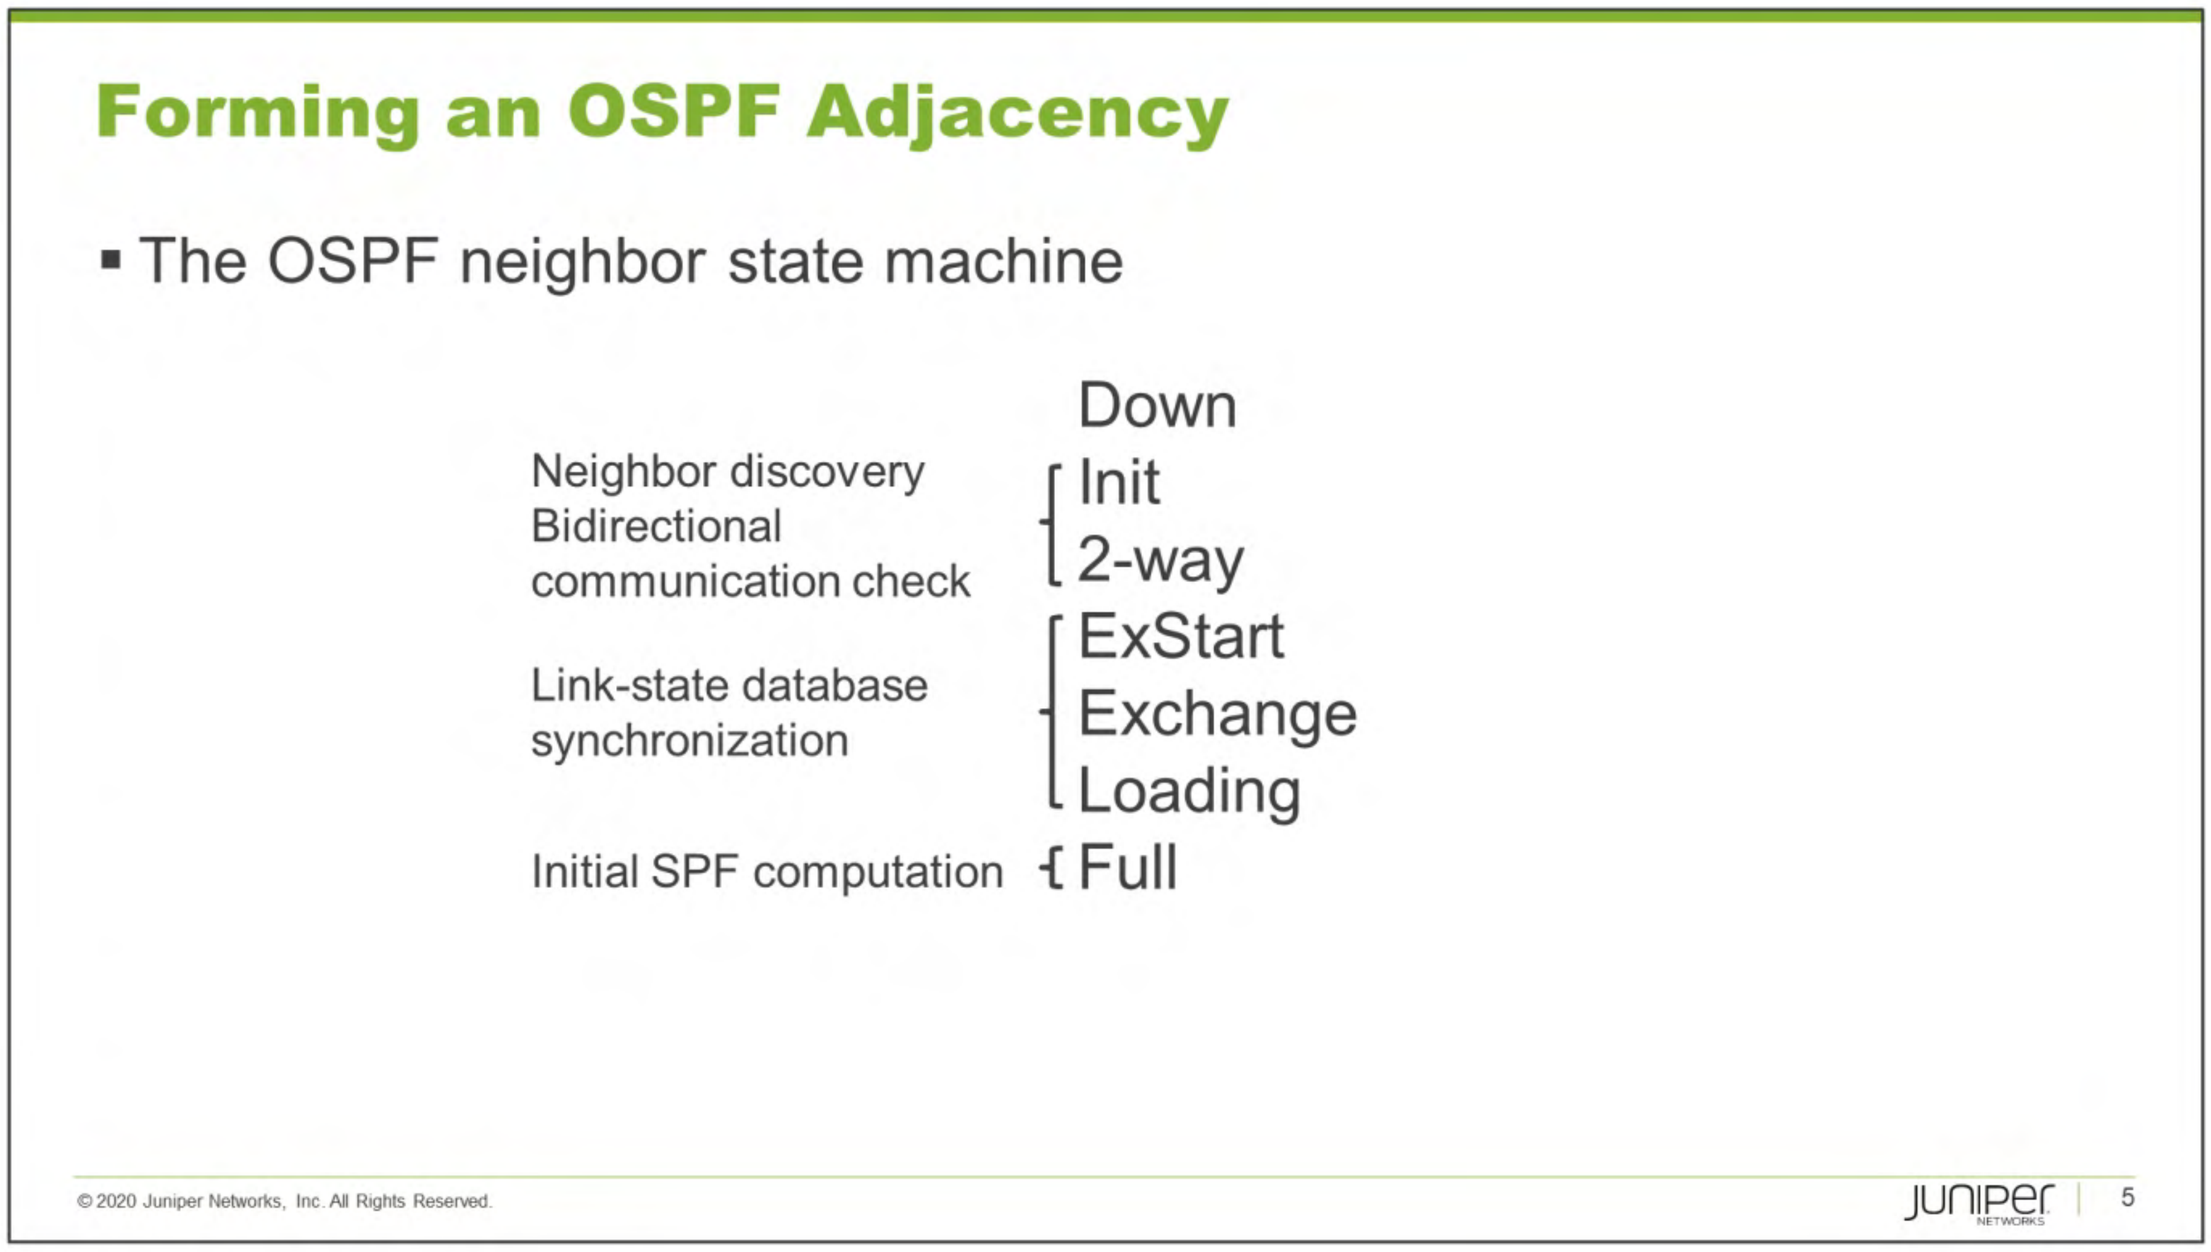

The OSPF Neighbor State Machine

Unlike protocols such as RIP, neighbors in OSPF are not stateless. For route computation, the protocol relies on a link-state database that must be maintained and synchronized between neighbors.

Each OSPF neighbor can be in any of these states:

- *Down*: Self-explanatory; no routes are exchanged;
- *Init*: The router is trying to detect neighbors behind an interface by sending OSPF hello packets. Those packets contain the list of neighbors seen on a given interface, and this information is use to confirm bidirectional connectivity;
- *Two-way*: The router detected hello packets from a neighbor, and can find himself in his neighbor's hello packet 'neighbors seen' list, which confirms that communication is bidirectional;
- *ExStart, Exchange, Loading*: In these states, the link-state database is being synchronized; and
- *Full:* The link-state database synchronization is completed, and the adjacency is maintained using hello packets.

Moving from Down to Init State (Interface issues)

```
[edit protocols]
lab@R1# show
ospf {
    area 0.0.0.0 {
        interface ge-0/0/1.0;
        interface lo0.0;
```

```
[edit protocols]
lab@R2# show
ospf {
    area 0.0.0.0 {
        interface ge-0/0/1.0;
        interface lo0.0;
```

```
lab@R2> show ospf interface
Interface          State   Area        DR ID        BDR ID
Nbrs
lo0.0              DR      0.0.0.0     192.168.2.1  0.0.0.0
```

Possible Interface Issues
1. Wrong interface configured.
2. Interface is down.
3. No IPv4 address configured on the interface.

Why is the ge-0/0/1 interface not included in the `show ospf interface` command?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Interface Issues.

When configuring OSPF, interfaces are configured to be part of a specific area. Once configured those interfaces should be displayed when the **show ospf interface** command is run. If, for some reason, it does not appear; What could be the reason?

There are three main reasons why this would happen:

1. The wrong interface was configured. OSPF uses the logical interface value. If the physical interface is entered with no logical interface designation, unit 0 is assumed. If a unit number is entered and the unit number used in OSPF doesn't match the unit configured under the `[interfaces]` hierarchy, then the interface will not be recognized by OSPF.
2. The interface is physically down or administratively set to be down
3. No IPv4 address is configured on the interface.

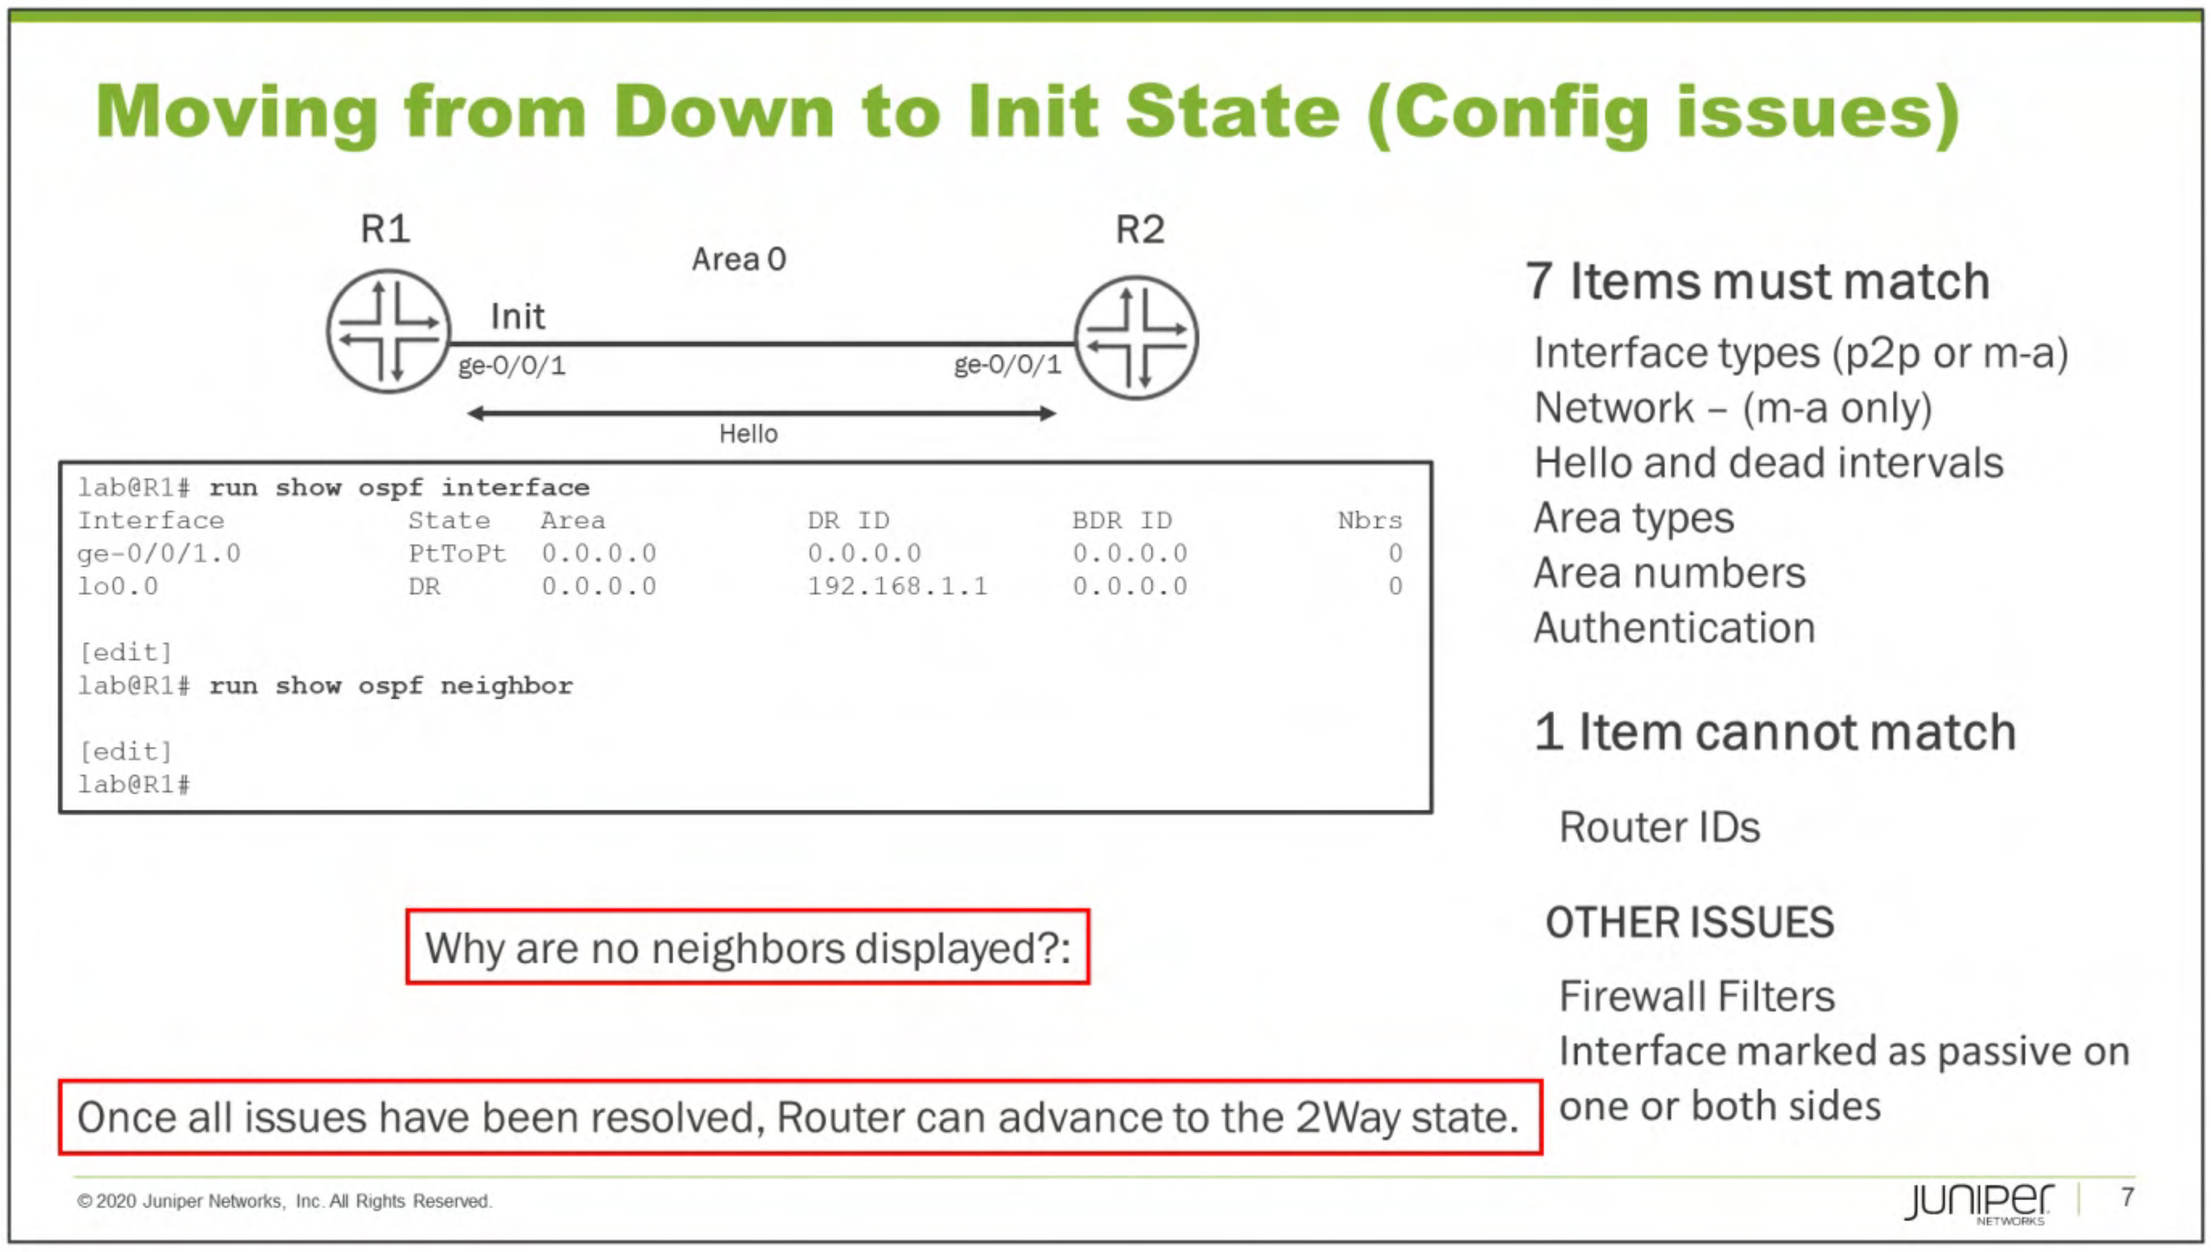

Configuration Issues

Once an interface is recognized by OSPF, then the next assumption is that it should appear when you execute the **show ospf neighbor** command. If the command does not return the expected neighbor, there are a large number of reasons that could cause the situation.

First there are 7 possible OSPF configuration options that need to match for OSPF neighbors to agree that they are able to form an adjacency and to move forward in the neighbor state machine process. These 7 possible items are:

1. Interface types must match. If one side is p2p, the other side must be p2p. If one side is multi-access the other side must be multi-access
2. Network mask must match if both sides are configured for multi-access
3. Hello interval values must match (default is 10)
4. Deat interval value must match (default is 40)
5. Area types must match. If one side is configured to be a Stub network, then the neighbor must be configured as a Stub. If one side is configured as an NSSA area, then the neighbor must also be configured to be an NSSA area.
6. Area numbers must match. An Area 0 link can only peer to a link that is part of Area 0.
7. Authentication type and keys must match.

In addition there is one item that cannot match. This is the Router-ID. The Router-ID is used to establish uniqueness in the OSPF network. If two connected routers have the same Router-ID, OSPF will not form an adjacency.

There are a couple of other issues that could come into play that would prevent two OSPF routers from advancing within the neighbor state machine process. They are as follows:

1. Firewall filters blocking OSPF
2. One or both of the interfaces trying to form the adjacency is marked as passive in OSPF. An interface in OSPF that is marked as passive will not send a hello message. Without a hello message a router cannot progress beyond the init state.

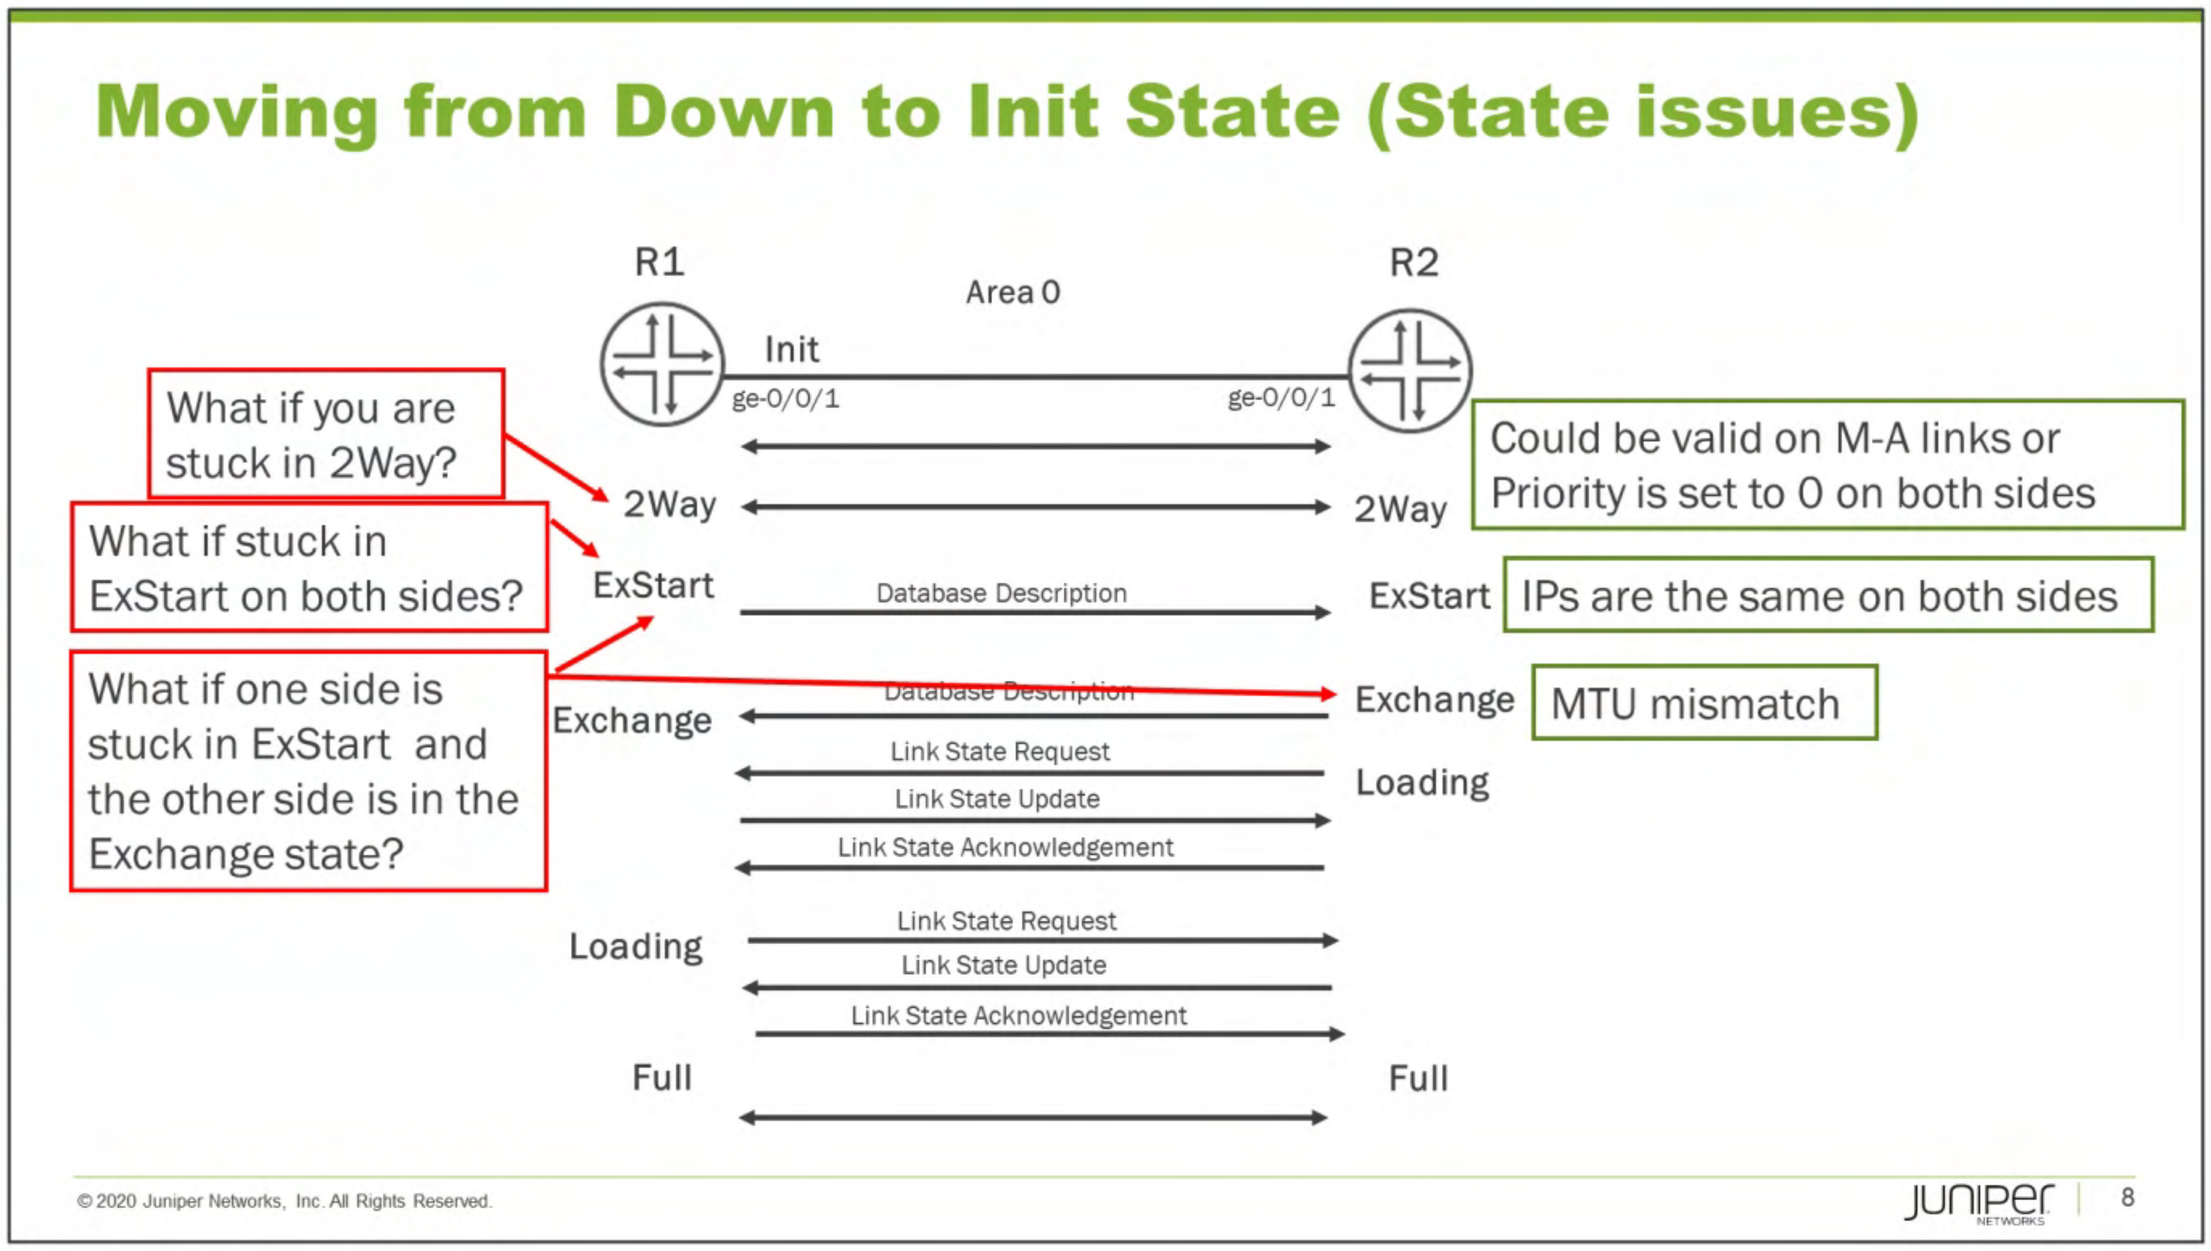

State Issues

Once both routers have exchanged hello messages they can advance to the 2Way state.

In a multi-access network with four or more routers sharing a network link, the 2Way state could be a valid end state. This is because `Full` adjacencies are only formed with the DR and BDR and not with other routers on the link. But if there are only two routers connected to a multi-access network segment, 2Way is not an expected state. What if you find the OSPF state stuck in 2Way in this scenario? Check to see if both routers have their priority set to 0. If they are both have a priority of zero, then no DR or BDR can be elected and 2Way will be the final state. To correct this issue, change the priority value of one or both sides or set both interfaces interface-type to p2p.

If you find that both sides are stuck in ExStart, the issue is that both interfaces have the same IP address. Before Database Description messages can be exchanged a negotiation takes place, based on the IP addresses of the interfaces. Since the addresses are identical, the negotiation cannot move forward, so both sides get stuck in ExStart.

Lastly if one side is stuck in ExStart and the other side is in the Exchange state, the issue is normally MTU size. The smaller MTU will be found on the side stuck in ExStart.

As you can see there are a large number of issues which can impact the ability for OSPF to achieve the `Full` state and start to forward traffic. The better you know the issues surrounding the forming of adjacencies the faster you can get your network up and running.

If you haven't memorized all of these options, there are many Junos commands that can be run that will help you identify adjacency issues.

OSPF Adjacency: Useful Commands

- **show ospf neighbor**

```
user@router> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
111.127.0.55      ge-0/0/1.0             Full      111.127.255.1    128    28
111.127.0.16      ge-0/0/2.0             2Way      111.127.255.5      0    22
111.127.0.42      ge-0/0/3.0           ExStart     111.127.255.8    128    26
```

- Used with the **detail** option the command shows the adjacency age, and also DR and BDR (on a LAN interface) .
- Remember that on a LAN, the state between two neighbors that are not DR or BDR is `2way`, not `Full`

© 2020 Juniper Networks, Inc. All Rights Reserved.

Useful Commands: show ospf neighbor

The **show ospf neighbor** command provides you with the state information of the neighbors. This command is useful for a summary check of OSPF status. If you include the **detail** option you can see the adjacency age, and also (on a LAN) DR and BDR.

Remember that on a LAN segment, the state between two neighbors that are not the DR or BDRis 2way, not Full. However, in the example above there is only one neighbor associated with the ge-0/0/2 interface. Seeing that the priority is 0 would give you a clue as to the potential problem.

For the ge-0/0/3 interface, based on the state the problem could be either MTU or IP address. It would be necessary to view the results of the connected routers **show ospf neighbor** command to determine the proper corrective action.

Remember both 2Way and ExStart are valid transition states. It may be that you ran the command and captured the results while the router was transiting that neighbor state. The command should be run a couple of times to ensure that the OSPF state is stuck at that state and not just passing through that state.

OSPF Adjacency: Useful Commands

- **show ospf statistics**
  - Only displays error counters

```
user@router> show ospf statistics | find errors
Receive errors:
  803 area mismatches
  1 netmask mismatches
  29 stub area mismatches
  15 nssa mismatches
  213 mtu mismatches
  2 Hellos received with our router ID
  9 Hellos received on point-to-point LAN with DR/BDR elected
```

  - Check for increasing counters
    - This will not tell you which interface errors are coming from, but will give you an idea about the problem
- Use the **clear ospf statistics** command to verify that the issue has been resolved

© 2020 Juniper Networks, Inc. All Rights Reserved.

Useful Commands: show ospf statistics

The **show ospf statistics** command shows a number of useful OSPF packet-related statistics. The most important information is at the end of the output, where you can find error counters aggregated per error type. Even if counters are aggregated, a quick check for increasing counters can immediately point troubleshooting in the right direction.

- *Area mismatch*: Local and remote routers do not agree on area number.
- *Netmask mismatches*: A hello packet received on an interface shows a different netmask compared to the one configured locally.
- *Stub and NSSA mismatch*: A check on the capability bits of a hello packet shows the remote and local routers do not agree on an area being NSSA or stub.
- *MTU mismatches*: Self-explanatory error, points to interface misconfiguration.
- *Hellos received with our router ID*: Symptom of a duplicate router-id.
- *Hellos received on point-to-point LAN with DR/BDR elected*: One side has configured the link as point-to-point, the other as LAN.

The error counters are aggregated, not separated on a per-interface basis. To find out the actual neighbor that is causing the problem you will need to enable traceoptions.

Since the counters could be displaying errors that have been previoiusly corrected, use the **clear ospf statistics** command before running the **show ospf statistics** command. In this way increasing numbers will highlight current issues.

OSPF Adjacency: Traceoptions (1 of 4)

- If error counters are increasing, you might need to enable traceoptions to find their cause
  - The OSPF statistic will only give you error counters

```
[edit protocols ospf]
user@router# show
traceoptions {
    file ospf.log size 10m files 3;
    flag error detail;
    flag hello detail;
}
```

    - Hellos are sent by default every 10 seconds on both a point-to-point and LAN interfaces
    - Start by tracing only errors, add hello only if needed.
  - We will now look at some sample of log entries

© 2020 Juniper Networks, Inc. All Rights Reserved.

Using OSPF Traceoptions

The OSPF statistic will only give you aggregated error counters. It won't tell you which neighbor is causing the problem, and often not even what the problem really is.

To find out, you might need to enable traceoptions; a good approach is to use **`flag error`** at first, and—especially in case of sporadic neighbor-down event – add **`flag hello`** when needed.

The latter is extremely useful in case of sporadic peer-down events due to a neighbor timing out; having **`flag hello`** enabled on both sides can show you if the loss is unidirectional or bidirectional, so you will be able to correlate it with transmission or interface errors.

OSPF Adjacency: Traceoptions (2 of 4)

- Protocol configuration errors

```
Feb 28 15:17:02.216071 OSPF packet ignored: area mismatch (0.0.0.1)
from 172.22.138.1 on intf ge-0/0/1.0 area 0.0.0.0
Feb 28 17:05:46.243212 OSPF packet ignored: area stubness mismatch
from 172.22.138.1 on intf ge-0/0/1.0 area 0.0.0.1
Feb 28 17:09:24.716352 OSPF packet ignored: area nssaness mismatch
from 172.22.138.1 on intf ge-0/0/1.0 area 0.0.0.1
Feb 28 17:11:30.583164 OSPF packet ignored: configuration mismatch
from 172.22.138.1 on intf ge-0/0/1.0 area 0.0.0.1
```

show log *ospf.log* | match mismatch

Not all entries are easy to interpret.

For example: The last log entry is an interface-type mismatch.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Protocol Configuration Errors

The slide highlights some examples of log entries due to protocol configuration errors. Respectively:

- Wrong area;
- Wrong area type (stub vs. non-stub);
- Wrong area type (nssa vs. non-nssa); and
- Wrong interface type (point-to-point versus LAN).

Other possible messages are related to authentication failure; they are rather self-explanatory, and the related problem is easy to correct.

OSPF Adjacency : Traceoptions (3 of 4)

- Interface configuration problems

```
Feb 28 15:07:45.302191 OSPF packet ignored: MTU mismatch from
172.22.138.1 on intf ge-1/1/4.0 area 0.0.0.0
```

```
Feb 28 17:30:09.113675 OSPF packet ignored: netmask 255.255.255.252
mismatch from 172.22.138.1 on intf ge-1/1/4.0 area 0.0.0.1
```

- In case of interface problems, error messages are clear and easy to interpret
- If you do not see anything in the logs, use **`traceoptions`** with **`flag hello send receive`**
  - Verify that you are actually receiving hello packets from your neighbor (to rule out communication issues)

© 2020 Juniper Networks, Inc. All Rights Reserved.

Interface Configuration Problems

This slide highlights some other possible problems that stem from configuration mistakes on the interface, such as:

- MTU mismatch between neighbors; and
- Different network mask (i.e. prefix length).

In either case, the error messages should be easy to interpret.

OSPF Adjacency : Traceoptions (4 of 4)

- OSPF traceoptions file naming
  - If you name the OSPF traceoptions file as *ospf*, take care not to confuse these two commands:
    - **show log ospf**: Displays the file named ospf from the /var/log directory
    - **show ospf log**: Is an internal command that displays some timing statistics about SPF runs; typically, it is not a very useful command

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Trace File Naming

It is a rather natural choice to name the OSPF traceoptions file simply *ospf*; if you do so, take care not to confuse these two commands:

- **show log ospf** displays the file called 'ospf' from the /var/log directory; and
- **show ospf log** is a completely different command which displays some timing statistics about SPF runs.

Typically, **show ospf log** is not very useful – it is more of a curiosity then a useful troubleshooting tool. Sample output follows:

```
user@router> show ospf log
Topology default SPF log:

   Last instance of each event type
When            Type            Elapsed
00:02:10        SPF             0.000863
00:02:10        Stub            0.000200
00:02:10        Interarea       0.000210
00:02:10        External        0.000051
00:02:10        NSSA            0.000040
```

OSPF Adjacency : Monitor Traffic Interface

- Fast and easy way to view OSPF adjacency messages.

```
lab@R2# run monitor traffic interface ge-0/0/1 detail no-resolve
…
18:18:48.441562 In IP (tos 0xc0, ttl   1, id 37619, offset 0, flags [none], proto: OSPF (89), length: 76)
172.20.66.1 > 224.0.0.5: OSPFv2, Hello, length 56 [len 44]
        Router-ID 192.168.1.1, Backbone Area, Authentication Type: none (0)
        Options [External, LLS]
         Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.0, Priority 128
…
18:18:55.792647 Out IP (tos 0xc0, ttl   1, id 55775, offset 0, flags [none], proto: OSPF (89), length: 80)
172.20.66.2 > 224.0.0.5: OSPFv2, Hello, length 60 [len 48]
        Router-ID 192.168.2.1, Backbone Area, Authentication Type: none (0)
        Options [External, LLS]
         Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.0, Priority 128
         Designated Router 172.20.66.2
         Neighbor List:
           192.168.1.1
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Monitor Traffic Interface.

You might also want to trace the protocol messages on a router interface by using the **monitor traffic interface *interface-name*** command. This command is essentially a front end to a well known tcpdump utility. You can use various filtering options in this command, for example, to capture only OSPF messages you can use the following filter:

```
user@srx> monitor traffic interface ge-0/0/1.0 matching "dst 224.0.0.5"
```

The 224.0.0.5 is the multicast address used by OSPF to send messages to all other OSPF routers.

Agenda: OSPF Troubleshooting

- Troubleshooting OSPF Adjacency Issues
- →Troubleshooting LSDB Consistency Issues
- Troubleshooting OSPF Routing issues
- Case Study – Adjacency Issues

© 2020 Juniper Networks, Inc. All Rights Reserved.

Troubleshooting LSDB Consistency Issues

The slide highlights the topic we discuss next.

OSPF LSDB Integrity issues

- Duplicated Router ID's
  - The router-id is part of the OSPF packet header so the problem is easily detected between directly-connected neighbors

```
Mar  5 10:43:20.192787 OSPF packet ignored: our router ID received
from 172.22.138.2 on intf ge-1/0/4.0 area 0.0.0.0
```

  - A more complex scenario is when the two routers with duplicated ID are not directly connected
    - The results are usually disastrous for the whole OSPF domain
  - Duplicate IDs can happen with anycast applications
    - Setting the router-id under routing-options is always a good idea
- Broken area 0.0.0.0 or some areas not connected to area 0.0.0.0
  - Can be resolved using virtual links
  - Virtual links are considered temporary solutions

© 2020 Juniper Networks, Inc. All Rights Reserved.

LSDB Integrity

One OSPF network problem that is often hard to troubleshoot is a duplicated Router ID. Two neighboring routers will not establish adjacency if they have an identical Router ID, however if the two routers with the same Router ID are not physically connected it presents a more difficult problem. You can only detect the failure by monitoring the LSDB and SPF and discovering regular routing flaps when some routes are continuously recalculated because the routers with the same Router ID advertise different pieces of routing information. A good indicator of this issue would be continuous SPF runs.

Broken OSPF backbone area or an area that is not connected to the backbone present another issue. OSPF implements an inter-area loop protection mechanism by ensuring that ABRs do not generate summary LSAs to the backbone area for the summaries they receive in a non-backbone area. The missing summary routes would break the inter-area communications in this scenario. To ensure that OSPF backbone area is contiguous and other areas are connected to it you should configure a virtual link. Virtual links were designed to solve a temporary issue. It is not meant to be a permanent solution. A direct connection should replace the virtual link as soon as possible.

Agenda: OSPF Troubleshooting

- Troubleshooting OSPF Adjacency Issues
- Troubleshooting LSDB Consistency Issues
- →Troubleshooting OSPF Routing issues
- Case Study – Adjacency Issues

© 2020 Juniper Networks, Inc. All Rights Reserved.

Troubleshooting OSPF Routing Issues

The slide highlights the topic we discuss next.

OSPF Routing Issues (1 of 2)

- OSPF routing issues
  - A broader class of problems where the protocol does not seem to compute routes according to the expected behavior
    - Missing routes or traffic blackholes
    - Unexpected routing decisions, suboptimal routing
    - Domain instability
    - Others
  - Routing issues are often the result of configuration or network design errors
    - Some design weaknesses go unnoticed until an outage or a network change exposes them

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Routing Issues

Under this general definition we can group a broader class of problems where the protocol does not seem to compute routes according to the expected behavior.

Some symptoms could be missing routes, routing loops, suboptimal or unexpected routing decisions and even domain-wide IGP instability.

Given the maturity of the OSPF protocol, most of these issues are not caused by software bugs, but by mistakes in design or configuration. In this regard, care must be taken when investigating routing issues; often a design weakness goes unnoticed until an outage or a configuration change triggers it.

The typical example is an incorrectly designed, sparsely connected area 0. Everything will appear to work until a circuit failure causes the two parts of the area to become disconnected, causing a split area 0 and an outage even in peripheral areas.

Investigating OSPF Routing Issues (2 of 2)

- Troubleshooting OSPF routing issues
  - Troubleshooting OSPF routing issues can be complex because the range of potential issues is very broad
    - Generally, there are no universal troubleshooting recipes
    - Too many possible causes make an exhaustive approach impossible
  - Still, it is possible to give the outline of a plan which can be used to troubleshoot most of the OSPF-specific routing issues
    - To do this, we must become familiar with some additional OSPF-specific commands

© 2020 Juniper Networks, Inc. All Rights Reserved.

Troubleshooting OSPF Routing Issues

Troubleshooting OSPF routing issues can be complex because the range of potential issues is very broad. As with any complex routing protocol, it is difficult to give a universal recipe that allows you to get to the root cause of any OSPF-specific routes; the variety of configuration options, of possible symptoms and causes and the dependency on specific topologies would make an exhaustive step-by-step plan too complex.

Still, it is possible to give the outline of a plan which can be used to troubleshoot most of the OSPF-specific routing issues. To do this, we will need to become familiar with some additional OSPF-specific commands and traceoptions.

OSPF Routing: Useful Commands (1 of 5)

- Some useful commands:
  - **`show route protocol ospf`**
    - Display OSPF-computed routes and their attributes
  - **`show ospf route`**
    - Will tell you the type (intra-area, inter-area, external type-1 and type-2, etc.) of each of the prefixes computed by OSPF
  - **`show ospf database`**
    - Allows you to check the content of the link-state database

© 2020 Juniper Networks, Inc. All Rights Reserved.

Useful OSPF Commands

The slide highlights useful commands that can help you troubleshoot OSPF routing issues. The order in which they are listed reflects the order in which they are typically used.

Summarization Issues

- Summarize the loopbacks of Area 1 so they appear as a single route on R1

```
lab@R2# show protocols ospf
area 0.0.0.0 {
interface ae1.0;
    interface ae0.0;
}
area 0.0.0.1 {
    area-range 192.168.0.0/22;
    interface ge-0/0/6.0;
    interface lo0.0;
}
```

```
lab@R3# show protocols ospf
area 0.0.0.0 {
interface ae2.0;
    interface ae0.0;
}
area 0.0.0.1 {
    area-range 192.168.0.0/22;
    interface ge-0/0/6.0;
    interface lo0.0;
```

- The task has not been achieved with this configuration. Why?

The area-range of 192.168.0.0/22 includes the Area 1 loopbacks 192.168.1.1 thru 192.168.3.1.but does not include the 192.168.4.1. This will cause the 192.168.4.1 summary LSA to be created on the ABR and passed to the R1 router.

Router	loopback	Area
R1	192.168.10.1	0
R2	192.168.1.1	1
R3	192.168.2.1	1
R4	192.168.3.1	1
R5	192.168.4.1	1

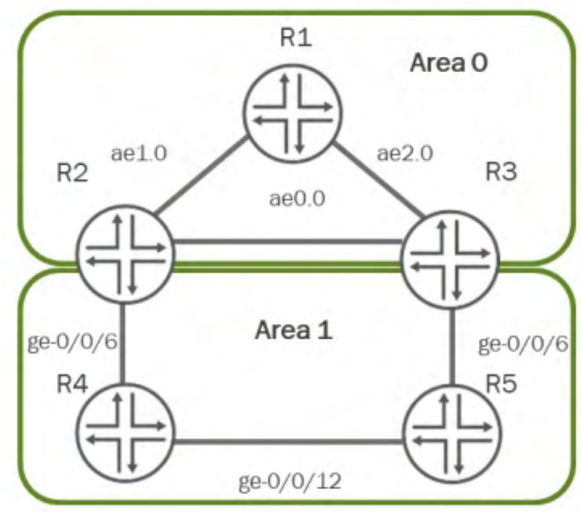

© 2020 Juniper Networks, Inc. All Rights Reserved.

Summarization Issues

Since configuration is the main reason for routing issues in OSPF, being able to view configurations and understand potential issues is important.

Reviewing the diagram and configuration snippets, the goal is to summarize the loopbacks of Area 1 and have them appear as a single summary route on R1. This task is not successful. Why?

First the area-range command does not include all of the Area 1 loopbacks. The summary includes those routes from 192.168.0.0 and 192.168.3.255. The 192.168.4.1 is not included. Changing the summary to a /21 would fix the problem, but could create other problems such as black holing addresses advertised by the summary but not owned by Area 1. such as 192.168.0.1 (If it were an active address on another device.).

Command Issues

- Summarize the loopbacks of area 1 so they appear as a single route on R1

```
lab@R2# show protocols ospf
area 0.0.0.0 {
interface ae1.0;
    interface ae0.0;
}
area 0.0.0.1 {
     nssa {
        area-range 192.168.4.0/22;
    };
    interface ge-0/0/6.0;
    interface lo0.0;
}
```

```
lab@R3# show protocols ospf
area 0.0.0.0 {
interface ae2.0;
    interface ae0.0;
}
area 0.0.0.1 {
     nssa {
        area-range 192.168.4.0/22;
     }
     interface ge-0/0/6.0;
     interface lo0.0;
```

Router	loopback	Area
R1	192.168.1.1	0
R2	192.168.4.1	1
R3	192.168.5.1	1
R4	192.168.6.1	1
R5	192.168.7.1	1

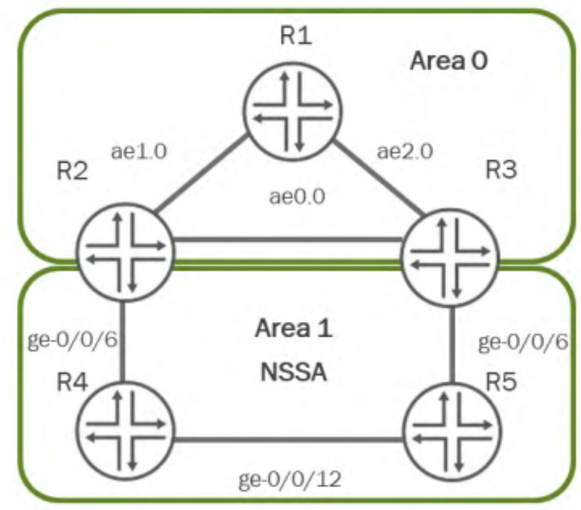

- All loopbacks show up on R1 as individual T3 LSAs. Why?

Using the area-range command after the nssa stanza will only summarize T7 LSAs. Loopbacks are T1 LSAs.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Command Placement Issues

In this situation, the goal is the same as before, advertise all the loopbacks from Area 1 as a single T3 LSA. This configuration results in all loopbacks appearing in the R1 database as T3 LSAs. Again the question is why.

This is another configuration issue. When the area-range command is configured after the nssa stanza, OSPF looks for T7 LSAs that fall within the area-range prefix. Since the loopbacks are T1 LSAs, they will not be summarized at all.

The area-range command must be configured to follow the area 1 stanza and not the nssa stanza for OSPF to summarize the loopbacks.

Functionality Issues

- R1 should be able to reach all loopbacks in Area 1 except for the loopback on R5

```
lab@R2# show protocols ospf
area 0.0.0.0 {
interface ae1.0;
    interface ae0.0;
}
area 0.0.0.1 {
    area-range 192.168.4.0/22;
    area-range 192.168.7.1/32 restrict;
    interface ge-0/0/6.0;
    interface lo0.0;
}
```

```
lab@R3# show protocols ospf
area 0.0.0.0 {
interface ae2.0;
    interface ae0.0;
}
area 0.0.0.1 {
    area-range 192.168.4.0/22;
    area-range 192.168.7.1/32 restrict;
    interface ge-0/0/6.0;
    interface lo0.0;
```

Router	loopback	Area
R1	192.168.1.1	0
R2	192.168.4.1	1
R3	192.168.5.1	1
R4	192.168.6.1	1
R5	192.168.7.1	1

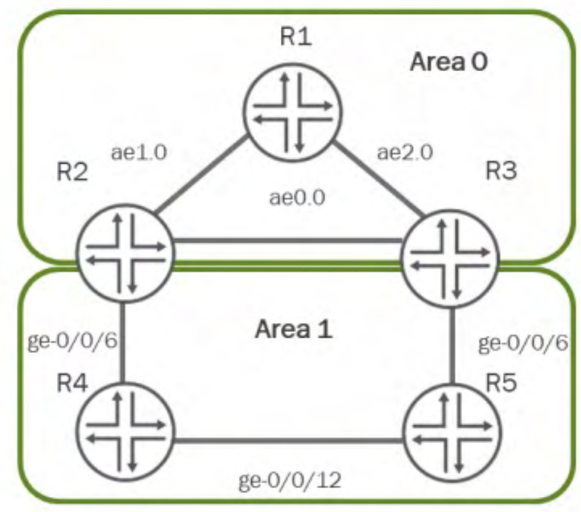

R1 can still reach R5. Why?

The restrict command does not override the summary area-range command.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Command Operation Issues

To address this issue, you will need to ensure that R1 can reach all loopbacks in Area 1 except for the loopback associated with R5. This solution created two area-range commands. One area-range for the block of loopbacks and one area-range to restrict the single loopback. The area-range command does not work this way and will ignore the **`area-range restrict`** command.

The way to resolve this issue is to remove the **`area-range`** command and only use the **`area-range restrict`** option.

Agenda: OSPF Troubleshooting

- Troubleshooting OSPF Adjacency Issues
- Troubleshooting LSDB Consistency Issues
- Troubleshooting OSPF Routing Issues

→Case Study – Adjacency Issues

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study - Adjacency Issues

The slide highlights the topic we discuss next.

OSPF Case Study: Adjacency Issues

- Topology, symptoms, and relevant information:
  - You recently made alterations to your existing OSPF network. You added a new router (R5) to Area 1. While you are verifying your network, you notice problems. You find that all adjacencies on R2 are no longer working and R5 cannot form adjacencies with neighbors. Physical connections are verified and working fine. These devices are not using security policies, so no traffic would be impacted.

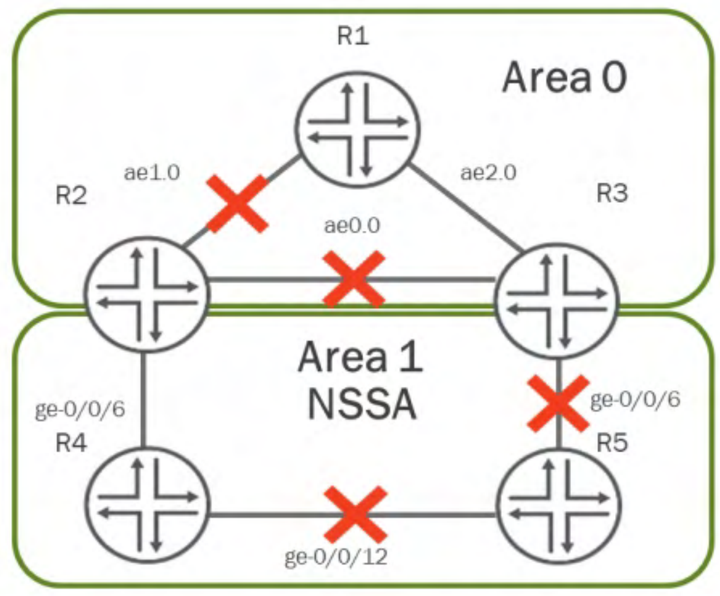

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Case Study: Adjacency Issues

The slide presents a network topology, reported symptoms and description of a case related to OSPF adjacency issues. In this case study you added a new router to Area 1. The verification showed that the new router does not establish adjacencies. While you are troubleshooting the issue you find that R2 is no longer reachable. The physical connections have been verified and showed no issues. You have ensured that the routers do not have any security policies or firewall filters blocking OSPF.

Troubleshooting R5 Neighborship (1 of 3)

- Check R5 neighborship state
  - The ExStart state indicates the most probable MTU issue

```
user@R5> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.222.0.21      ge-0/0/12.0            ExStart   10.222.1.4       128    32
10.222.0.13      ge-0/0/6.0             ExStart   10.222.1.3       128    37
```

- Compare OSPF interface parameters

```
user@R5> show ospf interface ge-0/0/12.0 detail
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/12.0        BDR     0.0.0.1         10.222.1.4      10.222.1.5         1
  Type: LAN, Address: 10.222.0.22, Mask: 255.255.255.252, MTU: 9178  Cost: 1
---(more)---
```

```
user@R4> show ospf interface ge-0/0/12.0 detail
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/12.0        DR      0.0.0.1         10.222.1.4      10.222.1.5         1
  Type: LAN, Address: 10.222.0.21, Mask: 255.255.255.252, MTU: 1500  Cost: 1
---(more)---
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Check R5 Neighborship State

You begin with checking the R5 neighbors. The **`show ospf neighbor`** command indicates that the two neighbors are stuck in exchange start state.

Verify OSPF Interface Parameters

Use the **`show ospf interface`** command to check the interface parameters—the **`detail`** option is useful. In this case, the command displays mismatched MTU values on the R4 and R5 interfaces.

Troubleshooting R5 Neighborship (2 of 3)

- Adjust the MTU settings and check the neighborship
  - ge-0/0/6 is still in ExStart

```
user@R5> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.222.0.21      ge-0/0/12.0            Full      10.222.1.4       128    34
10.222.0.13      ge-0/0/6.0             ExStart   10.222.1.3       128    31
```

- Check the interface parameters again

```
user@R5> show ospf interface ge-0/0/6.0 detail
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/6.0         BDR     0.0.0.1         10.222.1.3      10.222.1.5         1
  Type: LAN, Address: 10.222.0.13, Mask: 255.255.255.252, MTU: 1500  Cost: 1
```

```
user@R3> show ospf interface ge-0/0/6.0 detail
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/6.0         BDR     0.0.0.1         10.222.1.5      10.222.1.3         1
  Type: LAN, Address: 10.222.0.13, Mask: 255.255.255.252, MTU: 1500  Cost: 1
```

Do you see what's wrong?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Modify the MTU Settings

You modify the interface MTU settings to ensure that they match at all neighboring routers. Then you check the neighborship again. This time the command output shows that R4 and R5 became fully adjacent but R3 neighborship is still in an exchange start state.

Check the Interface Parameters

You check the interface parameters again. This time the MTU settings appear to be the same. A careful examination would reveal that the two neighbors are using the same IP address on their interfaces, but this problem is not always easily spotted.

Troubleshooting R5 Neighborship (3 of 3)

- Enable traceoptions

```
[edit protocols ospf]
user@R5# show
traceoptions {
    file ospf.log;
    flag packets;
    flag error detail;
}
```

- Duplicate IP addresses on ge-0/0/6 interface!

```
user@R5> show log ospf.log
Jan 31 18:25:56 exA-2 clear-log[10835]: logfile cleared
Jan 31 18:25:56.870847 OSPF hello from 10.222.0.13 (IFL 2147404756, area 0.0.0.1) absorbed
Jan 31 18:25:58.152391 OSPF periodic xmit from 10.222.0.13 to 224.0.0.5 (IFL 2147405268
area 0.0.0.1)
Jan 31 18:26:00.979655 OSPF resend last DBD to 10.222.0.13
Jan 31 18:26:00.979832 OSPF sent DbD 10.222.0.13 -> 10.222.0.13 (ge-0/0/6.0 IFL 69 area
0.0.0.1)
Jan 31 18:26:00.979862   Version 2, length 32, ID 10.222.1.5, area 0.0.0.1
Jan 31 18:26:00.979887   options 0x50, i 1, m 1, ms 1, r 0, seq 0xadf4b41, mtu 1500
Jan 31 18:26:00.980634 OSPF packet ignored: no matching interface from 10.222.0.13, IFL 0
---(more)---
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Using Traceoptions to Debug OSPF

For the sake of completeness you enable the traceoptions to debug the protocol operations.

The output of the **`show log ospf.log`** command clearly indicates that OSPF database description packet was ignored with the reason that the router does not have a matching interface.

Troubleshooting R2 Neighborship (1 of 3)

- Check OSPF interface parameters

```
user@R2> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ae0.0               DR      0.0.0.0         10.222.1.2      0.0.0.0            0
  Type: LAN, Address: 10.222.0.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 10.222.0.1, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 2, Start time: 2013 Mar 13 23:35:00 UTC
---(more)---
```

```
user@R3> show ospf interface detail
Interface           State   Area            DR ID           BDR ID          Nbrs
ae0.0               DR      0.0.0.0         10.222.1.3      0.0.0.0            0
  Type: LAN, Address: 10.222.0.2, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 10.222.0.2, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 2, Start time: 2013 Mar 13 23:35:00 UTC
---(more)---
```

  - Interfaces show matching parameters

© 2020 Juniper Networks, Inc. All Rights Reserved.

Check R2 OSPF Interfaces

Proceed to troubleshooting the R2 adjacencies. You start with inspecting the router and its neighbors' interfaces. The command on the slide shows the perfectly matching parameters. At the same time, you spot that the interfaces have MD5 authentication configured.

Troubleshooting R2 Neighborship (2 of 3)

- Enable traceoptions

```
[edit protocols ospf]
user@R2# show
traceoptions {
    file ospf.log;
    flag hello detail;
    flag error detail;
}
```

- OSPF log file indicates authentication failure

```
user@R2> show log ospf.log
Mar 13 23:51:28 srxA-1 clear-log[9348]: logfile cleared
Mar 13 23:51:30.029065 OSPF periodic xmit from 10.222.0.9 to 224.0.0.5 (IFL 78 area
0.0.0.0)
Mar 13 23:51:31.000899 OSPF hello from 10.222.0.6 (IFL 74, area 0.0.0.1) absorbed
Mar 13 23:51:31.429258 OSPF periodic xmit from 10.222.0.5 to 224.0.0.5 (IFL 74 area
0.0.0.1)
Mar 13 23:51:35.489116 OSPF packet ignored: authentication failure (bad cksum).
Mar 13 23:51:35.489914 OSPF packet ignored: authentication failure from 10.222.0.2
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Debugging OSPF at R2

You enable traceoptions to monitor OSPF packet flow. The **`show log`** ***`ospf.log`*** command output on the slide indicates that the router ignores OSPF Hello packets due to authentication failure.

Troubleshooting R2 Neighborship (3 of 3)

- Collected data
  - The previously active adjacencies failed all at one time
  - Debugging reveals authentication issues
  - Interfaces show a new authentication key rollover time
  - Authentication issues began after the new key rollover
- Solution
  - Make sure that all authentication keys match
  - Check neighborship

```
user@R2> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.222.0.2       ae0.0                  Full      10.222.1.3       128    37
10.222.0.10      ae1.0                  Full      10.222.1.1       128    36
10.222.0.6       ge-0/0/6.0             Full      10.222.1.4       128    39
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Collected Data

At this time, you know that there is an authentication issue between R2 and its neighbors. The detailed output of the **`show ospf interface`** command provides a clue that the adjacencies were affected at nearly the same time.

Solution

Given the information collected you can conclude that R2 OSPF adjacencies failed due to misconfigured authentication parameters that were activated by accident at the time you were troubleshooting the R5 adjacencies.

After you adjusted the authentication at R2, the new check with the **`show ospf neighbor`** command shows that all R2 neighbors are now in full adjacency.

Summary

- In this content, we:
  - Listed useful commands that are used to troubleshoot and verify OSPF
  - Isolated different OSPF issues

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Discussed:

- Useful commands that can be used to troubleshoot and verify correct operation of OSPF; and
- Isolating different OSPF issues.

Review Questions

1. What issue should be investigated if two routers are trying to form an adjacency and one router is in Exstart and the other is in the Exchange state?
2. Which OSPF parameters must match to ensure adjacency formation?
3. What OSPF parameter cannot match between two connected routers in order for them to form an adjacency?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review Questions

1.

2.

3.

Lab: Troubleshooting OSPF

- Monitor and troubleshoot the operation of an OSPF network with multiple areas.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab: Troubleshooting OSPF

The slide lists the objective for this lab.

Answers to Review Questions

1.

The MTU values on each side should be investigated for a mismatch.

2.

The following items must match for an OSPF adjacency to form between two connected routers: OSPF area, netmask, Hello and Dead intervals, OSPF options, and interface MTU must match for neighbors to become adjacent.

3.

The Router-ID cannot match on the two OSPF routers that want to form an adjacency.

Education Services

Advanced Junos Service Provider Routing

Chapter 6: IS-IS

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - Explain the concepts and operation of IS-IS
  - Describe various IS-IS link-state PDU types
  - List IS-IS adjacency rules and troubleshoot common adjacency issues
  - Configure and monitor IS-IS

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Will Discuss:

- IS-IS concepts and operation;
- IS-IS link-state protocol data unit (PDU) types;
- IS-IS adjacency rules and troubleshooting; and
- The configuration and monitoring of IS-IS in the Junos operating system.

Agenda: IS-IS

→Overview of IS-IS
- IS-IS PDUs
- Neighbors and Adjacencies
- Configuring and Monitoring IS-IS

© 2020 Juniper Networks, Inc. All Rights Reserved.

Overview of IS-IS

The slide lists the topics we will discuss. We discuss the highlighted topic first.

Overview of IS-IS

- An interior gateway protocol based on the SPF algorithm
  - Uses link-state information to make routing decisions
- Developed for routing ISO CLNP packets
  - IP was added later
  - Defined in ISO/IEC 10589, RFC 1142, RFC 1195, and RFC 2763

© 2020 Juniper Networks, Inc. All Rights Reserved.

The IS-IS Protocol

The IS-IS protocol is an IGP that uses link-state information to make routing decisions. It also uses the shortest-path-first (SPF) algorithm, similar to OSPF.

The protocol was developed by Digital Equipment Corporation for its DECnet Phase V. The International Organization for Standardization (ISO) standardized IS-IS to be the routing protocol for the ISO's Connectionless Network Protocol (CLNP) and is described in ISO 10589. The ISO was working on IS-IS at the same time as the Internet Advisory Board (IAB) was working on OSPF. ISO proposed that IS-IS be adopted as the routing protocol for TCP/IP in place of OSPF. This proposal was driven by the opinion that TCP/IP was an interim protocol suite that would eventually be replaced by the Open Systems Interconnection (OSI) suite.

Integrated IS-IS

- Implementation of IS-IS for routing IP in addition to CLNS, also called dual IS-IS
- All routers run a single routing algorithm
- SRX Series and MX Series platforms support CLNP/CLNS routing
  - M Series, and T Series platforms do not support native CLNP/CLNS routing
- IS-IS routers exchange LSPs
  - Similar to OSPF LSAs/link-state update packets
  - IS-IS PDUs are used to transmit the routing information
    - Sometimes called packets to conform with IP terminology
  - IP reachability information is included in the updates

© 2020 Juniper Networks, Inc. All Rights Reserved.

CNLS

Connectionless Network Service (CLNS) is the service that CLNP provides to route messages to their destinations. This service is independent of any other messages and can transmit data before a circuit is established. M and T Series routers only support IP routing with IS-IS. They do not support CLNP or CLNS routing. Full IS-IS CLNP routing (including End System–to–Intermediate System [ES-IS] routing) is supported on SRX Series Gateways and MX Series Routers.

Dual IS-IS

The Integrated IS-IS protocol was proposed to support the transition from TCP/IP to OSI. Integrated IS-IS, also known as dual IS-IS, provides a single routing protocol capable of routing both CLNS and IP. Integrated IS-IS was developed to operate in a CLNS, IP, or CLNS/IP setting.

Single Algorithm

Like all integrated routing protocols, Integrated IS-IS requires that all routers run a single routing algorithm. LSPs sent by routers running Integrated IS-IS include all destinations running either IP or CLNP Network Layer protocols. Routers running Integrated IS-IS still must support protocols such as the Address Resolution Protocol (ARP), the Internet Control Message Protocol (ICMP) for IP, and the ES-IS protocol for CLNP.

LSPs

Standard IS-IS packets must be modified to support multiple Network Layer protocols. IS-IS packet formats were designed to support the addition of new fields without a loss of compatibility with nonintegrated versions of IS-IS.

IS-IS adds type/length/value (TLV) objects (discussed further in the following pages) to support integrated routing. These TLVs tell intermediate systems (ISs) which Network Layer protocols are supported by other ISs and whether end stations running other protocols can be reached. They also include any other required Network Layer, protocol-specific information.

IS-IS Concepts

- IS-IS network is a single autonomous system
  - End systems: Network entities (hosts) that send and receive packets
  - Intermediate systems: Network entities (routers) that send, receive, and forward packets
  - PDUs: Protocol data units; term for IS-IS packets
- A single AS can be divided into smaller groups called areas, which are organized hierarchically
  - Level 1 intermediate systems route within an area or toward a Level 2 system
    - Attached bit
  - Level 2 intermediate systems route between areas and toward other ASs

© 2020 Juniper Networks, Inc. All Rights Reserved.

Operation of IS-IS

An IS-IS network is a single AS, also called a routing domain, that consists of *end systems* (ESs) and *intermediate systems* (ISs). End systems are network entities that send and receive traffic. Intermediate systems are routers which can send, receive, and forward traffic.

IS-IS Areas

In IS-IS, a single AS can be divided into smaller groups called *areas*. Routing between areas is organized hierarchically, allowing a domain to be divided administratively into smaller areas. IS-IS accomplishes this organization by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area, and Level 2 intermediate systems route between areas and toward other ASs. A Level 1/Level 2 system routes within an area on one interface and between areas on another.

A Level 1/Level 2 system sets the attached bit in the Level 1 PDUs that it generates into a Level 1 area to indicate that it is a Level 2-attached backbone router and that it can be used to reach prefixes outside the Level 1 area. Level 1 routers create a default route for interarea prefixes, which points to the closest (in terms of metrics) Level 1/Level 2-attached router.

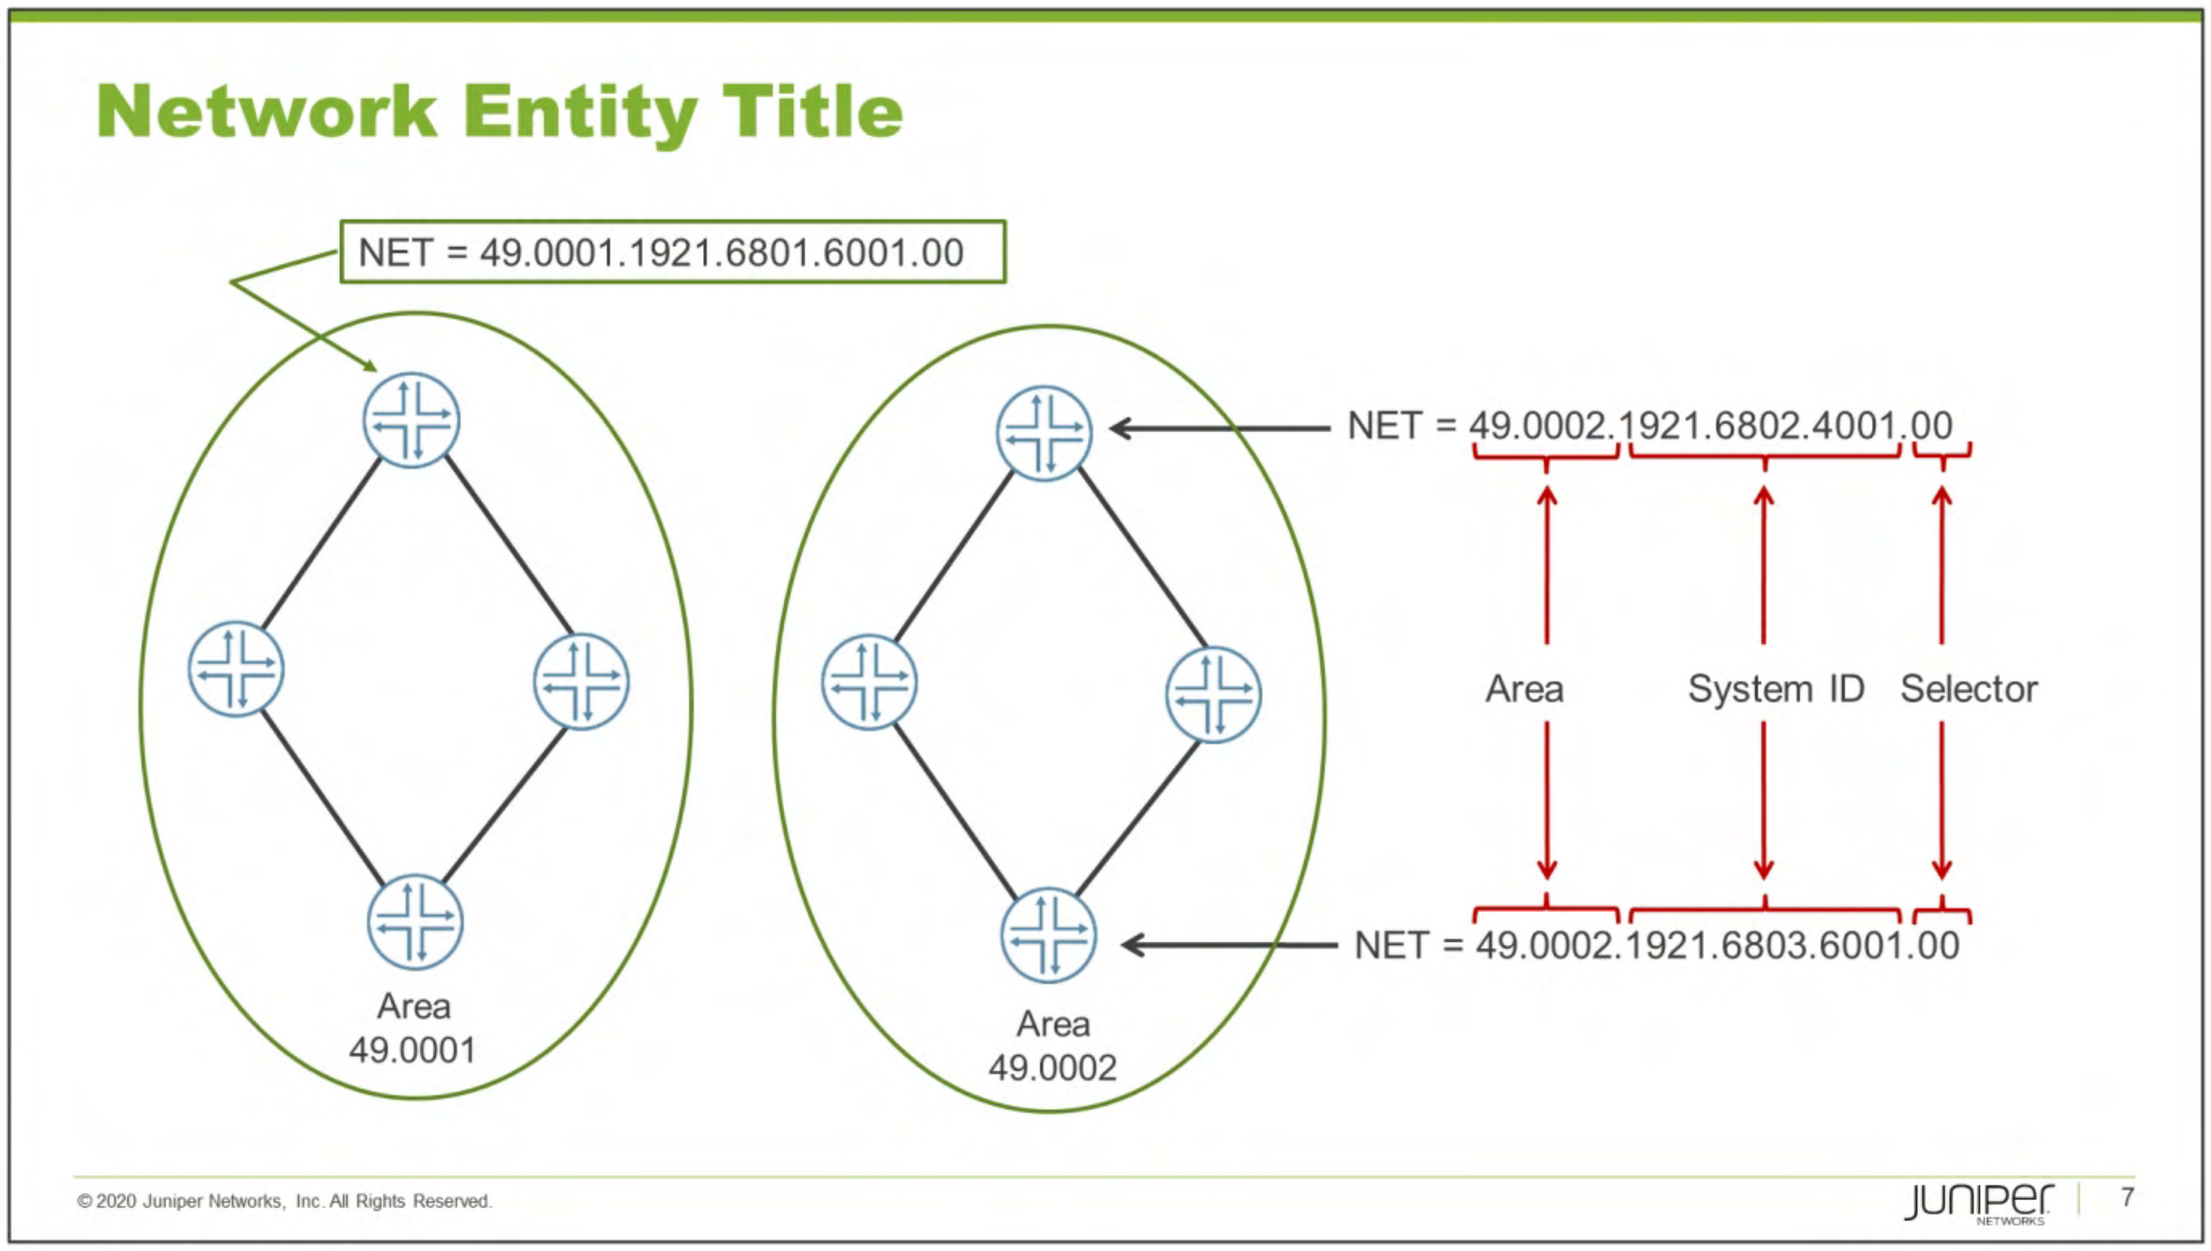

Network Entity Title

The slide displays a sample IS-IS network with two distinct areas configured—49.0001 and 49.0002. The area address can range from 1 to 13 bytes in length and is always the first component of the Network Entity Title (NET) address. The middle portion of the NET is the system ID of the router, which should be unique throughout the entire domain.

The system ID is always 6 bytes and can be any set of hexadecimal digits you want. One common convention is placing the IP address of the loopback interface into the system ID portion of the NET. You accomplish this convention by ensuring that the dotted-decimal IP address is written with all of the leading zeros filled in. On the slide, the router in the top left of the network has 192.168.16.1 configured as its loopback address. This address can also be written as 192.168.016.001. We now have an address with 12 digits in it, exactly the number we need for the system ID. We simply move the dots between the digits into a format used in the hexadecimal notation. This process results in a system ID of 1921.6801.6001.

The NSEL byte (or N-selector) is the last byte of the NET address. It identifies the destination network layer service that should receive the traffic. For a router, this part of the address will always be 00. The IS-IS protocol may assign this value to represent the pseudo-node.

IS-IS and OSPF Comparison (1 of 2)

Comparing:	OSPF	IS-IS
Developed by	IETF	ISO
Protocol used	TCP/IP	CLNP/CLNS
Format used	Packets	PDUs
Hierarchy	2 Levels BB — Area 0 NBB — any area that is not Area 0)	2 Levels BB — Connected Level 2 (any area) NBB — Connected Level 1 (same area)
Addressing	32 bit (4 Octets) Network/host	Net Address: (Area.System_ID.Sel)
Algorithm used	SPF	SPF
Naming Convention	Hosts/Routers	End-Systems/Intermediate Systems
Backbone	Area 0	Contiguous Connection of L2 ISs
Non-Backbone	Any area this is not Area 0	L1 routers in a common area
Border Routers	ABR — 1 router in multiple areas	L1/L2 — 1 router in a single area

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS and OSPF Comparisons (1 of 2)

The slide compares the OSPF and IS-IS protocols

IS-IS and OSPF Comparison (2 of 2)

Comparing:	OSPF	IS-IS
Advertise MA networks	DR – BDR	DIS – N/A
Link Info	LS Updates LS Advertisements	LS PDUs TLVs
Link State Messages	Database Description	CNSP
Link State Messages	LS Request LS Acknowledgement	PSNP
Link State Aging	1 hr max, 30m RFC, 50m-Junos Fixed	20m max, 15m, Configurable
Summarize	Automatic – area-range	Manual – aggregate route & policy
Policies	Import and Export	Export Only
IPv6 and Traffic Eng	IPv6 in OSPFv3 & TE optional	Supported by default

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS and OSPF Comparisons (2 of 2)

The slide continues comparing the OSPF and IS-IS protocols.

OSPF Areas

- A single OSPF router can be part of multiple areas
- An OSPF link can only form an adjacency with a link in the same area
- The backbone consists of all routers with at least one interface in area 0

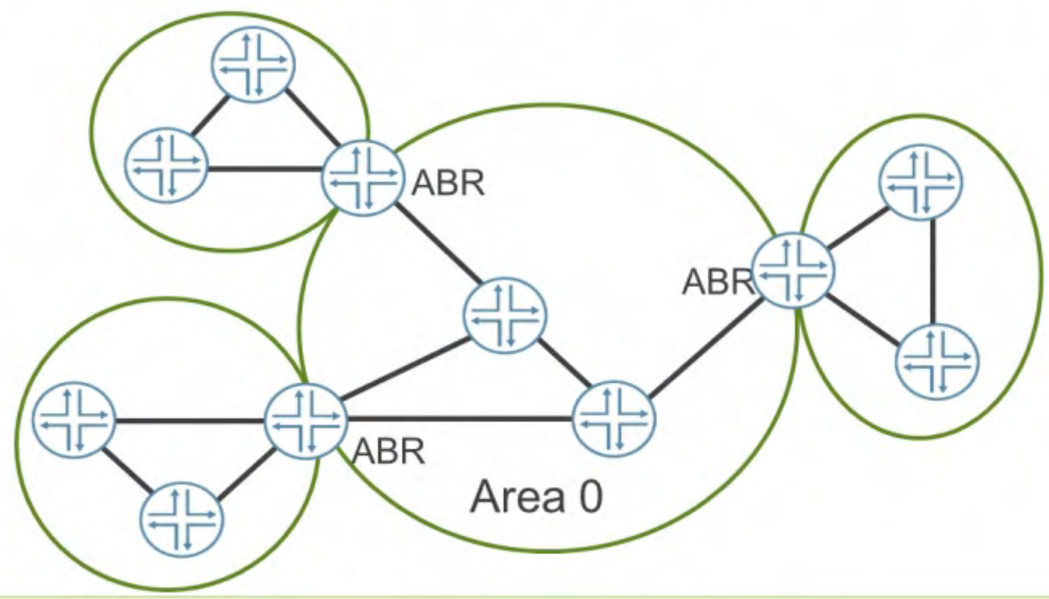

© 2020 Juniper Networks, Inc. All Rights Reserved.

OSPF Areas

IS-IS and OSPF are link-state routing protocols with many similarities. Differences exist between them as well. Recall that, under OSPF, routers separate areas. The slide shows how a typical OSPF network might be broken up into areas. Some interfaces are in one area, and other interfaces are in another area. When an OSPF router has interfaces in more than one area, it is an ABR.

IS-IS Areas

- A single IS-IS router can be in only one area per instance
- A L2 IS-IS interface can form an adjacency with a router in the same area or in a different area
- The backbone is a contiguous set of L2 links

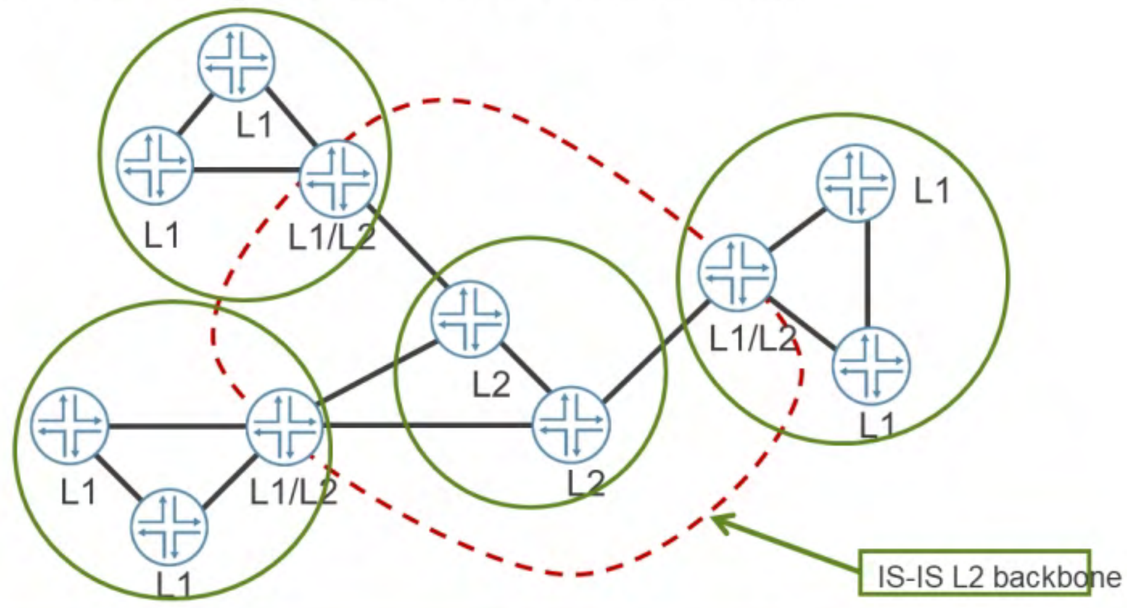

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Areas

In IS-IS areas, a Level 1/Level 2 router fulfills the same purpose as an area border router (ABR) in OSPF. Likewise, the collection of Level 2 routers in IS-IS is the backbone, while Area 0 is the backbone in OSPF. However, in IS-IS, all routers are completely within an area, and the area borders are on the links, not on the routers. The routers that connect areas are Level 2 routers, and routers that have no direct connectivity to another area are Level 1 routers. An intermediate system can be a Level 1 router, a Level 2 router, or both (an L1/L2 router).

Agenda: IS-IS

- Overview of IS-IS
- →IS-IS PDUs
- Neighbors and Adjacencies
- Configuring and Monitoring IS-IS

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS PDUs

The slide highlights the topic we discuss next.

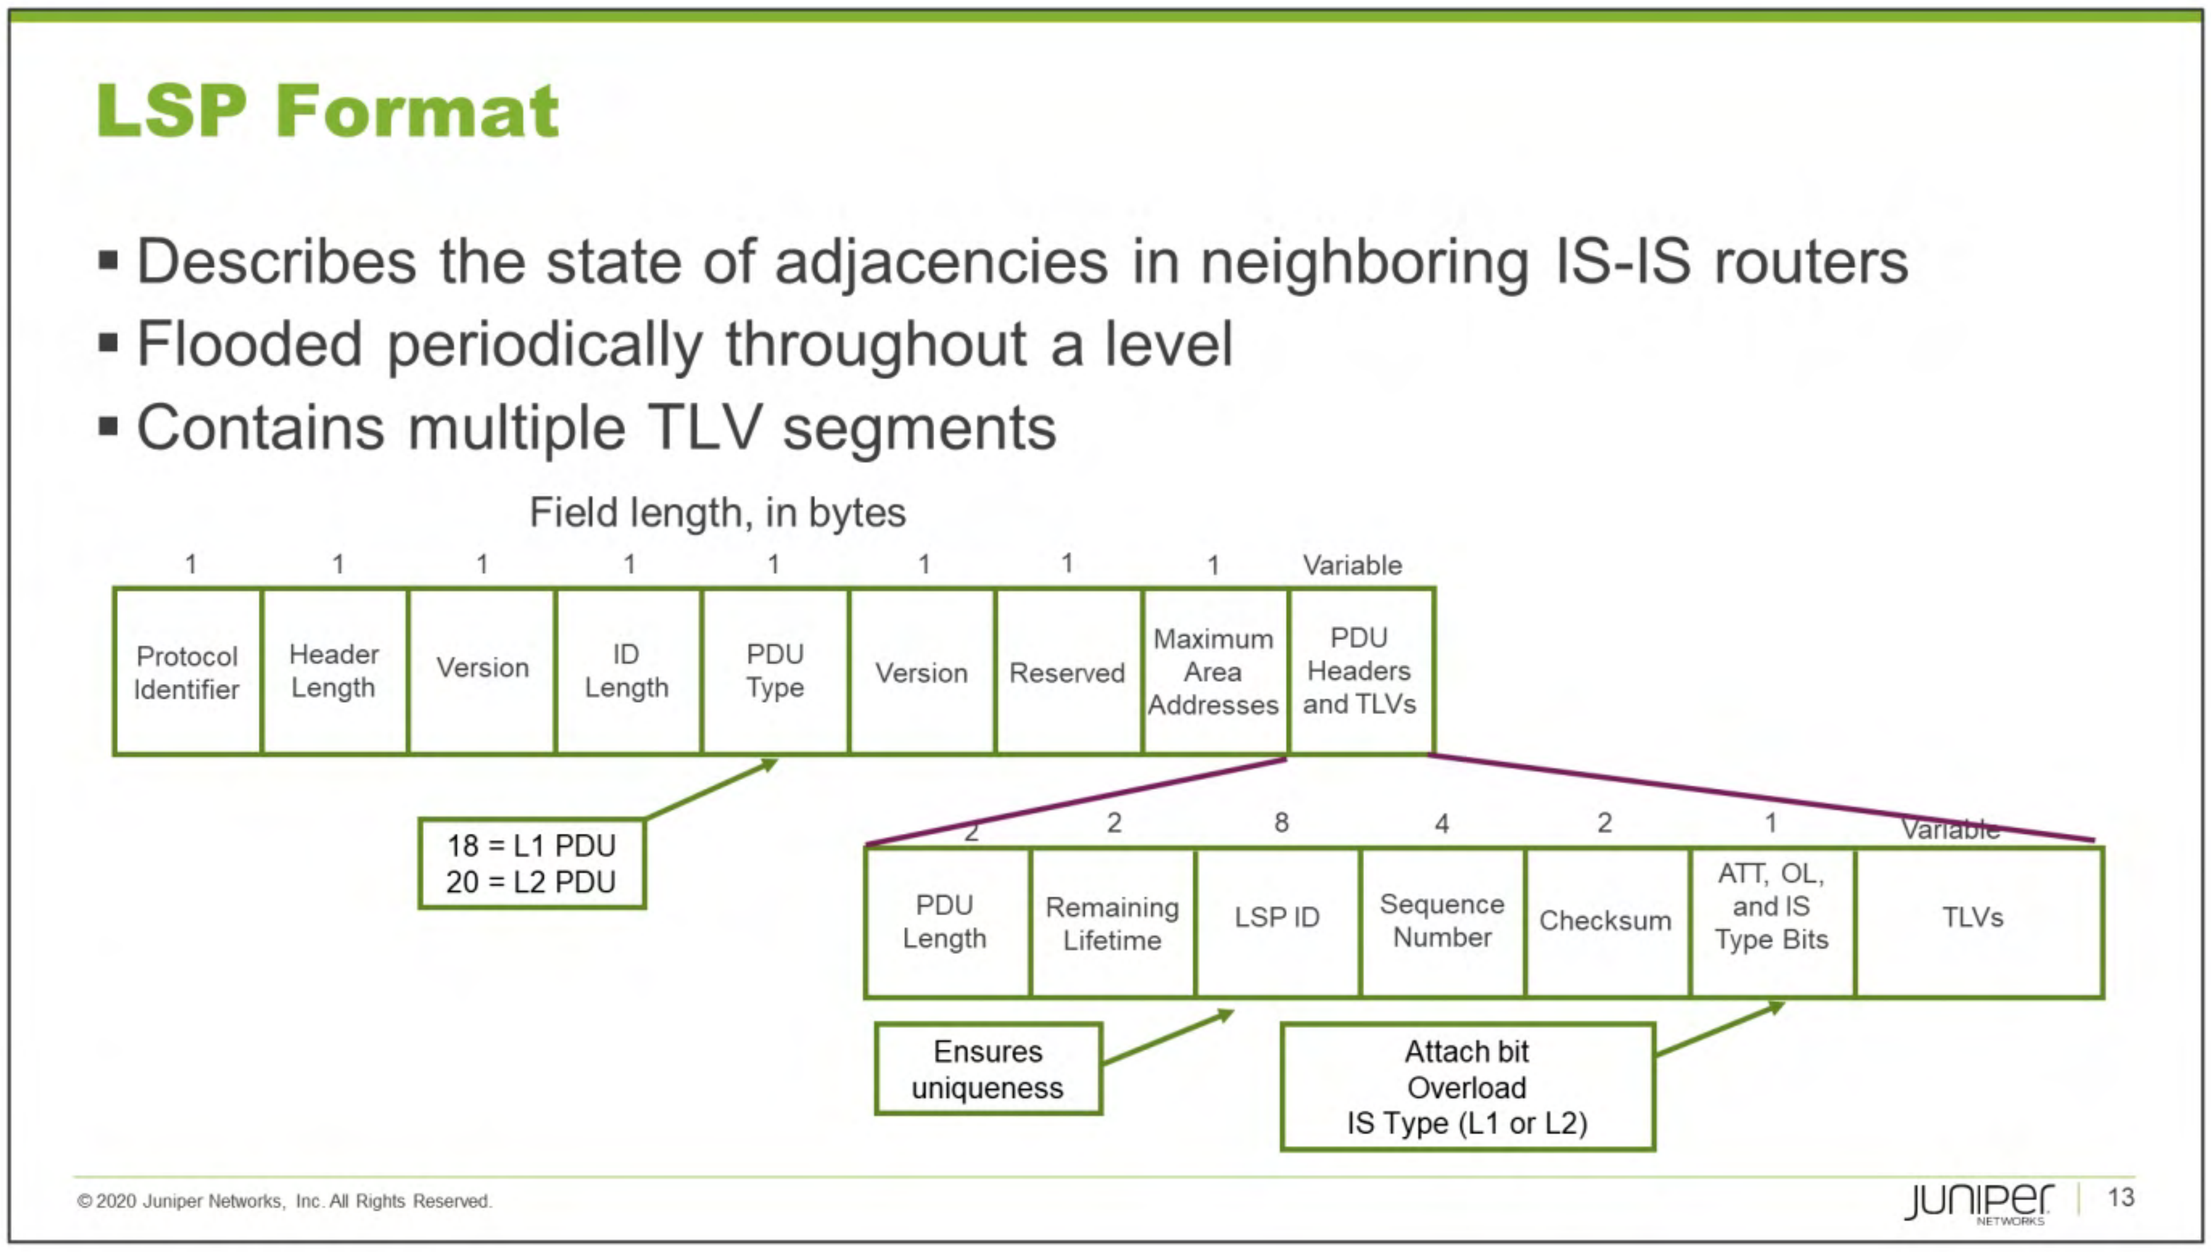

Describes Adjacency State

The LSP that each IS-IS router generates describes the state of the local router and its adjacencies in the network.

Some fields of interest in the LSP header include the ID length and the maximum area address, which are set to a constant value of 0x00. This value does not mean that Junos OS does not support their functionality. Instead, these fields are used for backward compatibility with older protocol implementations. The system ID is always 6 bytes in length, and the maximum area addresses supported is always 3 bytes. By setting these two fields to a value of 0, Junos OS is reporting that it supports the default values for these two settings.

You might also notice that the LSP header contains two version fields. The first of these fields, according to the original IS-IS specification, was designed as an extension of the protocol ID field. Most modern implementations, including Junos OS, do not support this function and instead place a constant value of 0x01 in the field to represent the version of the protocol. The second version field is the actual specified location for the protocol version, which is also set to a value of 0x01—the current version of the protocol.

Flooded Periodically

The IS-IS LSPs are flooded when either a network change occurs or often enough to keep the database from having stale information. Each LSP has a 2-byte field named **`remaining lifetime`**. Each IS-IS router initializes this field to a certain value when the LSP is created. By default, this timer value is set to 1200 seconds, or 20 minutes. Each router takes this value and begins a countdown toward 0. Before the timer expires (at approximately 317 seconds), the originating system regenerates the LSP and floods it to all its neighbors.

Contains TLV Segments

Each individual LSP contains multiple TLV segments that describe the originating router. Many possible TLVs can be sent, but each system only sends those that it needs to communicate to adjacent systems. The details of the TLV segments are discussed in future slides using the output of the **`show isis database extensive`** command. You can also view the received TLVs using the **`monitor traffic detail`** command. The following output shows a sample from this command:

Continued on the next page.

TLV Segments (Continued)

```
user@router> monitor traffic detail interface so-0/1/0 size 1514
Listening on so-0/1/0
11:55:48.470418  In ISIS(186), 30:30:30:30:30:30 > 30:30:30:30:30:30, hlen: 27, v: 1,
    sys-id-len: 6 (0), max-area: 3 (0), L2 LSP
    lsp-id: 1921.6804.8001.00-00, seq: 0x00000008, lifetime:  1189s
    chksum: 0x86c9 (correct), PDU length: 186, L1L2 IS
      Area address(es) TLV #1, length: 4
          Area address (3): 49.0001
      Protocols supported TLV #129, length: 2
          NLPID(s): IPv4, IPv6
      Traffic Engineering Router ID TLV #134, length: 4
          Traffic Engineering Router ID: 192.168.48.1
      IPv4 Interface address(es) TLV #132, length: 4
          IPv4 interface address: 192.168.48.1
      Hostname TLV #137, length: 8
          Hostname: SaoPaulo
      IPv4 Internal reachability TLV #128, length: 24
          IPv4 prefix: 192.168.48.1/32
            Default Metric: 00, Internal, Distribution: up
          IPv4 prefix: 10.222.60.0/24
            Default Metric: 10, Internal, Distribution: up
      Extended IPv4 reachability TLV #135, length: 17
          IPv4 prefix: 192.168.48.1/32
            Metric: 0, Distribution: up, no sub-TLVs present
          IPv4 prefix: 10.222.60.0/24
            Metric: 10, Distribution: up, no sub-TLVs present
      IPv4 External reachability TLV #130, length: 12
          IPv4 prefix: 192.168.49.0/24
            Default Metric: 00, Internal, Distribution: up
      Extended IPv4 reachability TLV #135, length: 8
          IPv4 prefix: 192.168.49.0/24
            Metric: 0, Distribution: up, no sub-TLVs present
      IS Reachability TLV #2, length: 12
          IsNotVirtual
          IS Neighbor: 1921.6805.2001.00, Default Metric: 10, Internal
      Extended IS Reachability TLV #22, length: 23
          IS Neighbor: 1921.6805.2001.00, Metric: 10, sub-TLVs present (12)
            IPv4 interface address: 10.222.60.2
            IPv4 neighbor address: 10.222.60.1
      Authentication TLV #10, length: 17
          HMAC-MD5 password: 00bb32fd7712bcea6003e516e2333077
```

LSP Notes

- PDU type field denotes a Level 1 or Level 2 PDU
  - Level 1 PDU = 18
  - Level 2 PDU = 20
- LSP ID field provides uniqueness in the domain
- Attached (ATT) bit is set if the IS is connected to another area
- Overload (OL) bit is set if the link-state database is overloaded
- IS type bits determine a Level 1 or Level 2 router (only 2 settings possible)
  - Level 1 router = 01 = 0x1
  - Level 1/2 router = 11 = 0x3

© 2020 Juniper Networks, Inc. All Rights Reserved.

PDU Type

The 1-byte PDU type field denotes whether the PDU is Level 1 or Level 2 PDU. A value of 18 represents a Level 1 PDU and a value of 20 is a Level 2 PDU. With this exception as well as different multicast destination addresses, the PDU contents and formats are identical to each other.

LSP ID Field

The ID of the LSP uniquely identifies it within the IS-IS domain. The 8-octet field is comprised of the router's system ID (6 bytes), the circuit ID (1 byte), and the LSP number (1 byte). The system ID is located within the NET address of the router and is used exactly as is. The LSP number field represents a fragmented LSP. The initial LSP receives a value of 0x00, and it is incremented by 1 for each following fragment. The circuit ID field helps distinguish LSPs advertised from a single router. By default, all LSPs representing the router as a node use a value of 0x00 for the circuit ID field. When a router is operating as a designated intermediate system (DIS), it places the value of the circuit ID into this field and generates a separate LSP for the broadcast segment. Junos OS uses a circuit ID value of 0x01 for the loopback interface as well as all point-to-point interfaces. Each broadcast segment receives a unique circuit ID value beginning at 0x02 and incrementing to a maximum value of 0xff.

Attached Bit

Any IS-IS router with a Level 2 adjacency to a system in another area sets the Attached bit in its Level 1 LSP. Level 1 routers can then forward data packets to that attached router for transport out of the local area.

Overload Bit

An IS-IS router can set the overload bit to inform the other routers in the network that it is incapable of reliably transmitting transit packets to other portions of the network. While the original intent of the bit was to ease memory problems on individual routers, modern-era routers do not have such constraints. In today's network environment, the bit is often used to take a router out of service for maintenance or to allow it to fully converge within the network before forwarding traffic.

IS Type Bits

The IS type bits inform other routers in the network about the capabilities of the local router. Two possible settings are available for these bits. When the bits are set to a value of 01 (0x01), the local router can support only Level 1 routing. A value of 11 (0x03), means that the local router can support both Level 1 and Level 2 routing.

IS-IS Messages

- Hello
- Link State PDUs (LSP)
- Partial Sequence Number PDUs (PSNP)
- Complete Sequence Number PDUs (CSNP)

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Messages

IS-IS uses a number of messages to form adjacencies and exchange information. These message types include:

- Hello messages,
- Link State PDUs (LSP),
- Partial sequence number PDUs (PSNP), and
- Complete sequence number PDUs (CSNP).

Hello PDUs

- The hello mechanism for neighbor discovery is used to build and maintain adjacencies
  - Similar to the hello mechanism in OSPF
- Separate hellos for:
  - Broadcast networks
    - Level-1 LAN hello (PDU type 15)
    - Level-2 LAN hello (PDU type 16)
  - Point-to-point hello (PDU type 17)
- Sent at regular intervals
  - DIS is sent at 1/3 the Hello timer
- Hello PDUs:
  - Identify the device
  - Describe its capabilities
  - Describe the parameters of the interface

© 2020 Juniper Networks, Inc. All Rights Reserved.

Purpose of IS-IS Hello PDUs

The purpose of the IS-IS hello PDU is to allow IS-IS routers to discover IS-IS neighbors on a link. Once the neighbors have been discovered and are adjacent, the hello PDU acts as a keepalive to maintain the adjacency and to inform the neighbors of any changes in the adjacency parameters.

IS-IS Hello PDU Types

Two kinds of IS-IS hellos PDUs exist: LAN hello PDUs and point-to-point hello PDUs. The LAN hello PDUs can be divided further into Level 1 and Level 2 hello PDUs. The format of the two types of LAN hello PDUs is identical. Note that on broadcast networks, IS-IS Level 1 and Level 2 hellos are coded with multicast address 01-80-C2-00-00-14 or 01-80-C2-00-00-15 respectively.

Hello Transmission

Routers send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. The hello interval field advertises this interval in the hello packet. By default, a DIS router sends hello packets every 3 seconds, and a non-DIS router sends hello packets every 9 seconds.

PDU Fields

The following list provides the details of the PDU fields:

- *Circuit type*: Defines the router as an Level 1, Level 2, or Level 1/Level 2 router;
- *Source ID*: Identifies the system ID of the router that originated the hello PDU;
- *Holding time*: Specifies the period a neighbor should wait to receive the next hello PDU before declaring the originating router dead;
- *PDU length*: Specifies the length of the entire PDU in octets;
- *Priority*: Provides a value between 0 and 127 used for DIS election; and
- *LAN ID*: Identifies the system ID or the DIS plus one more octet (the pseudo-node ID) to differentiate this LAN ID from another LAN ID that might have the same DIS.

LSPs

- LSPs
  - Used to build the link-state database
    - Similar to LSAs in OSPF
  - Separate LSPs for:
    - Level 1 systems (PDU Type 18)
    - Level 2 systems (PDU Type 20)
  - Sent as a result of:
    - Network change,
    - During adjacency formation, and
    - In response to a sequence number PDU (described on the next slide)
  - LSPs:
    - Identify an IS's adjacencies
    - Describe the state of its adjacencies
    - Describe its reachable address prefixes (routes)

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS LSPs

IS-IS sends LSPs at regular intervals and when an IS discovers any of the following:

- Its link to a neighbor is down;
- It has a new neighbor; or
- The cost of a link to an existing neighbor has changed.

Once LSPs are distributed appropriately, IS-IS runs the Dijkstra algorithm to compute optimal paths to each ES. This algorithm iterates on the length of a path, examining the LSPs of all ISs working outward from the host IS. At the end of the computation, IS-IS forms a connectivity tree yielding the shortest paths, including all intermediate hops, to each IS.

Sequence Number PDUs

- Partial sequence number PDU
  - Used to:
    - Maintain the link-state database synchronization
    - Acknowledge LSPs from a neighbor on a point-to-point network
    - Request a copy of a missing LSP on a broadcast network
  - Separate PDU types for Level 1 (26) and Level 2 (27) systems
  - Contains specific header information for the LSP being acknowledged or requested
- Complete sequence number PDUs
  - Used to maintain the link-state database synchronization
    - Sent periodically by all ISs on point-to-point networks
    - Only sent by DIS on broadcast networks
  - Separate PDU types for Level 1 (24) and Level 2 (25) systems
  - Contains header information for all LSPs in the IS's link-state database

© 2020 Juniper Networks, Inc. All Rights Reserved.

Partial Sequence Number PDUs

A receiver multicasts partial sequence number PDUs (PSNPs) when it detects that it is missing an LSP or when its LSP database is out of date. The receiver sends a PSNP to the system that transmitted the complete sequence number PDUs (CSNPs), effectively requesting that the missing LSP be transmitted. That router, in turn, forwards the missing LSP to the requesting router.

Complete Sequence Number PDUs

CSNPs contain a complete description of all LSPs in the IS-IS database. IS-IS sends CSNPs periodically on all links. The receiving systems use the information in the CSNP to update and synchronize their LSP databases. The DIS router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each LSP.

Type/Length/Values

- IS-IS information objects
  - Each piece of IS-IS information is defined as an object with three attributes:
    - *Object type*: The predefined code for the type of information contained in the object
    - *Object length*: The length of the information (allows for variable-length objects)
    - *Object value*: The actual information defined by the type attribute
  - TLVs are the building blocks of IS-IS PDUs, which are used for information exchange
    - Some TLVs are used in multiple PDUs
    - Some TLVs are PDU-specific
  - Similar to OSPF packet-specific and LSA-specific fields

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Information Objects

IS-IS PDUs use TLV encoding as the basic structure for all routing information. TLV encoding requires that the length of any field be defined explicitly when a PDU uses that field. IS-IS ignores all unknown TLVs, making the protocol easily extensible.

Some TLV Examples (1 of 2)

TLV Code	Defined in	Used for
1	ISO 10589	Area addresses
2	ISO 10589	IS neighbor metrics
6	ISO 10589	Neighbor LAN ID
8	ISO 10589	Padding
9	ISO 10589	LSP entries
10	ISO 10589	Authentication
22	RFC 5029	Extended IS reachability
128	RFC 1195	IP prefix, mask, and metrics
129	RFC 1195	Protocols supported

© 2020 Juniper Networks, Inc. All Rights Reserved.

TLV Variables, Part 1

The following list provides the details of several of the TLVs:

- *TLV 1 (Area address)*: Provides the area address encoded within the IS-IS NET on the loopback 0 (lo0) interface.
- *TLV 2 (IS reachability)*: Advertises the IS neighbors adjacent to the local router as well as the metric used to reach those neighbors.
- *TLV 10 (Authentication)*: Contains the authentication type and the configured password.
- *TLV 22 (Extended IS reachability)*: Advertises the IS neighbors adjacent to the local router and the routers that support traffic engineering capabilities. This TLV also contains multiple sub-TLVs that describe the user constraints placed on the router by the network administrator. This TLV also populates the traffic engineering database (TED).
- *TLV 128 (IP internal reachability)*: Advertises the IP address and subnet mask for each of the router's interfaces capable of supporting IP version 4 (IPv4) traffic.
- *TLV 129 (Protocols supported)*: Informs other routers in the network which Layer 3 protocols the local router supports. By default, Junos OS supports both IPv4 and IP version 6 (IPv6). On the MX Series Services Router and the SRX Series Services Gateways, you can also use Junos OS to support CLNS.

Some TLV Examples (2 of 2)

TLV Code	Defined in	Used for
130	RFC 1195	IP external information
132	RFC 1195	IP interface address
135	RFC 3784	Extended IP reachability
137	RFC 2763	Dynamic hostname mapping
222	RFC 5311	Multiple topologies (routing instances)
229	RFC 5311	Multiple topologies (routing instances)
232	RFC 5308	IPv6 support
235	RFC 3784	Multiple topologies (routing instances)
236	RFC 5308	IPv6 support

© 2020 Juniper Networks, Inc. All Rights Reserved.

TLV Variables, Part 2

The following list is a continuation of the details of several TLV variables:

- *TLV 130 (IP external reachability)*: Advertises the network and subnet mask for all routes advertised into IS-IS by using a policy.
- *TLV 132 (IP interface address)*: Advertises the host IP address for all router interfaces.
- *TLV 134 (Traffic engineering IP router ID)*: Advertises the 32-bit router ID (RID) of the local router.
- *TLV 135 (Extended IP reachability)*: Advertises the IP address and subnet mask for router interfaces that can support traffic engineering. This TLV also populates the TED.
- *TLV 137 (Dynamic hostname resolution)*: Advertises the ASCII hostname configured on the local router. Other IS systems use this TLV to resolve the hostname of the router for use in **`show`** command output and within certain TLVs.

Multiple Topology TLVs

- *TLV 222 (Multiple topology IS reachability)*: Advertises the IS neighbors adjacent to the local router and the routers that support multiple topologies of IS-IS.
- *TLV 229 (Multiple topologies supported)*: Advertises which multiple topologies of IS-IS the local router supports. Each topology is identified by a 12-bit ID field.
- *TLV 235 (Multiple topology IP reachability)*: Advertises the IP information for interfaces that support multiple topologies. This TLV contains multiple sub-TLVs, which define the actual information. Each set of sub-TLVs is accompanied by the 12-bit topology ID field.

IPv6 TLVs

The following list provides the details of the TLVs that support IPv6:

- *TLV 232 (IPv6 interface address)*: Advertises the IPv6 interface address for those interfaces that support IPv6 traffic.
- *TLV 236 (IPv6 reachability)*: Advertises information about the network link on which the IPv6 protocol is operating. This TLV contains multiple sub-TLVs that contain the actual metric information, among other things.

TLV 1—Area Address

- Advertises the area ID of the originating router
  - (1-byte) TLV type
  - (1-byte) TLV length
  - (1-byte) Area length
    - Repeated for each configured area
  - (Variable) Area ID
    - Ranges from 1 to 13 bytes
    - Repeated for each configured area

```
user@host> show isis database extensive | find TLV

TLVs:
    Area address: 49.0001 (3)
    Speaks: IP
    Speaks: IPv6
    IP router id: 192.168.48.1
    IP address: 192.168.48.1
    Hostname: SaoPaulo
```

TLV 1

© 2020 Juniper Networks, Inc. All Rights Reserved.

Area Address TLV

Each IS-IS router sends this TLV in both its Level 1 and its Level 2 LSPs. At times, during a network migration, for instance, when you might need more than one area configured on the router, up to three area addresses are configurable per system. The area ID field is repeated for each unique area address. The area address TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 1.
- *TLV length (1 byte)*: This field contains the length of the remaining fields in the TLV. Together, the TLV length and area length fields allow an IS-IS router to determine the number of area addresses encoded within the TLV.
- *Area length (1 byte)*: The length of the following advertised area is encoded in this field.
- *Area ID (variable)*: This field can range from 1 to 13 bytes. It contains the actual area address configured within the NET ID of the router.

TLV 2—IS Reachability

- Describes the IS neighbors of the local router
  - Metric cost to each neighbor is advertised
  - Metric and neighbor ID fields are repeated for each neighbor
    - (1-byte) TLV type
    - (1-byte) TLV length
    - (1-byte) Virtual flag
    - (1-byte) R (Reserved) bit, I/E bit, default metric
    - (1-byte) S (Supported) bit, I/E bit, delay metric
    - (1-byte) S (Supported) bit, I/E bit, expense metric
    - (1-byte) S (Supported) bit, I/E bit, error metric
    - (7-byte) Neighbor ID

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    ...
    IS neighbor: Sydney.00, Internal, Metric: default 10
```

TLV 2

Neighbor ID = System ID replaced with Hostname

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS Reachability TLV

Each IS-IS router advertises its adjacencies with neighboring systems through this TLV. The various metric fields as well as the neighbor ID field are repeated for each adjacent neighbor. The metric values are encoded in a 6-bit field resulting in possible metrics between 0 and 63. These values are sometimes referred to as *old-style* or *small* metrics. The IS reachability TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 2.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field. This length allows the router to determine the number of neighbors encoded within the TLV. Each set of neighbors and metrics occupies 11 octets of space. Therefore, the field length minus 1 (for the virtual flag) should be divisible by 11, resulting in the number of adjacent neighbors.
- *Virtual flag (1 byte)*: An IS-IS router sets this flag when the advertised information should be used to repair a nonadjacent Level 2 area. Junos OS does not support partition repair and this field is set to a constant value of 0x00.
- *Default metric (1 byte)*: The first bit in this field is a reserved bit and is set to a value of 0. The second bit in this field can be used to support internal and external metrics by indicating the metric type; internal and external metrics are not comparable. Because this TLV is never leaked, the I/E bit is always coded to a zero to indicate an internal type. The remaining 6 bits are used to encode the metric cost to reach the adjacent neighbor.
- *Delay metric (1 byte)*: Junos OS does not support the use of type of service (ToS) metrics within IS-IS. The S bit is set to a constant value of 1 (not supported), while the I/E and metric bits are all set to a constant value of 0.
- *Expense metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit is set to a constant value of 1 (not supported), while the I/E and metric bits are all set to a constant value of 0.
- *Error metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit is set to a constant value of 1 (not supported), while the I/E and metric bits are all set to a constant value of 0.
- *Neighbor ID (7 bytes)*: The ID of the adjacent neighbor is encoded in this field. It is comprised of the 6-byte system ID and the 1-byte circuit ID of the neighbor.

TLV 10—Authentication

- Encodes authentication data to ensure that only trusted information is placed into the link-state database
  - (1-byte) TLV type
  - (1-byte) TLV length
  - (1-byte) Authentication type
  - (Variable) Password

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    ...
    IS neighbor: Sydney.00, Internal, Metric: default 10
    IS extended neighbor: Sydney.00, Metric: default 10
      IP address: 10.222.60.2
      Neighbor's IP address: 10.222.60.1
    ...
    Authentication data: 17 bytes        <-- TLV 10
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Authentication TLV

If configured to support authentication, an IS-IS router includes this TLV in all advertised link-state PDUs. The authentication TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 10.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field.
- *Authentication type (1 byte)*: The specific form of authentication is encoded in this field. Plain-text authentication uses a value of 1, while Message Digest 5 (MD5) authentication uses a value of 54.
- *Password (Variable)*: The actual authentication data is stored in this field. When MD5 is used, the size of this field is always 16 bytes.

TLV 22—Extended IS Reachability

- Advertises additional capabilities and information about IS neighbors
  - Larger metric values
  - Traffic engineering parameters
  - Information repeated for each neighbor (system ID)
    - (1-byte) TLV type
    - (1-byte) TLV length
    - (7-byte) System ID
    - (3-byte) Wide metric
    - (1-byte) Sub-TLV length
    - (Variable) Sub-TLVs

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    ...
    IS neighbor: Sydney.00, Internal, Metric: default 10
    IS extended neighbor: Sydney.00, Metric: default 10
      IP address: 10.222.60.2                    subTLV 6
      Neighbor's IP address: 10.222.60.1         subTLV 8
    ...
    Authentication data: 17 bytes
```

TLV 22 *wide* metrics

© 2020 Juniper Networks, Inc. All Rights Reserved.

Extended IS Reachability TLV

Each IS-IS router also advertises its adjacencies with neighboring systems through this TLV. The system ID, metric field, and optional sub-TLV fields are repeated for each adjacent neighbor. The main use of TLV 22 for *regular* IS-IS is the larger metric values it supports. The extended IS reachability TLV uses a 24-bit field that results in possible metrics between 0 and 16,777,215. This range of metric values is often referred to as *new-style* or *wide* metrics. The extended IS reachability TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 22.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field.
- *System ID (7 bytes)*: The ID of the adjacent neighbor is encoded in this field. It is comprised of the 6-byte system ID and the 1-byte circuit ID of the neighbor.
- *Wide metric (3 bytes)*: The metric cost to reach the adjacent neighbor is encoded in this field.

The following list is a continuation of the fields contained in the extended IS reachability TLV:

- *Sub-TLV length (1 byte)*: The length of any optional sub-TLVs is encoded in this field. If no sub-TLVs are present, the field is set to a value of 0.
- *Sub-TLVs (Variable)*: Additional traffic engineering information is advertised in these fields, each with a separate type code, length, and value portion. Each of the fields is placed within the TED on the router. In addition, some sub-TLVs are placed into the LSDB on the router. The possible sub-TLVs include the following:
  - Administrative group (Type 3);
  - IPv4 interface address (Type 6);
  - IPv4 neighbor address (Type 8);
  - Maximum link bandwidth (Type 9);
  - Reservable link bandwidth (Type 10); and
  - Unreserved bandwidth (Type 11).

TLV 128—IP Internal Reachability

- Describes the internal IP information of the local router
  - IP prefix and metric advertised
  - Information repeated for each prefix
    - (1-byte) TLV type
    - (1-byte) TLV length
    - (1-byte) U/D bit, I/E bit, default metric
    - (1-byte) S (Supported) bit, R (Reserved) bit, delay metric
    - (1-byte) S (Supported) bit, R (Reserved) bit, expense metric
    - (1-byte) S (Supported) bit, R (Reserved) bit, error metric
    - (4-byte) IP address
    - (4-byte) Subnet mask

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    ...
    IP prefix: 192.168.48.1/32, Internal, Metric: default 0, Up
    IP prefix: 10.222.60.0/24, Internal, Metric: default 10, Up
```

TLV 128 *small* metrics

© 2020 Juniper Networks, Inc. All Rights Reserved.

IP Internal Reachability TLV

Each IS-IS router advertises its locally connected IP prefixes with the IP internal reachability TLV. The metric fields, the IP address, and the subnet mask are repeated for each advertised prefix. As with the IS reachability TLV, the metric values are encoded in a 6-bit field resulting in possible metrics between 0 and 63 (*small* metrics). The IP internal reachability TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field, which holds a constant value of 128.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field. This length allows the router to determine the number of prefixes advertised within the TLV. Each set of metrics and addresses occupies 12 octets of space.
- *Default metric (1 byte)*: The first bit in this field is known as the Up/Down (U/D) bit. It is used to provide information to IS-IS routers in a multilevel network to allow for prefix advertisements across a level boundary and to prevent routing loops. The second bit in this field can be used to indicate whether a given pair of metrics are comparable by indicating a metric type of either internal or external. Although some Junos OS versions made use of the I/E bit for TLV 128, the current release ignores this bit upon reception and treats all prefixes as internal. The final 6 bits represent the metric cost to reach the advertised prefix.

The following list details the remaining fields in the IP internal reachability TLV:

- *Delay metric (1 byte)*: unos OS does not support the use of ToS metrics within IS-IS. The S bit, the R bit, and the metric bits are all set to a constant value of 0.
- *Expense metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit, the R bit, and the metric bits are all set to a constant value of 0.
- *Error metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit, the R bit, and the metric bits are all set to a constant value of 0.

TLV 129—Protocols Supported

- Describes the Layer 3 protocols supported by the local router
  - (1-byte) TLV type
  - (1-byte) TLV length
  - (1-byte) Network Layer Protocol ID
    - Repeated for each supported protocol

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    Speaks: IP
    Speaks: IPv6
    IP router id: 192.168.48.1
    IP address: 192.168.48.1
    Hostname: SaoPaulo
```

TLV 129

© 2020 Juniper Networks, Inc. All Rights Reserved.

Protocols Supported TLV

Every IS-IS router sends this TLV in all LSPs. It lists the Layer 3 protocols supported by the local router, making IS-IS a true multiprotocol routing protocol. The network layer protocol identifier (NLPID) is repeated for each supported protocol. The TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field. A constant value of 129 is placed in this field.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field. This length allows a router to determine the number of NLPIDs encoded within the TLV.
- *Network Layer protocol ID (1 byte)*: The 8-bit NLPID is placed within this field. By default, Junos OS supports both IPv4 (0xCC) and IPv6 (0x8E) and encodes those values in this TLV. On MX Series routers and SRX Series devices, you can also configure Junos OS to support the CLNS protocol by configuring **`clns-routing`** at the `[edit protocols isis]` configuration hierarchy.

TLV 130—IP External Reachability

- Describes the external IP information of the local router
  - IP prefix and metric advertised
  - Information repeated for each prefix
    - (1-byte) TLV type
    - (1-byte) TLV length
    - (1-byte) U/D bit, I/E bit, default metric
    - (1-byte) S (Supported) bit, R (Reserved) bit, delay metric
    - (1-byte) S (Supported) bit, R (Reserved) bit, expense metric
    - (1-byte) S (Supported) bit, R (Reserved) bit, error metric
    - (4-byte) IP address
    - (4-byte) Subnet mask

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    ...
    IP external prefix: 192.168.49.0/24, Internal, Metric: default 0, Up
    ...
    Authentication data: 17 bytes
```

TLV 130

© 2020 Juniper Networks, Inc. All Rights Reserved.

IP External Reachability TLV

Each IS-IS router advertises prefixes that are natively external to IS-IS with the IP external reachability TLV. As with the IP internal reachability TLV, the metric, IP address, and subnet mask fields are repeated for each advertised prefix. The metric values are again encoded in a 6-bit *small* metrics field. The IP external reachability TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 130.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field. This length allows the router to determine the number of prefixes advertised within the TLV. Each set of metrics and addresses occupies 12 octets of space.
- *Default metric (1 byte)*: The first bit in this field is the U/D bit. It is used by IS-IS routers in a multilevel network to allow for prefix advertisements across a level boundary and to prevent routing loops. The second bit in this field can be used to indicate whether a given pair of metrics is comparable by indicating a metric type of either internal or external. Although some Junos OS versions made use of the I/E bit for TLV 130, the current release ignores this bit upon reception and treats all prefixes as external by virtue of their being carried within TLV 130. The final 6 bits represent the metric cost to reach the advertised prefix.

The following list details the fields remaining in the IP external reachability TLV:

- *Delay metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit, the R bit, and the metric bits are all set to a constant value of 0.
- *Expense metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit, the R bit, and the metric bits are all set to a constant value of 0.
- *Error metric (1 byte)*: Junos OS does not support the use of ToS metrics within IS-IS. The S bit, the R bit, and the metric bits are all set to a constant value of 0.
- *IP address (4 bytes)*: The IPv4 prefix being advertised in the TLV is encoded in this field.
- *Subnet mask (4 bytes)*: The subnet mask associated with the advertised prefix is stored in this field.

TLV 132—IP Interface Address

- Advertises the IP address of the local router's interface
  - IPv4 address field is repeated for each advertised address
    - (1-byte) TLV type
    - (1-byte) TLV length
    - (4-byte) IPv4 address

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    Speaks: IP
    Speaks: IPv6
    IP router id: 192.168.48.1
    IP address: 192.168.48.1      <- TLV 132
    Hostname: SaoPaulo
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

IP Interface Address TLV

IS-IS routers send this TLV in all LSPs, which lists IP addresses configured on the originating router. At least one interface address must be advertised while each router interface can be advertised. By default, Junos OS encodes the RID of the local system in this TLV. This RID is often the same as the primary address of the router's `lo0.0` interface. The IP interface address TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 132.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field. This length allows a router to determine the number of addresses encoded within the TLV.
- *IPv4 address (4 bytes)*: The RID is placed in this field.

TLV 134—TE IP Router ID

- Advertises the traffic engineering router ID of the local router
  - (1-byte) TLV type
  - (1-byte) TLV length
  - (4-byte) Router ID

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    Speaks: IP
    Speaks: IPv6
    IP router id: 192.168.48.1        <-- TLV 134
    IP address: 192.168.48.1
    Hostname: SaoPaulo
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

TE IP RID TLV

All routers configured to support traffic engineering, which is the Junos OS default setting, generate this TLV. The RID of the local router is placed in this field for use in the TED. The TE IP RID TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 134;
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field with a constant value of 4; and
- *IPv4 address (4 bytes)*: The RID of the local router is placed in this field.

TLV 135—Extended IP Reachability

- Advertises information about the local router's IP reachability
  - Larger metric values
  - Information repeated for each prefix
    - (1-byte) TLV type
    - (1-byte) TLV length
    - (4-byte) Metric
    - (1-byte) Up/Down bit, Sub bit, and prefix length
    - (Variable) Prefix
    - (1-byte) Optional sub-TLV type
    - (1-byte) Optional sub-TLV length
    - (Variable) Optional sub-TLV

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    ...
    IP extended prefix: 192.168.48.1/32 metric 0 up
    IP extended prefix: 10.222.60.0/24 metric 10 up
    ...
   Authentication data: 17 bytes
```

TLV 135 *wide* metrics

© 2020 Juniper Networks, Inc. All Rights Reserved.

Extended IP Reachability TLV

Each IS-IS router configured to support traffic engineering advertises its local IP prefix information using this TLV. Both locally connected and nonnative IS-IS routes use this TLV for reachability information, with no concept of internal or external metrics. The metric, prefix length, prefix, and sub-TLV fields are repeated for each advertised address.

The metric field uses 32 bits (*wide* metrics) for each advertised prefix, which results in possible metrics between 0 and 4,294,967,295. The extended IP reachability TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 135.
- *TLV length (1 byte)*: The length of the remaining fields in the TLV is placed in this field. This length allows the router to determine the number of prefixes advertised within the TLV.
- *Metric (4 bytes)*: The metric of the advertised prefix is encoded in this field.

The following list details the remaining fields of the extended IP reachability TLV:

- *Prefix length (1 byte)*: The first bit in this field is the U/D bit, which is used to allow prefixes to be advertised across a level boundary and to prevent routing loops. The second bit in this field is the sub bit, which denotes if any optional sub-TLVs are associated with the advertised prefix. The final 6 bits represent the length of the advertised prefix.
- *Prefix (variable)*: The advertised prefix is placed in this field.
- *Optional sub-TLV type* (1 byte): The type of the sub-TLV is encoded in this field. Junos OS currently supports only one sub-TLV type, which is a 32-bit route tag with a type code of 1.
- *Optional sub-TLV length (1 byte)*: The length of the remaining fields in the sub-TLV is placed in this field.
- *Optional sub-TLV (variable)*: For the supported route tag sub-TLV, the 32-bit tag value is placed in this field. IS-IS route tagging offers the same administrative filtering capabilities as the OSPF protocol. IS-IS traffic engineering extensions, which enable TLV 135, must be enabled to support route tagging. TE extensions are enabled for IS-IS by default.

TLV 137—Dynamic Hostname

- Advertises the ASCII hostname of the router to the network
  - (1-byte) TLV type
  - (1-byte) TLV length
  - (Variable) Hostname

```
user@host> show isis database extensive | find TLV
TLVs:
    Area address: 49.0001 (3)
    Speaks: IP
    Speaks: IPv6
    IP router id: 192.168.48.1
    IP address: 192.168.48.1
    Hostname: SaoPaulo
```

TLV 137

© 2020 Juniper Networks, Inc. All Rights Reserved.

Dynamic Host Name TLV

All IS-IS routers advertise their configured hostname into the network using this TLV. This advertisement allows network operators to view information about the network routers using a symbolic name instead of the hexadecimal system ID in various **show** commands. The dynamic hostname TLV contains the following fields:

- *TLV type (1 byte)*: The type of the TLV is encoded in this field with a constant value of 137.
- *TLV length (1 byte)*: The length of the hostname field is placed in this field. Possible values range from 1 to 255 octets.
- *Hostname (Variable)*: The locally configured ASCII hostname of the router is placed in this field.

Agenda: IS-IS

- Overview of IS-IS
- IS-IS PDUs
- →Neighbors and Adjacencies
- Configuring and Monitoring IS-IS

© 2020 Juniper Networks, Inc. All Rights Reserved.

Neighbors and Adjacencies

The slide highlights the topic we discuss next.

Neighbors and Adjacencies

- IS-IS adjacency rules:
  - Level 1 routers never form an adjacency with a Level 2 router
    - The reverse is also true
  - For Level 1 adjacencies, area IDs must be the same
  - For Level 2 adjacencies, area IDs can be different

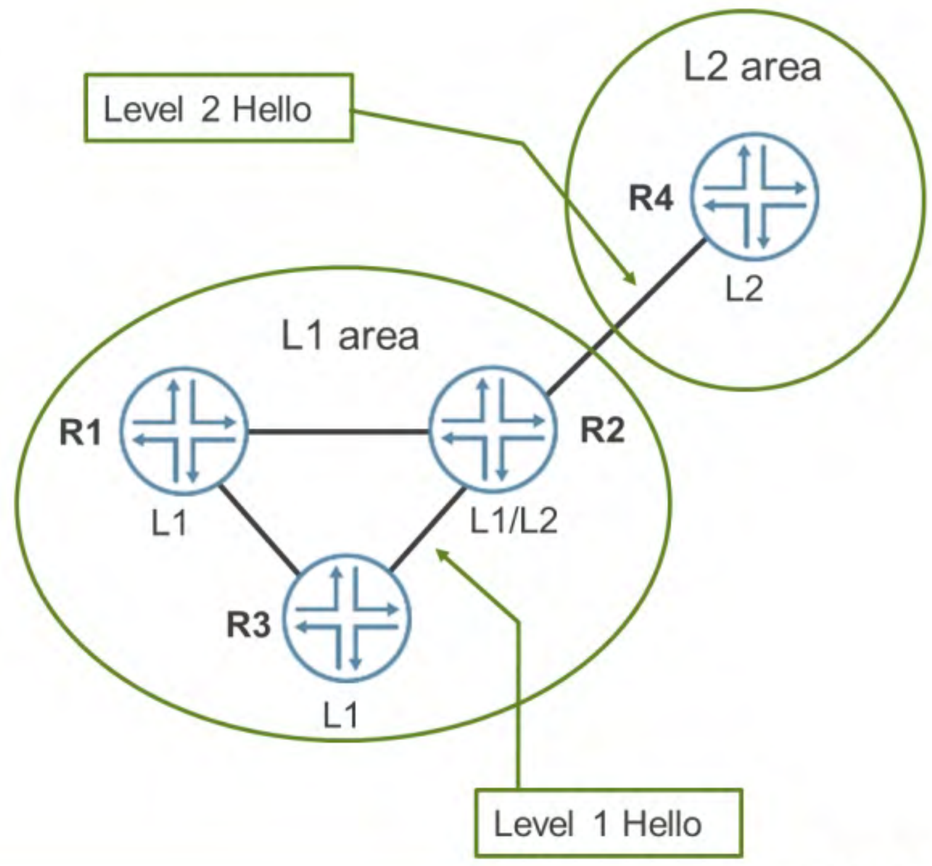

© 2020 Juniper Networks, Inc. All Rights Reserved.

Neighbors and Adjacencies

The slide lists the rules Level 1 and Level 2 routers must follow when forming an adjacency.

DIS

- IS-IS elects a DIS on broadcast and multiaccess networks
  - In a priority tie, the system with the highest SNPA or MAC address wins the DIS election
  - Separate DIS is elected for L1 and L2 (could be the same router)
- DIS characteristics:
  - DIS acts as the representative of the pseudo-node and advertises the pseudo-node to all attached routers
  - No backup DIS in IS-IS
  - Election is deterministic (unlike OSPF)

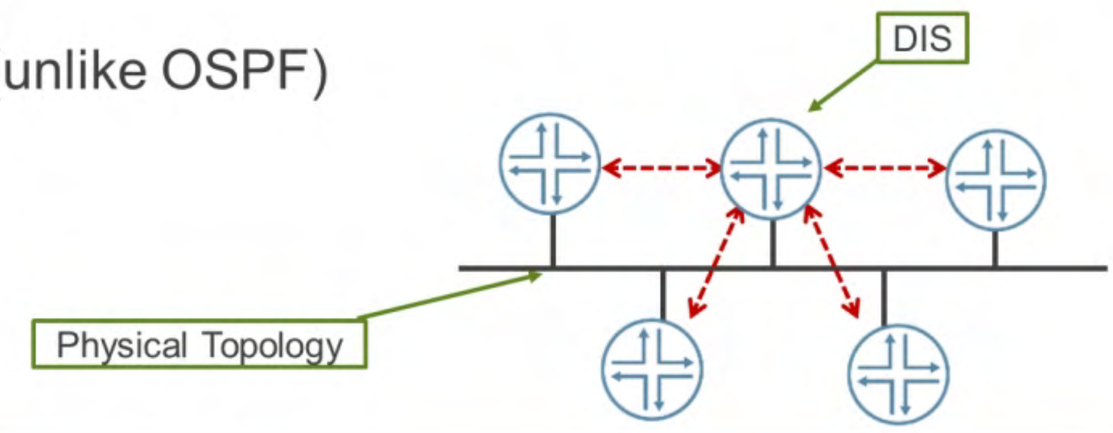

© 2020 Juniper Networks, Inc. All Rights Reserved.

DIS Election

The IS-IS DIS election process is achieved by assigning a Level 1 priority and a Level 2 priority on every IS-IS router interface, with a range of 0 through 127. Junos OS uses a default priority of 64 for both levels. The router advertises its priority in the hello PDUs sent from each interface. The L1 priority is advertised in L1 hello PDUs, and the L2 priority is advertised in L2 hello PDUs. If the priority is 0, the router is ineligible to become the DIS. Interfaces to nonbroadcast networks automatically have a priority of 0. The router with the higher priority value becomes the DIS. In the event of a tie, the router with the numerically highest subnetwork point of attachment (SNPA), which is a fancy name for a MAC address, wins the election.

DIS Characteristics

Unlike OSPF, however, an IS-IS router attached to a broadcast, multiaccess network establishes adjacencies with all of its neighbors on the network, not just the DIS. Each router multicasts its LSPs to all of its neighbors, and the DIS uses a system of PDUs—called sequence number PDUs—to ensure that the flooding of LSPs is reliable. Also, unlike OSPF, no election of a backup designated router occurs in IS-IS. If the IS-IS DIS fails, a new DIS is elected. Another characteristic is that if a new router with a higher priority than the existing DIS becomes active, or if the new router has an equal priority and a higher subnetwork point of attachment (MAC address), the new router becomes the DIS. When the DIS changes, a new set of LSPs must be flooded.

Pseudo-Node

- An IS-IS network running on a broadcast segment is considered a single router (called a pseudo-node)
  - Each router advertises a single link to the pseudo-node, including the DIS
  - Each router forms an adjacency with each of its neighbors on a broadcast, multi-access network

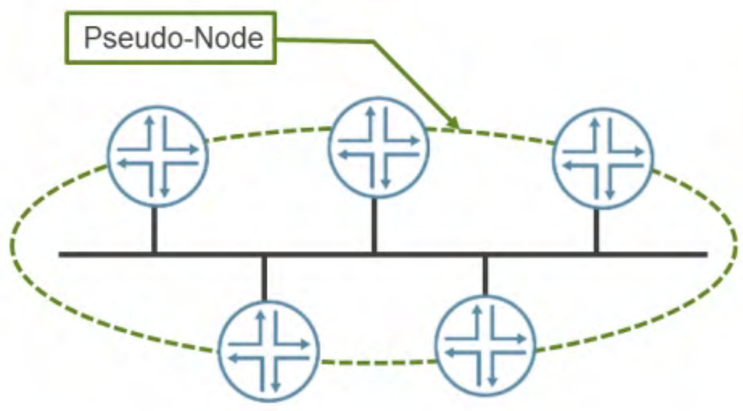

The pseudo-node is considered one router with regard to SPF calculations. The pseudo-node has a cost of zero and is **not** considered an extra hop in the SPF calculation.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Pseudo-Node

IS-IS elects a DIS on broadcast and multiaccess networks for the same reason as OSPF. Rather than having each router connected to the LAN advertise an adjacency with every other router connected to the LAN, the network itself is considered a router—a pseudo-node. Each router, including the DIS, advertises a single link to the pseudo-node. The DIS also advertises, as the representative of the pseudo-node, a link to all of the attached routers.

Troubleshooting IS-IS Adjacencies

- If no adjacency exists, check for the following:
  - Physical Layer and Data Link Layer connectivity
  - Mismatched areas (if L1 router) and levels
  - Failure to support minimum protocol MTU of 1492
  - Lack of IP configuration on interfaces
  - Lack of, or malformed, ISO-NET
    - No NET configured

© 2020 Juniper Networks, Inc. All Rights Reserved.

IP Configuration Necessary

You can check several things while troubleshooting IS-IS adjacency problems. When establishing adjacencies in IS-IS, all routers on a link must agree upon the IP subnet to which they belong, except on point-to-point links, which can be unnumbered or use /32 addressing. Junos OS needs the IP portion of the network to function so that the IS-IS adjacency will work.

Agenda: IS-IS

- Overview of IS-IS
- IS-IS PDUs
- Neighbors and Adjacencies
- →Configuring and Monitoring IS-IS

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configuring and Monitoring IS-IS

The slide highlights the topic we discuss next.

Configuring IS-IS (1 of 2)

- You must include the ISO family on all interfaces on which you want to run IS-IS and a NET on one of the router's interfaces (usually `lo0`)

```
[edit]
   user@router# show interfaces
   ge-0/0/1 {
       unit 0 {
           family inet {
               address 10.0.24.1/24;
           }
           family iso;
       }
   }
   lo0 {
       unit 0 {
           family inet {
               address 192.168.2.1/32;
           }
           family iso {
               address 49.0001.0192.0168.0201.00;
           }
       }
   }
```

Necessary to run IS-IS on an interface even though a NET address is not configured for this interface

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configuring IS-IS, Part 1

For IS-IS to run on the router, you must enable IS-IS on the router, configure a NET on one of the router's interfaces (preferably the loopback interface, `lo0`), and configure the ISO family on all interfaces on which you want IS-IS to run.

Junos OS supports the assignment of multiple ISO NETs to the router's loopback interface. Such a configuration might prove helpful when migrating two previously independent IS-IS domains into a single routing domain.

Configuring IS-IS (2 of 2)

```
[edit protocols]
user@router# set isis interface ge-0/0/1.0 level 1 disable
[edit protocols]
user@router# set isis interface ge-0/0/1.0 point-to-point
[edit protocols]
user@router# set isis interface lo0.0 level 2 disable

[edit protocols]
user@router# show
isis {
    interface ge-0/0/1.0 {
        point-to-point;
        level 1 disable;
    }
    interface lo0.0 {
        level 2 disable;
    }
}
```

You can use the point-to-point statement to configure a LAN interface to act like a point-to-point interface for IS-IS. It has no effect if configured on an interface that is already point-to-point.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Configuring IS-IS, Part 2

By default, all IS-IS interfaces are Level 1 and Level 2 interfaces. You might need to disable a particular level on a given interface. Notice that configuring IS-IS on Junos devices requires minimal configuration.

Reference Bandwidth

- You can change the interface cost to use the formula *reference-bandwidth/bandwidth*
  - Automatically alters the cost of interfaces
  - Allows for a consistent change across all interfaces
- Use the **`reference-bandwidth`** command within the `[edit protocols isis]` hierarchy

```
[edit protocols isis]
user@router# set reference-bandwidth 1g

[edit protocols isis]
user@router# show
isis {
    reference-bandwidth 1g;
    interface ge-0/0/0.0;
    interface ge-0/1/0.0 {
        level 2 metric 10;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Change the Cost Calculation

Although you can statically assign each interface with a cost, this process often proves to be an administrative hassle in all but the smallest networking environments. However, an option is available to administrators who want an automatic calculation of metrics in their networks.

The solution is to enable a bandwidth calculation similar to that used in OSPF. This calculation takes a supplied value—the reference bandwidth—and divides it by the bandwidth of the physical interface. This solution not only allows for an automatic cost assignment, but it also maintains a consistent ratio across all the router interfaces.

Reference Bandwidth

To enable the calculation, use the **`reference-bandwidth`** command within the `[edit protocols isis]` configuration hierarchy. The value entered has a value in bits per second. Use the standard Junos shortcut notations of *k* (kilobits per second), *m* (megabits), and *g* (gigabits). The entered value becomes the numerator value in the bandwidth calculation.

IS-IS Metric

- Metric of an interface indicates the overhead required to send packets out that interface
- Default IS-IS metric for all links is 10
  - Includes passive interfaces
  - Exception is the loopback interface
    - Default metric is 0
- Can set metric on a per-interface basis
  - Each level on an interface can also have a different metric

```
[edit protocols]
isis {
    interface ge-0/0/0.0 {
        level 2 metric 10;
    }
    interface ge-1/0/0.0 {
        level 1 metric 5;
    }
    interface lo0.0 {
        level 1 metric 5;
    }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Metric Cost of an Interface

Each IS-IS router must determine the metric (or cost) associated with that network. Often referred to as simply the metric, the cost is simply what the router views as the *overhead* associated with transmitting a packet on that interface. Each router advertises these cost values in its LSPs; each router can determine the total cost (or metric) to reach any destination in the network.

IS-IS Default Metric

Each router in as IS-IS network uses a default metric of 10 for all interfaces on the router. The notable exception to this default setting is the loopback interface of the router, which has a default metric of 0.

When you configure an IS-IS interface to operate in passive mode, the default metric is still set to a value of 10. This default is different from the operation of other router vendors, so be careful in a multivendor environment to ensure a consistent metric calculation.

Set on a Per-Interface Basis

If you want to alter the metrics in the network, you can set the cost for each interface. Within the interface portion of the `[edit protocols isis]` configuration hierarchy, the **`metric`** command assigns a permanent cost to that interface for a particular level. Each interface on the router can have a separate cost assigned to it.

Monitoring IS-IS Operation

- Various **show** commands exist to provide detailed information on the operation of IS-IS
  - **show isis interface**
  - **show isis adjacency**
  - **show isis spf log**
  - **show isis statistics**
  - **show isis route**
  - **show isis database extensive**

© 2020 Juniper Networks, Inc. All Rights Reserved.

Monitoring IS-IS Operation

This slide lists the **show** commands discussed in the following pages.

Displaying IS-IS Interface Status

- Use the **`show isis interface`** command to display the IS-IS parameters associated with an interface

```
user@R2> show isis interface ?
Possible completions:
  <[Enter]>            Execute this command
  <system-id>          System ID of intermediate system
  brief                Display brief output (default)
  detail               Display detailed output
  extensive            Display extensive output
  instance             Name of IS-IS instance
  logical-system       Name of logical system, or 'all'
```

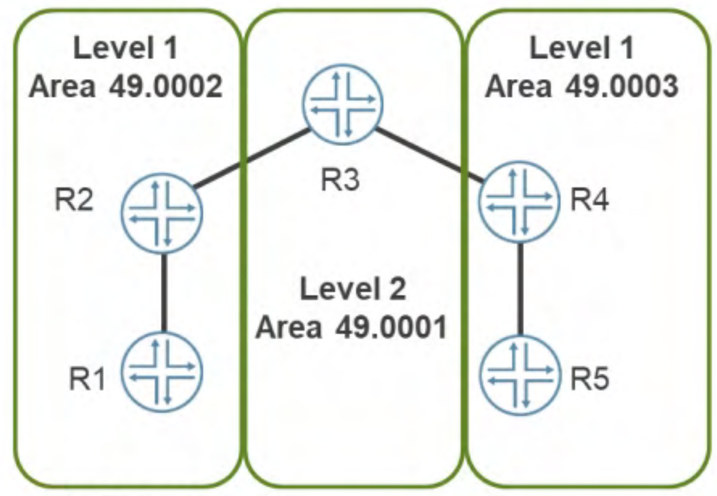


```
user@R2> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR  Level 2 DR        L1/L2 Metric
ge-1/0/5.0            3   0x1 R2.00       R3.02                 10/10
ge-1/1/4.0            1   0x2 R2.02       Disabled              10/10
lo0.2                 0   0x1 Passive     Disabled               0/0
```

"3" indicates both a Level 1 and Level 2 interface. This one is between the R2 and R3 routers.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying Interface Status

The **`show isis interface`** command displays the status of an interface. Use this command to ensure that IS-IS is configured correctly on the router. The output of this command includes the following fields:

- *`interface-name`* (detail output only): Displays the name of the interface.
- `Index` (detail output only): Displays the interface index assigned by the Junos kernel.
- `State` (detail output only): Displays the internal implementation information.
- `Circuit ID` (detail output only): Displays the circuit identifier.
- `Circuit type` (detail output only): Displays the circuit type, which can be `1`—Level 1 only, `2`—Level 2 only, or `3`—Level 1 and Level 2.
- `LSP interval` (detail output only): Displays the interface's link-state PDU interval.
- `Sysid` (detail output only): Displays the system identifier.
- `Interface` (brief output only): Displays the interface through which the adjacency is made.
- `Level 1 DR/Level 2 DR` (brief output only): Displays the Level 1 or Level 2 designated intermediate system.
- `L1/L2 Metric`: Displays the interface's metric for Level 1 and Level 2. If no information exists, the metric is 0.
- `Adjacencies` (detail output only): Displays the number of adjacencies established on this interface.
- `Priority` (detail output only): Displays the priority value for this interface.
- `Metric` (detail output only): Displays the metric value for this interface.
- `Hello (s)` (detail output only): Displays the interface's hello interval.
- `Hold (s)` (detail output only): Displays the interface's hold time.

Displaying IS-IS Adjacency Status

- Use the **`show isis adjacency`** command to display IS-IS parameters adjacency status

```
user@R2> show isis adjacency ?
Possible completions:
  <[Enter]>            Execute this command
  <system-id>          System ID of intermediate system
  brief                Display brief output (default)
  detail               Display detailed output
  extensive            Display extensive output
  instance             Name of IS-IS instance
  logical-system       Name of logical system, or 'all'
```

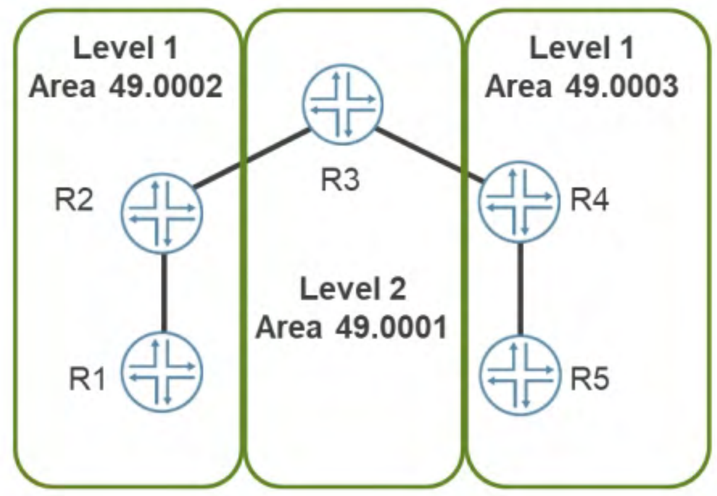


```
user@R2> show isis adjacency
IS-IS interface database:
Interface             System   L  State      Hold (secs) SNPA
ge-1/0/5.0            R3       2  Up                8    80:71:1f:c3:3:7d
ge-1/1/4.0            R1       1  Up               20    80:71:1f:c3:3:64
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying IS-IS Adjacency Status

The **`show isis adjacency`** command displays the status of IS-IS adjacencies. The output of this command includes the following fields:

- `Interface`: Displays the interface through which the neighbor is reachable.
- `System` (brief output only): Displays the system identifier (sysid), printed as a name if possible.
- `L`: Displays the level, which can be `1`—Level 1 only; `2`—Level 2 only; or `3`—Level 1 and Level 2. An exclamation point (`!`) preceding the level number indicates that the adjacency is missing an IP address.
- `State`: Displays the state of the adjacency. It can be `Up`, `Down`, `New`, `One-way`, `Initializing`, or `Rejected`.
- `Hold (secs)` (brief/standard output only): Displays the remaining hold time of the adjacency. Note that the **`show isis adjacency`** command returns brief output by default.
- `SNPA` (brief output only): Displays the subnetwork point of attachment (MAC address of the next hop).

Displaying Detailed Adjacency Information

- Display detailed IS-IS adjacency information with the `detail` switch

```
user@R2> show isis adjacency detail
IS-IS adjacency database:
R3
  Interface: ge-1/0/5.0, Level: 2, State: Up, Expires in 7 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 00:35:44 ago
  Circuit type: 3, Speaks: IP, IPv6, MAC address: 80:71:1f:c3:3:7d
  Topologies: Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: R3.02, IP addresses: 10.0.2.2

R1
  Interface: ge-1/1/4.0, Level: 1, State: Up, Expires in 22 secs
  Priority: 64, Up/Down transitions: 1, Last transition: 00:35:42 ago
  Circuit type: 1, Speaks: IP, IPv6, MAC address: 80:71:1f:c3:3:64
  Topologies: Unicast
  Restart capable: Yes, Adjacency advertisement: Advertise
  LAN id: R2.02, IP addresses: 10.0.1.1
```

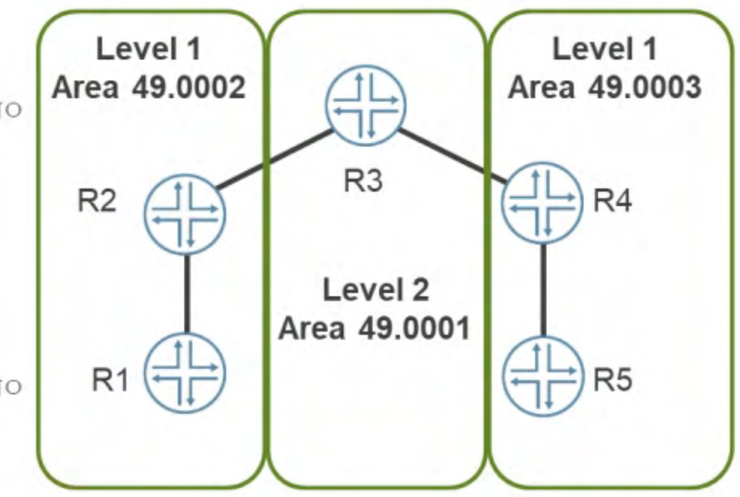

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying Detailed Adjacency Information

The **`show isis adjacency detail`** command displays detailed IS-IS adjacency information. The output of this command includes the following fields:

- `Expires in` (detail output only): Displays how long until the adjacency expires, in seconds;
- `Priority` (detail output only): Displays the priority to become the DIS;
- `Up/Down transitions` (detail output only): Displays the count of adjacency status changes from up to down or from down to up;
- `Last transition` (detail output only): Displays the time of the last up/down transition;
- `Circuit type` (detail output only): Displays the bit mask of levels on this interface, which can be `1`—Level 1 router, `2`—Level 2 router, or `1/2`—both Level 1 and Level 2 routers;
- `Speaks` (detail output only): Displays the protocols supported by this neighbor; and
- `IP addresses` (detail output only): Displays the IP address of this neighbor.

Displaying SPF Log

- Display information about SPF operation with the **`show isis spf log`** command

```
user@R2> show isis spf log
IS-IS level 1 SPF log:
Start time          Elapsed (secs) Count Reason
Mon Aug 15 23:06:42        0.003862     3 Reconfig
Mon Aug 15 23:06:42        0.000559     2 Updated LSP R2.00-00
Mon Aug 15 23:06:57        0.000046     2 Multi area attachment change
Mon Aug 15 23:06:58        0.000087     5 New adjacency R1 on ge-1/1/4.0
Mon Aug 15 23:17:59        0.000102     1 Reconfig
Mon Aug 15 23:20:32        0.000107     1 Reconfig
Mon Aug 15 23:21:12        0.000090     1 Reconfig
Mon Aug 15 23:22:17        0.000089     1 Reconfig
Mon Aug 15 23:36:12        0.000103     1 Periodic SPF
Mon Aug 15 23:48:27        0.000113     1 Periodic SPF
Tue Aug 16 00:00:56        0.000110     1 Periodic SPF
```

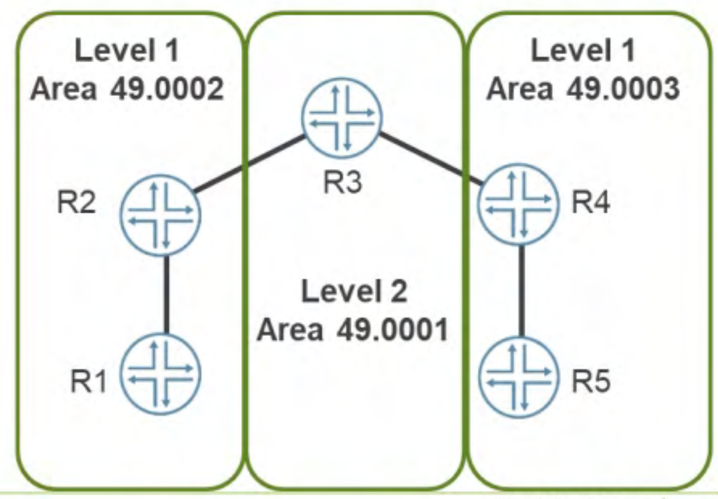

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying SPF Operation

The **`show isis spf log`** command displays the elapsed time to perform SPF calculations and the reasons why they were triggered. The output fields of this command are the following:

- `Node`: Displays the sysid of a node;
- `Metric`: Displays the metric to the node;
- `Interface`: Displays the interface of the next hop;
- `Via`: Displays the sysid of the next hop;
- `SNPA`: Displays the subnetwork point of attachment (MAC address of the next hop);
- `Start time` (log output only): Displays the time that the SPF computation started;
- `Elapsed time` (log output only): Displays the length of time required to complete the SPF computation in seconds;
- `Count` (log output only): Displays the number of times the SPF was triggered; and
- `Reason` (log output only): Displays the reason that the SPF computation was completed.

Displaying IS-IS Statistics

- Display various IS-IS statistics information with the **`show isis statistics`** command

```
user@R2> show isis statistics
IS-IS statistics for R2:
PDU type        Received    Processed        Drops          Sent        Rexmit
LSP                   38           38            0            25             0
IIH                 1935         1935            0          2456             0
CSNP                 432          432            0           428             0
PSNP                   0            0            0             0             0
Unknown                0            0            0             0             0
Totals                 0            0            0             0             0

Total packets received: 2405 Sent: 2882

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 24
Fragments rebuilt: 28
LSP regenerations: 9
Purges initiated: 0
```

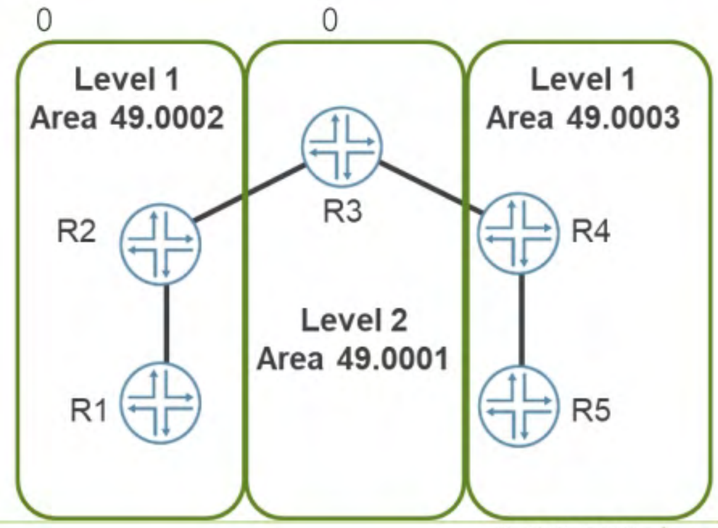

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying IS-IS Statistics

The **`show isis statistics`** command displays statistics about IS-IS traffic. The output fields of this command are the following:

- `PDU type`: Displays the protocol data unit type.
- `Received`: Displays the number of PDUs received since IS-IS started or since the statistics were zeroed.
- `Processed`: Displays the number of PDUs received less the number dropped.
- `Drops`: Displays the number of PDUs dropped.
- `Sent`: Displays the number of PDUs transmitted since IS-IS started or since the statistics were zeroed.
- `Rexmit`: Displays the number of PDUs retransmitted since IS-IS started or since the statistics were zeroed.
- `Total packets received/sent`: Displays the total number of PDUs received and transmitted since IS-IS started or since the statistics were zeroed.
- `SNP queue length`: Displays the number of CSNPs and PSNPs sitting on the SNP queue waiting for processing. This value is almost always 0.
- `LSP queue length`: Displays the number of LSPs sitting on the link-state PDU queue waiting for processing. This value is almost always 0.
- `SPF runs`: Displays the number of SPF calculations performed. If this number is incrementing rapidly, it indicates that the network is unstable.
- `Fragments rebuilt`: Displays the number of link-state PDU fragments that the local system has computed.
- `LSP regenerations`: Displays the number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
- `Purges initiated`: Displays the number of purges that the system initiated. A purge is initiated if the software decides that an link-state PDU must be removed from the network.

Displaying IS-IS Routes

```
user@R2> show isis route
IS-IS routing table                Current version: L1: 11 L2: 13
IPv4/IPv6 Routes
----------------
Prefix              L Version   Metric Type Interface        NH   Via
10.0.3.0/30         2      13       20 int  ge-1/0/5.0       IPV4 R3
10.0.4.0/30         2      13       30 int  ge-1/0/5.0       IPV4 R3
192.168.1.1/32      1      11       10 int  ge-1/1/4.0       IPV4 R1
. . .
user@R2> show route protocol isis

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.3.0/30         *[IS-IS/18] 01:05:37, metric 20
                     > to 10.0.2.2 via ge-1/0/5.0
10.0.4.0/30         *[IS-IS/18] 00:51:45, metric 30
                     > to 10.0.2.2 via ge-1/0/5.0
192.168.1.1/32      *[IS-IS/15] 01:05:35, metric 10
                     > to 10.0.1.1 via ge-1/1/4.0
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying IS-IS Routes

The **`show isis route`** command displays the routes in the IS-IS routing table. The output of this command includes the following fields:

- `Current version`: Displays the number of the current version of the IS-IS routing table.
- `L1`: Displays the version of Level 1 SPF that was run.
- `L2`: Displays the version of Level 2 SPF that was run.
- `Prefix`: Displays the destination of the route.
- `L`: Displays the level, which can be `1`—Level 1 only; `2`—Level 2 only; and `3`—Level 1 and Level 2.
- `Version`: Displays the version (or run) of SPF that generated the route.
- `Metric`: Displays the metric value associated with the route.
- `Type`: Displays the metric type. It can be `int` (internal) or `ext` (external).
- `Interface`: Displays the interface to the next hop.
- `Via`: Displays the system identifier of the next hop, displayed as a name if possible.

Route Preference Review:

- L1 - IS-IS Internal - 15
- L2 - IS-IS Internal - 18
- L1 - IS-IS External - 160
- L2 - IS-IS External - 165

Displaying Detailed IS-IS Database Information

```
user@R2> show isis database extensive
IS-IS level 1 link-state database:

R1.00-00 Sequence: 0x9, Checksum: 0xf332, Lifetime: 403 secs
   IS neighbor: R2.02                    Metric:       10
     Two-way fragment: R2.02-00, Two-way first fragment: R2.02-00
   IP prefix: 10.0.1.0/30                      Metric:       10 Internal Up
   IP prefix: 192.168.1.1/32                   Metric:        0 Internal Up

  Header: LSP ID: R1.00-00, Length: 148 bytes
    Allocated length: 284 bytes, Router ID: 192.168.1.1
    Remaining lifetime: 403 secs, Level: 1, Interface: -1610691560
    Estimated free bytes: 136, Actual free bytes: 136
    Aging timer expires in: 403 secs
    Protocols: IP, IPv6

  Packet: LSP ID: R1.00-00, Length: 148 bytes, Lifetime : 1198 secs
    Checksum: 0xf332, Sequence: 0x9, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 18, Packet version: 1, Max area: 0

  TLVs:
    Area address: 49.0002 (3)
    Speaks: IP
    Speaks: IPV6
    IP router id: 192.168.1.1
    IP address: 192.168.1.1
    Hostname: R1
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Displaying Detailed IS-IS Database Information

The **`show isis database extensive`** command provides detailed output for the contents of the IS-IS LSDB. The output of this command includes the following fields:

- `LSP ID`: Displays the LSP identifier.
- `Sequence`: Displays the sequence number of the LSP.
- `Checksum`: Displays the checksum value of the LSP.
- `Lifetime` (in seconds): Displays the remaining lifetime of the LSP, in seconds.
- `IP prefix` (detail and extensive output only): Displays the prefix advertised by this LSP.
- `IS neighbor` (detail output only): Displays an IS-IS neighbor of the advertising system.
- `Metric` (detail and extensive output only): Displays the metric of the prefix or neighbor.

The command output then displays detailed packet content information.

Summary

- In this content, we:
  - Explained the concepts and operation of IS-IS
  - Described various IS-IS LSP types
  - Listed IS-IS adjacency rules and how to troubleshoot common adjacency issues
  - Configured and monitored IS-IS

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Discussed:

- IS-IS concepts and operation;
- IS-IS LSP types;
- IS-IS adjacency rules and troubleshooting; and
- The configuration and monitoring of IS-IS in Junos OS.

Review Questions

1. What is a PDU and what is its purpose?
2. The LSP contains multiple segments. What are they called and what do they do?
3. What happens if an IS-IS enabled router with a higher priority than all other routers is added to a broadcast network?
4. How would you configure only a Level 1 adjacency on a Junos device's interface?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review Questions

1.

2.

3.

4.

Lab: IS-IS Configuration and Monitoring

- Configure, monitor, and troubleshoot the operation of an IS-IS network using multiple levels.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab: IS-IS Configuration and Monitoring

The slide provides the objective for this lab.

Answers to Review Questions

1.

A PDU is a protocol data unit. It is used to send information about the IS-IS configuration between network devices.

2.

The segments are called type/length/value (TLVs). They describe the originating router and the various pieces of information that can be transmitted.

3.

Because IS-IS is deterministic, the added IS will assume the role of the DIS immediately and flood out new PDUs to all other ISs on that segment.

4.

Set `family iso` and `net` on the physical and loopback interfaces. On the loopback interface, add an IP address and an ISO address. Ensure that the Area Address portion of the NET is the same. Under `[protocols isis]`, disable Level 2.

Education Services

Advanced Junos Service Provider Routing

Chapter 7: Advanced IS-IS Operations and Configuration Options

Engineering Simplicity

Objectives

- After successfully completing this content, you will be able to:
  - Display and interpret the IS-IS LSDB
  - Perform advanced IS-IS configuration options
  - Implement IS-IS routing policy

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Will Discuss:

- The display and interpretation of the IS-IS link-state database (LSDB);
- Advanced IS-IS configuration options; and
- The implementation of IS-IS routing policy.

Agenda: Advanced IS-IS Operations and Configurations Options

→IS-IS Operations

- IS-IS Configuration Options
- IS-IS Routing Policy

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Operations

The slide lists the topics we will discuss. We discuss the highlighted topic first.

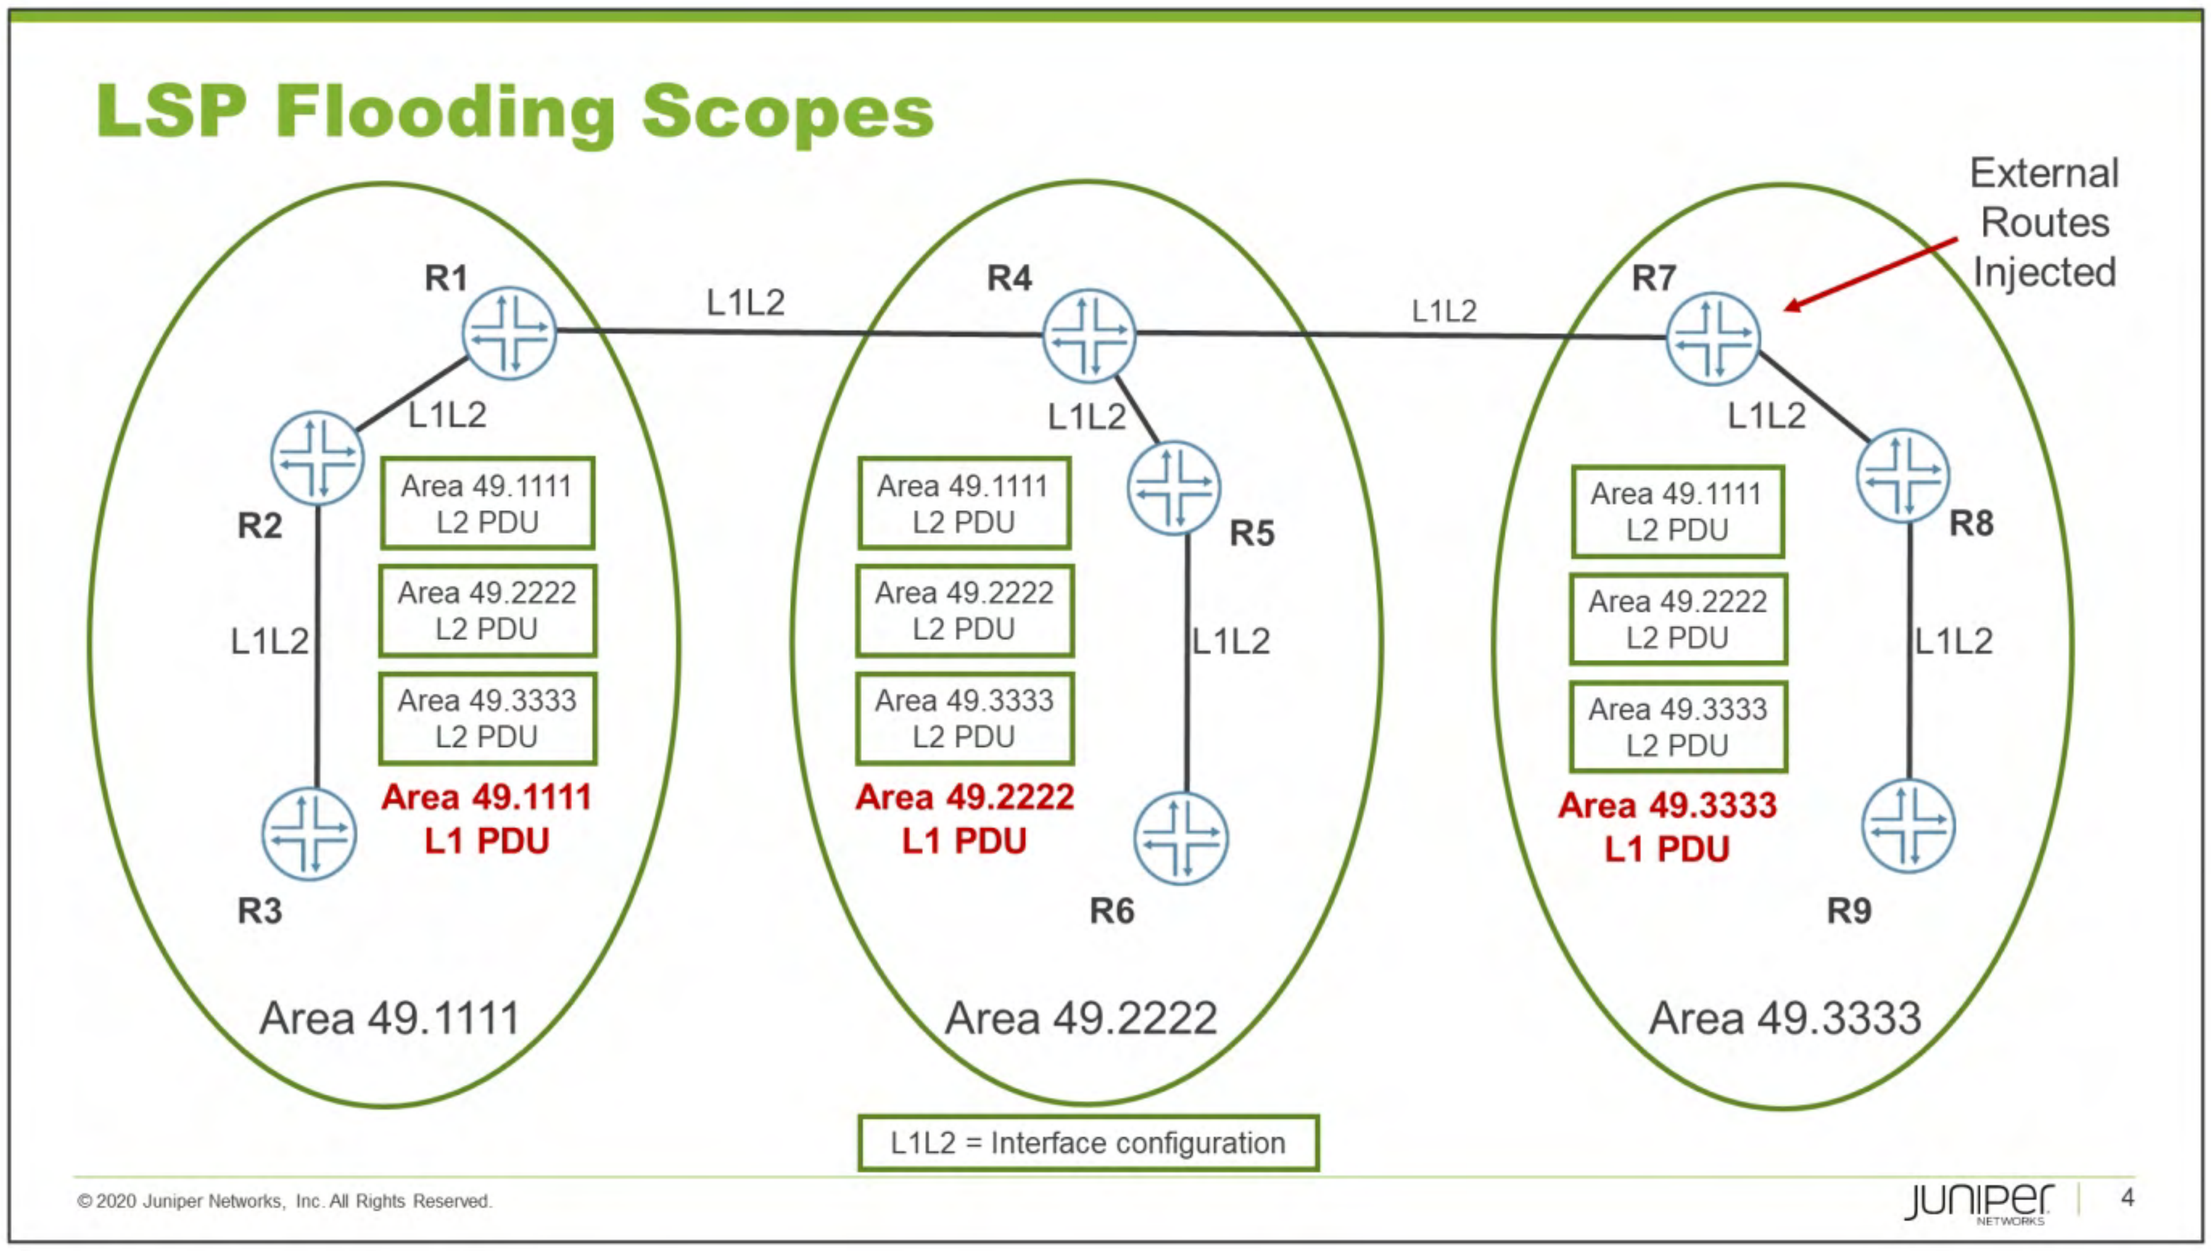

Link-State PDU Flooding Scopes

Each link-state protocol data unit (PDU) in the IS-IS protocol has a specific scope within which it can be flooded. The slide graphically displays these flooding scopes.

Note that the Level 1 link-state PDUs (LSPs) are generated within each area. Because these LSPs have a Level 1 flooding scope, they remain within their own particular area and are not seen in other areas. The L1/L2 router at the edge of the area places the routing information contained within the LSP into a Level 2 LSP and forwards it across the area boundary.

All Level 2 LSPs are flooded across every contiguous Level 2 area. This flooding results in Level 2 LSPs within every area that represents all IS-IS routes. Within Area 49.1111, for example, Level 2 LSPs exist that represent Area 49.2222 and Area 49.3333.

Sample IS-IS Database

```
user@R1> show isis database

IS-IS level 1 link-state database:
LSP ID                      Sequence       Checksum Lifetime Attributes
R1.00-00                  0x8              0xcc42      957 L1 L2 Attached
R2.00-00                  0x9              0xbdfa     1055 L1
R3.00-00                  0x7              0x54d2      500 L1
R3.02-00                  0x4              0xdddb      677 L1
  4 LSPs

IS-IS level 2 link-state database:
LSP ID                     Sequence   Checksum Lifetime Attributes
R1.00-00                   0x6              0xa5a1     1102 L1 L2
R4.00-00                   0x9              0xc92f      909 L1 L2
R5.00-00                   0x6              0xd7d2     1109 L1 L2
  3 LSPs
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Sample IS-IS Database

The example on the slide shows the details of the IS-IS database. To this point, we have been examining small portions of the database through the use of the **`extensive`** keyword. Now, we take a step back and examine the database as a whole. You can gather details by examining the database closely:

- This router has both Level 1 and Level 2 adjacencies. You can determine this fact by the existence of two databases on the router.
- Within the Level 1 area, R1 is the router that can communicate with other Level 2 systems. You can determine this fact by the `Attached` keyword in its Level 1 LSP.
- The R2 and R3 routers are configured to operate at Level 1 only because their attributes are set to 0x01 (not shown in the output). Notice the L1 designation and the absence of a L2 notation within the `Attributes` field.
- The R3 router is the designated intermediate system (DIS) for one of its links in the network. You can determine this fact by the extra LSP advertised as R3.02-00. The circuit ID of the interface upon which it is the DIS is 0x02.

Shortest-Path-First Algorithm

- Based on the Dijkstra algorithm
  - LSDB
  - Candidate database
  - Tree database
- Runs on a per-level basis on each router
  - Independent calculation of the topology
- Results are passed to the routing table
  - Decision as to whether the route is marked active is made by the routing engine.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Dijkstra Algorithm

After a router receives a new LSP and places it into the LSDB, the router runs an algorithm known as the Dijkstra algorithm (or shortest-path-first [SPF] algorithm). This computation uses the database as a data source and results in a loop-free network topology using the best metric from the local router to all nodes in the network.

During the course of this calculation, the algorithm uses three databases—the LSDB, the candidate database, and the tree database. Remember that the LSDB is the total compilation of routing knowledge in the network. Conceptually, it consists of multiple tuples in the form of router ID (RID), neighbor ID, and cost, which describe each link in the network.

When the algorithm operates, the local router moves its own local tuple into the tree database and all tuples for its links into the candidate database. It then performs the following steps until the candidate database is empty:

1. For each new entry in the candidate database, the router determines the cost from root to each neighbor ID. The SPF algorithm moves the tuple with the lowest cost from the candidate database into the tree database. If multiple tuples exist with an equal cost, one is chosen randomly.
2. If a new neighbor ID appears in the tree database, any tuples in the LSDB with a router ID equal to the new tree entry's neighbor ID are moved into the candidate database.
3. Each tuple in the candidate database is evaluated and any tuples whose neighbor ID is currently in the tree database and whose cost to the root is greater than the entry currently in the tree database are deleted. Step 1 is repeated until the candidate database is empty.

Continued on the next page.

Run on a Per-Level Basis (Contd.)

The Dijkstra algorithm is calculated for each LSDB present on the local router. Recall that each IS-IS level maintains a separate copy of the database. These separate copies allow each level to have a separate forwarding topology independent of another level.

Results Are Passed to the Routing Engine

Once the algorithm is run, the routing table on the Routing Engine receives the results in the tree database. At this point, the Routing Engine controls the determination of whether to use any individual IS-IS route to forward traffic. The preference value assigned to each route often handles this choice.

The action of calculating the best IS-IS route prior to the route being placed into the routing table has a consequence in regard to the filtering of routing knowledge. An individual filter (or policy) operates on IP routes that enter the router as IP routes and are placed into the routing table. This form of data flow is not present in IS-IS because the routing information enters the router as an LSP and is only placed into the routing table *after* the router performs the Dijkstra algorithm. Hence, the only method of filtering IS-IS routing knowledge is to keep that information from being placed into the database (on a per-level basis) in the first place.

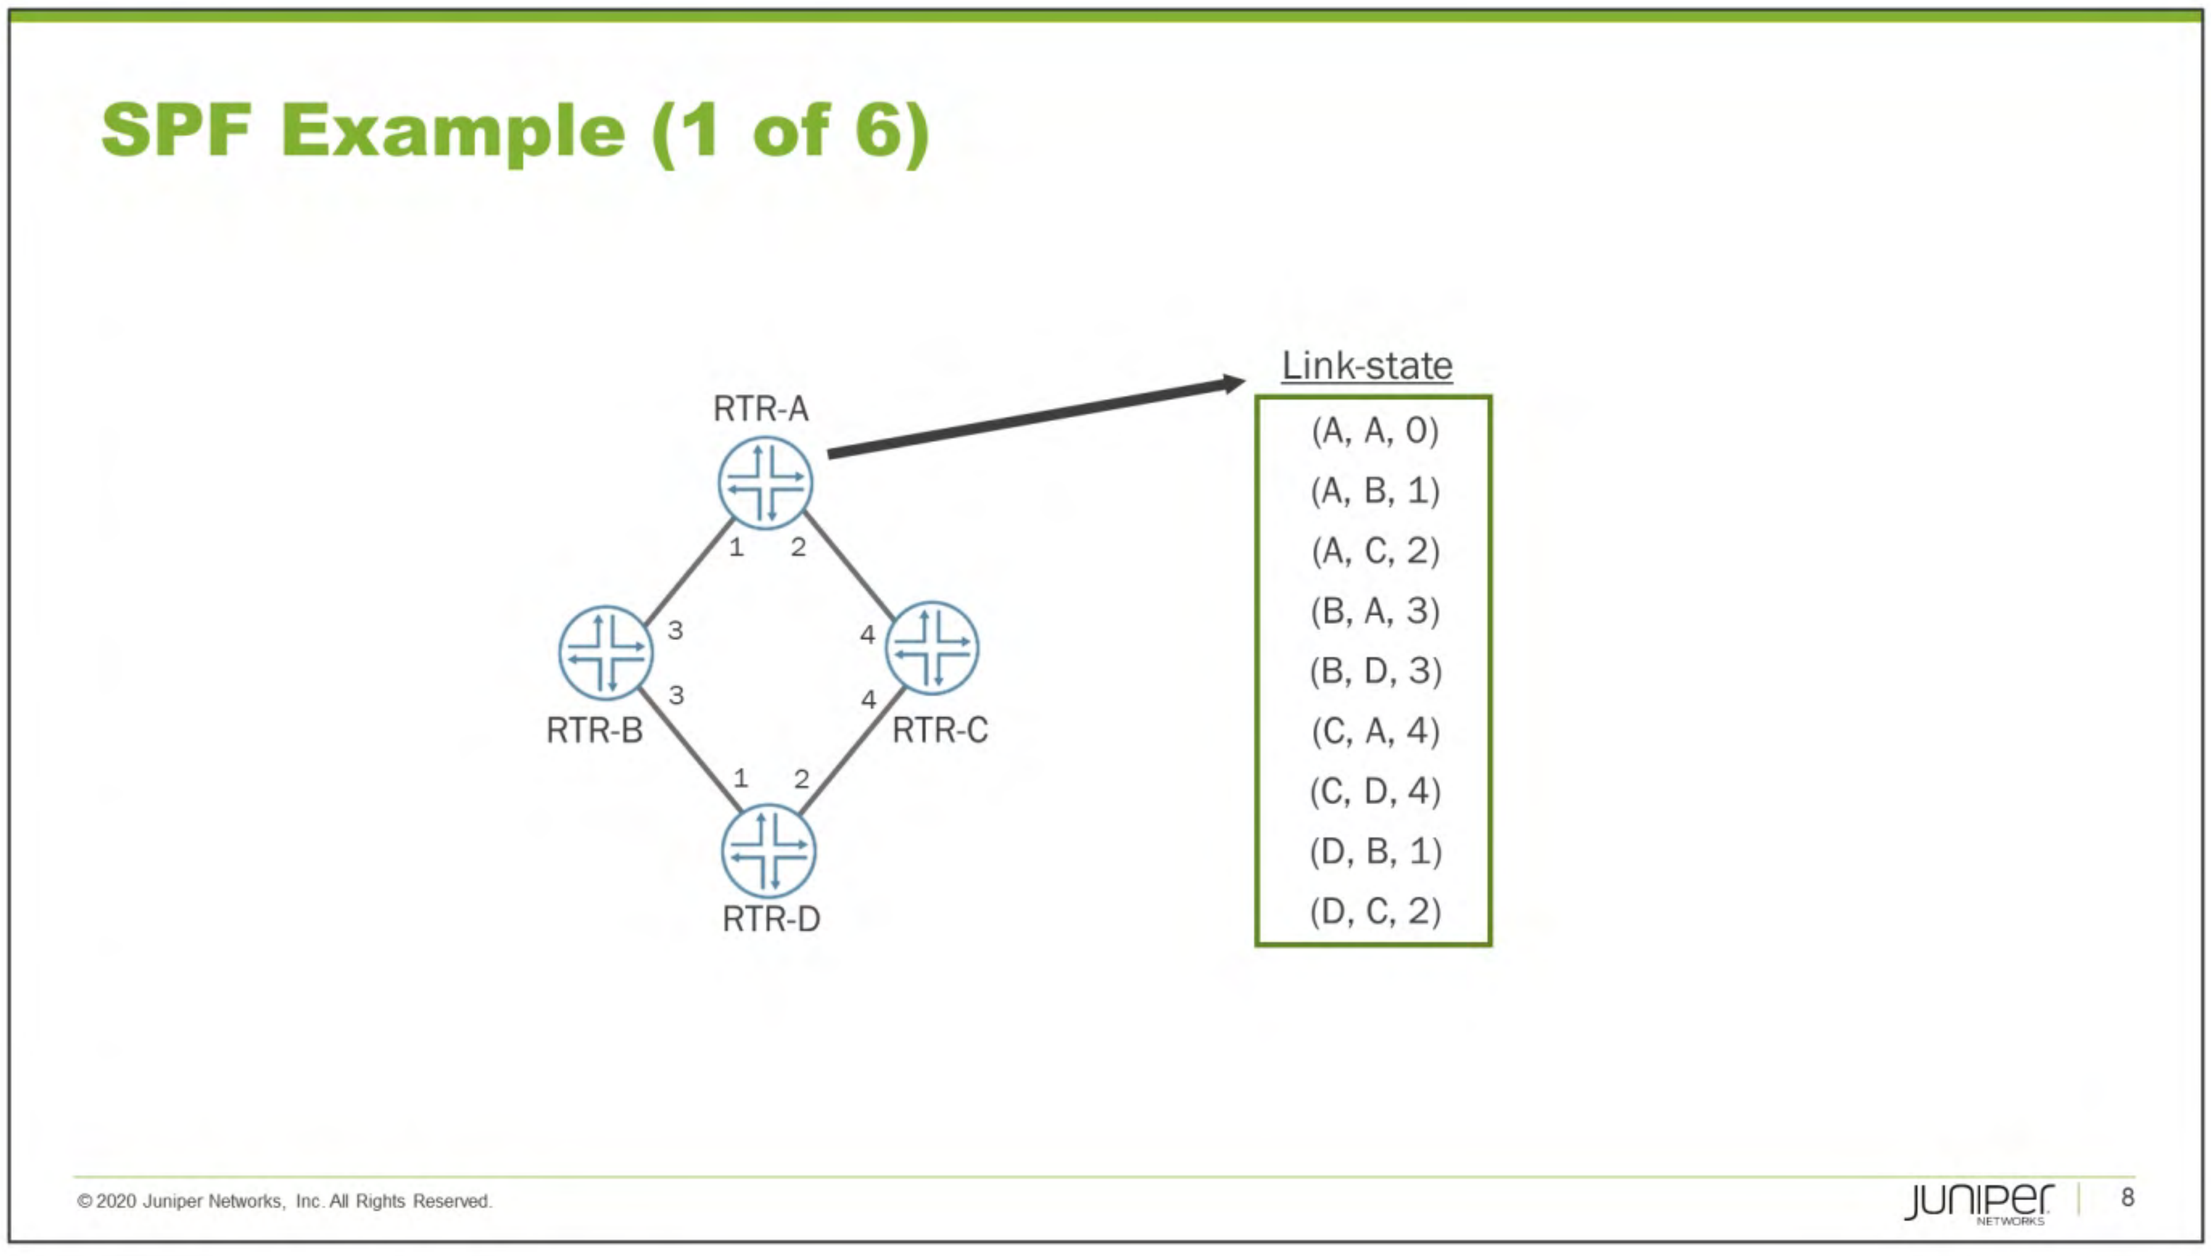

SPF Calculation Example: Part 1

In the following slides, an example SPF calculation is displayed. This graphic shows the beginning state of the network including the routers involved, the configured link metrics, and the LSDB. The network and the LSDB have recently converged and the local router, RTR-A, is running an SPF calculation to determine the shortest path to each node in the network.

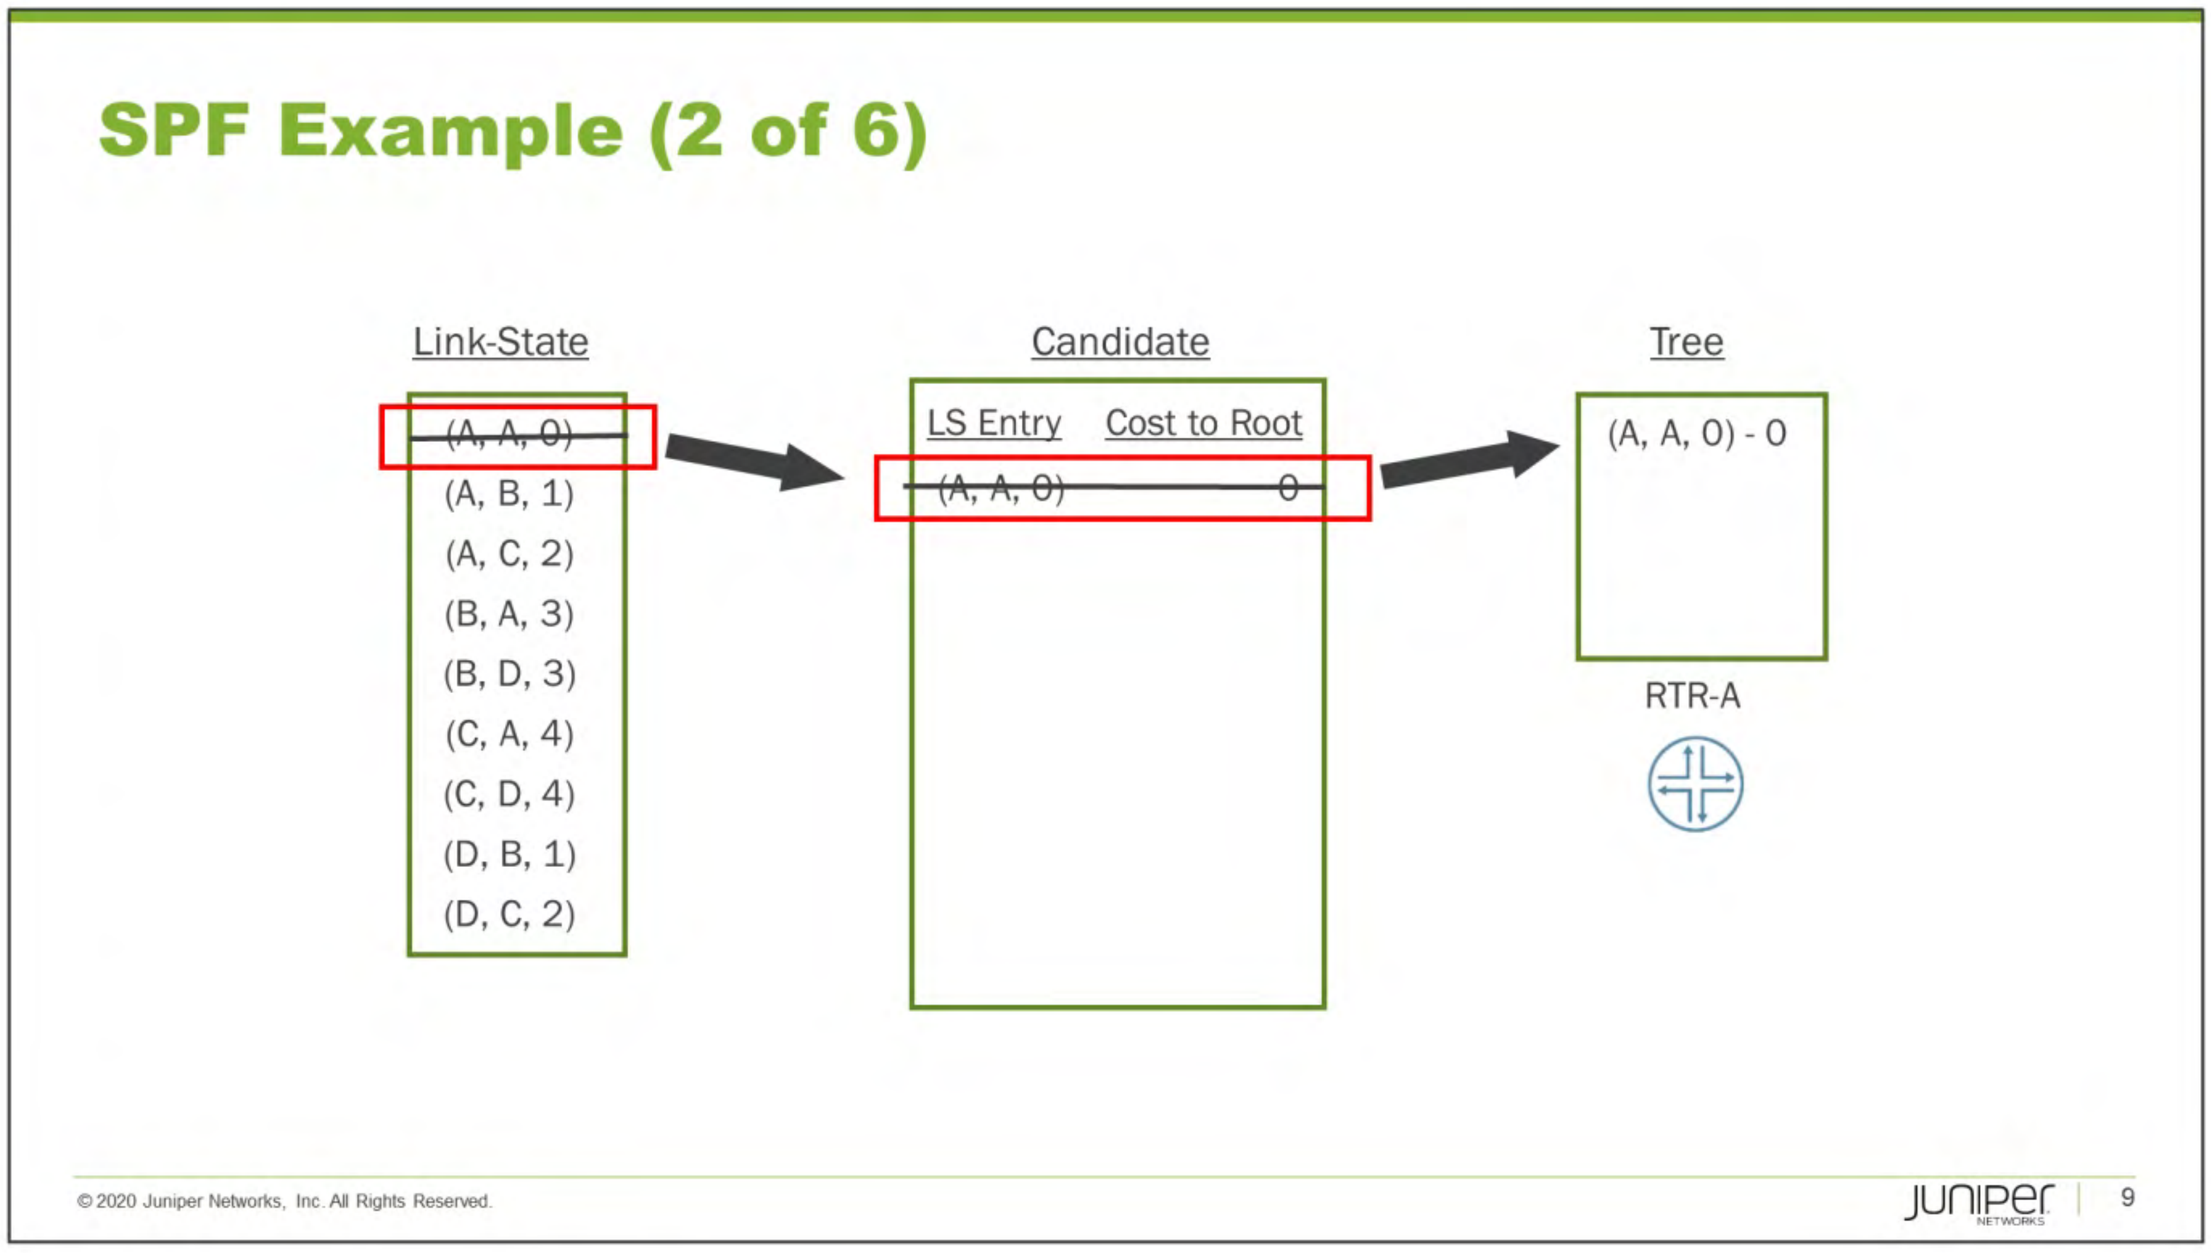

SPF Calculation Example: Part 2

RTR-A begins by moving its own local database tuple (A, A, 0) into the candidate database. The total cost from the neighbor ID to the root is calculated, which results in a 0 value. In other words, RTR-A is directly connected to itself!

The lowest, and only, tuple in the candidate database is moved to the tree database and RTR-A places itself on the network map.

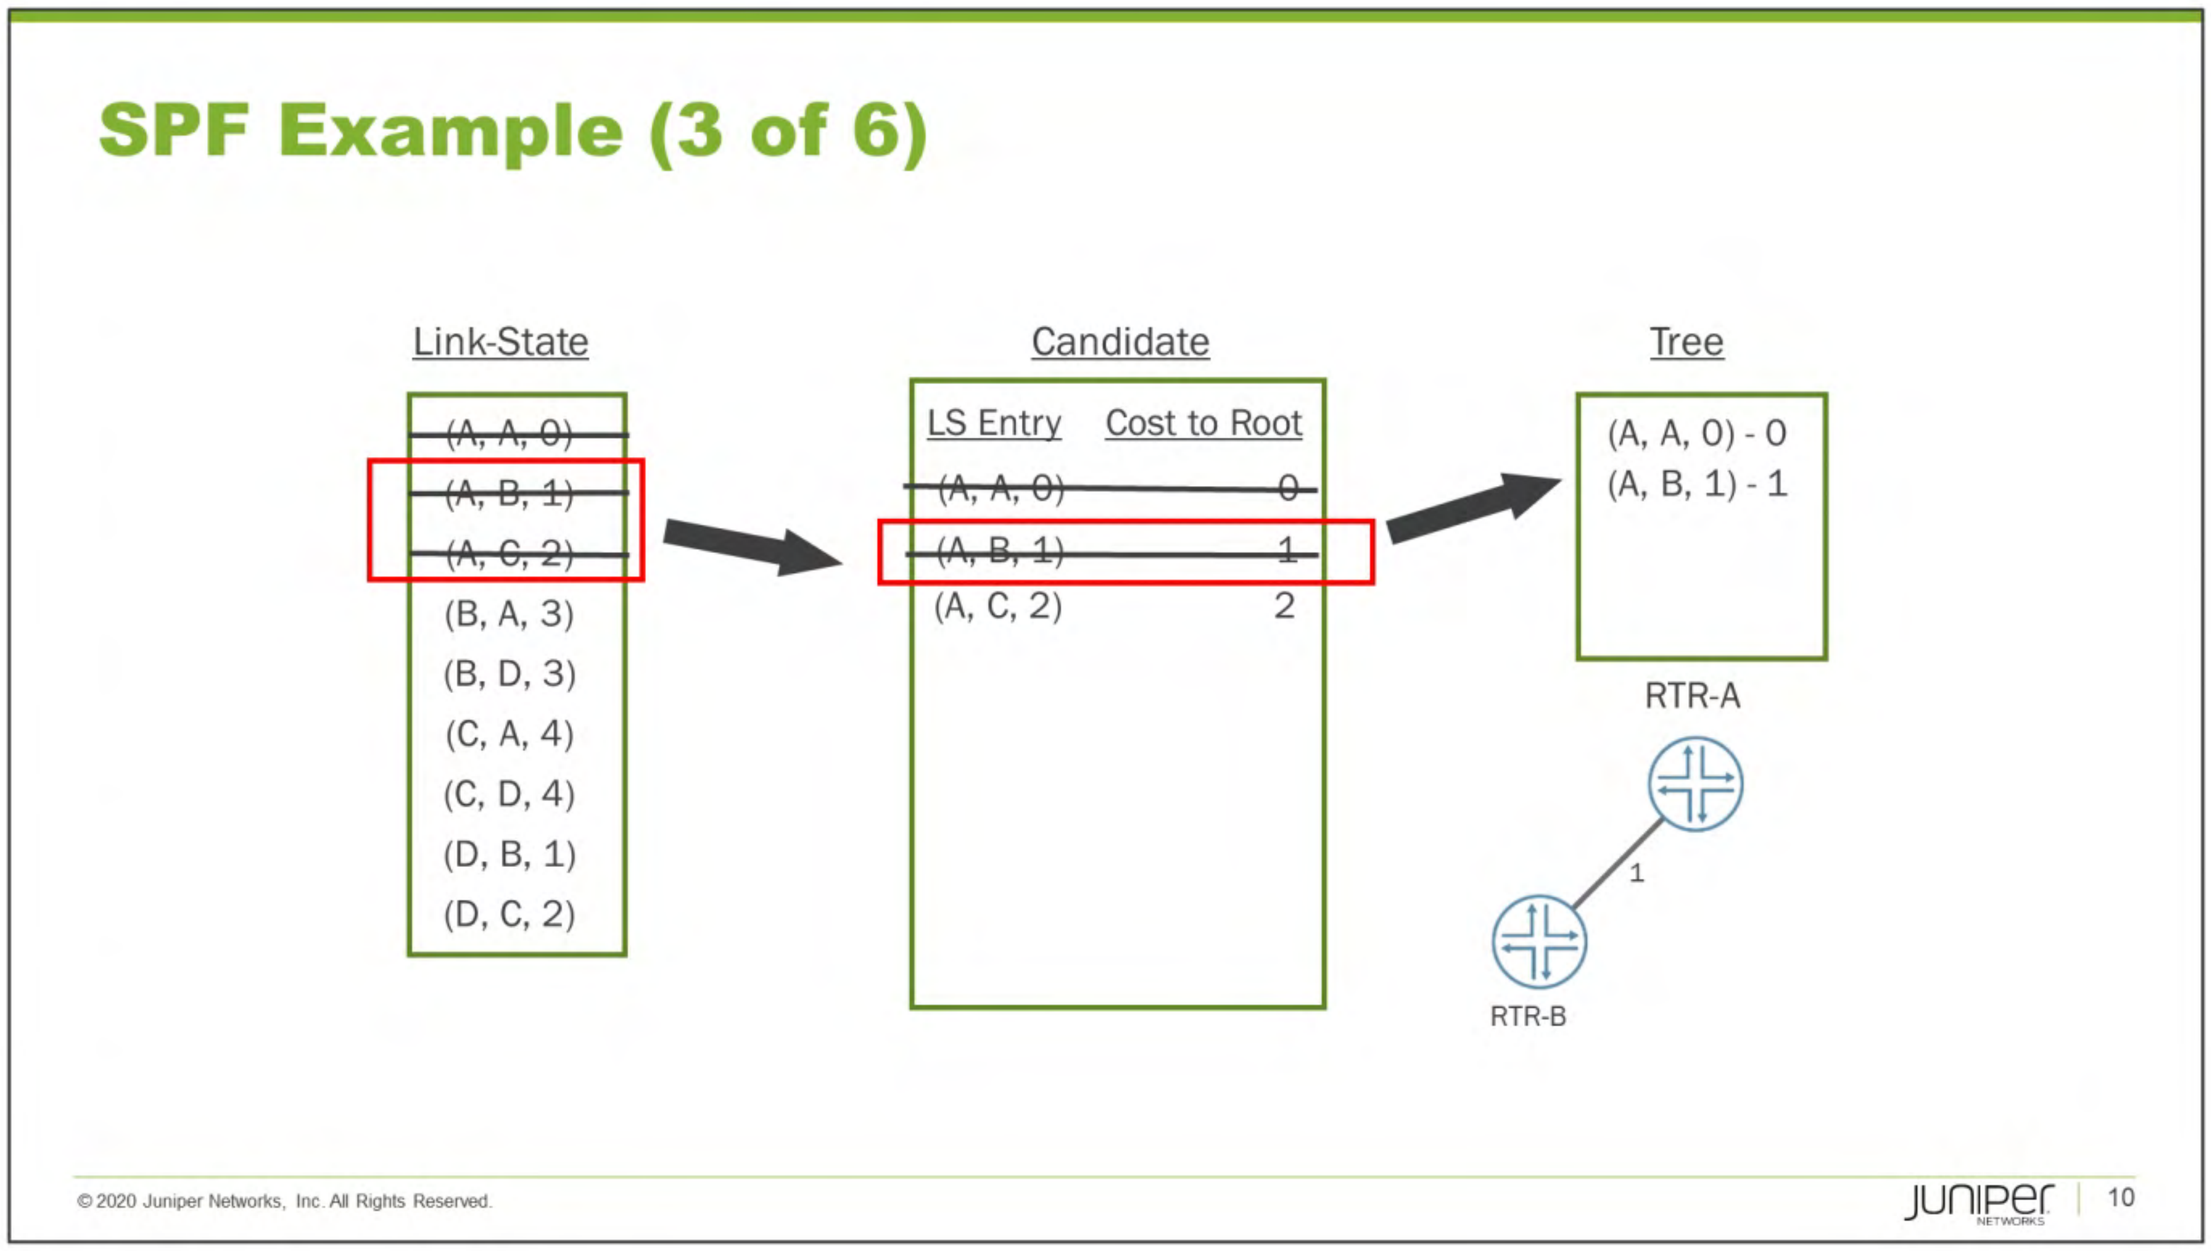

SPF Calculation Example: Part 3

All tuples from the most recent node added to the tree database are now added to the candidate database. Because RTR-A is the most recent entry to the tree database, all of RTR-A's tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are removed from the candidate, of which there are none. (For example, if B were already in the tree database, the tuple (A, B, 1) would be eliminated.)

The cost to each neighbor ID from the root is then calculated. It costs RTR-A 0 to reach itself and 1 to reach RTR-B, so the total cost to RTR-B is 1. The same calculation is done for RTR-C, and the total cost of 2 is placed into the candidate database.

The lowest cost tuple in the candidate database, (A, B, 1), is now moved to the tree database, and RTR-B is placed on the network map.

The candidate database is not empty, so the algorithm continues.

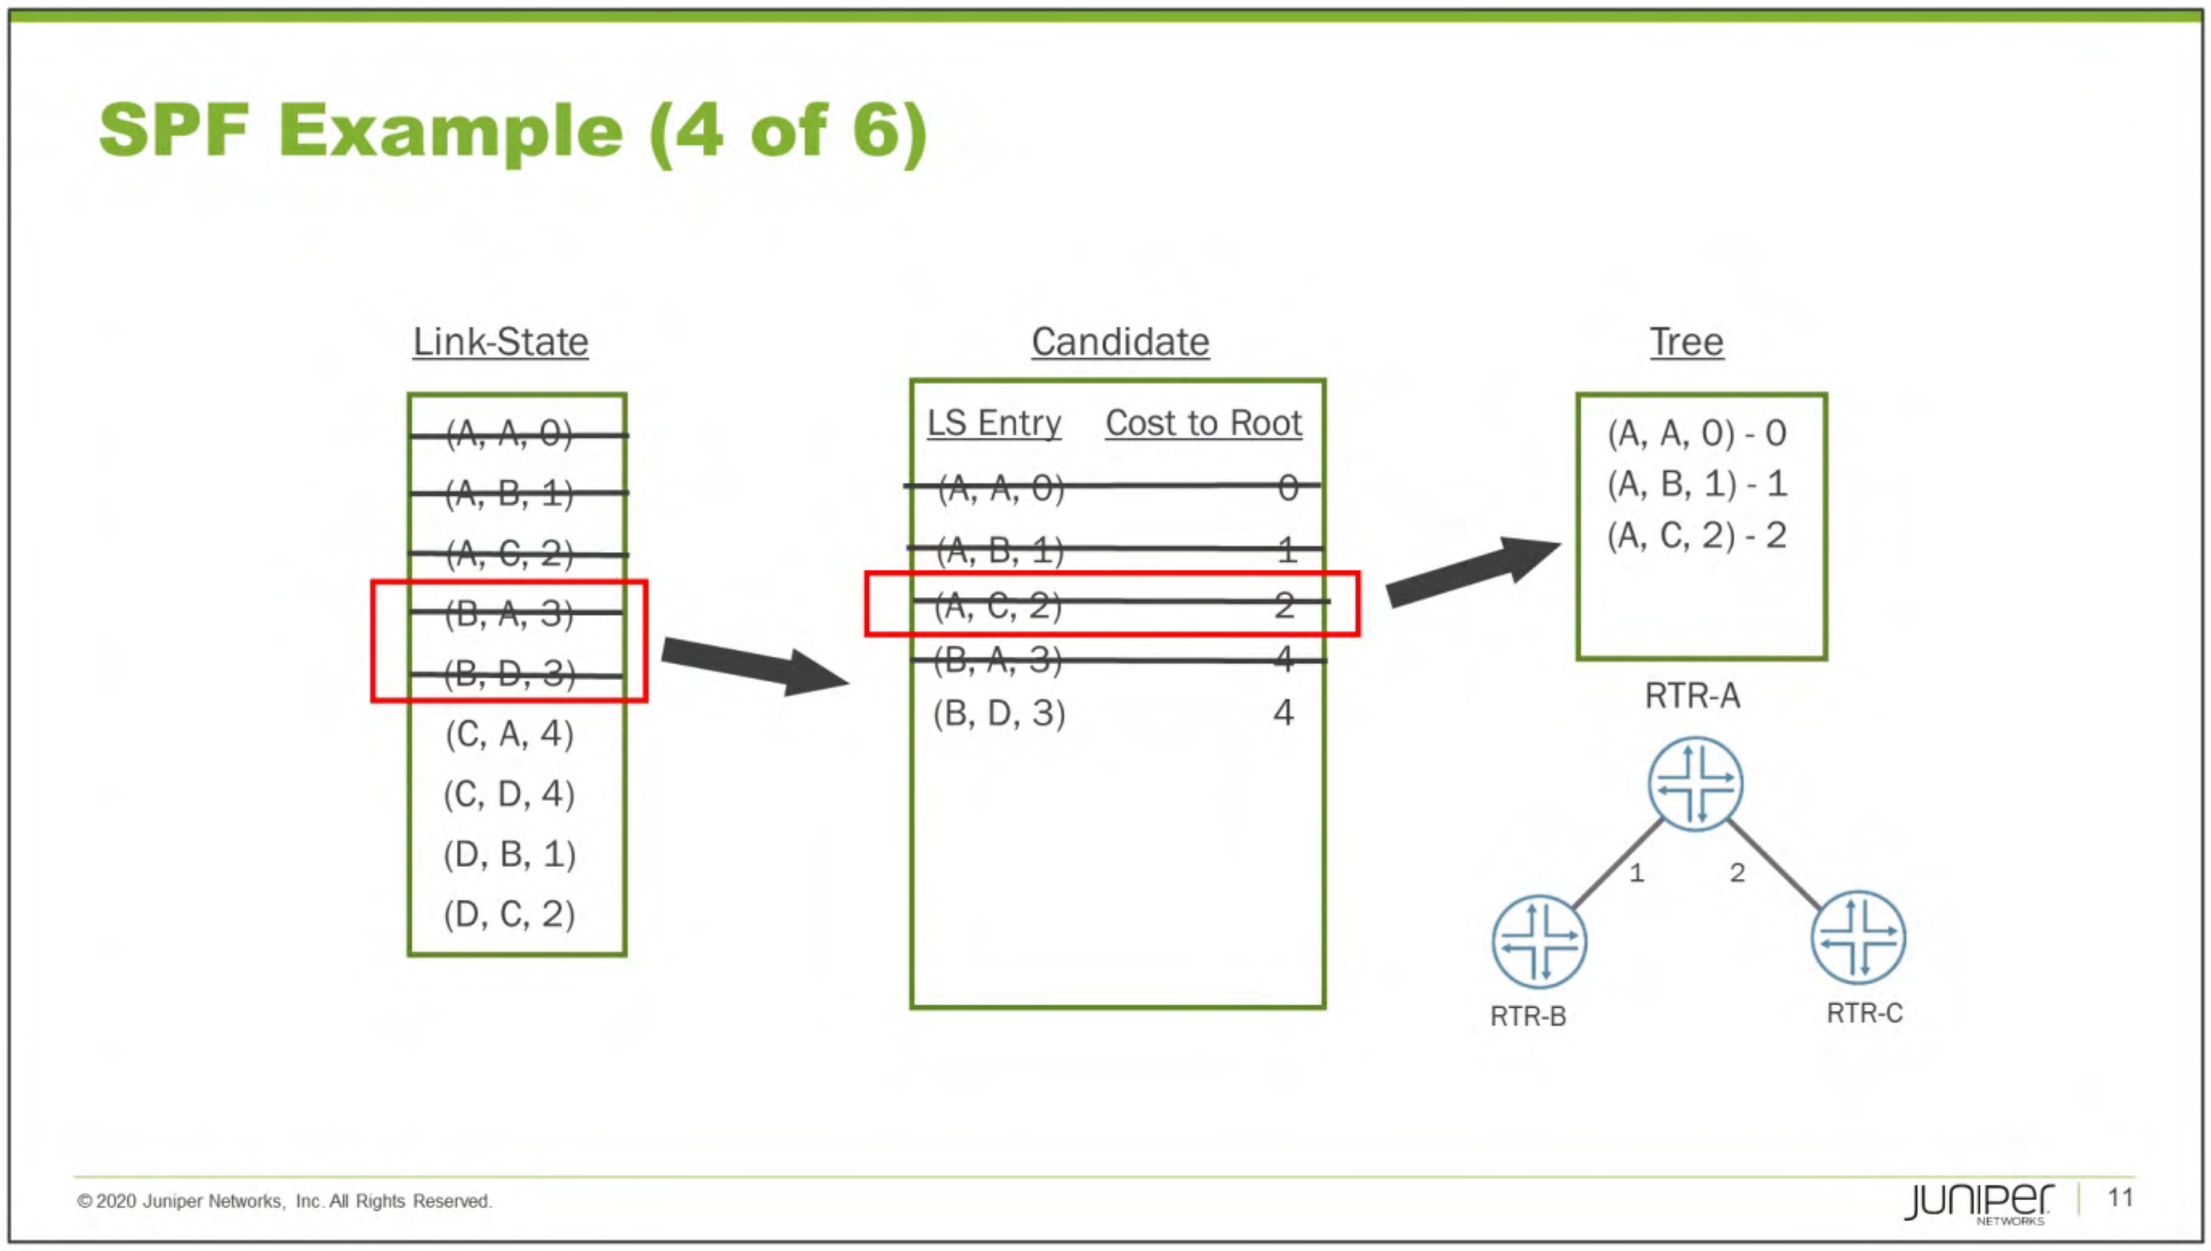

SPF Calculation Example: Part 4

RTR-B is the most recent entry to the tree database, so all RTR-B's tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are then removed from the candidate. Thus, the (B, A, 3) tuple is removed because RTR-A already has the shortest path to RTR-A.

The cost to each neighbor ID from the root is then calculated. It costs RTR-B 3 to reach RTR-D, and it costs 1 to reach RTR-B from the root. So the total cost to reach RTR-D from the root through RTR-B is 4.

The lowest cost tuple in the candidate database, (A, C, 2), is now moved over to the tree database, and RTR-C is placed on the network map.

The candidate database is not empty, so the algorithm continues.

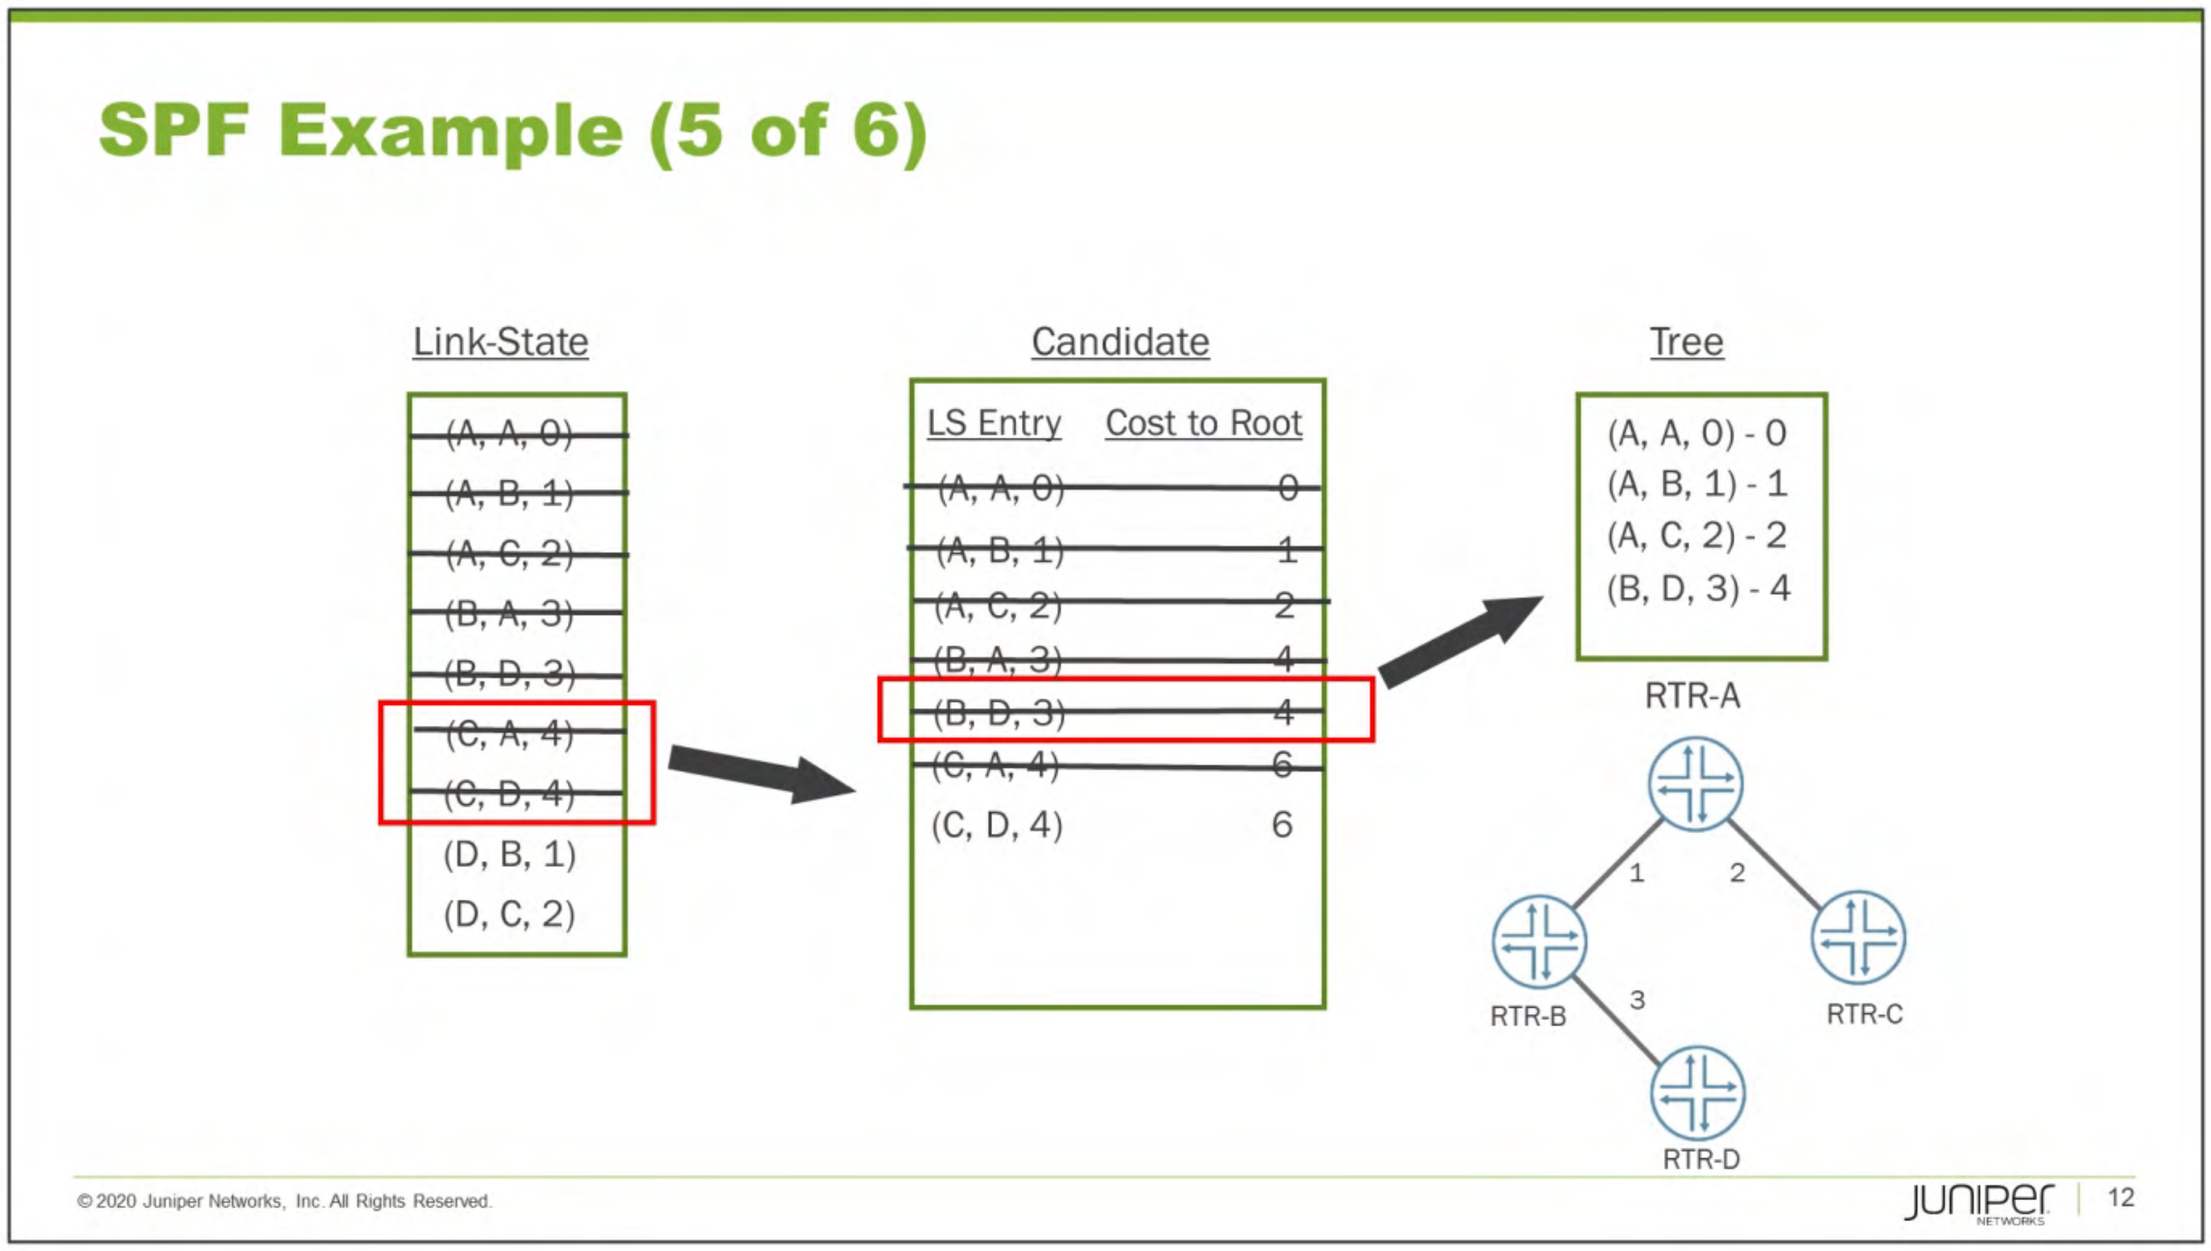

SPF Calculation Example: Part 5

Because RTR-C is the most recent entry to the tree database, its tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are then removed from the candidate. Thus, the (C, A, 4) tuple is removed because RTR-A already has the shortest path to RTR-A.

The cost to each neighbor ID from the root is then calculated. It costs RTR-C 4 to reach RTR-D, and it costs 2 to reach RTR-C from the root. So the total cost to reach RTR-A through RTR-C is 6.

The lowest cost tuple in the candidate database, (B, D, 3), is moved to the tree database, and RTR-D is placed on the network map.

The candidate database is not empty, so the algorithm continues.

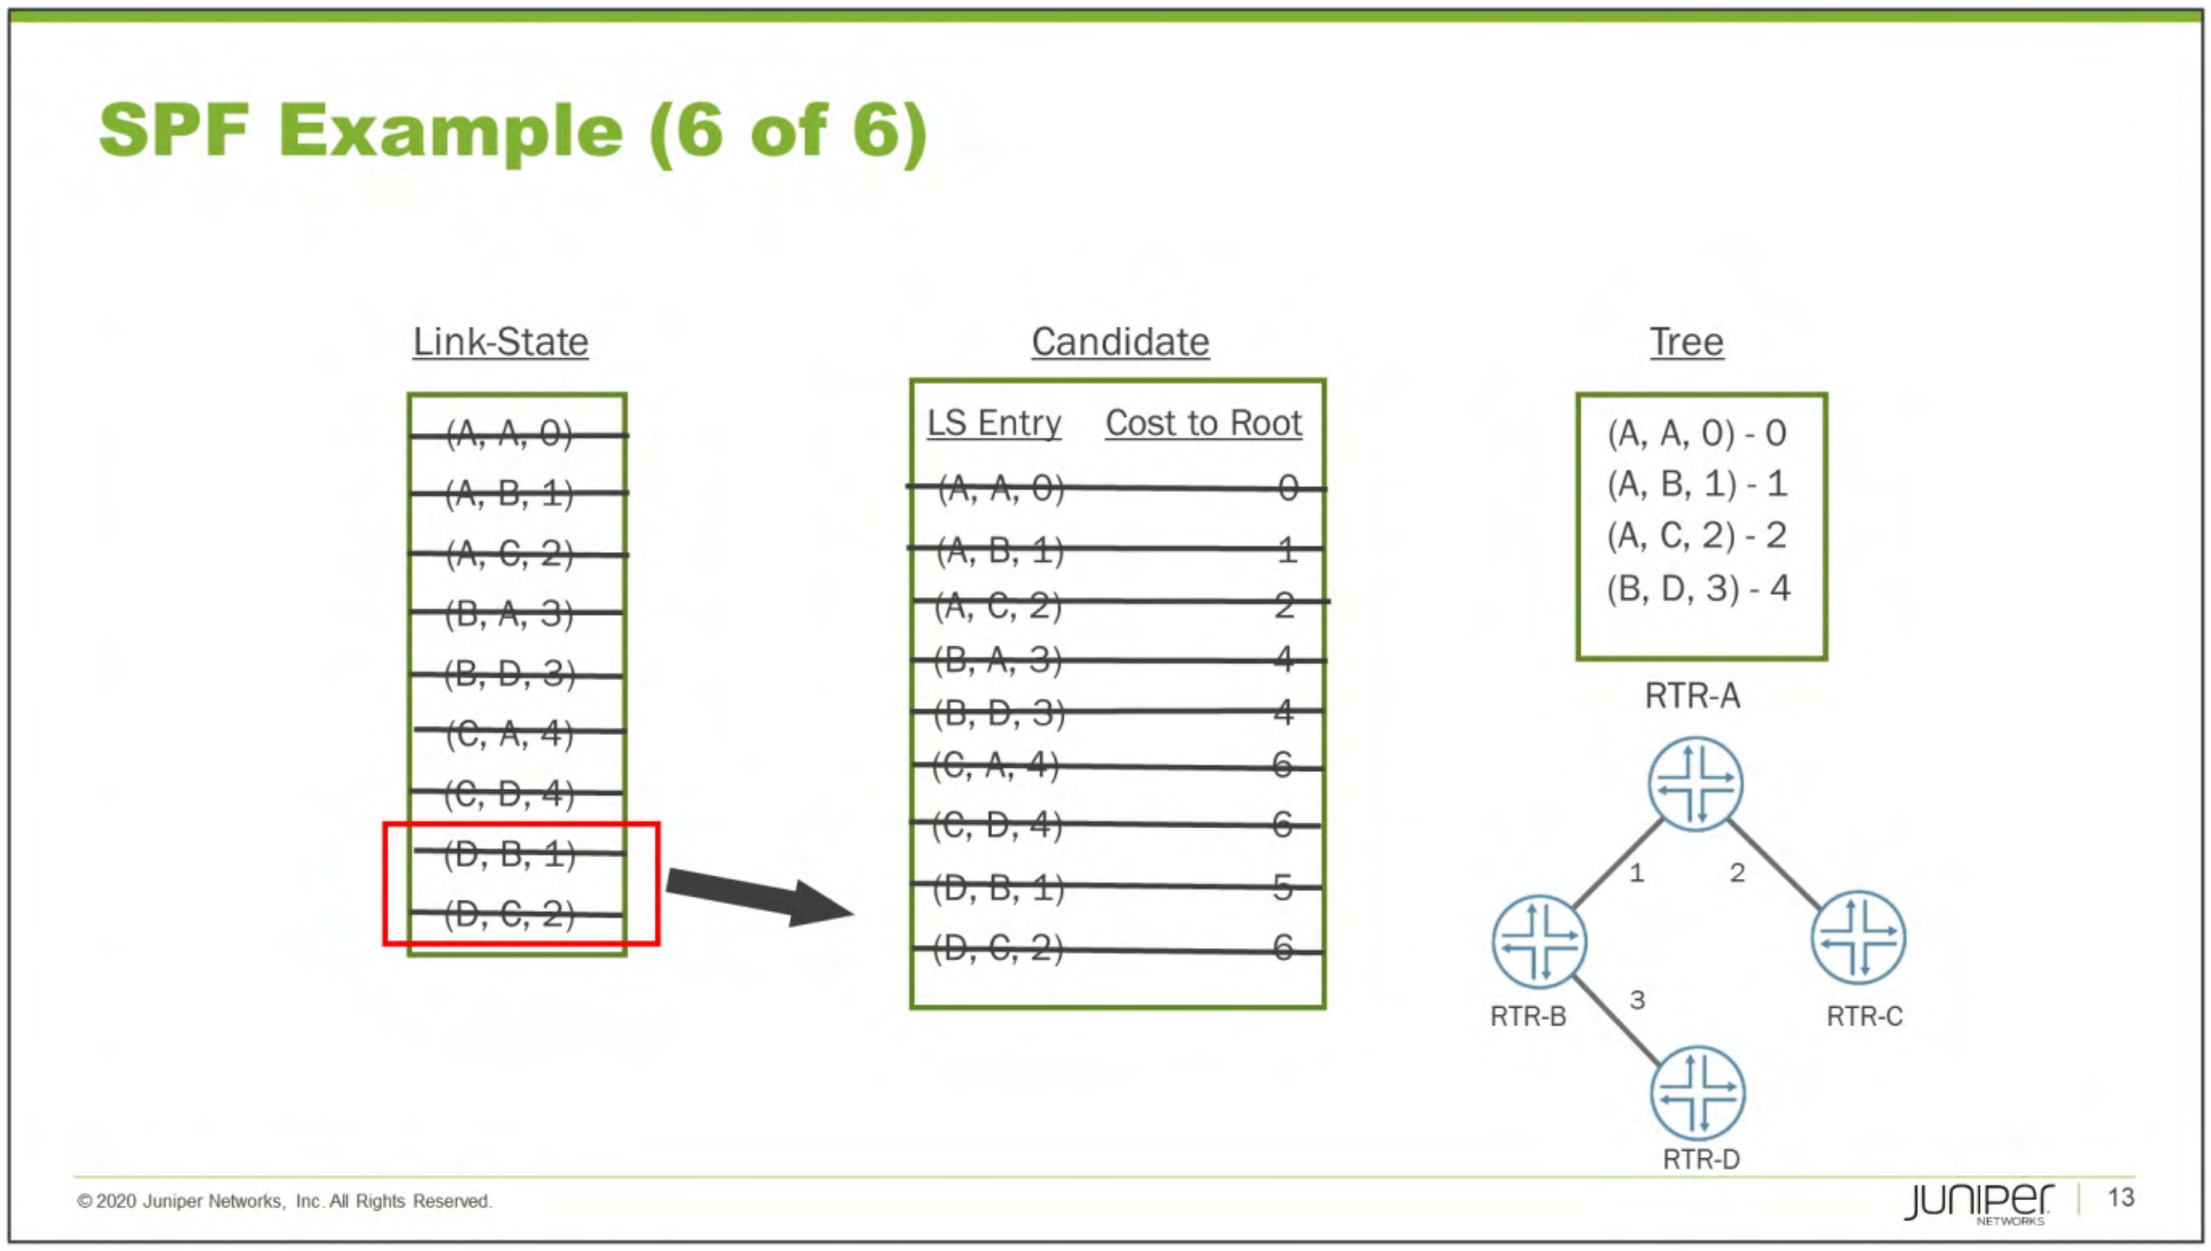

SPF Calculation Example: Part 6

RTR-D, through its link to RTR-B, is the most recent entry to the tree database. Therefore, its tuples are moved from the LSDB into the candidate database.

All known nodes in the tree database are then removed from the candidate. Thus, the (C, D, 4), (D, B, 1), and (D, C, 2) tuples are removed because RTR-A already has paths to RTR-B, RTR-C, and RTR-D. The candidate database is now empty of all tuples, so the algorithm stops.

RTR-A now has a complete network map built with the total cost to each node calculated. This information is then passed to the routing table for its use.

Controlling SPF Calculations

- Three consecutive SPF runs can occur before a mandatory hold-down occurs
  - Keeps the network stable during change
  - Holddown timer is configurable (Default is 5000 milliseconds = 5 seconds)

```
[edit protocols isis]
user@router# show spf-options ?
Possible completions:
  ...
  delay                Time to wait before running an SPF (50..1000 milliseconds)
  holddown             Time to hold down before running an SPF (2000..10000 milliseconds)
  rapid-runs           Number of rapid SPF runs before SPF holddown (1..5)
```

- A 200-millisecond delay is preconfigured between the back-to-back SPFs
  - Altered with the **spf-delay** parameter
  - Supports values that range from 50 to 1000 ms

```
[edit protocols isis]
user@router# set spf-delay 100
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Back-to-Back SPF Calculations

To enhance the convergence time of IS-IS during topology changes, Junos enables the SPF calculation to be run three times in a back-to-back fashion before encountering a mandatory hold-down period. The hold-down timer is configurable and supports values between 2000 and 10000 milliseconds. The timer ensures that the network can continue to forward packets with potentially incorrect routing knowledge during the convergence period. The timer also allows the routing process to not control the CPU at the expense of other routing functions.

Controlling the Delay Between Calculations

A configurable timer slightly delays consecutive SPF calculations. The default timer value is 200 milliseconds (ms), which you can alter with the **spf-delay** statement. The **spf-delay** statement is supported both at the global IS-IS configuration hierarchy and within an IS-IS routing instance. Delay values ranging from 50 ms to 1000 ms (1 second) are supported.

The mandatory 5-second hold-down timer is still brought to bear after the third consecutive rapid SPF calculation, regardless of the `spf-delay` setting.

As a best practice, we recommend you set the `spf-delay` value slightly larger than the worst possible propagation delay found in your network. For example, if the delay across the network is 200 ms, then you might set the `spf-delay` to 250 ms. This setting allows time for all of the duplicate LSPs to arrive at all routers before the SPF calculation begins.

Partial Route Calculation

- Full SPF calculation is run in two stages:
  - Build a shortest-path tree to each IS in the network
  - Calculate the best path to the IP reachability information advertised by each IS
- Only recalculates the IP reachability information
  - Received LSPs are examined for changes
- Automatically enabled and cannot be disabled

© 2020 Juniper Networks, Inc. All Rights Reserved.

Full SPF Is Actually Two Steps

The IS-IS routing protocol actually performs two separate calculations during each SPF computation. The first stage of this full SPF is the generation of a tree representing all IS nodes in the network. The second stage then maps the advertised IP prefixes onto the IS tree for determining the shortest path to each router. These routes are then placed into the routing table on the router.

Only Recalculate IP Reachability

If an IS-IS router begins advertising a prefix (or removes a prefix from the network), but does not alter its IS reachability, no need exists to recalculate the IS tree for the network. This partial route calculation (PRC) reruns only the IP reachability portion of the SPF algorithm. Each router in the network makes the decision to run a full or partial SPF independently based on the contents of the received LSP. Each field is examined to determine the extent of the changes (if any). Should a router find that only IP information is changed, a PRC is scheduled and run, after which the results are passed to the routing table on the router.

Automatically Enabled

The functionality of the PRC is automatically enabled when you configure IS-IS within Junos—you do not have to do anything to benefit from the enhancement. Conversely, no configuration option is available to disable this feature.

Agenda: Advanced IS-IS Operations and Configuration Options

- IS-IS Operations
- →IS-IS Configuration Options
- IS-IS Routing Options

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Configuration Options

The slide highlights the topic we discuss next.

IS-IS Wide Metrics

- TLVs 2, 128, and 130 use 6 bits for metric values
  - Maximum metric for an interface is 63
  - Configured metrics larger than the maximum value are advertised as 63
  - These TLVs are sent by default
- TLVs 22 and 135 use larger metric spaces
  - Maximum value for an interface is now 16,777,215
  - Both TLVs are sent by default
- Router can be configured to only use the wide metrics of TLVs 22 and 135
  - Configured for an entire level
  - Only these TLVs are sent in the LSP
    - Removes external route distinction; results in automatic leaking of redistributed routes from Level 1 to Level 2

```
[edit protocols isis]
user@router# set level 2 wide-metrics-only
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Small Metric TLVs

The IS-IS protocol uses 6 bits in type/length/values (TLVs)—2, 128, and 130—to advertise the metric of an individual interface or external route respectively. Therefore, the largest possible metric that an interface can be assigned is 63. Any larger value configured on the interface, or calculated using the reference-bandwidth, is advertised as a 63 in those TLVs in all LSPs. This value does not affect the total metric to reach a network as determined by the SPF algorithm. The algorithm adds multiple 63 values together to reach a total cost. The largest value possible for this total metric cost is 1023 to any destination in the network.

Wide Metric TLVs

With the creation of the traffic engineering TLVs (22 and 135), the concept of limiting an interface cost to six bits was changed to allow for greater granularity and scalability. The new TLVs use 24 bits to advertise the metric, for a maximum link cost of 16,777,215 and the field for the total cost is expanded to 32 bits.

Only Advertise Wide Metrics

The default operation of IS-IS is to advertise both the small and wide metric TLVs in all LSPs. You can inform the local router to use only the wide metric TLVs with the **`wide-metrics-only`** command. This command is applied on a per-level basis and allows the local router to advertise larger metrics using only the extended TLVs. Note that wide metrics cannot differentiate between internal and external prefixes. As a result, the use of only wide metrics results in the automatic leaking of external prefixes from Level 1 areas into Level 2 areas. You can use routing policy to block the leaking of external prefixes into the backbone area if needed.

Wide Metrics Operation (1 of 3)

```
[edit protocols isis]
user@R1# show
interface ge-1/1/0.0 {
    level 1 disable;
}
interface lo0.0 {
    level 1 disable;
}
```

```
user@R1> show route 192.168.48.1 terse
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination          P Prf   Metric 1   Metric 2  Next hop        AS path
* 192.168.48.1/32      I  18        10             >10.222.60.2

user@R1> show isis database R1 level 2 extensive | find TLVs | match prefix
    IP prefix: 10.222.60.0/24, Internal, Metric: default 10, Up
    IP prefix: 192.168.52.1/32, Internal, Metric: default 0, Up
    IP extended prefix: 10.222.60.0/24 metric 10 up
    IP extended prefix: 192.168.52.1/32 metric 0 up
```

TLV 135

TLV 128

© 2020 Juniper Networks, Inc. All Rights Reserved.

Wide Metrics Operation: Part 1

The example on the slide (as well as the following slides) shows a small IS-IS network consisting of 2 routers. Each of the routers is running only Level 2 on all of its interfaces and is using all of the default settings.

We see that the current cost to the R2 router's loopback address of 192.168.48.1/32 is 10. This cost accounts for the cost to reach R2 over the ge-1/1/0.0 interface (10.222.60.0/24) on R1 in addition to the metric advertised by R2 for its loopback address (0 by default). A look at the LSDB information advertised by R1 shows that both TLV 128 (IP prefix) and TLV 135 (IP extended prefix) are being sent to announce the IP prefixes available on R1.

Wide Metrics Operation (2 of 3)

```
[edit protocols isis]
user@R1# show
interface ge-1/1/0.0 {
    level 1 disable;
    level 2 metric 10000;
}
interface lo0.0 {
    level 1 disable;
}
```

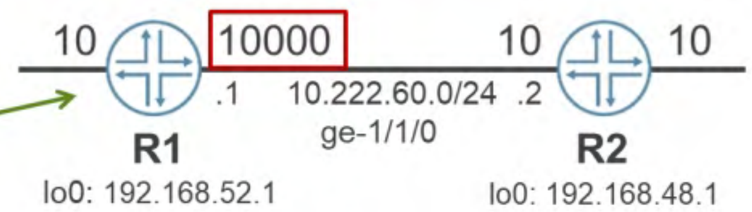


```
user@R1> show route 192.168.48.1 terse
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination          P Prf   Metric 1   Metric 2  Next hop        AS path
* 192.168.48.1/32      I  18        63             >10.222.60.2

user@R1> show isis database R1 level 2 extensive | find TLVs | match prefix
    IP prefix: 10.222.60.0/24, Internal, Metric: default 63, Up
    IP prefix: 192.168.52.1/32, Internal, Metric: default 0, Up
    IP extended prefix: 10.222.60.0/24 metric 63 up
    IP extended prefix: 192.168.52.1/32 metric 0 up
```

TLV 135

TLV 128

© 2020 Juniper Networks, Inc. All Rights Reserved.

Wide Metrics Operation: Part 2

The R1 router has configured a metric value of 10,000 on its interface toward R2. After committing the configuration, we see that the metric cost to reach R2's loopback is now 63. The LSDB shows both TLV 128 and TLV 135 being advertised with a metric of 63, the maximum metric available to the IP internal reachability TLV.

Wide Metrics Operation (3 of 3)

```
[edit protocols isis]
user@R1# show
level 2 wide-metrics-only;
interface ge-1/1/0.0 {
    level 1 disable;
    level 2 metric 10000;
}
interface lo0.0 {
    level 1 disable;
}
```

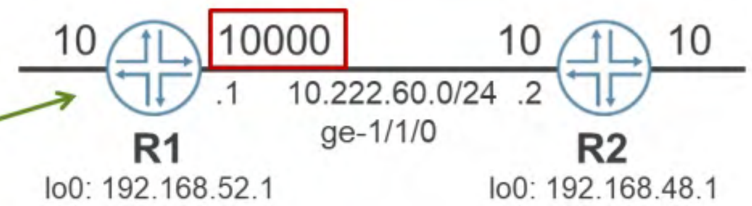


```
user@R1> show route 192.168.48.1 terse
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination          P Prf   Metric 1   Metric 2  Next hop        AS path
* 192.168.48.1/32      I  18     10000             >10.222.60.2

user@R1> show isis database R1 level 2 extensive | find TLVs | match prefix
  IP extended prefix: 10.222.60.0/24 metric 10000 up
  IP extended prefix: 192.168.52.1/32 metric 0 up
```

TLV 135

© 2020 Juniper Networks, Inc. All Rights Reserved.

Wide Metrics Operation: Part 3

Sydney has now configured its Level 2 operation to use only wide metrics through the **`wide-metrics-only`** command. The metric cost to R2's loopback address is now listed as 10,000 in the routing table. The LSDB now shows that only TLV 135 is being advertised with a metric value of 10,000 assigned to it.

Effects of Altering Metrics

- The TLV values within an LSP advertise the metric values, which populate the LSDB
- As each router runs the SPF algorithm, each LSA is examined individually for the cost of the outgoing interface
  - The final metric calculation uses that cost
- Routers can disagree about the cost on a network link
  - Can result in asymmetric routing in the network
  - R1 sees a cost of 20 to reach the R4 router through R2
  - R4 sees a cost of 20 to reach the R1 router through R3

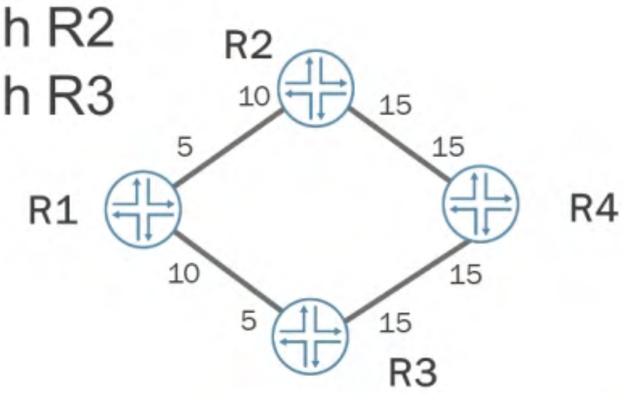

© 2020 Juniper Networks, Inc. All Rights Reserved.

Advertisement of Metric Values

After the IS-IS process on a router decides which metric to assign to each interface, that information is flooded into the domain within either a Level 1 LSP or a Level 2 LSP.

SPF Calculations

After receiving a new LSP from another router, the local router performs an SPF calculation using all the values contained in the LSPs in the database. As mentioned on a previous slide, the cost is calculated from the root node to every other node in the network using the metric cost of the outgoing interfaces.

Routers Can Disagree

Because each router performs a separate SPF calculation, two routers on either side of a point-to-point link can disagree on the metric of that link. The example on the slide shows that each router has determined a different metric value for its attached links. This difference results in the R1 router calculating a total cost of 45 (5+15+25=45) to reach the R4 router, while the R4 router calculates a total cost of 60 (30+20+10=60) to reach the R1 router.

Although the configuration of different metrics for a single link does not affect the operation of IS-IS, asynchronous routing can occur within the network, which might cause response delays for some end users and make troubleshooting the network challenging for network administrators.

IS-IS Authentication

- Authentication can occur within multiple places
  - Level 1
  - Level 2
  - Interface
- Three authentication types are supported
  - None (default)
  - Simple
  - MD5
- MD5 includes an encrypted checksum with all packets
  - Provides better security than `simple` type

© 2020 Juniper Networks, Inc. All Rights Reserved.

Authentication Configured in Various Places

Within the IS-IS protocol, you can enable authentication in multiple places within the configuration hierarchy. Each specific location encrypts certain IS-IS PDU packets. Any authentication configured at either Level 1 or Level 2 secures all hello PDUs, LSPs, and sequence number PDUs sent by the local router for that specific level.

Authentication at the interface level secures only hello PDUs. Again, this configuration occurs for either Level 1 or Level 2.

Authentication Types

Junos supports three different forms of authentication for the IS-IS protocol. These types are none, simple, and Message Digest 5 (MD5). The type of authentication used must match on all routers within a level.

MD5 Authentication Checksum

For better security in an IS-IS network, we recommend you use authentication type MD5. MD5 includes an encrypted checksum in all IS-IS packets—not a simple text password. Each IS-IS router uses the same MD5 algorithm to calculate the checksum value, so it virtually guarantees interoperability and a correct result.

Authentication Configuration

- Level authentication affects all IS-IS PDUs
  - Link-state, sequence number, and hello
- Per-interface authentication affects hello PDUs only and takes precedence over per-level settings

```
[edit protocols isis]
user@router# show
level 1 {
    authentication-key "$9$bssYomPQ69pkq39puhc8X7V2a"; # SECRET-DATA
    authentication-type md5;
}
level 2 {
    authentication-key "$9$dXVYoDjqQ39gomTz6CAvW8X-ViHmFnCDi1h"; # SECRET-DATA
    authentication-type simple;
}
interface fe-0/0/0.0 {
    level 2 {
        hello-authentication-key "$9$1sEEclw4JH-d2oGq.Ctp01h7NbgaU"; # SECRET-DATA
        hello-authentication-type md5;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Level Authentication

Configuration of any authentication at either Level 1 or Level 2 secures hello, link-state, and sequence number PDUs.

In the example on the slide, MD5 authentication is enabled for Level 1, and simple authentication is configured for Level 2.

Note

Take care when configuring authentication on a point-to-point interface. IS-IS uses a single hello PDU on these interfaces (as opposed to a broadcast interface having a Level 1 and a Level 2 hello). Thus, the local router uses the authentication configured at the lowest level for the hello PDUs on this interface. Potential problems might arise if one side of the interface is operating both Level 1 and Level 2 while the other side is operating only Level 2.

Per-Interface Authentication

As is a common practice within the Junos configuration hierarchy, IS-IS authentication options configured at an interface level supercede any options configured at a higher level.

Interface fe-0/0/0.0 on the slide has MD5 authentication configured for hello PDUs sent at Level 2 using the **`hello-authentication`** commands. This hello authentication will be used only on that specific interface.

Authentication Control

- Level 1 or Level 2 authentication can be disabled for specific PDUs
  - Hello PDUs
    - `no-hello-authentication`
  - Complete sequence number PDUs
    - `no-csnp-authentication`
  - Partial sequence number PDUs
    - `no-psnp-authentication`
- Stop verifying authentication on all PDUs with the **`no-authentication-check`** command
  - Useful for migration purposes

© 2020 Juniper Networks, Inc. All Rights Reserved.

Selectively Disable Authentication

For easier interoperability with implementations from other vendors as well as for more control by network operators, Junos can selectively disable authentication for certain PDUs. This ability stops both the securing of transmitted PDUs as well as the verification of received PDUs. In essence, the router operates as if no authentication was ever configured within IS-IS for the specific PDUs configured.

No Authentication Verification on Received PDUs

To aid you during a migration period or for troubleshooting an authentication problem, you can ignore the verification on received packets with the **`no-authentication-check`** command. Transmitted packets are still secured, but all received packets are used, regardless of any authentication information contained in them.

Mesh Groups

- IS-IS floods LSPs to all neighbors by default
- Certain physical topologies make this flooding unnecessary
  - R4 router will receive three copies of the same LSP
- Once configured, the group members do not reflood LSPs within the group

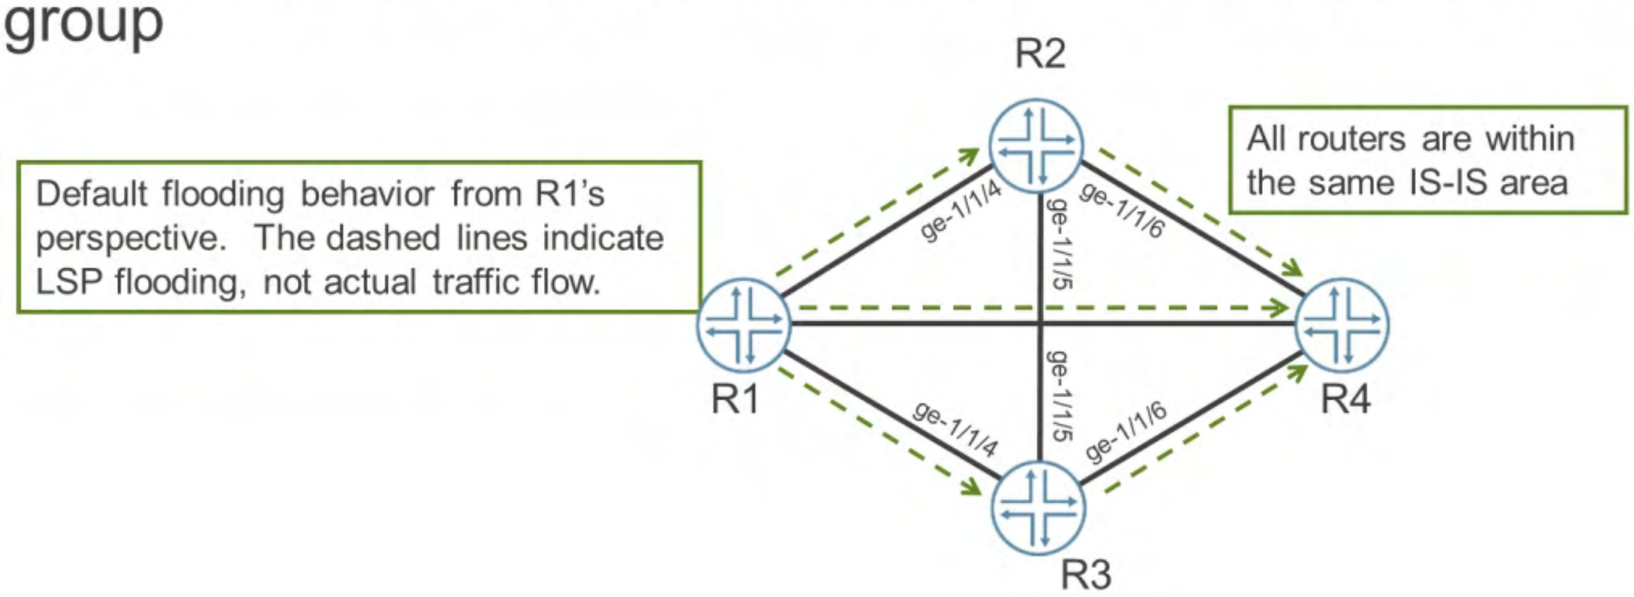

© 2020 Juniper Networks, Inc. All Rights Reserved.

Default IS-IS Flooding

As mentioned on previous slides, the default IS-IS operation is to flood all LSPs to any adjacent router for a configured level.

Full-Mesh Topologies

At times, the default flooding mechanism might be a disadvantage. Such is the case with a full-mesh physical or logical topology. In this type of an environment, each IS-IS router receives extra copies of the same LSP. These extra copies are not needed and can be discarded safely.

In the example on the slide, the R4 router receives an LSP generated from the R1 router three times.

Mesh Group Members

The IS-IS protocol has the option of configuring a mesh group for certain peers. Once configured, the mesh group members do not re-flood LSPs within their group. Only LSPs received from outside the group membership are flooded within the group.

If a mesh group were configured in the example on the slide, the R4 router would receive only a single copy of the LSP from the R1 router.

Mesh Group Configuration

- Each interface is configured with a group number
  - 32-bit numbers can be different on separate interfaces
- To prevent an interface from flooding any LSPs, you can use the **blocked** keyword

```
[edit protocols]
user@R2# show
isis {
    interface ge-1/1/4.0 {
        mesh-group 1;
    }
    interface ge-1/1/5.0 {
        mesh-group 1;
    }
    interface ge-1/1/6.0 {
        mesh-group 1;
    }
}
```

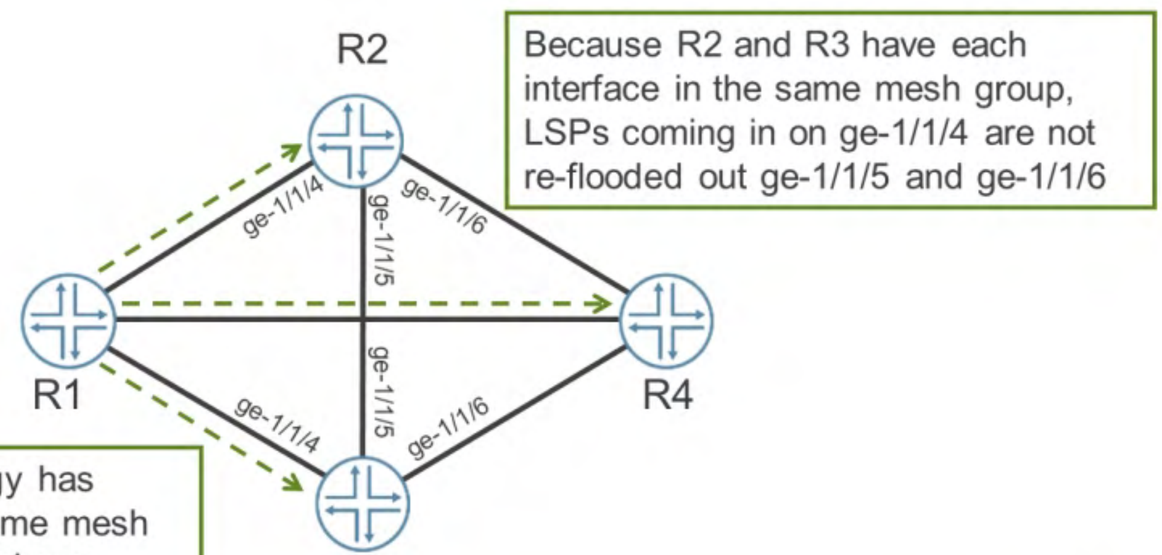

Because R2 and R3 have each interface in the same mesh group, LSPs coming in on ge-1/1/4 are not re-flooded out ge-1/1/5 and ge-1/1/6

Each router in the example topology has each interface configured in the same mesh group. The configuration for R2 is shown.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Group Numbers

To configure a mesh group within Junos, assign each interface within IS-IS a group number by using the **mesh-group** command. This number can be any 32-bit value. Within each local router, any LSPs received on an interface assigned to a mesh group are not flooded out any interfaces within that same mesh group.

Prevent All Flooding

Because the **mesh-group** command alters the default IS-IS flooding, you can disable all flooding out a particular interface by using the **mesh-group blocked** command. This command might be useful in an environment where a local IS-IS router might want to receive LSPs from an adjacent neighbor but not send any LSPs to that neighbor.

```
[edit protocols isis]
user@router# show
interface ge-0/0/1.0 {
    mesh-group blocked;
}
```

Overload Bit

- Allows advertisement of routing information to neighbors while indicating the node should not be used for transit traffic
  - Other routers ignore the LSP during SPF calculation
  - Turns off the attached bit
- Can be set permanently or with a timeout value
  - Timer is between 60 and 1800 seconds
  - Timer only runs after rpd starts

```
[edit protocols]
user@router# show
isis {
    overload;
    interface so-0/0/0.0;
    interface ge-0/1/0.0;
}

user@router> show isis database
IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
Router-1.00-00             0x36      0x8cf7     1007 L1 L2
Router-2.00-00             0x37      0x4c3a     1067 L1 L2 Overload
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Avoid Transit Traffic

The concept of the overload bit was first available within the IS-IS protocol. Put simply, its function is to advertise information into the IS-IS network, but to prevent transit traffic on the router. This functionality can be useful in two situations—first, when a router must be taken out of the network for maintenance, and second, when the router has many BGP peers. In the first case, traffic should avoid the router because its ability to forward traffic can be compromised. In the second case, a network administrator might want the router to fully establish its BGP peering sessions before handling transit traffic.

When an IS-IS router is configured for overload, a bit is set to 1 within the attributes field in the LSP header. This configuration is then visible to all other routers in the level within the LSDB.

In the example on the slide, the Router-2.00-00 LSP is advertised with the overload bit set.

Overload Settings

You can turn the `overload` setting on or off as a permanent value, or you can associate a timer with it. If the timer is omitted, then the overload bit is set once you commit the configuration. The bit remains until you remove the overload setting from the configuration, and the configuration is committed once again. However, if you assign a timer value, the bit is not set automatically. The timer associated with the overload bit initializes only when the routing protocol daemon (rpd) also initializes. This timer can run from between 60 and 1800 seconds. Once the timer expires, the bit is removed from the LSPs, but the configuration still contains the **`overload`** command.

CSNP Interval

- DIS router sends CSNP packets on a LAN interface every 10 seconds
- Can be altered on a per-interface basis
  - Value can be between 1 and 65,535 seconds

```
[edit]
user@router# run show isis interface detail
IS-IS interface database:
ge-0/2/0.0
  Index: 3, State: 0x6, Circuit id: 0x2, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    2             1       64     10         3        9 R1.02 (us)
[edit]
user@router# set protocols isis interface ge-0/2/0 csnp-interval 40
[edit]
user@router# commit
[edit]
user@router# run show isis interface detail
IS-IS interface database:
ge-0/2/0.0
  Index: 3, State: 0x6, Circuit id: 0x2, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 40 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    2             1       64     10         3        9 R1.02 (us)
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

CSNP Packets on a LAN Interface

By default, the DIS router on a LAN interface advertises a complete sequence number PDU (CSNP) every 10 seconds. This advertisement allows other routers on the link to ensure that their databases are complete. Perhaps more importantly, it allows the other IS-IS routers to know when the DIS is no longer available. This fast CSNP timer allows IS-IS to not require the election of a backup DIS.

Alter the Timers

If you do not need the quick timer of the CSNP, you can change it. One possible situation where this change might be useful is on a broadcast link with only two routers. If the DIS is not available any longer, the other router becomes aware of this unavailability from a lost adjacency or downed network link.

To lengthen or shorten the timer value, you can use the **`csnp-interval`** command.

On a per-interface basis, you can set the timer value as short as 1 second or as long as 65,535 seconds.

Agenda: Advanced IS-IS Operations and Configurations Options

- IS-IS Operations
- IS-IS Configuration Options
- →IS-IS Routing Policy

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Routing Policy

The slide highlights the topic we discuss next.

IS-IS Default Policies

- User-defined import policies are not allowed
- IS-IS route information in an LSP is populated from the configuration within the `[edit protocols isis]` hierarchy
  - Subnets are placed into an LSP for advertisement into the network
  - Can block internal and external subnets in export policy
    - Different behavior than OSPF

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Import Policies

Like OSPF, IS-IS is a link-state protocol. The routing table is populated by the results of the SPF algorithm. Unlike OSPF, however, Junos does not allow user-defined import policies in IS-IS.

Default Export Policy for IS-IS

Remember that policies affect routes into and out of the routing table. The information in an IS-IS LSP is not gathered from inet.0. Instead, the LSP is populated through the configuration of IS-IS within the `[edit protocols isis]` configuration hierarchy. However, for IS-IS you can block internal subnets as well as external subnets from entering the LSP in an export policy, essentially overriding the inherit first term. This behavior is different from OSPF, where export policy affects only external subnets.

Redistributing Routes in IS-IS

- Route redistribution requires an export policy at the global IS-IS level
  - Routes from other IGPs (RIP or OSPF)
  - Routes from other protocols (static or aggregate)
- Beware of routing loops
  - When multiple redistribution points exist
    - Route table preference values can cause redistributed routes to be preferred

```
[edit policy-options]
user@router# show
policy-statement redistribute-routes-into-isis {
   term find-other-igp-routes {
      from protocol [rip ospf];
      then accept;
   }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Route Redistribution in IS-IS

You can add routes to an advertised LSP through an export policy. A new TLV is added to the update for each route accepted by the policy. These types of policies often match and accept all routes from a particular protocol. For example, the `redistribute-routes-into-isis` policy on the slide advertises all RIP and OSPF routes into the network.

Routing Loop Concerns

Again, you must take care when redistributing routes between protocols to avoid a routing loop.

IS-IS Route Attributes on Export

- Change metric values for IS-IS routes
  - Use the **metric** keyword in a policy
  - Configured value placed in appropriate LSP for that level
- Administrative marking for policy matching
  - Use the **tag** keyword in a policy
    - Allows other routers to match on the defined value
    - Useful for selective redistribution

```
[edit policy-options]
policy-statement example-isis-policy-1 {
   term set-static-metric {
      from protocol static;
      then {
         metric 40;
         tag 68;
         accept;
      }
   }
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

IS-IS Export Attributes

IS-IS can change the advertised metric value for routes redistributed into the protocol. A policy action of **then metric *value*** changes the value in the advertised LSP.

IS-IS Route Tagging

IS-IS route tagging offers the same administrative filtering capabilities as the OSPF protocol. Route tagging for IS-IS requires traffic engineering extensions, which are enabled by default.

The example on the slide shows an IS-IS export policy that matches all active static routes in inet.0 and advertises them into IS-IS with a metric of 40 and a route tag of 68.

Case Study: External Routes (1 of 5)

- Case study topology

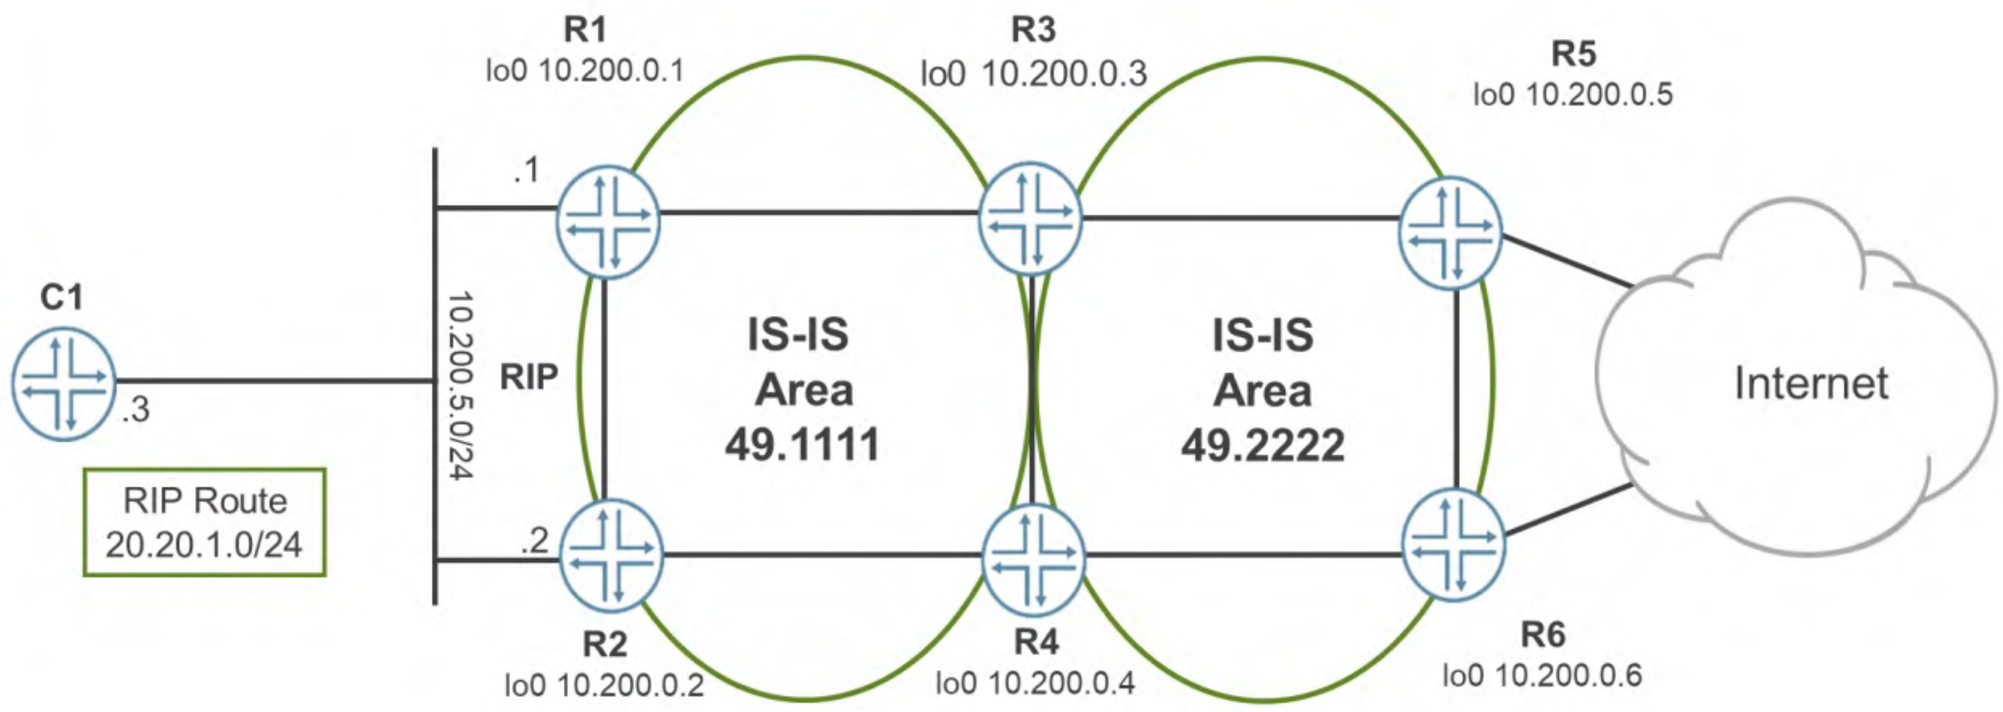

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study Topology

The slide displays the case study topology that will be used in subsequent slides. R1, R2, and C1 all run RIP.

Case Study: External Routes (2 of 5)

- Case study background:
  - The RIP router is advertising routes
    - Redistribute a single RIP route into IS-IS—Advertise the 20.20.1.0/24 RIP route
  - The R1 and R2 routers have an IS-IS default route
    - Default route is generated because of the attached bit
    - Redistribute the default route into RIP

© 2020 Juniper Networks, Inc. All Rights Reserved.

Case Study Background

The slide describes the goal of the case study. The goal is to advertise a single RIP route into IS-IS as well as send a default route to RIP.

Case Study: External Routes (3 of 5)

- Redistribute the RIP route into IS-IS

Both polices are defined and applied on the R1 and R2 routers.

- Redistribute the IS-IS route into RIP

```
[edit]
user@R1# show policy-statement redistribute-rip
term 1 {
     from {
         protocol rip;
         route-filter 20.20.1.0/24 exact;
     }
     then accept;
}
[edit]
user@R1# show protocols
isis {
    export redistribute-rip;
}
```

```
[edit]
user@R2# show policy-statement redistribute-default
term 1 {
    from {
        protocol isis;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}
[edit]
user@R2# show protocols
rip {
    export redistribute-default;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Create a Policy for the RIP Route

The first step is to create a policy and apply the policy to IS-IS. In this case, two match conditions were used, creating a logical AND. This policy will then be applied under the `[edit protocols isis]` hierarchy by performing a **`set protocols isis export redistribute-rip`** command.

Redistribution of Default Route

The next step is to redistribute the default IS-IS route into RIP using an export policy under RIP. We will use two match conditions again. This policy will be applied under `[edit protocols rip]`.

Case Study: External Routes (4 of 5)

- Redistribute the default route from IS-IS into RIP
  - High preference 160 to low preference 100

```
user@R2> show route 0/0

r2.inet.0: 20 destinations, 22 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[RIP/100] 00:04:33, metric 2
                    > to 10.200.5.1 via ge-1/0/6.110
                    [ISIS/160] 00:00:06, metric 0
                    > to 10.203.5.2 via ge-1/0/9.2
```

Toward the C1 RIP router

Toward the R4 IS-IS router

  - The RIP route is the active route
  - Modify the IS-IS external preference to 90

```
[edit policy-options]
user@R2# show policy-statement isis-preference
    term 1 {
        from {
            protocol rip;
            route-filter 0.0.0.0/0 exact;
        }
        then {
            preference 90;
            accept;
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Redistribution of Default Route

By default, RIP has a lower preference than IS-IS external routes. Because RIP has a better preference, the default route for RIP is preferred. In this sample network, this preference creates a loop, because the IS-IS routers point to the RIP router as its gateway while the RIP router points to the IS-IS routers.

To fix the loop, modify the IS-IS route preference to a lower value than the RIP route.

Case Study: External Routes (5 of 5)

- Mutual redistribution
  - The default route has been fixed
  - IS-IS now has a lower preference than RIP and is the preferred route.

```
user@R2> show route 0/0

r2.inet.0: 20 destinations, 22 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[ISIS/90] 00:00:06, metric 0
                    > to 10.203.5.2 via ge-1/0/9.2
                    [RIP/100] 00:04:33, metric 2
                    > to 10.200.5.1 via ge-1/0/6.110
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

The Result

The result of the preference change is now a default that points properly to the IS-IS router.

Prefix Limits for External Routes

- Junos is built to handle a large number of external routes
  - Internet routes are **not** normally imported into IS-IS
  - Usually occurs because of a configuration mistake
  - Can leave a portion of your network unusable
- Limit can be placed on the number of routes allowed using a routing policy
  - When the limit is reached:
    - External routing information no longer transmitted in LSPs
    - Overload state initiated
  - Requires a manual step to *fix* the problem

```
[edit protocols isis]
user@router# show
level 1{
    prefix-export-limit 400;
}
level 2{
    prefix-export-limit 600;
}
```

© 2020 Juniper Networks, Inc. All Rights Reserved.

Junos Supports Large Numbers of Routes

IS-IS implementations can encounter problems when large numbers of external routes are injected into the LSDB. Junos has great protocol stability and handles these external routes without failing. However, a configuration mistake could make a portion of your network unusable, because only the Juniper Networks routers would be left operating.

Limit External Routes

To help network administrators when a configuration mistake occurs, Junos allows a limit to be placed on the number of external routes exported into IS-IS. The **`prefix-export-limit`** command informs the router how many routes to accept using a routing policy configuration. You configure this option at a specific level using a 32-bit value, which provides a range of routes from 1 to 4,294,967,295. Once the route limit is reached, the router transitions into an overload state where the bit is set in all LSPs. The external routing information is no longer transmitted in the LSPs as well. The local router remains in this state until the number of external routes returns to a level below the configured limit. This level reduction requires the administrator to manually change the existing configuration; either the number of advertised routes must be reduced or the routing policy must be changed.

Summary

- In this content, we:
  - Displayed and interpreted the IS-IS LSDB
  - Performed advanced IS-IS configuration options
  - Implemented IS-IS routing policy

© 2020 Juniper Networks, Inc. All Rights Reserved.

We Discussed:

- The display and interpretation of the IS-IS LSDB;
- Advanced IS-IS configuration options; and
- The implementation of IS-IS routing policy.

Review Questions

1. What is the maximum default IS-IS link metric? Can this be changed? How?
2. What are the different forms of IS-IS authentication?
3. What are the default policies for IS-IS?

© 2020 Juniper Networks, Inc. All Rights Reserved.

Review Questions

1.

2.

3.

Lab: Advanced IS-IS Configuration Options and Routing Policy

- Configure, monitor, and troubleshoot the operation of an IS-IS network. Additionally, configure interface metrics, wide metrics, and authentication.
- Configure routing policies to redistribute routes between interior gateway protocols.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab: Advanced IS-IS Configuration Options and Routing Policy

The slide provides the objectives for this lab.

Answers to Review Questions

1.

The maximum default link metric is 63. By default, Junos supports the sending and receiving of wide metrics. To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the **`wide-metrics-only`** statement.

2.

None, simple, and MD5 are the three forms of IS-IS authentication.

3.

Import polices are not allowed for IS-IS, because they could lead to inconsistencies in the LSDB. The default export policy for IS-IS is to reject everything. IS-IS floods LSPs to announce local routes and any routes learned by the protocol. Exporting can be used only to announce information from other protocols.

Acronym List

ABR . . . . . area border router
ARP . . . . . Address Resolution Protocol
AS . . . . . autonomous system
ASBR . . . . . autonomous system boundary router
BFD . . . . . Bidirectional Forwarding Detection
BGP . . . . . Border Gateway Protocol
BGP4 . . . . . Border Gateway Protocol version 4
CLI . . . . . command-line interface
CLNP . . . . . Connectionless Network Protocol
CLNS . . . . . Connectionless Network Service
CSNP . . . . . complete sequence number PDU
DIS . . . . . designated intermediate system
DoS . . . . . denial-of-service
DR . . . . . designated router
EGP . . . . . exterior gateway protocol
ES . . . . . end system
ES-IS . . . . . End System–to–Intermediate System
GUI . . . . . graphical user interface
IAB . . . . . Internet Advisory Board
IANA . . . . . Internet Assigned Numbers Authority
IBGP . . . . . internal BGP
ICMP . . . . . Internet Control Message Protocol
IGP . . . . . interior gateway protocol
IPv4 . . . . . IP version 4
IPv6 . . . . . IP version 6
IS . . . . . Intermediate System
ISO . . . . . International Organization for Standardization
ISP . . . . . Internet service provider
JNCP . . . . . Juniper Networks Certification Program
LSA . . . . . link-state advertisement
LSDB . . . . . link-state database
LSP . . . . . link-state PDU
MD5 . . . . . Message Digest 5
MED . . . . . multiple exit discriminator
MOSPF . . . . . Multicast Open Shortest Path First
NET . . . . . network entity title
NLPID . . . . . network layer protocol identifier
NLRI . . . . . network layer reachability information
NSSA . . . . . not-so-stubby area
OSI . . . . . Open Systems Interconnection
PDU . . . . . protocol data unit
PFE . . . . . Packet Forwarding Engine
PSNP . . . . . partial sequence number PDU
RIB . . . . . Routing Information Base
RID . . . . . router ID
rpd . . . . . routing protocol daemon
SNPA . . . . . subnetwork point of attachment
SPF . . . . . shortest path first
TED . . . . . traffic engineering database
TLV . . . . . type/length/value
ToS . . . . . type of service
TTL . . . . . time-to-live

Education Services

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1110 PZ SCHIPHOL-RIJK
Amsterdam, Netherlands
Phone: 31.0.207.125.700
Fax: 31.0.207.125.701

Copyright 2020
Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, services marks, registered marks, or registered services marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Engineering
Simplicity

EDU-JUN-AJSPR, Revision V19A